



**Hewlett Packard
Enterprise**



AI from Edge to Cloud – It's real

S9948 - AI from Edge to Cloud: How HPE and Seagate deliver
Quality and Efficiency across the Manufacturing Supply Chain

Bharath Ramesh, Head – WW IoT Product Management & Marketing, HPE
Bruce King, Data Science Technologist, Seagate Technology


18th March 2019

Edge is the NEW “Off Prem”



Edge – The NEW “Off Prem”

More sensors
= More data
= More analytics
= Smarter Operations



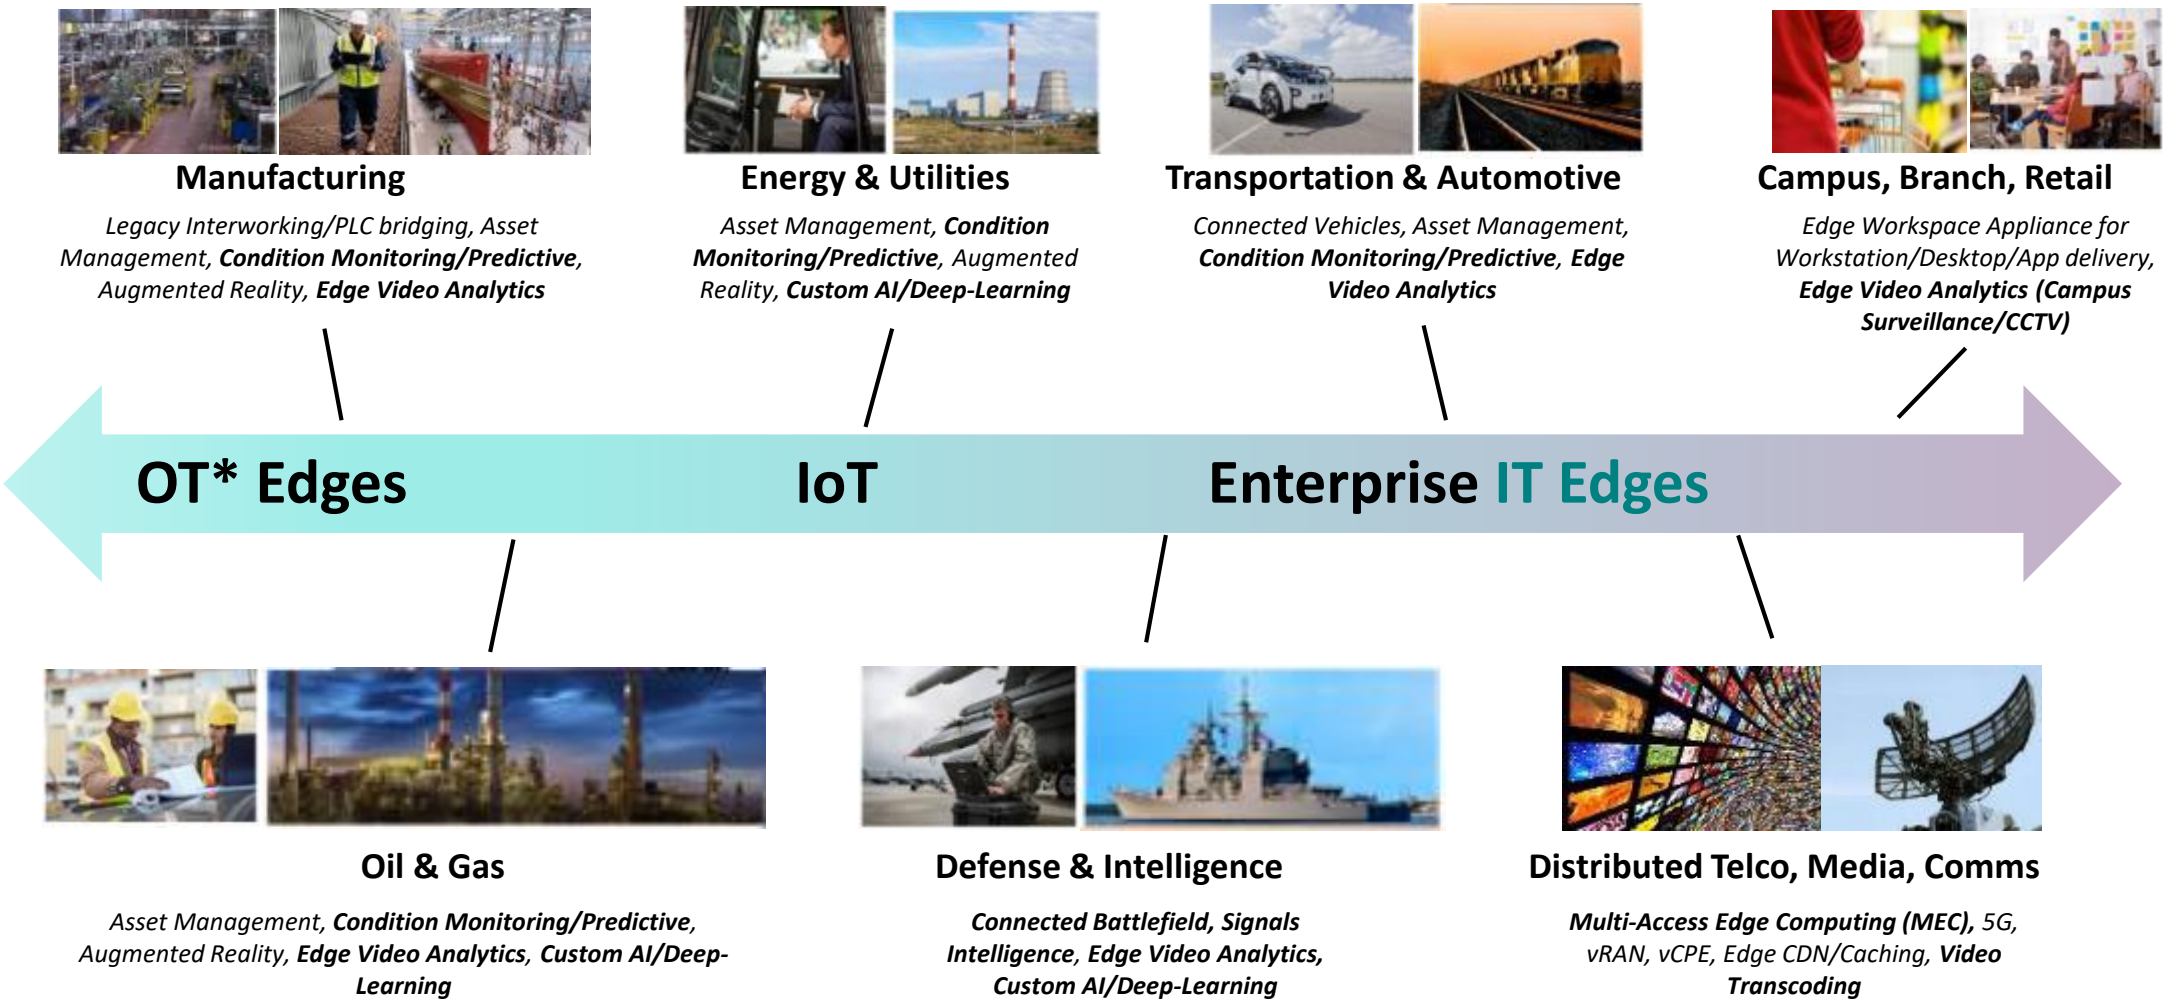
Datacenter – The “On Prem”

