

EVENT STREAM PROCESSING EXPLAINED

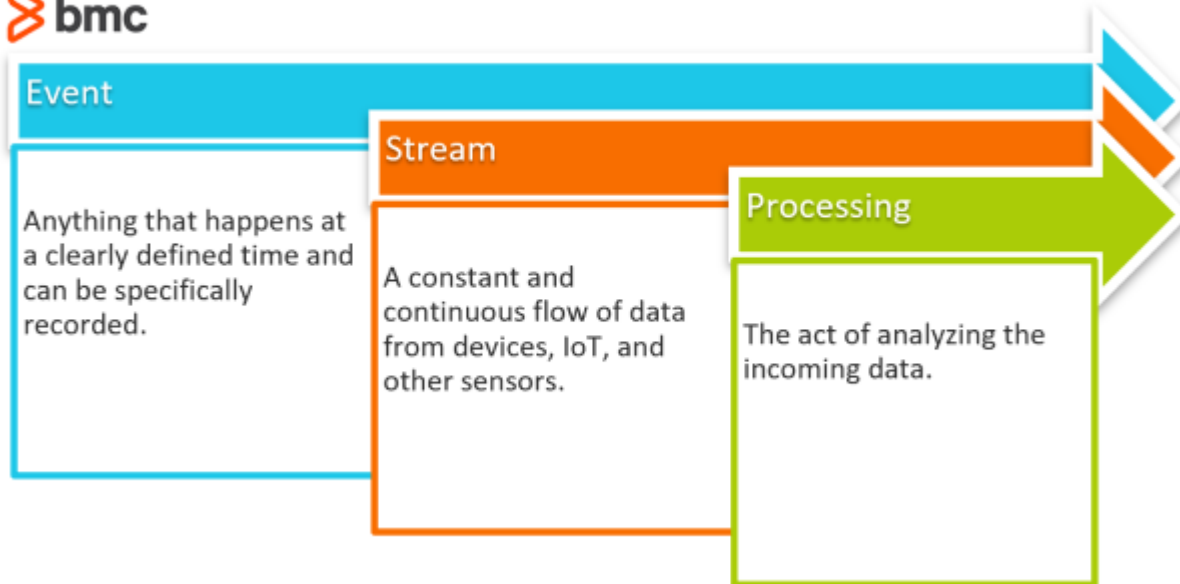


Event stream processing is a reactive programming technique that filters, analyzes, and processes data from a [data streaming source](#) as the data comes through the pipe. It is used for a number of different scenarios in real-time applications.

As we rely more and more on data generated from our [phones, tablets, thermostats, and even cars](#), the need for it to be analyzed *while still streaming* only increases. The aim is two-fold:

1. To process data in real-time.
2. To act on those data signals in as close to real-time as possible.

One way the Internet of Things (IoT) data can be evaluated while it is in motion is with event stream processing. The IT community is adopting this technique at greater rates, and it is because of all the advantages and applications event stream processing offers to the community.



What is event stream processing?

Sometimes shortened to ESP, event stream processing is composed of three simple terms: event + stream + processing.

Event

An **event** is anything that happens at a clearly defined time and can be specifically recorded. The possible types of events can be vast! Events can be created from an event source. These could be systems, business processes, sensors, data streams.

On an application, there could be a login event. On a thermometer, there could be a "too hot" event. On a fire alarm, there is a smoke-detected event. On databases, there can be data uploaded events.

Stream

A **stream** is a constant and continuous flow of objects that navigate into and around companies from thousands of connected devices, IoT, and any other sensors. An event stream is a sequence of events ordered by time. Enterprises typically have three different kinds of event streams:

- **Business transactions** like customer orders, bank deposits, and invoices
- **Information reports** like social media updates, market data, and weather reports
- **IoT data** like GPS-based location information, signals from SCADA systems, and temperature from sensors

Processing

Processing is the final act of analyzing the incoming data. How the data is processed is dependent on the function of the processor. It could trigger alerts that a sensor needs attention. It could define the next path the data is supposed to travel down along the pipe. It could classify the data and tag it with metadata.

Goal of event stream processing

Event stream processing, then, is a form of reacting to a stream of event data and processing the data in a near, real-time manner.

The ultimate goal of ESP deals with identifying meaningful patterns or relationships within all of these streams in order to detect things like event correlation, causality, or timing.

ESP is a useful real-time processing technique for events that either/both:

- Occur frequently and close together in time
- Require immediate attention

Strong examples of when ESP would be necessary are in the areas of e-commerce, fraud detection, [cybersecurity](#), financial trading, and any other type of interaction where the response should be immediate.

How event stream processing works

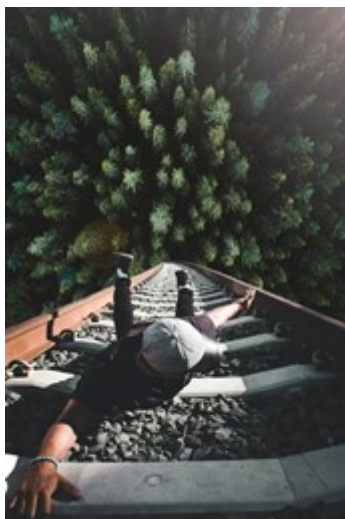
Steps to event processing:

1. Event Source
2. Event Processing
3. Event Consumer

Event stream processing requires the existence of a streaming data source. There is nothing to process if there is no data source.

Once the source is created, the source must emit events to the processor. Likewise, the processor needs a way to listen and receive the output of the event source. This can occur through an [API](#).

And, finally, like the sound of a man screaming help in a forest, what is the point of processing data if there is no audience to hear it? An event processor needs a consumer to [output](#) its processes to. It can be a data dashboard, another database, a [user analytics report](#).



