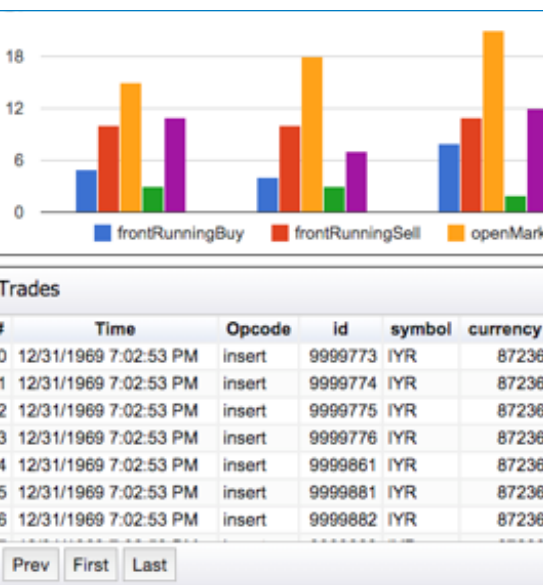


SAS® Event Stream Processing

Act on data while it's in motion to keep a real-time pulse on your business



What does SAS® Event Stream Processing do?

SAS Event Stream Processing analyzes and understands millions of events per second, detecting patterns of interest as they occur. The results show the correct actions to take, what alerts to issue, which data to store and which events to ignore.

Why is SAS® Event Stream Processing important?

Get value out of streaming data from operational transactions, sensors, devices, transmissions, etc. – while the data is in motion. With submillisecond response times, the software assesses data streams using a suite of prebuilt operators, functions, routines and advanced analytics. So you can derive insight into events and take appropriate actions.

For whom is SAS® Event Stream Processing designed?

A configurable, embeddable engine, the software is designed for application developers, IT engineers and system architects. The web-based interface makes it easy to manipulate data and build projects, while a dashboard delivers information visually so you can test and validate results.

Staying competitive in a data-driven world means working and making decisions incredibly fast. You need to assess continuously changing conditions, monitor for events of interest as they happen and follow through with contextually relevant action. With event stream processing, the data doesn't stop. Continuous queries constantly analyze live data to detect patterns of interest so you can respond in real time.

SAS Event Stream Processing helps you understand events while they are in motion (known as event streams). Instead of the traditional "stream, score and store" model, queries are stored and data is streamed through them. Data is analyzed continually as it's received, updating the situational intelligence as new events take place. Value, which otherwise would be lost through information lag, is captured immediately. The solution is designed for optimal performance with a flexible processing model that provides submillisecond response for high-volume throughput.

Benefits

- **Get instant information, take immediate action.** Streaming data from operations, transactions, existing systems, and Internet of Things (IoT) sensors and devices is valuable when you can act upon key insights in real time. Regardless of data or format, SAS immediately enhances the value of streaming data for information-driven applications, including customer experience, asset and IT performance, fraud, risk, security and more.
- **Store the right data.** Faster, better and more powerful stream processing ensures data corrections are made before high-volume throughput data is stored – reducing data management processing costs. Continuous queries run faster in the flexible threaded-processing model, which natively supports dynamic updates, deletions and insertions while data is in motion.
- **Make sound decisions based on sound analyses.** Having accurate data at your disposal means you can take the right action. SAS Event Stream Processing provides in-stream, prebuilt data quality operators, natural language text-processing functions and a wealth of advanced analytics algorithms to detect even the most complex event stream patterns.
- **Take full control and adapt to changes quickly.** No matter how fast your data moves, how much data you have, or how many data sources you're connecting to, it's all streamlined and under your control from a single, intuitive interface. Consolidated event stream data instructions defined in SAS and other languages ensure streamlined, governed pattern detection, which makes it easier to evolve with changing data conditions.

Overview

SAS Event Stream Processing ingests large volumes of streaming data quickly - millions of events per second - so you can understand events in the data while it's in motion. No data stream is too big or fast. The solution's processing speed is bounded only by the hardware environment's limitations.

Incoming data is read through adapters and connectors, which are part of a publish-and-subscribe architecture. Event data publishes

into a source window of an event stream processor. A visual interface makes it easy to define the windows, procedures and operators. In turn, it's simple to define continuous queries through which the data will stream. Streaming data is examined for patterns and can be intelligently filtered to store anomalies that demand deeper investigation. Or, if no relevancy is detected, the data can be discarded. Downstream applications subscribe to receive streaming analysis results with prescribed actions.

