



www.ibm.com

1 New Orchard Road, Armonk
New York 10504-1722
Tel: 914 499 1900

IBM Informix time-series and spatial capabilities

“ We were looking for an enterprise-class database that was truly tailored for use in an IoT environment, and IBM Informix fit the bill perfectly. With Informix, we get excellent scalability and strong security, and can manage both structured and unstructured data effectively. What’s more, the platform has a very small footprint, making it ideal as an embedded database. ”
Petrosoft

The company

IBM needs no introduction. Informix is an object-relational database originally developed by Informix Corporation, which was founded in 1980. In 1996, the company acquired

Illustra Technologies, upon whose time-series capabilities, Informix’s capabilities are based.

The Informix database was acquired by IBM in 2001 and in 2017 IBM and

HCL Technologies (Software Division) agreed to a long-term, 15-year partnership to co-develop, support, and market the product. The database is available both on premises and in the cloud. A pre-configured, ready to use Docker container of the

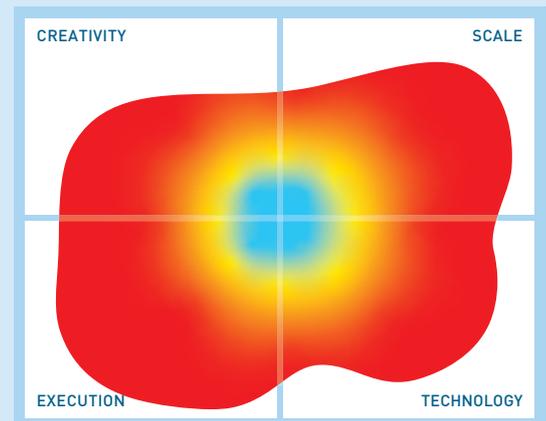
free Informix Developer Edition is available on www.dockerhub.com.

What is it?

IBM Informix is an object-relational database with native support for both time-series and geo-

spatial data. After its acquisition of Illustra in the mid-1990s its initial focus was on irregular time-series as generated within capital markets (stock ticks). It subsequently extended this to support regular time-series of the type that might derive from sensors. In 2015 the product was extended – leveraging its existing geo-spatial capabilities – to support spatiotemporal queries for moving objects and, even more recently, support for high frequency (sub-second) time-series, sometimes known as Hertz data, has been added.

The product is positioned as a database for Internet of Things (IoT) environments and highly distributed applications. It has a small enough footprint (it will run on a Raspberry Pi Zero, though



The image in this Mutable Quadrant is derived from 13 high level metrics, the more the image covers a section the better. Execution metrics relate to the company, Technology to the product, Creativity to both technical and business innovation and Scale covers the potential business and market impact.

preferably a Pi2 or later), together with fire-and-forget functionality, to make it practical to deploy at the edge. Conversely, it is equally capable of running in mainframe environments. The product is also suitable for – and there are clients using it for – environments where the emphasis is on monitoring and metrics rather than analytics.

How does it work?

Informix supports a native TimeSeries datatype with relevant data stored as illustrated in **Figure 1**. Note that this is much more efficient than would otherwise be the case for a relational database, which would typically have multiple rows for each Meterid: one for each time stamp. Compared to Informix’s approach, which does not require indexes, a relational storage model would require these on the meterID (or more), thereby adding significant overhead to the storage model. Specialised compression capabilities are provided. In addition, if it is a regular time-series Informix allows you to skip the storage of the time stamps altogether. If readings are static – subject to intervals you define – you can compress the results or filter them out. A variety of windowing data analytics functions – moving, static, rolling and so on – are provided.

Ingestion is via the Informix TimeSeries Loader, which is parallelised to ensure high performance.

Meterid	Reading
1001	2016-01-01,daily,(12.34, 12567),(12.56,9000),(12.35,55567),...
2002	2016-05-05,daily,(199.08,6780),(198.55,3405),(198.12,250),...
2011	2016-09-01,daily,(9.34,8067),(9.56,9000),(9.40,10780),...