The landscape of available tooling to process data streams is large and growing, and the best solutions, as usual, depend on the use case, your tech stack, your budget, and your team's skill levels. Possible cloud-based solutions are:

- Amazon MSK

- Amazon Kinesis
- Apache Kafka
- Azure Stream Analytics
- Google Pub/Sub

Benefits of event stream processing

Event stream processing processes events at the moment they occur, completely changing the order of the traditional analytics procedure. It allows for a faster reaction time and creates the opportunity for proactive measures to be taken before a situation is over.

Processing data in this way is advantageous because it:

- Creates opportunity for a real-time response time. For example, a potential new business service.
- Requires less memory of the service to process individual data points rather than whole datasets.
- Extends a company's ability to process data to a greater number of potential data sources, such as IoT devices, edge devices, and distributed event sources

Event stream processing is a smart solution to many different challenges and it gives you the ability to:

- Analyze high-velocity big data while it is still in motion, allowing you to filter, categorize, aggregate, and cleanse before it is even stored
- Process massive amounts of streaming events
- Respond in real-time to changing market conditions
- Continuously monitor data and interactions
- Scale according to data volumes
- Remain agile and handle issues as they arise
- Detect interesting relationships and patterns

When to use event stream processing

Event stream processing can be found in virtually every industry where stream data is generated, whether it be from people, sensors, or machines. As IoT continues to expand its technologies, it will continue to see dramatic increases in the real-world applications of stream processing.

Systems are growing more varied and dynamic. Applications consist of an ever-growing number of micro-decisions, processing growing numbers of events sources and events.

AI has entered the picture, too, and *it* is becoming the consumer for event processes—instead of solely people—and AI is taking on the burden of responsibility to address real-time events, reaching far and beyond the limits of the human attention span, care to detail, and ability to react to limited number of problems.

(Learn about [AI/human augmentation](#).)

Instances where event stream processing can solve business problems include:

- Ecommerce

- [Fraud detection](#)
- Network monitoring
- Financial trading markets
- [Risk management](#)
- Intelligence and surveillance
- Marketing
- Healthcare
- Pricing and analytics
- Logistics
- Retail optimization

Though far from exhaustive, this list begins to show the huge variety of uses for event stream processing as well as how wide-reaching it can go.

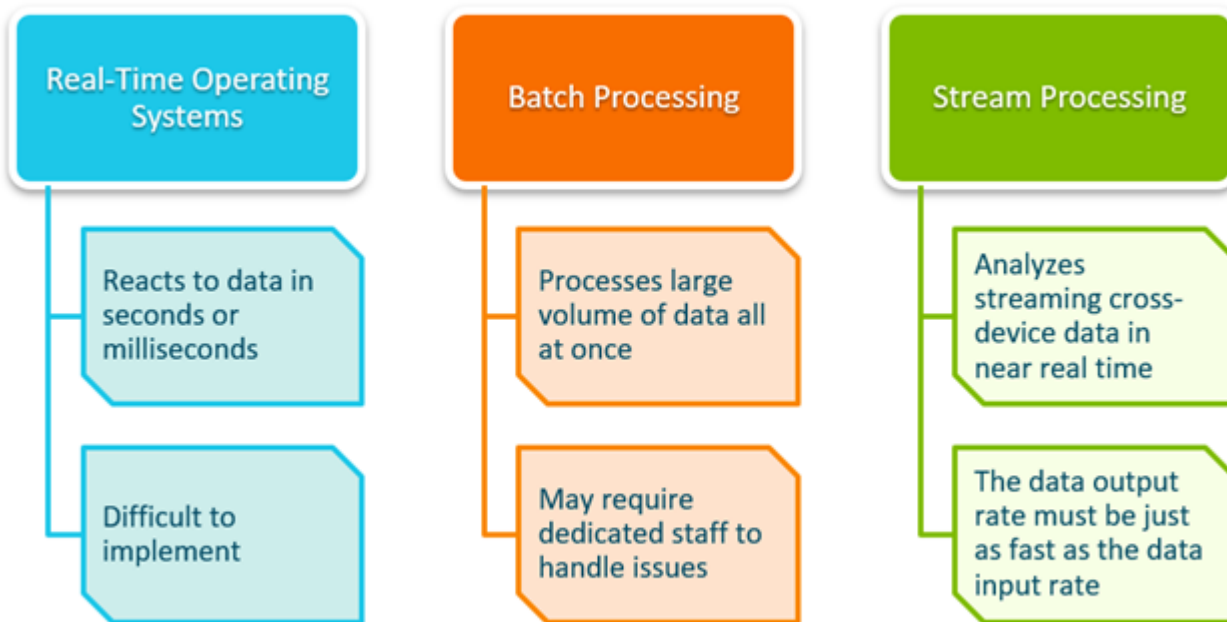
Event stream processing here to stay

Overall, event stream processing is here to stay and will only prove itself more crucial as the demand by consumers, desires the data to be computed and interpreted in real-time. ESP is a powerful tool which allows people to get closer to their customers, their companies, and their events, in order to take their analytics to the next level. As devices become more and more connected, having the ability to continuously stream and analyze big data will become vital.

Will your organization be ready?

Related reading

- [BMC Machine Learning & Big Data Blog](#)
- [BMC AIOps Blog](#)
- [Real Time vs Batch Processing vs Stream Processing](#)
- [Creating & Using Snowflake Streams](#)
- [What Is Pub/Sub? Publish/Subscribe Messaging Explained](#)
- [Machine Learning for Data Management](#)



Benefits & drawbacks of common data processing types