



**Hewlett Packard
Enterprise**



**OPEN
PLATINUM™**

Cloudline Autonomous Driving Solutions

Accelerating insights through a new generation of data
and analytics using open design principles

March, 2019



The 4G / Edge Autonomous Vehicle Computing Stack

Current Generation: 4G and Cloud Data Centers

/Autonomous
/Sensing
/Communication
/Battery
/Navigation
/Mirrorless
/Ecology

A car can travel **4 feet**
in the amount of time it takes
for data to travel round trip at
80ms latency

500 Miles / ~80ms Latency

Autonomous Vehicles

A future with IoT at its center requires fast computing solutions that current infrastructure does not support. An autonomous car is said to require approximately **4,000 gigabytes per day**. [Mashable](#)



The 4G Network

This is the current generation of cellular networks that exist in the US today



Cloud Data Centers

Not suitable for the low latency real-time predictive analytics autonomous vehicles will require. These could be used to store the large amounts of un-structured data for future data-mining and analytics. **A car would have traveled 4 feet with 80ms of data latency.**



The 5G / Edge Autonomous Vehicle Computing Stack

Next Generation: 5G and Micro Data Centers

/Autonomous
/Sensing
/Communication
/Battery
/Navigation
/Mirrorless
/Ecology

100m

48
mph

A car can travel **4 inches**
in the amount of time it takes
for data to travel round trip at
5ms latency

5 Miles / ~5ms Latency

Autonomous Vehicles

A future with IoT at its center requires fast computing solutions that current infrastructure does not support. An autonomous car is said to require approximately **4,000 gigabytes per day**. [Mashable](#)




The 5G Network

The speed and latency requirements of autonomous vehicles will require 5G technology. The carriers are spending \$100's billions in order to upgrade their infrastructure to handle the 5G demand



Micro Data Centers

With autonomous vehicles constantly on the move a requirement for data centers on the edge will evolve. Companies like Vapor IO are designing self-enclosed micro data centers that will be installed under cell towers.



HPE Cloudline Building Blocks for Big Data Streaming

Purpose-built with Open Standards Based Hardware

Different requirements along the data pipeline stages demand different node geometries