Consume and connect streaming data

Is your data center prepared to handle the rapid growth projected for IoT data? Most are not. Even data streams generated from current web transactions, sensor devices and operational systems can be challenging to consolidate and fully exploit. SAS Event Stream Processing can consume a wide array of data stream formats, and can easily integrate results into existing systems and applications. An extensive suite of prebuilt adapters and connectors enables you to source and manipulate live data streams to detect, filter, aggregate, correlate and delineate patterns of interest. Actions are applied to both structured and unstructured text data, with results sent to other systems to improve efficiency and lower costs. Your organization can consume the data sources it needs today with assurance that the same technology will work tomorrow.

Flexible, visual development environment

SAS Event Stream Processing provides a highly visual, interactive interface for building modular, continuous queries that use SAS advanced analytics algorithms and rules to pinpoint event relevance. System architects and application developers have point-and-click access to the palette of windows and connectors, making it easy to design complex streaming analytical models. An interactive test mode lets users evaluate logic and validate results before deployment. Events can be shared across projects and historic activity can be retained and compared to current events - all from the visual drag-and-drop environment. This means data in motion can be easily defined, updated and revised to address new and quickly emerging events.

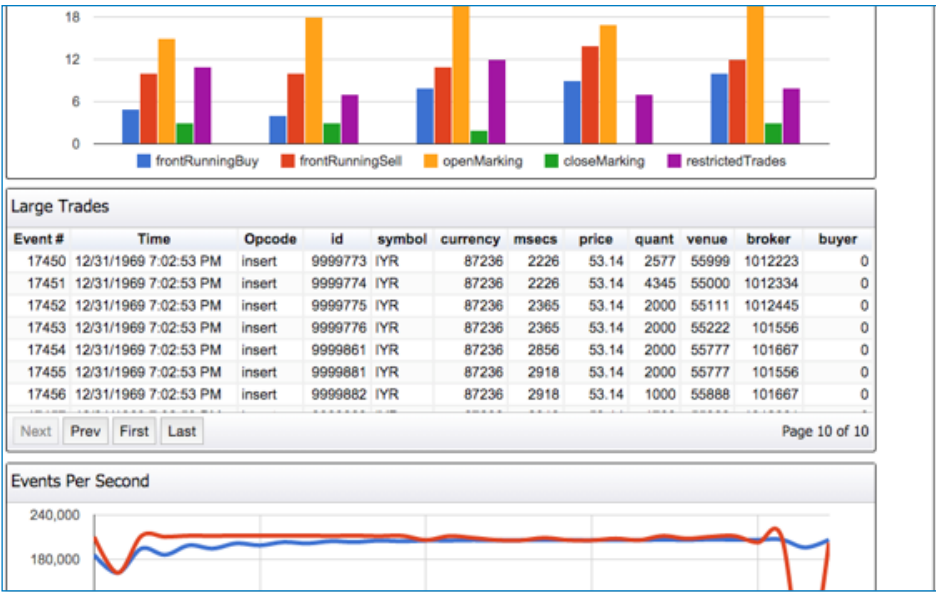


Figure 1: Test and refine event stream models in an easy-to-use HTML5 interface.

