

January 2015

# WHITE PAPER

## CLOSING THE REAL-TIME INTELLIGENCE GAP

### BACKGROUND AND OBJECTIVES

The volume, velocity and variety of data ingested by today's data-intensive analytic platforms presents a significant challenge for decision makers. It is no longer feasible to rely solely on historical and batch analysis to make well-informed decisions that are timely and actionable. Instead, organizations must begin pushing analytic and alerting capabilities as far upstream toward the data collection point as possible in order to elicit truly real-time insights from those data streams.

There are a significant number of tools on the market that address the problem of downstream analytics in a voluminous environment; however, the ability for these tools to perform real-time analysis and alerting is limited by the performance of today's solutions used to Extract, Transform, and Load (ETL) data into downstream systems due to the latency they add between data collection and data analysis. Moreover, the volume and variety of data being ingested make it impossible for analysts and decision makers to find all of the data they need across the various analysis platforms upon which it lands.

This paper will illustrate how the use of Napatech accelerators can overcome this gap and, when coupled with open source messaging solutions, provide near real-time alerting and extraction of data for ingest into decision-making tools.

### VALUE TO CLIENT

The use of Napatech accelerators can immediately improve real-time analysis capabilities by pushing intelligence to the point of data ingress and providing:

#### REAL-TIME ALERTING

The ability to know what data is entering your system in real time, before it reaches decision making tools, provides intelligent alerts to stakeholders informing them of the presence of new data that is of interest for their Area of Responsibility (AOR).

#### IN-LINE ANALYTICS

Making use of perishable insights, that is data whose value declines rapidly over time, requires that organizations begin to analyze data at the very moment it is received. Doing so ensures that an organization can begin acting on what is happening immediately.

#### INTELLIGENT DATA FLOW

By inspecting data immediately upon ingress, data flow decisions can be made to direct data to downstream consumers at line rate. This minimizes the unnecessary flow of data through downstream brokers and processing engines.

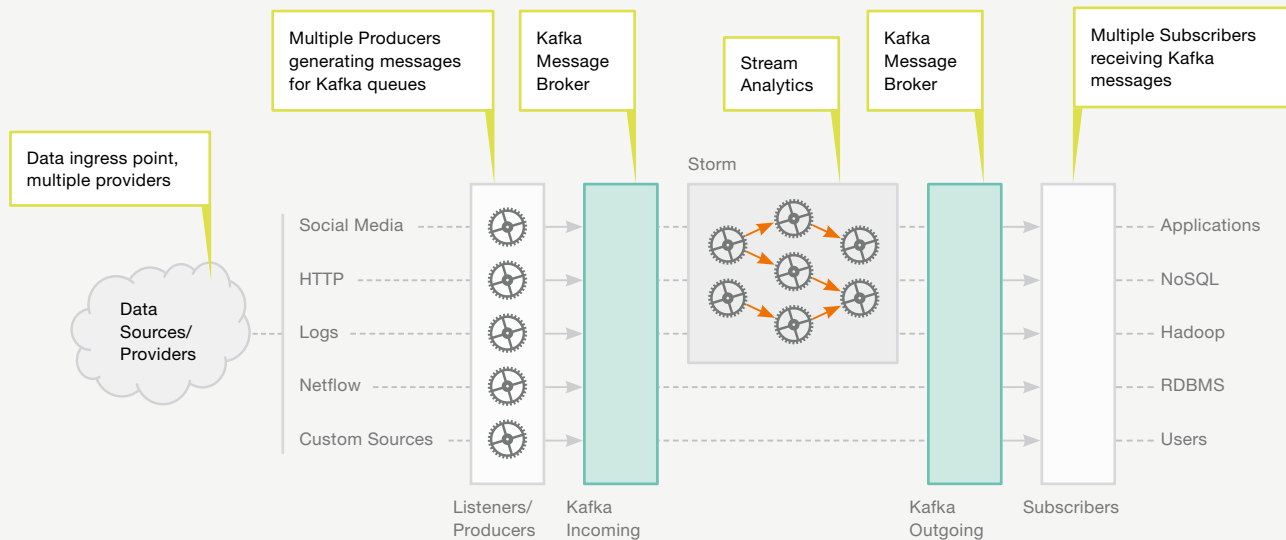
#### TECHNICAL APPROACH

Conventional approaches to streaming analytics involve downstream software tools that inspect and analyze data before forwarding it further downstream to systems and applications for consumption by end users. A typical streaming analytic architecture is pictured in Figure 1 below.

