

Product Data Sheet: TimeBase

Time Series Database

TimeBase is a high performance event-oriented time-series database engine and messaging middleware. TimeBase is designed explicitly for very fast population and retrieval of massive volumes of time-series data and delivering that data for subsequent use by both QuantOffice and third-party software. High volumes data such as fundamental data, news data, daily, bar, tick and Level II (depth of market data) are handled with superior speed. Processing speeds are measured in millions of messages per second on terabytes of data. TimeBase can be populated with historical data and receive real-time data feeds. TimeBase has a sophisticated time-series join engine, capable of efficient on-the-fly merging of multiple data streams with arbitrary temporal characteristics into a unified query response. This capability is central to the TimeBase architecture and is available regardless of the nature of the data source.

Data Feeds

TimeBase supports digitized news, fundamental data, Level I and Level II data. Many market data vendors and exchanges are supported.

Multiple data feeds can populate TimeBase simultaneously, even for the same instrument.

Instruments

Equities, options, futures, currencies, baskets and custom synthetic instruments are supported. Synthetic instruments range from baskets, to the output of strategies themselves.

Data Periodicities and Regularity

Any periodicity of data can be handled: down to 1 millisecond intervals such as tick and Level II data. Data can be both *regular* and *irregular*. For example, irregular tick data can be maintained and converted into bar data (regular).

Bar Creation

When recording tick or intra-day data, TimeBase can create and maintain time-based bars automatically. In QuantOffice, users can create bars using more complex techniques, which are then stored in TimeBase. For example, users can define logic for “by equal volume” or “by equal trade number” bar generation. Such bars are stored in TimeBase in real-time and are accessible to QuantOffice in real-time.

Data Population

TimeBase supports historical population, backfilling and real-time data feeds. Deployed with QuantOffice, TimeBase can be used for *both back-testing and managing strategies in production*.

User Interaction

TimeBase has a Graphical User Interface (GUI), called TimeBase Administrator, which provides system and data configuration and graphing capabilities which are capable of displaying years' worth of tick data across thousands of instruments in seconds. With the GUI, the user can zoom into the graphs to inspect data at the most granular level. All graphs and the raw data can be downloaded.

TimeBase Administrator - Data View

Time	deltix.qsrv.hf.pub.TradeMessage		deltix.qsrv.hf.pub.BestBidOfferMessage			
	Price	Size	Bid Price	Bid Size	Offer Price	Offer Size
01/02/2009 09:30:31.000			0.000000	0.000000	84.660000	1.000000
01/02/2009 09:30:31.000			84.610000	2.000000	0.000000	0.000000
01/02/2009 09:30:31.000	84.630000	74,300.000000				
01/02/2009 09:30:31.000			84.470000	4.000000	0.000000	0.000000
01/02/2009 09:30:31.000			84.400000	50.000000	0.000000	0.000000
01/02/2009 09:30:32.000			0.000000	0.000000	84.660000	2.000000
01/02/2009 09:30:32.000			84.450000	1.000000	0.000000	0.000000
01/02/2009 09:30:32.000			84.450000	2.000000	0.000000	0.000000
01/02/2009 09:30:32.000			84.400000	50.000000	0.000000	0.000000
01/02/2009 09:30:32.000			84.480000	20.000000	0.000000	0.000000
01/02/2009 09:30:32.000			84.410000	1.000000	0.000000	0.000000
01/02/2009 09:30:32.000			84.410000	3.000000	0.000000	0.000000
01/02/2009 09:30:32.000			84.520000	1.000000	0.000000	0.000000
01/02/2009 09:30:32.000			84.540000	1.000000	0.000000	0.000000
01/02/2009 09:30:34.000	84.650000	200.000000				
01/02/2009 09:30:34.000			84.650000	2.000000	0.000000	0.000000
01/02/2009 09:30:34.000	84.660000	100.000000				
01/02/2009 09:30:34.000			84.540000	1.000000	0.000000	0.000000
01/02/2009 09:30:41.000	84.630000	100.000000				
01/02/2009 09:30:41.000			84.630000	1.000000	0.000000	0.000000
01/02/2009 09:30:41.000			84.540000	1.000000	0.000000	0.000000
01/02/2009 09:30:42.000			0.000000	0.000000	84.660000	1.000000
01/02/2009 09:30:42.000	84.660000	100.000000				
01/02/2009 09:30:42.000			84.630000	1.000000	0.000000	0.000000
01/02/2009 09:30:42.000			0.000000	0.000000	84.660000	2.000000
01/02/2009 09:30:42.000			0.000000	0.000000	84.709900	2.000000
01/02/2009 09:30:43.000	84.709900	100.000000				
01/02/2009 09:30:43.000			0.000000	0.000000	84.709900	1.000000
01/02/2009 09:30:45.000	84.709900	100.000000				
01/02/2009 09:30:47.000			0.000000	0.000000	84.709900	2.000000
01/02/2009 09:30:58.000			84.560000	1.000000	0.000000	0.000000
01/02/2009 09:30:58.000	84.630000	100.000000				
01/02/2009 09:30:58.000			0.000000	0.000000	84.650000	16.000000
01/02/2009 09:30:58.000			84.540000	1.000000	0.000000	0.000000
01/02/2009 09:30:58.000			0.000000	0.000000	84.650000	17.000000
01/02/2009 09:30:58.000			0.000000	0.000000	84.650000	18.000000
01/02/2009 09:30:59.000	84.630000	100.000000				