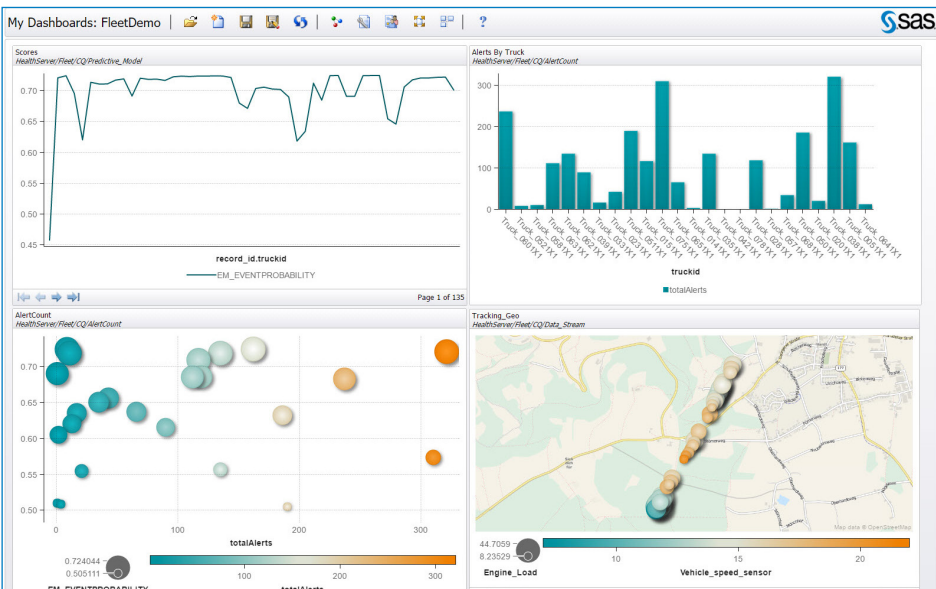


Figure 2: A real-time dashboard lets you monitor live events from a single location.

Advanced pattern-matching algorithms

SAS Event Stream Processing provides parsing, filters, joins, field calculations and pattern-matching functions available right out of the box. This includes prebuilt functions specifically designed for routine data management tasks, like transformations, normalization, matching, identification and more. Unstructured text processing functions provide natural language processing (NLP) extractions of concepts, entities and facts, as well as classification of text and sentiment identification. In-stream clustering builds homogeneous groups of events on the fly while predictive algorithms estimate future events, creating advanced patterns of interest. You can examine multiple events in a single query, including sequential and time-based (temporal) events.

Dashboard view of live event streams

The streamviewer in SAS Event Stream Processing is a real-time dashboard that lets you monitor live event streams from a single location. Use the dashboard to produce SAS graph visualizations, view multiple models across different SAS Event Stream Processing servers, share dashboards and create embeddable dashboards.

Key Features

In-stream learning model windows

- In-stream learning model windows allow you to use different window types to specify data stream input sources, patterns of interest and derived output actions. Learning model windows include:
 - Train. Develop an advanced analytical model in stream and pass the resulting model updates to a score window.
 - Score. Apply the trained model to current events in stream to produce score output.
 - Calculate. Use with data normalization and transformation methods, as well as learning models that bundle training and scoring together.
 - Model supervisor. Control what model to deploy, and when and where to deploy it (for example, to the score window).
 - Model reader. Integrate offline ASTORE models and publish a model to another streaming analytics window, such as the score window.
- Unified project and server management via SAS Event Stream Manager provides:
 - An easy-to-use interface for constructing and managing repeatable deployment plans for projects executing on SAS Event Stream Processing servers.
 - Ability to monitor memory and disk usage across the SAS Event Stream Processing server environment.
 - Easy addition of new SAS Event Stream Processing servers for improved monitoring.
 - Autodiscovery of new SAS Event Stream Processing servers using SAS Event Stream Processing agent technology for automated monitoring.

Ability to consume and connect streaming data

- Extensive suite of data adapters and connectors for publishing and subscribing to live data streams of both structured and unstructured data.
- Predefined adapters include read and write (i.e., publish and subscribe): Hadoop (HDFS Yarn), Apache Camel, OS/soft PI, Axeda, RabbitMQ, Solace, Tervela Data Fabric, XML/JSON File Socket Adapter, SAS® LASR™ Analytic Server, IBM DB2, IBM Netezza, IBM WebSphere MQ, SAP Sybase ASE, Tibco Rendezvous, JMS, File/Socket, Database ODBC and SAS data sets.
- Customized publish/subscribe APIs can also be written in C or Java.
- Publish only from the following: Twitter, Log Sniffers (Oracle, Greenplum), Network Sniffer, Google Protobuf Sniffer, SYSLOG, HTTP RESTful interfaces; subscribe only to SOAP and SMTP.

Adaptable, in-stream analytics and data manipulation

- Machine-learning streaming algorithm support lets you create scoring and learning procedural windows for various learning algorithms to include density-based clustering (DBSCAN) and K-means clustering.
- DBSCAN lets you identify clusters in an unsupervised manner. A combination of train and score windows are used to periodically update the DBSCAN model.
- A new geofence window type allows you to track the location of objects relative to borders of a geofence. You can also alert when an object approaches, enters or leaves the defined geofence boundaries.
- Model training on historical data (for accurate model development) complements high-performance analytics for at-rest data.
- Integration with SAS model management capabilities provides faster integration of analytical models with SAS Event Stream Processing projects.