Typically these architectures include multiple stream producers that categorize data flows and submit them to queues on message brokers for delivery to downstream systems and stream processing engines. These message brokers and data processing engines handle the processing and flow of data and can be deployed in a robust and scalable manner; however, two compounding latency factors limit their effectiveness when it comes to enabling true real-time analysis:

**1. Upstream latency** – This is the delay introduced by systems that the data must pass through before it reaches the broker or processing engine. This includes but is not limited to the data provider/source, listeners or aggregation points, and the network or transport layer itself.

**2. Software latency** – This is the delay introduced by the interaction of software with the hardware components of the system and includes but is not limited to interrupts, DMA and context switches.



**FIGURE 1**  
Conventional Streams Processing

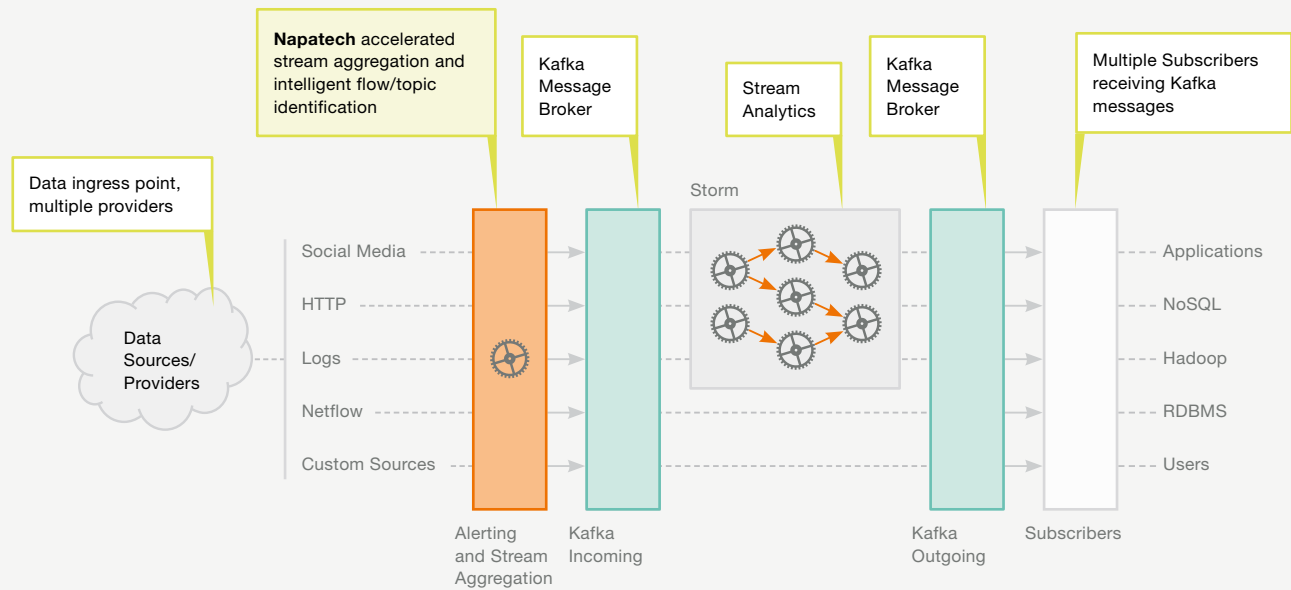
Mitigating these issues requires placing the beginning of your analytic pipeline as close to the point of ingress as possible and leveraging hardware acceleration for initial data analysis and triage. Napatech acceleration technology can be leveraged to solve both of these problems and accelerate the streaming analytics pipeline.

Napatech accelerators use hardware, Field Programmable Gate Arrays (FPGAs), to perform in-line event processing and line-rate packet analysis at 1/10/40/100 Gbps speeds. By placing these devices near the edge of your network, directly behind your security perimeter, it is possible to aggregate, triage and process dozens of incoming data streams simultaneously.