“IoT Event Producers”
Edge Processing of data in motion




HPE Cloudline 2800



HPE Edgeline Products



“Fast Data”
Core Processing of data in motion



HPE Cloudline 3100 with NVME




“Big Data”
Analysis of data at rest




HPE Cloudline 3100 with LFF HDDs




“Data Lake”
Object Storage- Warm/Cold



Streaming
Analytics Storage

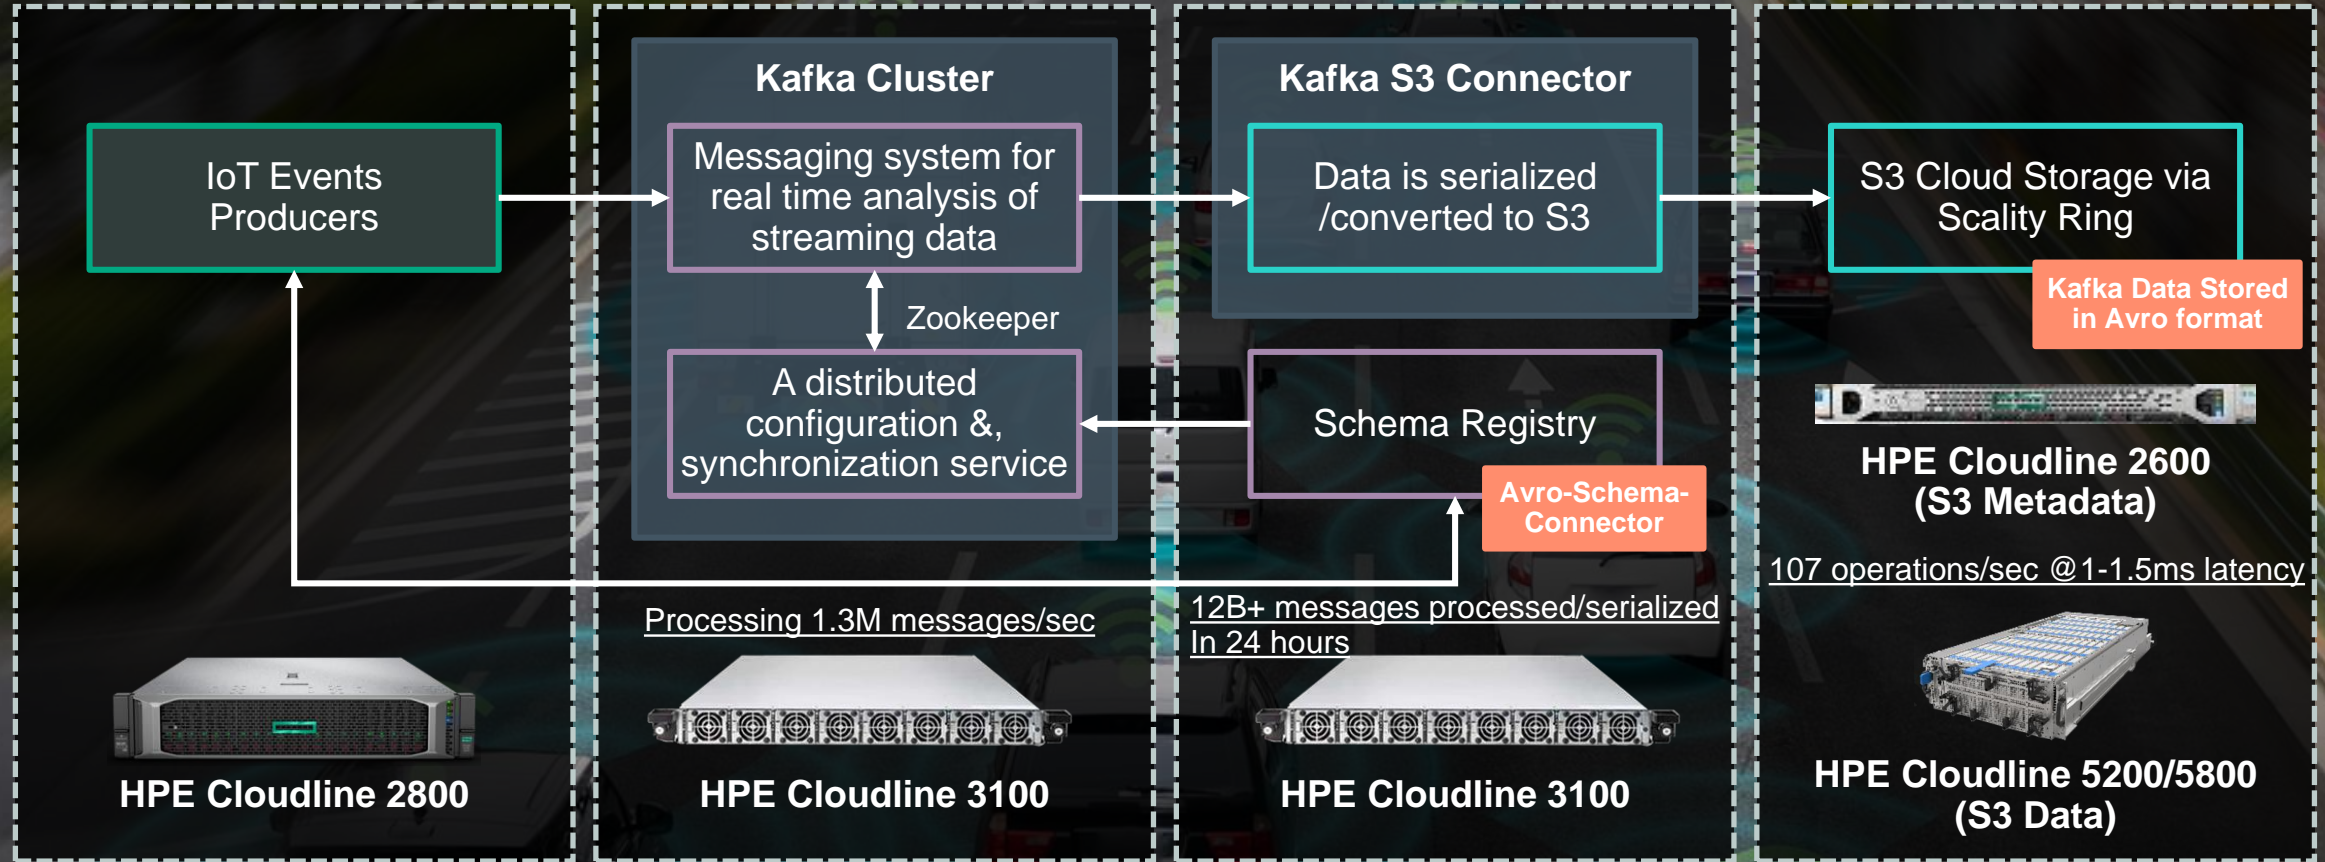


HPE Cloudline 5200 or 5800



Streaming
Analytics Storage

HPE Cloudline for Autonomous Driving Solutions



Data Ingest

Analytics

Object Storage

Edge to Cloud

The SMACK Stack for Big Data / AI / ML



Apache Zookeeper™

An open source Apache project that provides centralized infrastructure and services that enable synchronization across a Hadoop cluster. ZooKeeper maintains common objects needed in large cluster environments. Examples of these objects include configuration information, hierarchical naming space, and so on. Applications leverage these services to coordinate distributed processing across large clusters.