TimeBase Administrator – Graph View



Deltix Product Suite

Together with the QuantOffice products, TimeBase is a core component of the Deltix Product Suite that addresses all stages of the Alpha Generation Process. As such, TimeBase is either deployed as part of a QuantServer and QuantOffice deployment, or stand-alone, integrated with a client's in-house or third-party applications.

Architecture and Design

TimeBase provides diverse ways of specifying data types, as well as hints for encoding values for transmission and persistence. One benefit of such diversity is a more native representation of user's data model. For example, TimeBase natively supports enumerations. As an example, consider a data item whose value is 'SPOT', 'FORWARD', 'SWAP' or 'N/A', in TimeBase, one can define an enumeration, and everyone will always know what the allowed values are. Additionally, this data item can map to native Java or C# enumerations for programming clarity.

The use of rich data types enhances system performance. This is achieved by TimeBase knowing the detailed data type of every value, and the user's ability to define at a very fine-grained level how the data should be encoded. TimeBase can then allocate less resource (memory, network bandwidth, disk space) and process data quicker. In the above example of an enumeration, each enumerated data

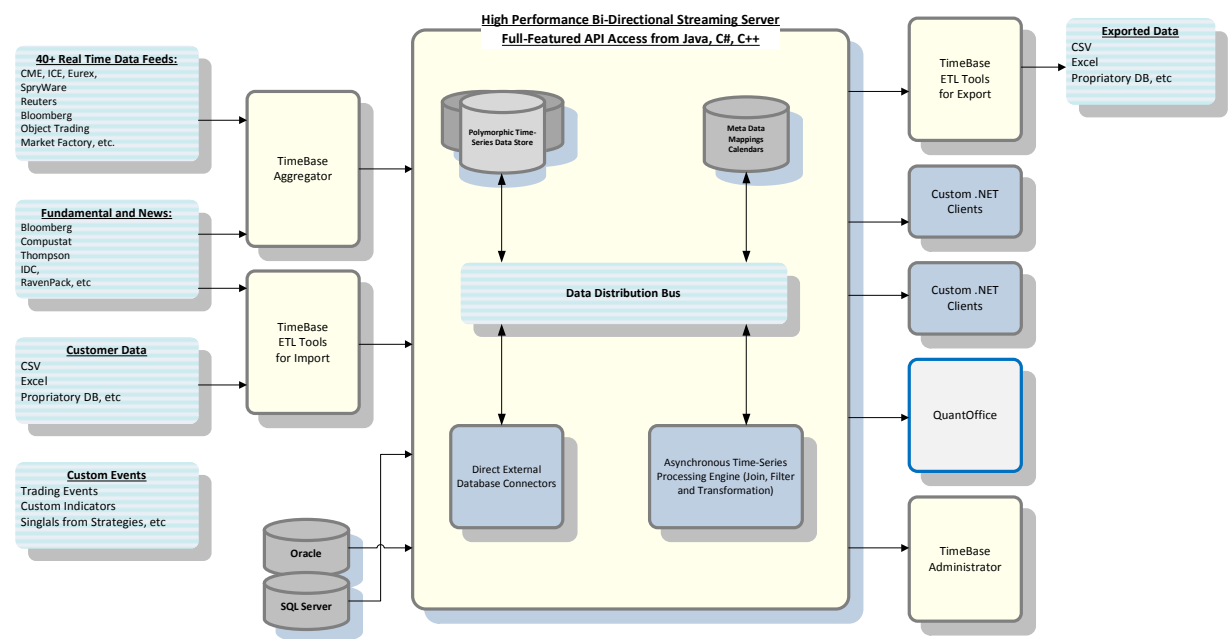
item only requires 2 bits to store its value (there are only 4 possible values), even if it appears to the user as a text constant.

TimeBase supports special compression for small alphanumeric codes which allows the data item to appear as text, while it is internally stored in extremely compressed form. Examples of alphanumeric codes include exchange codes, currency codes, liquidity provider codes, or any other codes the user desires to define.

TimeBase supports so-called relative encoding. In a price bar, it makes sense to store close as a number, and open, high, low as a delta from close. This saves space and increases speed. In TimeBase the relative encoding is a generic part of the type system and can be used with any user-defined data structure, not only price bars.

