



www.sas.com

Wittington House, Henley Road, Medmenham,
Marlow, Buckinghamshire SL7 2EB, UK
Tel: 01628 486 933

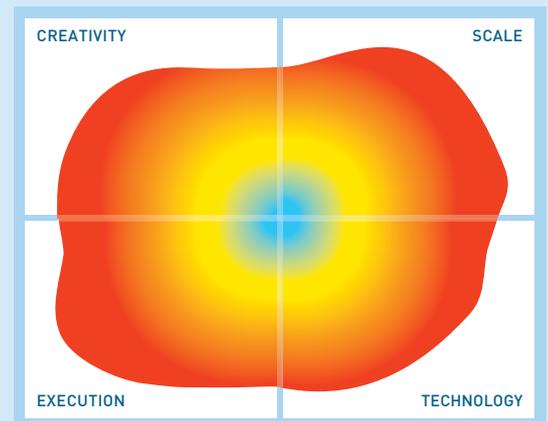
SAS Event Stream Processing

The company

SAS is headquartered in North Carolina and was founded in 1976. The company started life by offering a 4GL-based statistical package for financial and economic analysis but has expanded to offer a multi-faceted product set, providing organisations with an information delivery system that spans all aspects of analytics and associated technologies. This includes business analysis, analytics and data science as well as practical applications thereof (customer intelligence and so forth). In addition, complementary capabilities such as data quality and data integration are also offered.

Since its foundation SAS has expanded primarily in an organic manner and is now a multi-billion-dollar company that has offices and partners worldwide. It remains privately owned. A notable point is that SAS consistently invests a significant proportion of its revenues into R&D, typically at around 25%. This emphasis on R&D, plus the company's refusal to listen to the blandishments of the stock market (where short-term trends are more important than long-term profits), has meant that SAS is, by some margin, the largest pure-play vendor in the analytics market.

“Subsea surveillance of oil platforms to help avoid unplanned production interruptions, resulted in the prediction of a submersible pump failure that saved the company \$3 million.”
Major energy company



The image in this Mutable Quadrant is derived from 13 high level metrics, the more the image covers a section the better. Execution metrics relate to the company, Technology to the product, Creativity to both technical and business innovation and Scale covers the potential business and market impact.

What is it?

SAS has been offering stream processing for a number of years, especially for fraud applications. However, a major driver is now the Internet of Things (IoT) and applications such as predictive maintenance and smart “anything”. The basic principle is to capture and apply analytics to data in motion.

In terms of the products themselves, SAS Event Stream Processing (ESP) comes with a variety of options and components, as illustrated in **Figure 1**. The plus signs indicate additional license options. CAS stands for Cloud Analytic Services. The ESP for Edge option will run on something as small as a Raspberry Pi and supports analytics on those small but powerful platforms. Both event-based and window-based (time sliced) options are available.

What does it do?

Perhaps the biggest differentiator for SAS in this space is its analytic capabilities, though it's also notable for its performance. The options available are illustrated in **Figure 2**, many of which have only been introduced in the latest release (5.1). A key feature is the continuous improvement of in-stream models, using machine learning.

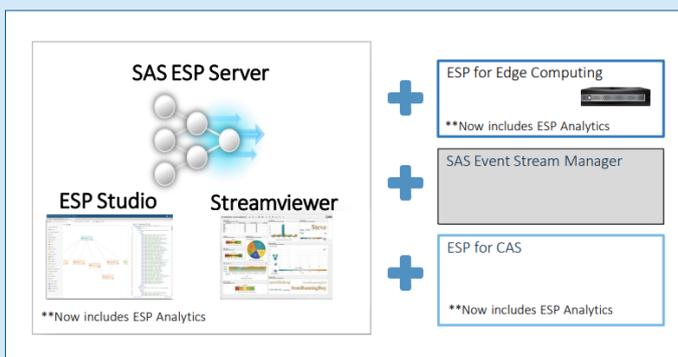


Figure 1

Analytics & modelling	★★★★★	Development	★★★★★
Architecture	★★★★★	Integration	★★★★★
Connectivity	★★★★★	Self-service	★★★★★
Deployment	★★★★★	Streaming functionality	★★★★★

