



Powerful columnar database for time series and unstructured data

*The ideal solution for your time
series and database requirements*

The ultimate in modern analytic database systems: Kerf runs on most operating systems and produces timely queries and analytic results.

Kerf is competitively priced and uses the latest technologies for high performance columnar databases.

Kerf is in use by cutting edge data providers, data scientists, quants and academic users worldwide.

Kerf is suited to daily data, ticks, quotes, telemetry data and log files.

Sample Use cases

- Real time tick data processing to support multi strategy trading groups
- Backtesting historical data to support rapid trade idea review and testing.
- Risk calculations and compliance checking tools
- Trading cost analysis tool
- Logfile analysis
- NoSQL stores
- Smartmeter/electrical load data
- ETL for data scientists
- Complex event processing

IF you have time series problems and want robust high performance software which is efficient, easily maintained and scales to fit your data problems, Kerf might be the answer for you. Kerf uses the latest in high performance techniques to make time series joins and aggregates blazingly fast, processing billions of rows per second. Kerf is a high performance programming language in addition to high performance C code to Kerf is a high performance scalable columnar database and integrated data programming environment particularly suited to working with timeseries. Kerf's priorities are speed, maintainability and ease-of-use. It solves complicated time-series problems with a shorter learning curve and faster development time than other systems. Kerf is well suited to building ticker plants, processing log files, analysis of telemetry data and financial time series.

- Kerf is a revolutionary high speed scalable columnar database and integrated data programming environment particularly suited to working with timeseries.
- Kerf is fast: billions of rows per second can be processed in single threaded mode.
- Kerf is programmable: new aggregates and complex algorithms can be written in the Kerf programming language.
- Kerf is simple to use: queries are constructed using the SQL syntax familiar to most programmers. Data is represented in a JSON-like format, and can be presented to other applications in standard JSON.
- Kerf is flexible: optimized NoSQL queries are available to the end user.
- Kerf is scalable: server mode, multi-threaded queries and concurrency are available in Kerf.
- Kerf is suited to real time and historical databases.
- Kerf will save you money: flexible licensing, familiar and maintainable source code, and supercharged queries make for a low total cost of ownership over competing solutions

Kerf licenses and support are available from

Briarcliff Hall LLC	Kerf Software
261 Madison Avenue	2946 MLK Way
New York, NY 10016	Berkeley CA 94703
sales@briarcliff-hall.com	support@getkerf.com
1 (201) 572-2766	

More information can be found at
<http://www.kerfsoftware.com/>

Get in touch with the Kerf team (see addresses in panel). We love challenging data problems, and would be excited to hear about yours!

Kerf

Kerf – next generation Time-Series database platform for very large volumes of data to rapidly support your trading, research, risk, operational or reporting requirements. Bring your data to life using the scale, speed, flexibility and agility of Kerf. Key advantages are:

- ❑ Simple, powerful programming syntax for ease of use and maintenance
- ❑ Speed: Kerf's first priority is speed and leads the speed race with ability to exceed a billion rows a second
- ❑ High performance joins and aggregations on timeseries data
- ❑ Ease of use: familiar SQL queries and JSON data formatting
- ❑ All standard API's, no need to learn an extensive proprietary language
- ❑ Time is a first class data type
- ❑ Very light weight installation and can scale instantly
- ❑ Cost Mitigation / Efficiency: Kerf is compact and very efficient with numeric time series data so can reduce the hardware and hosting costs for your applications.
- ❑ Server mode
- ❑ Production high frequency trading algorithm calculation engine
- ❑ Back testing tool for new strategies
- ❑ Real time tick store
- ❑ Kerf can be used as a market data store and on-demand calculation engine
- ❑ Rapid pre-processing for risk or other applications
- ❑ Flexibility: tables in Kerf have the option of being unstructured, or schemaless, for dealing with highly irregular data
- ❑ Prebuilt loaders for Bloomberg, Thomson Reuters, Quandl, Xignite with more coming

What is Kerf

Kerf is a high performance scalable columnar database and integrated data programming environment particularly suited to working with timeseries. Kerf's priorities are speed, maintainability and ease-of-use. It solves complicated time-series problems with a shorter learning curve and faster development time than other systems. Kerf is well suited to building ticker plants, processing log files, NoSQL data stores, analysis of telemetry data and complex analytics of financial time series.

Power

Kerf's simple syntax makes your engineers productive immediately. Nanosecond TimeStamps are a first class data type in Kerf, and time based queries are highly performant. Kerf queries are composed in a standard SQL syntax for simple and maintainable development. There are optimized core primitives for time series data processing, and new aggregations and time based functions can be written in Kerf and applied to Kerf databases. Kerf's data syntax is similar to JSON, and Kerf tables can be output as reports in standard ASCII JSON. Kerf can also organize data in a schemaless NoSQL manner, with optimized key/value retrieval. Kerf's query and aggregation execution speed is such that expensive distributed solutions can be replaced with one Kerf node.

Capabilities

Kerf has a server mode, and a high level API for interprocess communication and parallel processing. Kerf can be called by other programming languages via a C api, or in server mode via a networked JSON client. Kerf can also do NoSQL queries on unstructured data using the atlas structure. Kerf has feed handler and loaders for popular financial data feeds such as Bloomberg, Quandl, Reuters, Xignite and TAQ data.

Technology

Kerf has its roots in the APL family of programming languages. Founder Kevin Lawler wrote the first open source K interpreter, Kona, before writing Kerf. Kerf is designed to be easily usable for non-APL experts, removing the unusual syntax and opacity of this family of programming languages. Kerf uses a familiar syntax for function definition, SQL queries, and represents data in JSON-like form. In addition to saving developer time, the simplicity and legibility of Kerf code make it easily maintainable. It retains the powerful APL abstractions for array data processing, and the analytic performance capabilities of an APL while generating maintainable code in a powerful rapid prototyping language.

History

Kerf is the team's third major programming language release, and a fifth generation language and database system. It draws on lessons from over twenty years of programming. It is mature technology.

The team was first introduced to array languages in 2006 by Dennis Shasha at NYU. This led to work with Arthur Whitney's and Kx System's kdb+ family of languages at the investment banks Cantor Fitzgerald and Merrill Lynch.

A precursor language to Kerf called Kona was started around 2009. Kona was open-sourced in Summer 2010 and improved upon with the community for the following four years. Lessons from Kona would inspire Kerf. The first lines of code for Kerf were written in Summer 2014. Kerf officially launched as a product in Spring 2015.