TimeBase has a built-in message broker. All interactions with the database, either via Deltix products, user or third-party applications, are handled by this message broker.

TimeBase Logical Design



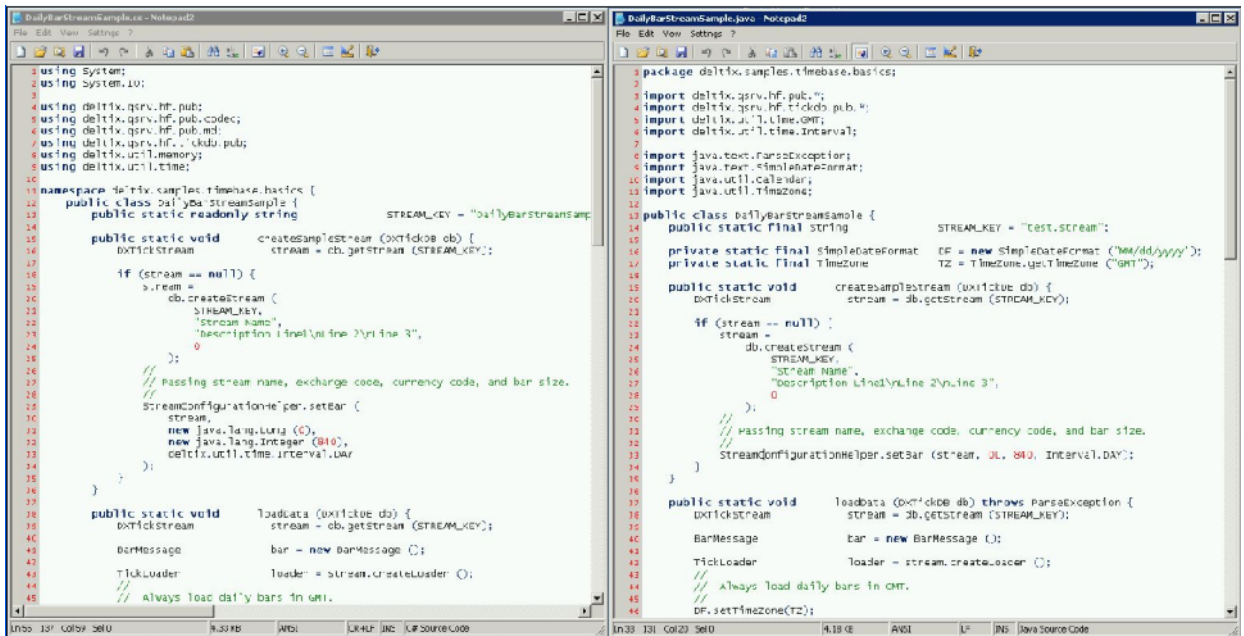
Hardware

It is a common misconception that high performance requires either accelerated or very high end (that is, expensive) hardware. It does not. But, high performance does require superior software engineering: a skill often forgotten in a world of ever more powerful computers. TimeBase is written in Java, and so runs on Windows and Linux platforms.

API

TimeBase has a rich API, against which C#, C++ and Java code can be written. Examples are shown below:

Interfacing with the TimeBase[™] API: C# and Java code examples



The image displays two side-by-side screenshots of code editors, illustrating C# and Java code examples for interfacing with the TimeBase API.

Left Editor (C#): The code is in a file named `DailyBarStreamSample.cs`. It uses the `System` and `System.IO` namespaces. It defines a `namespace deltax.samples.timebase.basics` containing a `public class DailyBarStreamSample`. The class has a `public static readonly string STREAM_KEY = "dailybarstreamsample";`. It includes a `createSampleStream` method that takes a `DateTime` object and returns a `Stream`. It also has a `loadData` method that takes a `DateTime` object and returns a `BarMessage` object. The code uses `StreamConfiguration` to set the stream name, exchange code, currency code, and bar size.

Right Editor (Java): The code is in a file named `DailyBarStreamSample.java`. It uses the `deltax.samples.timebase.basics` package. It imports `deltax.grv.hf.pub`, `deltax.grv.hf.tickdb.pub`, `deltax.jl.timegmt`, `deltax.jl.timeinterval`, `java.text.ParseException`, `java.text.SimpleDateFormat`, `java.util.Calendar`, and `java.util.TimeZone`. It defines a `public class DailyBarStreamSample`. The class has a `public static final String STREAM_KEY = "test.stream";`. It includes a `createSampleStream` method that takes a `DateTime` object and returns a `Stream`. It also has a `loadData` method that takes a `DateTime` object and returns a `BarMessage` object. The code uses `StreamConfiguration` to set the stream name, exchange code, currency code, and bar size.