“ Our transportation assets are equipped with edge devices that manage hundreds of data elements per second to optimize their operation. We use SAS Event Stream Processing in our edge devices to enable our customers with analytical insights on their data in real time to minimize operating costs, identify predictive maintenance issues before they become a problem. ”

Major rail transportation manufacturer

As far as **Figure 2** is concerned, DS2 is a SAS specific language and ASTORE, also proprietary, is used to support algorithms that are trained offline but which are scored in-stream. Model Manager, which is a separate product, supports PMML (predictive modelling mark-up language) and will convert supported model types (more than a dozen currently) into SAS code for deployment on the ESP Server. There is also support for RESTful APIs to run other models. Python notebooks are supported to drive the ESP engine, publish events to ESP, and display results, and models written in Python and C are also supported, though not Java-based models. R will be supported later during 2018. Additional facilities in the ESP Server that are

any downtime on the server, and similarly, add new ESP servers as required. It will automatically discover new servers using ESP agent technology for automated monitoring.

SAS provides connectivity to over 300 end points and supports a variety of standard protocols including MQTT, BACNet publisher connector and adaptor (for smart homes), OPC-UA connector and adaptor (for machine-to-machine communications), a UVC connector (Video4Linux), a WebSocket publisher connector, and a URL publisher connector (for RSS news feeds, JSON from a weather service or News from an HTML page). There are also facilities provided so that you can write your own connectors. In this context it is worth mentioning the new SAS ESP Community (communities.sas.com/IOT), which is moderated by SAS but user-driven.

Why should you care?

Machine learning and the Internet of Things are top of mind for many chief data officers and SAS has significant expertise in both of these areas. It is the world's largest business intelligence and analytics company. It should therefore not come as any great surprise that the company is exceptionally strong when it comes to the breadth of analytic capabilities it provides, regardless of whether this is for data at rest or data in motion. Put this together with the fact that Mckinsey has predicted that the Internet of Things will generate \$11 trillion of economic impact by 2025 and you have significant opportunities for SAS to exploit its strength in analytics, across a range of IoT-based deployments. The company is also actively researching analytics for Blockchain, another environment which is set to have a major impact.

In-Stream Analytic Model Deployment
 SAS® DS2, Python, C
 SAS® ASTORE Scoring support
 SAS® Model Manager Integration

Streaming Algorithms & Machine Learning

Streaming Summary - Univariate Statistics	Streaming K-Means	Support Vector Machines*
Streaming Pearson's Correlation	Streaming DBSCAN	Streaming Support Vector Machines
Streaming Segmented Correlation	Random Forest*	Streaming Linear Regression
Weibull Distribution Fitting	Gradient Boosting Tree*	Streaming Logistic Regression
Short Time Fourier Transform	Factorization Machine*	Streaming Fit Statistics
Streaming Text Tokenization	Support Vector Data Description*	Streaming Receiver Operating Characteristic (ROC)
Streaming Text Vectorization	Deep Neural Network*	Streaming Histogram
Moving Relative Range	Convolutional Neural Network*	Image Processing
	Bayesian Network*	Robust Principle Components Analysis*

*Out-of-Stream Training & In-Stream Scoring

Figure 2

worth mentioning include in-stream geofencing, text analytics, in-stream time pattern recognition (including time-series similarity analysis and time-series clustering), and the ability to build data quality rules into the streaming process.

Of the other elements within the ESP environment, ESP Studio provides an environment for constructing visual models, designed for use by non-technical personnel, while Streamviewer provides a visual analytic dashboard environment that lets you combine real-time and historic data. Event Stream Manager is used to update or deploy analytic algorithms without requiring

The Bottom Line

As a company, SAS is long-established and has a reputation for enterprise-level capability. At the same time it has been a leading light in the analytics market for 40 years. In our view, that makes SAS a major contender, if not a leader, in the market for streaming analytics platforms.

[FOR FURTHER INFORMATION AND RESEARCH CLICK HERE](#)