Optimized processing that scales

With event stream processing, systems need to be online, all the time. And processing must happen very fast, which demands optimal performance. SAS has patented, instantaneous 1+N way failover, guaranteed delivery without persistence and dynamic updates to ensure consistent, relevant processing. The low latency, high-volume SAS Event Stream Processing software outperforms other stream processing engines with commodity hardware. And the distributed, in-memory grid processing scales linearly as your data grows, optimizing hardware investments.

Key Features (continued)

- Flexible, modular, window-driven architecture to define complex continuous queries:
 - Based on an extensive suite of interchangeable window types and operators, to detect an unlimited number of patterns, correlations, computations and aggregations.
 - Prebuilt, common data quality routines are used to cleanse, standardize and filter live-stream data before it's stored, reducing downstream processing.
 - Patterns of interest can include nearly unlimited advanced analytics calculations with in-stream, machine-learning K-means clusters and live-stream analytical scoring. Analytical models include all SAS DATA step, DS2 and third-party code snippets.
- Highly embeddable to gateways, devices and any existing C++ application (with dedicated thread-pool processing).
- Streaming algorithms for DBSCAN, regression, decision tree and naive Bayesian.

Expanded deployments and open data support

- Python publish/subscribe API:
 - Publish events and subscribe to SAS Event Stream Processing using Python.
- Use languages in SAS Event Stream Processing (in-process event stream handlers):
 - DS2 (current), Python and C/C++.
- BOSH Cloud Foundry: for large-scale distributed services.
- Data stream support:
 - Hortonworks DataFlow (HDF) NiFi integration: SAS and HDF can provide immediate, streaming and deep intelligence.
 - MapR streams support.
 - Connectors and adapters for IoT (Twitter, Kafka, Flume, MQTT, Pico/SpryWare, gateways, devices, Cassandra (adapter only)), Boardreader.
- SAS Event Stream Processing for SAS® Viya™ and SAS Cloud Analytic Services (CAS):
 - Deploy SAS Viya data management/machine-learning models to SAS Event Stream Processing for streaming analytics.
 - Stream, analyze, filter, identify patterns and score data while it moves in/out of CAS.

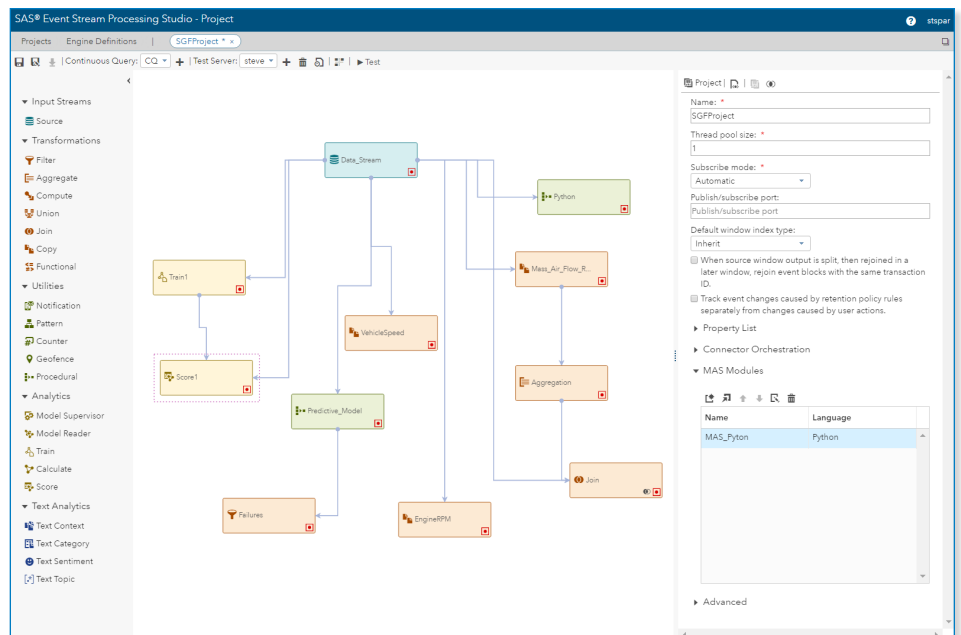


Figure 3: Design, test and share streaming models with SAS Event Stream Processing.

