kx | it's about time

Streaming Analytics - Change the Game

Serverless kdb+

Instant Processing Capacity, As You Need It. Automatically.

Serverless computing is a cloud-computing model by which your infrastructure team or your cloud service provider will run and manage physical infrastructure, dynamically allocating these resources by demand of the application itself.

Kdb+, with its small footprint, high performance profile, and elegant design makes it ideal for offloading certain aspects of time-series analysis to a serverless solution. Infrastructure is allocated at runtime, according to the amount of resources required by kdb+, rather than needing to pre-define units of capacity that may lie idle or dormant for long periods of time.

Benefits of Serverless kdb+ Include:

- No hardware to provision - your kdb+ application will run without requiring you to provision or manage servers.
- High Availability: Serverless kdb+ processes scale automatically with built-in high-availability and fault tolerance.
- Continuous scaling: your application scales automatically and appropriately in sync with the size of changing workloads.
- Cost per execution. Costs based on kdb+ execution time per function. There is no cost when kdb+ isn't running.

Serverless allows kdb+ developers to focus on the application logic. Building "serverless" means that kdb+ developers can focus on the core product instead of having to manage and operate virtual machines, physical machines or associated infrastructure.

The reduced overhead lets kdb+ developers reclaim time and energy that can be spent on developing services, especially where those applications have to dynamically scale and have demanding high-availability requirements.

At a Glance

Kdb+ is now available on AWS Lambda enabling users to avail of the benefits of serverless computing: Instant availability, automatic hardware provisioning, pay-as-you-use processing.

The Kx Advantage

- Complex Event Processing for real-time data capture and analysis across multiple data sources
- In-memory computing supporting "compute then store" for instant results
- Low latency messaging for high-speed data capture and distribution
- Lambda/HTAP/Translytical for combined analysis of real-time and historical data
- In-memory data grid for high performance, distributed computing
- Optimized for time-series data using columnar storage and built-in temporal data types
- Single Integrated software stack for quicker, easier implementation and lower TCO
- Integrated real-time visualization and enabling self-service Business Intelligence
- Built-in resilience and fault tolerance capabilities

Kdb+ is now available in serverless mode on AWS Lambda

AWS Lambda is an event-driven, serverless computing platform. It runs code in response to events and automatically manages the computing resources required by that code.

Lambda can be described as a type of serverless Function-as-a-Service (FaaS). FaaS is one approach to building event-driven computing systems. It relies on functions as the unit of deployment and execution and provides provision-free scalability and built-in reliability.

Lambda functions can be triggered by a variety of events that occur on AWS or on supporting third-party services. They enable you to build reactive, event-driven systems. When there are multiple, simultaneous events to respond to, Lambda runs more copies of the function in parallel and scales with the size of the workload.

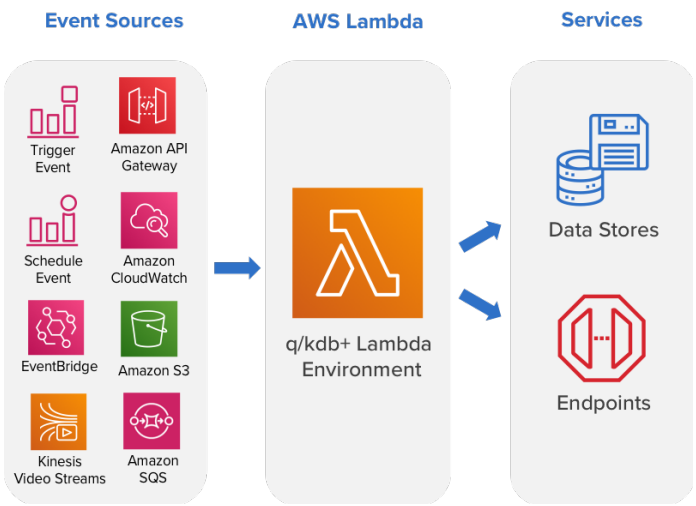

Optimal Processing

Parallel capable "on-demand", workloads written for kdb+/q can easily fit into serverless. A business process that has to run in either intense batch load or requires an extract-translate-load process suits this model. This can allow kdb+ to offload services such as foreign data ingest into a serverless process that exists only to handle that data translation, then shuts down.

On-Demand Availability

Serverless is a programmatically friendly extension of the existing "on-demand" licensing model for kdb+. Only pay for what you use for a functional workload, but invoke workloads from the application logic in kdb+/q, and not from an infrastructure layer with fixed licenses.

Licensing

Commercial licensing and support for kdb+ running in a serverless environment is now available. Whether or not you are an existing customer of Kx, you may use serverless kdb+.

For more information please visit serverless.kx.com

About Kx

Kx is a division of First Derivatives, a global technology provider with more than 20 years of experience working with some of the world's largest finance, technology, automotive, utility, manufacturing and energy institutions. Kx technology, incorporating the kdb+ time-series database, is a leader in high-performance, in-memory computing, streaming analytics and operational intelligence. Kx delivers the best possible performance and flexibility for high-volume, data-intensive analytics and applications across multiple industries. The Group operates from 15 offices across Europe, North America and Asia Pacific, including its headquarters in Newry, and employs more than 2,400 people worldwide. For more information about Kx please visit www.kx.com. For general enquiries, write to info@kx.com. For press inquiries, write to pr@firstderivatives.com.

Head Office
3 Canal Quay,
Newry,
BT35 6BP
N. Ireland
+44 (0)283 025 2242

London
Cannon Green Building,
27 Bush Lane,
EC4R 0AN
United Kingdom
+44 (0)207 337 1210

New York
45 Broadway,
NY 10006
USA
+1 (212) 447 6700

Singapore
One Raffles Quay,
North Tower,
#30-03,
048583
Singapore
+65 6592 1960

Sydney
22 Pitt Street,
NSW 2000
Australia
+61 (02) 9236 700

Toronto
31 Lakeshore Road
East, Suite 201
Mississauga, On,
L5G 4V5
Canada
+1 289 329 0636

Tokyo
20th Floor,
Shin-Marunouchi
Center Building,
1-6-2, Marunouchi,
Chiyoda-ku,
Tokyo, Japan
+81 (0)36 634 9799

Contact us at sales@kx.com or visit www.kx.com