



it's about time

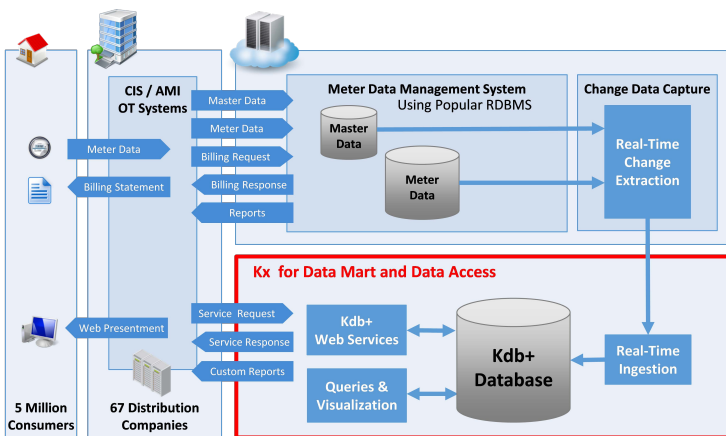
Kx powers Data Mart for Ontario's Smart Meter Data

Case Study: Independent Electricity System Operator Smart Meter Data Mart

The Independent Electricity System Operator (IESO) works at the heart of Ontario, Canada's power system – ensuring there is enough power to meet the province's energy needs in real time while also planning and securing energy for the future. The IESO, as the Smart Metering Entity, operates the provincial smart meter data processing service (the MDM/R) for 67 distribution system operators and over 4 million customers. This service processes over 100 million meter reads per day for billing customers on time-of-use rates.

Implementation

The IESO implemented a Data Mart using Kx technology to support the current volumes of data access requests and anticipated growth, 24x7x365 availability, and current and future demands for advanced analytics in Ontario's energy sector.



Quick Facts
5 million smart meters
9 years of history
300 billion sensor reads
26.3 TB storage (tiered)
130M sensor reads / day
Delivers query results over 100 times faster than before
One 16-core server with a hot standby

In order to get data out from its existing source database system into the Data Mart, the IESO implemented a change data capture tool to extract data from the source database system and deliver this data as flat files.

Challenge

The IESO was faced with a large database of smart meter data (over 250 billion meter reads and growing at over 100 million per day), growing volume and variety of data retrieval and analytics requests, and the need for providing processed data to utilities in a timely manner. The existing transaction system and its relational database system were not originally designed to handle these increased requirements, and needed to be offloaded so as not to impact the core operation of the MDM/R.

Approach

The IESO needed a technology that could handle large volumes of time series smart meter data efficiently and at the lowest total cost. The solution needed to handle simultaneous data ingestion, transformation, and data retrieval requests, and be synchronized in near real-time with the source database system.

After conducting market research, evaluations and proof of concept testing, the IESO found that Kx kdb+ and its q programming language, delivered very high levels of performance at the lowest total cost of ownership. The IESO's testing of technologies involved simulated data representing 260 billion meter reads for 5 million meters and over 5 years. IESO conducted over 25 tests on this data including data ingestion, on-demand retrievals, aggregations and research analytics.

As soon as the data is available from the source database, Kx kdb+ ingests the data, performs a number of validation checks, inserts this data into a new data model, and makes it immediately available for data retrieval and analytics.

Implementation Overview

- In-memory and columnar database, and programming system
- High performance initial data load and intra-day data ingestion, transformation and loading
- High performance data access and analytics
- Integration with internal and external systems
- Secure web-interface and logging
- High availability and service level management
- Low infrastructure and support requirements

As kdb+ is a time-series database it facilitated the storing and making available all of historical data and versions for archival and analytics purposes. Most recent data requests are retrieved using kdb+'s in-memory database, while all historical and versioned data is maintained across two tiers of storage. As needed, older data is compressed to further save on storage and infrastructure costs.

Integrated with kdb+ is a high-performance internet facing web-services interface for fulfilling data retrieval requests from utilities and authorized parties. This interface receives and returns responses to XML requests using data stored in kdb+. This solution handles hundreds of web service requests per second with the ability to scale to more as needed, and fulfills requests spanning a full year of data in milliseconds.

Kx also provides a Graphical User Interface for business analytics and data scientists to perform ad-hoc analytics on live and historical data, which is useful for operational investigations and exploratory queries.

As a high performance analytics engine, Kx supports ad-hoc aggregation and reporting, including custom extracts. For example, with kdb+ IESO is able to aggregate an entire year of data for all meters in a few minutes, which supports analysis of load profiles, consumption behavior and different rate models.

The solution was implemented with security and privacy by design. It is housed in secure data centers, enforces authorization and authentication rules and logs all requests and responses, and was subjected to extensive testing by the IESO, utilities and an independent consultant.

The solution was implemented for high availability, with no-single points of failure, automated process fault detection and recovery, and service level monitoring and reporting.

About Kx

Kx has been the software leader for complex analytics on massive-scale streaming data for over two decades. The Kx technology is an established and trusted standard for trading, surveillance and research in financial services. Kx has emerged as the elegant, integrated solution to the analytic needs of the Pharma, Retail, Utilities and High Tech Manufacturing industries scaling effortlessly to the data challenges of the Internet of Things.

Kx is a division of First Derivatives plc. Listed on the London Stock Exchange [FDP:LN] First Derivatives is a specialist software and consulting organization with a uninterrupted track record of customer success and profitable growth since it was listed in 2002.

Volumes at a Glance

- 67 distribution utilities and their authorized service providers
- 40 billion measurements per year and over 4.2 million active metering points
- 7+ years of historical data; thousands of tables
- Ingests and makes data available for analytics at 300K to millions of records per second
- Supports hundreds of web-service requests per second
- Fulfills data requests in milliseconds
- Two Kx kdb+ data servers in a high availability configuration

Results

With the Kx powered Data Mart for smart meter data the IESO achieved the following results:

Improved performance and responsiveness: Improved response time to small and large data requests, delivering results over 100 times faster than before.

Improved service availability: Provided 24x7 web service access to data for utilities and their customers, enabling access to data even when the source meter data management is on maintenance periods.

Enhanced scope of services: Made more data available (365 days vs 100 days) for data retrieval through web-services interface. Enhanced on-line access to data from 27 months to 10+ years.

Enhanced analytics: Enabled ad-hoc investigation and analysis that was not possible before, including analyzing data from live data processing and historical master and meter data to improve system operations. Supports future demands for advanced analytics in Ontario's energy sector with the ability to rapidly aggregate and perform statistical analysis on all historical data.