



## THE TREND

# Transforming the customer experience with near real-time insights

As a recognized industry innovator, HPE is paving the path for faster, smarter data center infrastructure solutions by adding near real-time insights to the InfoSight predictive analytics platform.

This means that Jeff Dutton, Data Platform Architect for HPE InfoSight, is busier than ever these days.

“I’m swamped with phone calls,” Dutton says excitedly. “Everyone wants to know more about getting value from their data in real-time.”

This surge in demand for real-time analytics across HPE products is coming not only from InfoSight customers; it’s coming from other business units inside HPE. Customers, both external and internal, are hungry to learn how to leverage streaming edge-to-core-to cloud data pipelines for their uses cases.

Streaming data pipelines allow almost any organization to deliver higher value, faster, to customers. In this case, InfoSight is transforming the customer experience with AI for data center to predict and prevent problems.

Delivering on the real-time element of the customer experience, however, necessitated a major overhaul to the InfoSight architecture, which was one of the most critical lessons learned by Dutton and his team.

“Moving to a near real-time user experience requires an infrastructure that self heals, scales massively and is always available to process streaming workloads, no matter what,” says Dutton. “This is fast becoming new table stakes for every business.”

# A worthy challenge: delivering customer value faster

HPE's vision for helping customers is to slash the time it takes to turn ideas into value, which is transforming industries, markets and lives. For InfoSight, delivering on this vision means delivering customer value in seconds and minutes, rather than hours or days.

InfoSight was already recognized as a category leader based on its ability to monitor infrastructure, predict possible problems and recommend ways to enhance performance.

It works by receiving infrastructure information from “call home” sensors and running analytics against the massive amount of usage data it has accumulated over the years as well as incoming sensor data. Located in the storage devices themselves, these sensors also collect network, compute, and hypervisor data.

HPE has over 20 billion sensors deployed in data centers all around the globe sending trillions of metrics each day to InfoSight, providing analytics on petabytes of telemetry data.

“We’re watching what’s going on in the environment and putting the data through expert systems using various AI and machine learning techniques,” says Dutton. “The goal is to tell the customer what to do to prevent problems and optimize their experience.”

The next challenge: delivering insights on all of this, faster.



**Billions**  
of Global Sensors

**Trillions**  
of Metrics

**Petabytes**  
of Data



**“The streaming part of the journey  
has been the most challenging.”**

**- Jeff Dutton**

## THE IT CHALLENGE

# Introducing stream processing: closing the gap between data and insight

To deliver value faster, InfoSight needed to evolve beyond its classic, batch mode, big data architecture.

The speed by which Dutton and the team could deliver insights to customers was gated by the time it took to run batch processes against very large, extremely complex data sets.

As an initial experiment, the team migrated from an enterprise database to Apache Spark to see if faster batch was the answer. They quickly understood they would need to transition to work with the data in real time to realize their business vision. This meant introducing stream processing.

“Batch analytics is something nearly every developer and architect understands fundamentally, regardless of their tech stack,” notes Dutton. “The streaming part of the journey has been the most challenging.”

Stream processing solutions process data in motion, computing on data directly as it is produced or received. This turns around the batch-mode paradigm of storing data in a database, a file system, or other form of mass storage where applications query or compute over the data at rest. With stream processing the application logic, analytics, and queries exist continuously, and data flows through them continuously.

“Batch continues to play an important role in the InfoSight architecture. It’s leveraging several petabytes of historical data,” says Dutton. “The source for near real-time insight, however, becomes the stream.”

# The next leap forward: big data at speed

HPE is collaborating with Lightbend to design, implement, and operate a reference architecture based on InfoSight, which is fast becoming a blueprint for other HPE product lines.

Delivering insights on streaming data sets necessitated a new approach for Dutton and the team. Unlike batch, data services in a dynamic, streaming application must run forever and be able to scale up, as well as down, on demand. Here, up and down are measures of size. The actual scaling is horizontal, across multiple nodes, rather than vertically, within a single machine.

In addition to scaling elastically, the streaming application must be able to recover from failures quickly. A logical data service may failover to backup instances. Failed service instances may be restarted automatically. This means that failure must be a first class, expected concept in the software. Failures are normal, anticipated, handled, and, ideally, self heal within the application.

Dutton turned to the Lightbend Platform for the elasticity and resilient self healing required to deliver big data at speed.

Lightbend Platform includes microservices frameworks for processing continuous application logic; multiple streaming engines for handling tradeoffs between data

latency, volume, transformation, and integration; machine learning and deep learning tools for applying algorithms from the data science team; and intelligent management and monitoring tools for reducing the risk of running this always-on system in production.

“I have been working with Lightbend since 2014 and knew they could deliver on the elasticity and resilience we required,” shares Dutton. “When they packaged their Reactive stack with Apache Kafka, Apache Spark, and a multitude of other technologies into a de facto Fast Data distribution, it made good business sense to use their platform for the next generation of InfoSight.”

Back to Dutton’s phone ringing off the hook. Lightbend Platform and Dutton’s approach are fast becoming a reference for other HPE business units looking to deliver similar real-time customer value. It is making such sound business sense that HPE Storage and HPE Pointnext are teaming to bring the streaming solution to HPE’s top customers.

## A BRIEF LOOK INSIDE HPE INFOSIGHT, POWERED BY LIGHTBEND PLATFORM

Lightbend Platform provides HPE InfoSight with a complete streaming, Fast Data architecture, including support for microservices, multiple streaming engines, and an end-to-end console for cluster management, monitoring, and security. To ensure system responsiveness and developer happiness, HPE InfoSight benefits from on-site and direct portal access to Lightbend experts and 24/7 production support.

### Microservices Architecture

Lightbend Platform  
(Akka, Lagom, Play, Scala)

### Streaming Engines

Akka Streams  
Apache Spark

### Data Backplane

Apache Kafka

### Machine Learning

Spark MLlib

### Infrastructure Orchestration

Mesosphere DC/OS

### Monitoring, Management, and Security

Lightbend Console  
Fortify SCA for Scala

### Lightbend Subscription Services and Support

Embedded Lightbend Engineer (ELE)  
Developer Assist  
Production Support 24/7  
Legal Indemnification & Warranties