In accelerated streams processing architectures these devices would be configured to work in concert with a messaging system like Apache Kafka and a streams processing system like Apache Storm, to identify and direct traffic flows while also issuing alerts about incoming data to users and applications downstream. As multiple streams of data enter the device they are merged into a single stream that can be interrogated for initial insights and classified into flows or topics for Kafka. All of this is accomplished within the hardware of the accelerator at line rate. Once initial triage and topical flow identification is complete, the data is handed off to Kafka for publishing to subscribers downstream including Storm for further event processing. A notional architecture for this solution is depicted in Figure 2 below.

Through the use of hardware acceleration technology we are able to aggregate Kafka stream production and begin analyzing data immediately upon entry into our environment. This not only simplifies the event processing architecture, but also accelerates time to insight. Furthermore, the in-hardware implementation of initial triage and flow identification algorithms eliminates latency that would be introduced by a software instantiation of the same. Lastly, future generations of the Napatech architecture are expected to include support for external, DMA addressable, co-processors such as GPUs, Xeon Phi and FPGAs. The ability to implement Kafka or Storm engines within a co-processor will allow even more analytics to be pushed to the edge and empower a new breed of low latency, custom developed in-line analytic capabilities.

Today, there is an almost unlimited amount of data available to analysts upon which to gain insights and inform the decision making process. Data originating from the Internet of Things, sensors, markets and the Web, among other sources, represents a treasure trove of useful information. However, the pace at which this data is generated often causes its value to decline rapidly over time. Real-time systems attempt to mitigate this by performing analytics on an organization's most recent data, but the performance of these systems is limited by the latency introduced to the data flow by upstream systems. Enabling truly real-time analytics necessitates pushing the beginning of your analytic pipeline as close to edge of your network as possible.



**FIGURE 2**  
Accelerated Streams Processing

The use of intelligent network accelerators, coupled with open source message brokers and streams processing engines, will empower organizations to begin acting upon information and events as they occur.

### DISCOVER THE POWER OF NAPATECH

Napatech accelerators are designed to handle the maximum theoretical throughput of data for a given port speed. Napatech offers a range of accelerators supporting speeds from 10 Mbps to 100 Gbps. A single, common Application Programming Interface (API) allows application software to be developed once and used with a broad range of Napatech accelerators. This allows combinations of different accelerators with different port speeds to be installed in the same server. Additional features include:

- Napatech accelerators can identify, filter and distribute flows to up to 32 CPU cores
- Data merging functionality allows flows from different ports on different accelerators to be merged for analysis
- Data sharing functionality allows multiple applications to access the same data at the same time
- All of this can be performed with very low server CPU load

### CHOOSE THE MARKET LEADER

Napatech is the market leading provider of accelerators for network management and security applications. Napatech provides global sales and support from local offices in all major continents, which is included in the price of the

accelerator. This means that our highly experienced support resources are available for design and integration support, as well as field support without extra charge.

Napatech accelerators are manufactured to the highest standards by outsourced manufacturers in Switzerland and the USA supporting all major certifications including NEBS for telecom applications.

### COMPANY PROFILE

Napatech is the world leader in accelerating network management and security applications. As data volume and complexity grow, the performance of these applications needs to stay ahead of the speed of networks in order to do their jobs. We make this possible, for even the most demanding financial, telecom, corporate and government networks.

Now and in the future, we enable our customers' applications to run faster than the networks they need to manage and protect.

### Napatech. FASTER THAN THE FUTURE





APRIL 2010

**EUROPE, MIDDLE EAST  
AND AFRICA**

Napatech A/S  
Copenhagen, Denmark

Tel. +45 4596 1500  
ntemeasales@napatech.com  
www.napatech.com

**NORTH AMERICA**

Napatech Inc.  
Boston, Massachusetts  
Mountain View, California  
Washington D.C.

Tel. +1 888 318 8288  
ntamericassales@napatech.com  
www.napatech.com

**SOUTH AMERICA**

Napatech Brasil LTDA  
São Paulo, Brazil

Tel. +55 11 2127 0782  
ntsouthamericasales@napatech.com  
www.napatech.com

**APAC**

Napatech Japan K.K.  
Tokyo, Japan  
Tel. +81 3 5326 3374

Napatech Korea  
Seoul, South Korea  
Tel. +82 2 6001 3545

ntapacsales@napatech.com  
www.napatech.com