Enterprise administration and management

Centralized administration and management of event stream processing activities are key to enforcing IT command and control of project handling, processing optimization and maintenance. The solution provides transparent governance that aligns to business requirements and can integrate stream processors and algorithms from SAS and other sources. An optional caching store is also available if data sizes exceed the available memory.

Extend other SAS® solutions with streaming analytics

SAS Event Stream Processing enhances other SAS solutions such as SAS Customer Intelligence, SAS Asset Performance Analytics, SAS Security Intelligence and more. Enabling these solutions with event streaming means anomalies can be flagged and new patterns of interest can be uncovered and monitored while data is in motion. Combined with other SAS solutions, SAS Event Stream Processing can provide deeper insights – at faster rates.

Key Features (continued)

Visual monitoring of event streams

- Configure the dashboard for customized testing of streaming activity.
- Interactively filter and query live-stream activity to examine specific behavior of elements.
- Define and customize notifications by SMS, email and other alerts as part of event stream model workflow.
- Compare historical activity with current processing using graphical representations.
- Monitor stream processing detail by subscribing to events of interest.

In-memory, distributed and optimized processing that scales

- High-volume processing of millions of events per second. Low-latency response times (millisecond, submillisecond).
- Retained and aggregated data kept in memory.
- Ability to take advantage of distributed grid architectures.
- Processing speeds can be customized with flexible thread-pool sizing, caching stores and more.
- Includes patented, instantaneous 1+N way failover, native failover, guaranteed delivery without the use of persistence and other fault-tolerance functions to ensure successful event stream processing activity.
- Full and open access to all event metadata.

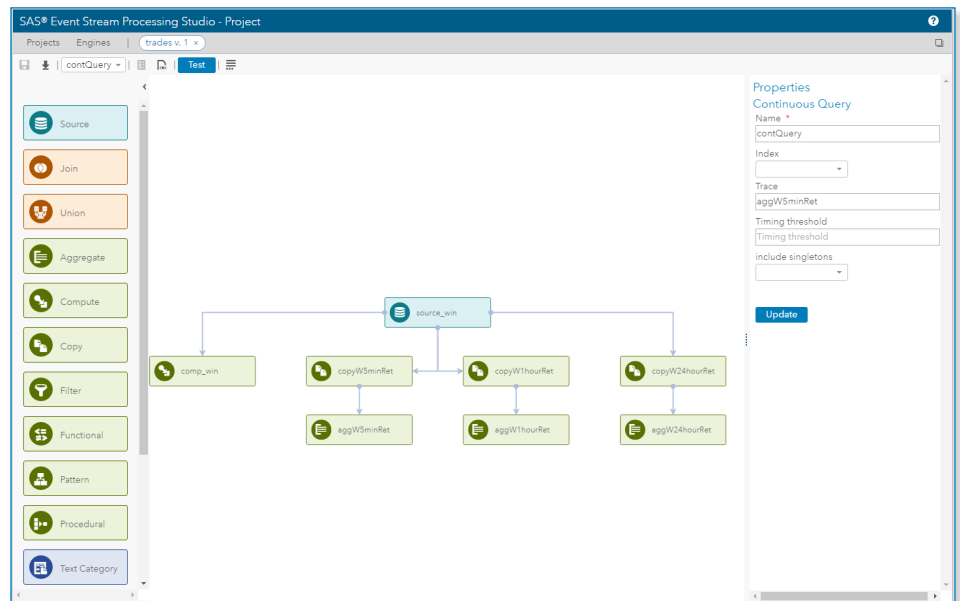


Figure 4: Use different window types to put together a process flow for streaming data.

TO LEARN MORE »

To learn more about SAS Event Stream Processing system requirements and see other related material, please visit: sas.com/ESP.

Key Features (continued)

Enterprise administration and management

- RESTful interface to XML factory server makes it easy to introduce multiple projects, validate XML syntax, control analysis windows and define project retention policies.
- Centralized multiproject management eases administrative burdens, setting controls to dynamically load projects, start/stop/remove and persist projects.
- Can include authentication for TCP/IP client for both pub/sub and HTTP connections.
- Configuration files ease initial setup protocol specifications.
- Customizable retention strategy based on size, volume or time options associated with policy definitions.
- Enhanced error handling to increase the resilience of handling run-time errors.
- Localized log messages.
- Ability to investigate processing activity with queries into live event streams of interest.
- Authentication and encryption for events streaming in and out to other systems.

To contact your local SAS office, please visit: sas.com/offices