Figure 1 – Support for native time-series datatype

Time-series architecture	★★★★★
General architecture	★★★★★½
Time-series processing	★★★★★½
Tools	★★★★½

Performance and scale	★★★★★½
Geo-spatial	★★★★★½
Embedding	★★★★★
Other functionality	★★★★★½

“ We estimate that using Informix works out at around a tenth of the cost of a traditional relational database, and the time saved enables us to cut our time-to-market for new services.”
Hildebrand

You can access TimeSeries data using SQL (for which there are more than 100 pre-defined functions), either via a virtual table interface, which makes the time-series data look relational; or you can create user-defined routines (UDRs) using stored procedures. C (for which further functions are provided) is also available, as is Java. These UDRs are registered with the database optimiser and used for query and other CRUD purposes. While HCL/IBM do not have partnerships with any third party visualisation vendors, there are users employing tools such as Grafana and Tableau (via ODBC/JDBC connectivity) together with Informix time-series. Of course, the product integrates with IBM Cognos. Sophisticated analytics, such as pattern matching, are available, as shown in **Figure 2**, though there is no facility to discover patterns in the first place. There are specific TimeSeries API capabilities to support anomaly detection and other advanced analytics.

Pattern Match: IoT Scenario – An Example Facility Management

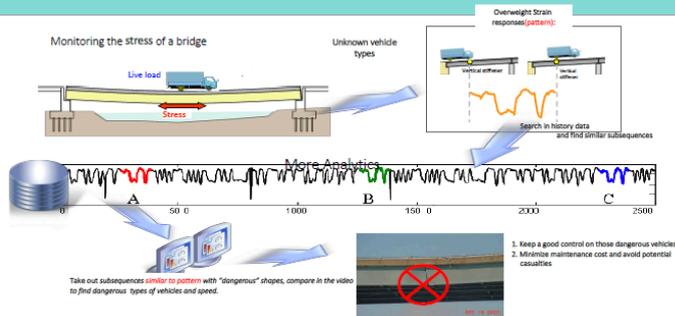


Figure 2 – Example of pattern matching

Geo-spatial capabilities represent an important complement to time-series in many IoT environments and Informix supports around 20 different co-ordinate systems with built-in indexing and native support for spatial datatypes. There is tight integration with ESRI and support for GeoJSON. It is also worth commenting on the utility of the spatiotemporal capabilities provided by Informix. These are not limited to traditional logistics examples, but also allow you to answer questions like “what was the location of bus number 3435 at 2019-03-01 15:30?”, “when was a delivery truck within 100 meters of the Good Food Diner between February 2-4, 2019?”, “when was



Figure 3 – InformixHQ

taxi number 324 at the Superlative Hotel?”, “how long do buses stay in our repair shop?”, or “which taxi driver witnessed an accident by finding which taxi was nearest to the location of the accident at 9:00 am?”.

Finally, it is worth mentioning InformixHQ, which is illustrated in **Figure 3**. This is a monitoring and administration tool that has been developed for the Informix environment itself. However, it does have features that will allow you to construct your own dashboards, though these probably need further development before the product will be suitable to support monitoring of third-party environments.

Why should you care?

Databases that offer time-series capabilities come in a wide range of shapes and sizes. Some are purely focused on analytics and/or monitoring, while others – including Informix – support both analytics and operational uses cases, as well as hybrid requirements that need both. Then again, many products other than Informix that support time-series have limited, if any, geo-spatial capabilities, which means they simply are not suitable for many IoT applications. Moreover, Informix is unique, as far as we know, in offering spatiotemporal search.

The Bottom Line

Informix has been offering TimeSeries and geo-spatial data longer than any other database on the market. In many cases for longer than competitive offerings have been in existence. The experience that comes with this longevity is clear in the depth and breadth of the capabilities offered.

FOR FURTHER INFORMATION AND RESEARCH [CLICK HERE](#)