A fast, in-memory data processing engine with elegant and expressive development APIs to allow data workers to efficiently execute streaming, machine learning or SQL workloads that require fast iterative access to datasets.



Apache
MESOS™

A cluster manager that provides efficient resource isolation and sharing across distributed applications or frameworks



A set of open-source libraries for designing scalable, resilient systems that span processor cores and networks



cassandra

A distributed database for managing large amounts of structured data across many commodity servers, while providing highly available service and no single point of failure.



A fast, scalable, durable, and fault-tolerant publish-subscribe messaging system. Kafka works in combination with Apache Storm, Apache HBase and Apache Spark for real-time analysis and rendering of streaming data

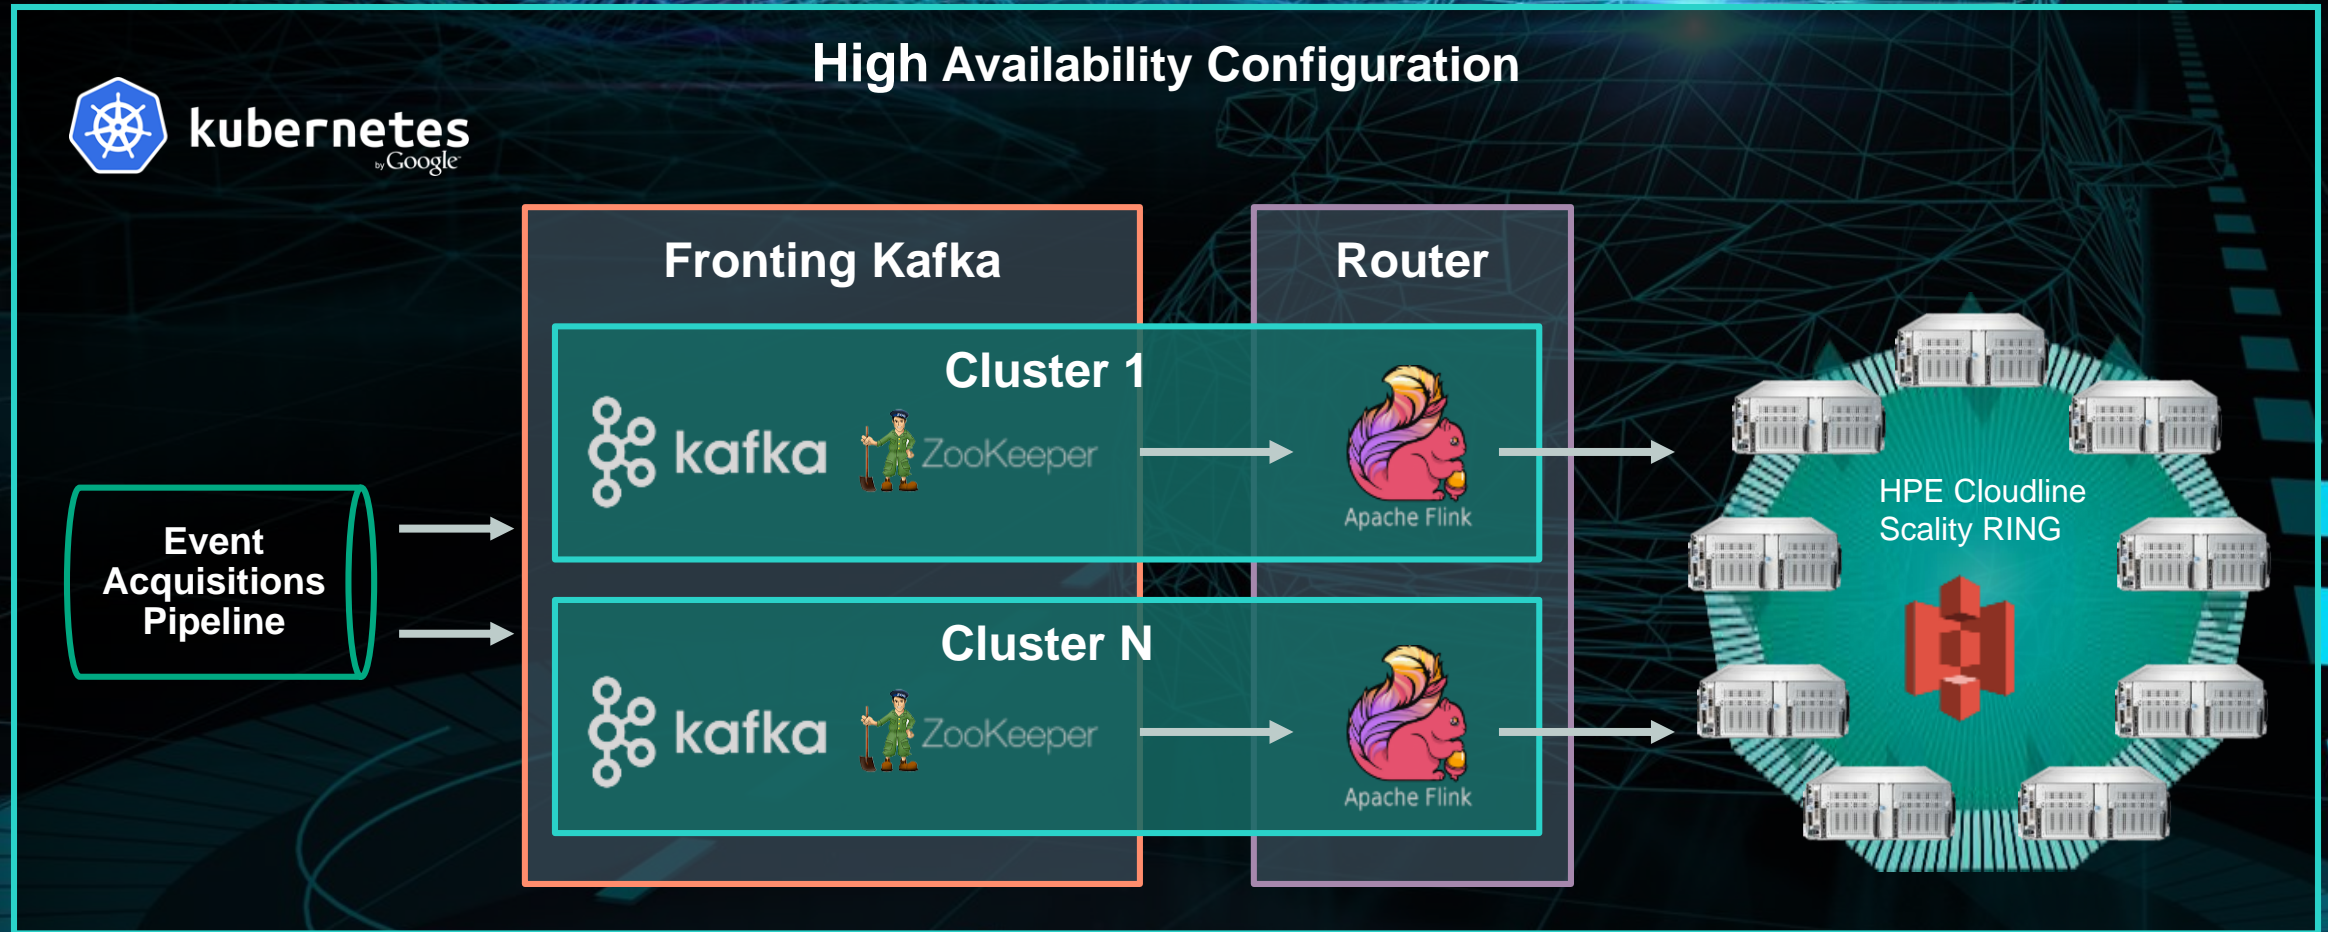
APACHE
HBASE



Flink

An open-source stream processing framework for distributed, high-performing, always-available, and accurate data streaming applications

Autonomous Driving Software Eco-System (Example)



HPE Cloudline Portfolio

General Purpose Servers



HPE CL2600 Gen10
1U 2S, 8SFF

The 1U open standard in an HPE flagship design



HPE CL2800 Gen10
2U 2S, 24SFF or 12LFF or 16 NVMe

Expandable, open system in a 2U HPE flagship design

Compute Servers



HPE CL3100 Gen10
1U 2S, 12LFF or 20 NVMe

Cloud storage server for Hadoop, Cassandra, real-time analytics



HPE CL4100 Gen10
Dual Nodes
Each 1U 2S, 8SFF or 2 NVMe

Optimized for high-density compute and NFV workloads

Storage Servers




HPE CL5200 Gen9
4U 2S, 80LFF

Very high density storage with and ultra-low \$/GB price point



HPE CL5800 Gen9
4U 2S, 100LFF

Optimized for cloud storage and big data workloads



Thank you!

For more information please visit us at OCP booth C6