

Fast, intuitive in-memory platform for stateful stream processing

ScaleOut StreamServer™ combines a scalable, stream-processing compute engine with an integrated, in-memory data grid (IMDG) into a powerful, unified software platform for stateful stream processing. Now you can perform lightning-fast event analysis using sophisticated in-memory state tracking to provide deep introspection and precise real-time feedback. Ideal for a wide range of applications, including the Internet of Things (IoT), manufacturing, logistics, and financial-services. ScaleOut StreamServer introduces breakthrough technology for the next generation in stream processing.

STATEFUL STREAM PROCESSING FOR DEEP INTROSPECTION

Live systems generate streams of incoming events that need to be tracked, correlated, and analyzed to identify patterns and trends — and then generate immediate feedback and alerts to steer operations.

With today’s ever more complex real-time systems, it’s not enough to just analyze patterns within data streams using conventional techniques. Applications need deeper introspection to extract full value from the telemetry they receive. They need to build dynamic models of data sources that they can continuously update and analyze. Called **stateful stream processing** and popularized as the “digital twin” by Gartner, this breakthrough approach can harness machine learning, neural networks, and other advanced techniques to enable deep introspection and provide precise, timely feedback for live systems.

By integrating a fast, scalable stream-processing engine with an in-memory data grid, ScaleOut Software has created a unified software platform for the next-generation of stream processing. Unlike mainstream platforms such as Apache Flink, Spark, and Storm, ScaleOut StreamServer enables applications to implement object-oriented models of data sources. It can host large populations of data objects in memory on a cluster of commodity servers and dispatch incoming streaming events to these objects for analysis. Applications now can process incoming data streams in a rich context of evolving state, enabling the use of sophisticated algorithms while delivering blazingly fast event handling.

MOVE STREAM PROCESSING TO WHERE THE DATA LIVES

ScaleOut StreamServer’s innovative architecture delivers both breakthrough capabilities and peak performance for stateful stream processing. It processes incoming data streams within an in-memory data grid — where the data lives — ensuring minimum latency and peak throughput. Other platforms need to pull state information from remote data stores, such as database servers and distributed caches; this creates delays and network bottlenecks. Instead, ScaleOut StreamServer delivers streamed events directly to their associated state data, enabling immediate, fully contextual processing. Its transparently scalable platform minimizes the latency required for event tracking and analysis, ensuring timely feedback and/or alerts for the largest workloads.

Here are just a few applications:



Ex: Portfolio tracking, wire-fraud detection, stock back-testing



Ex: Device tracking for manufacturing, vehicles, mobile devices



Ex: Real-time patient monitoring and alerting, health device tracking



Ex: Real-time inventory reconciliation, manufacturing flow optimization

...and much more



KEY FEATURES AND CAPABILITIES

ScaleOut StreamServer unleashes the power of stateful stream processing. Its seamless integration of a scalable stream processing engine and in-memory data grid creates a powerful, unified platform for building digital twin models and performing deep introspection on streaming data — with blazing performance and built-in high availability. These capabilities are delivered as an intuitive, easy to use SDK that makes application development in C# and Java simple and straightforward. Automatic code shipping to the in-memory data grid simplifies application deployment and helps ensure fast startup times.

Key Features	Benefits
Integrated IMDG & streaming engine	Enables dynamic digital twin models.
Stream processing runs in the grid	Eliminates network bottlenecks.
Automatic code shipping	Simplifies app. design and deployment.
Support for event posting with Reactive Extensions APIs	Minimizes latency for immediate feedback and is easy to use.
Kafka integration (connector & producer)	Enables connectivity to existing Kafka pipelines.
Scalable Kafka connections	Transparently scales throughput.
Comprehensive time window libraries	Enables easy integration of time windowing into digital twin models.
Data-parallel APIs (e.g., MapReduce)	Allows aggregate analysis of live data.
Automatic event routing to associated state (digital twin objects)	Simplifies design and allows seamless scaling for large workloads.
Object-oriented design in C# and Java	Separates application from platform orchestration; maximizes ease of use.

SCALABLE KAFKA INTEGRATION AND MORE

ScaleOut StreamServer was designed to seamlessly integrate into existing Kafka streaming data pipelines as both a consumer and producer of streaming data. It takes full advantage of the in-memory data grid's architecture to automatically scale the number of Kafka connections for large workloads and avoid bottlenecks.


Application developers can use familiar APIs from the popular Reactive Extensions (RX) library to implement extremely fast, lightweight event processing in both C# and Java. These APIs have been integrated with Kafka to allow posting of incoming Kafka messages to RX observables associated with grid objects for stateful stream processing. Likewise, RX observables can be associated with Kafka producers to quickly send outgoing messages.

Unique Advantages for Streaming Data

Traditional CEP and stream processing platforms, such as Apache Flink and Spark Streaming, focus on analyzing incoming data streams without regard to the context in which the data was created. The next generation of stream processing tracks the dynamic state of data sources as “digital twins,” offering a basis for much deeper introspection and more effective alerting. ScaleOut StreamServer's unique architecture, which executes streaming operations within an in-memory data grid, creates a breakthrough that enables stateful stream processing with digital twin models.

FAST, SCALABLE, POWERFUL — AND EASY-TO-USE

By integrating stream processing within an in-memory data grid, ScaleOut StreamServer offers capabilities not found in other platforms and delivers on the promise of stateful stream processing. Its familiar, object-oriented data storage and computing model in C# and Java supports advanced analysis algorithms and ensures a clean separation between application-specific code and the platform's orchestration of event processing.



Example of stateful stream processing for heart-rate monitoring devices

To illustrate the power of stateful stream processing with ScaleOut StreamServer's architecture, consider a heart-rate monitoring application which receives telemetry from wearable devices. ScaleOut StreamServer can route millions of incoming events to dynamic, in-memory models (“digital twins”) which track each patient's unique medical history and current condition, analyzing events and generating timely alerts to medical professionals when needed.

Now is the time to let ScaleOut StreamServer unlock the potential of stateful stream processing for your applications and give your business a powerful competitive edge.



Try ScaleOut for free.

Experience the power of in-memory computing on Windows or Linux.

www.scaleoutsoftware.com/try-for-free