Cloud – The current “Off Prem”

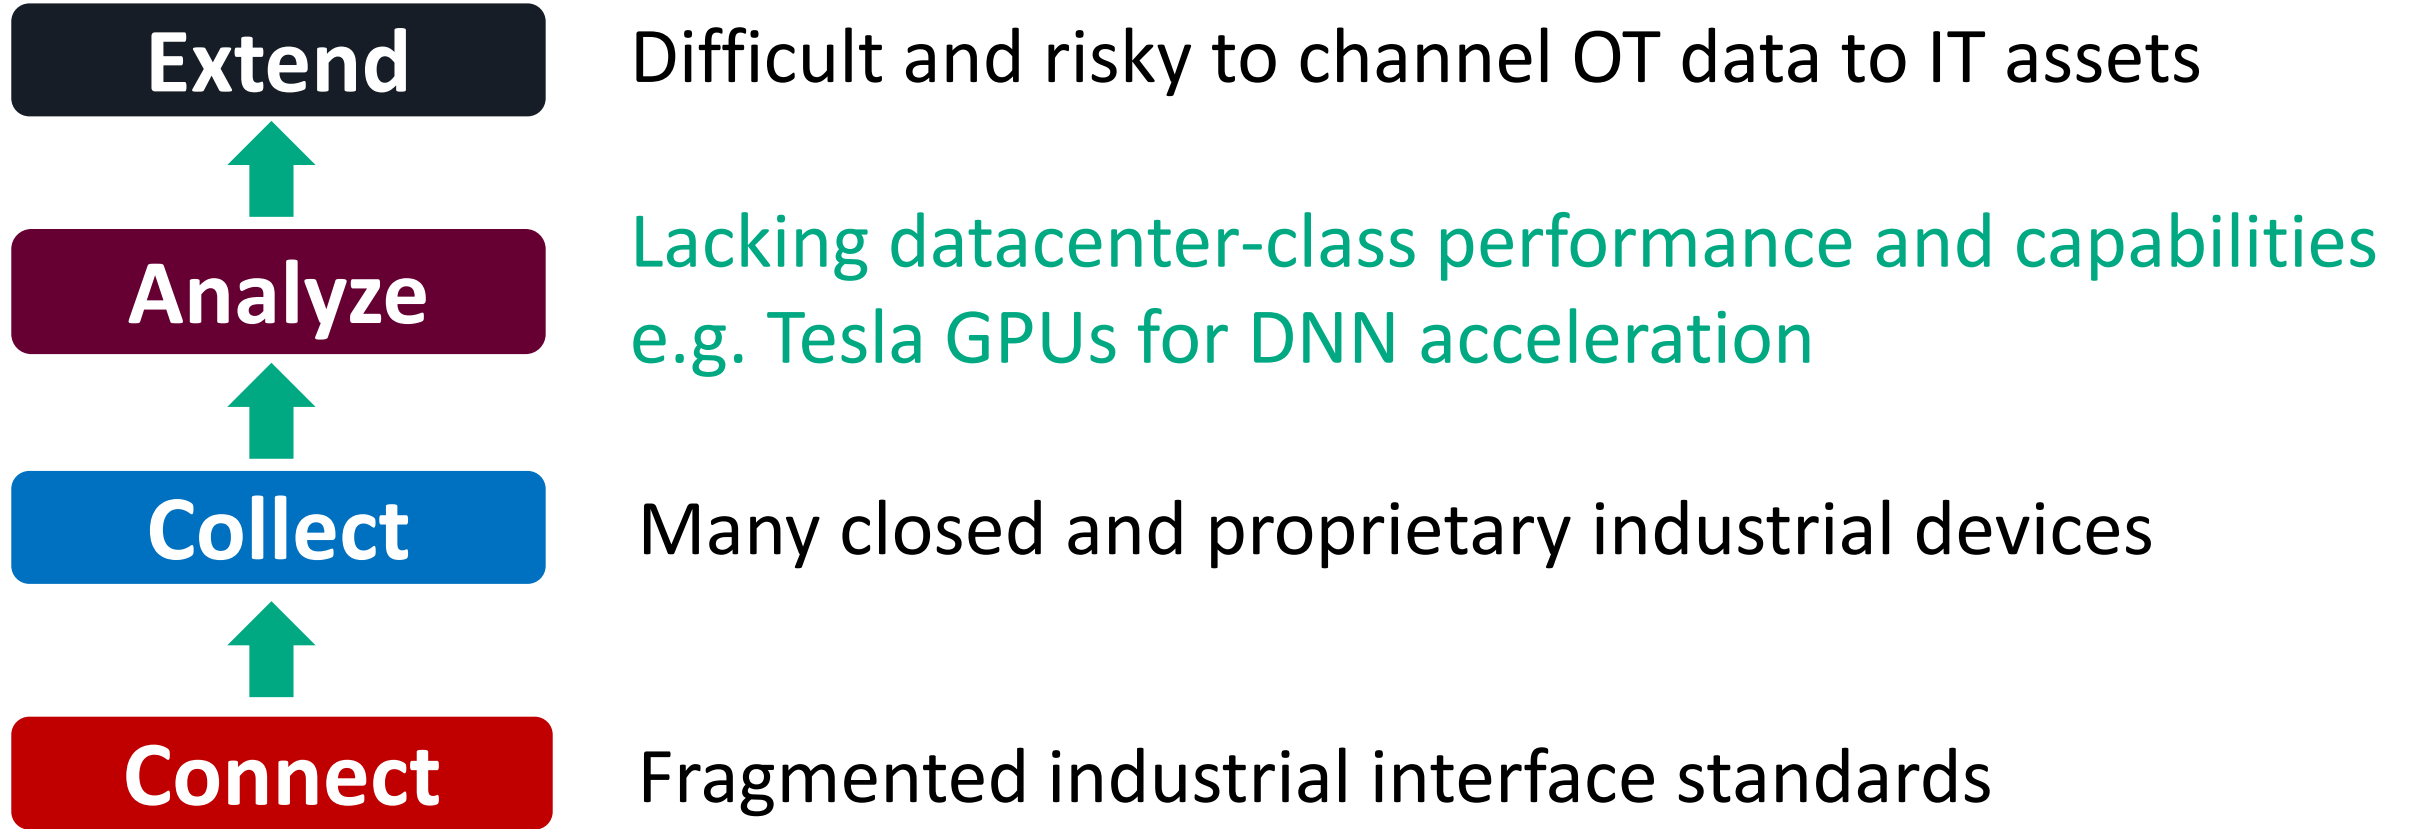


The Edge is a spectrum of use cases that benefit from Deep Learning

* OT = Operational Technology

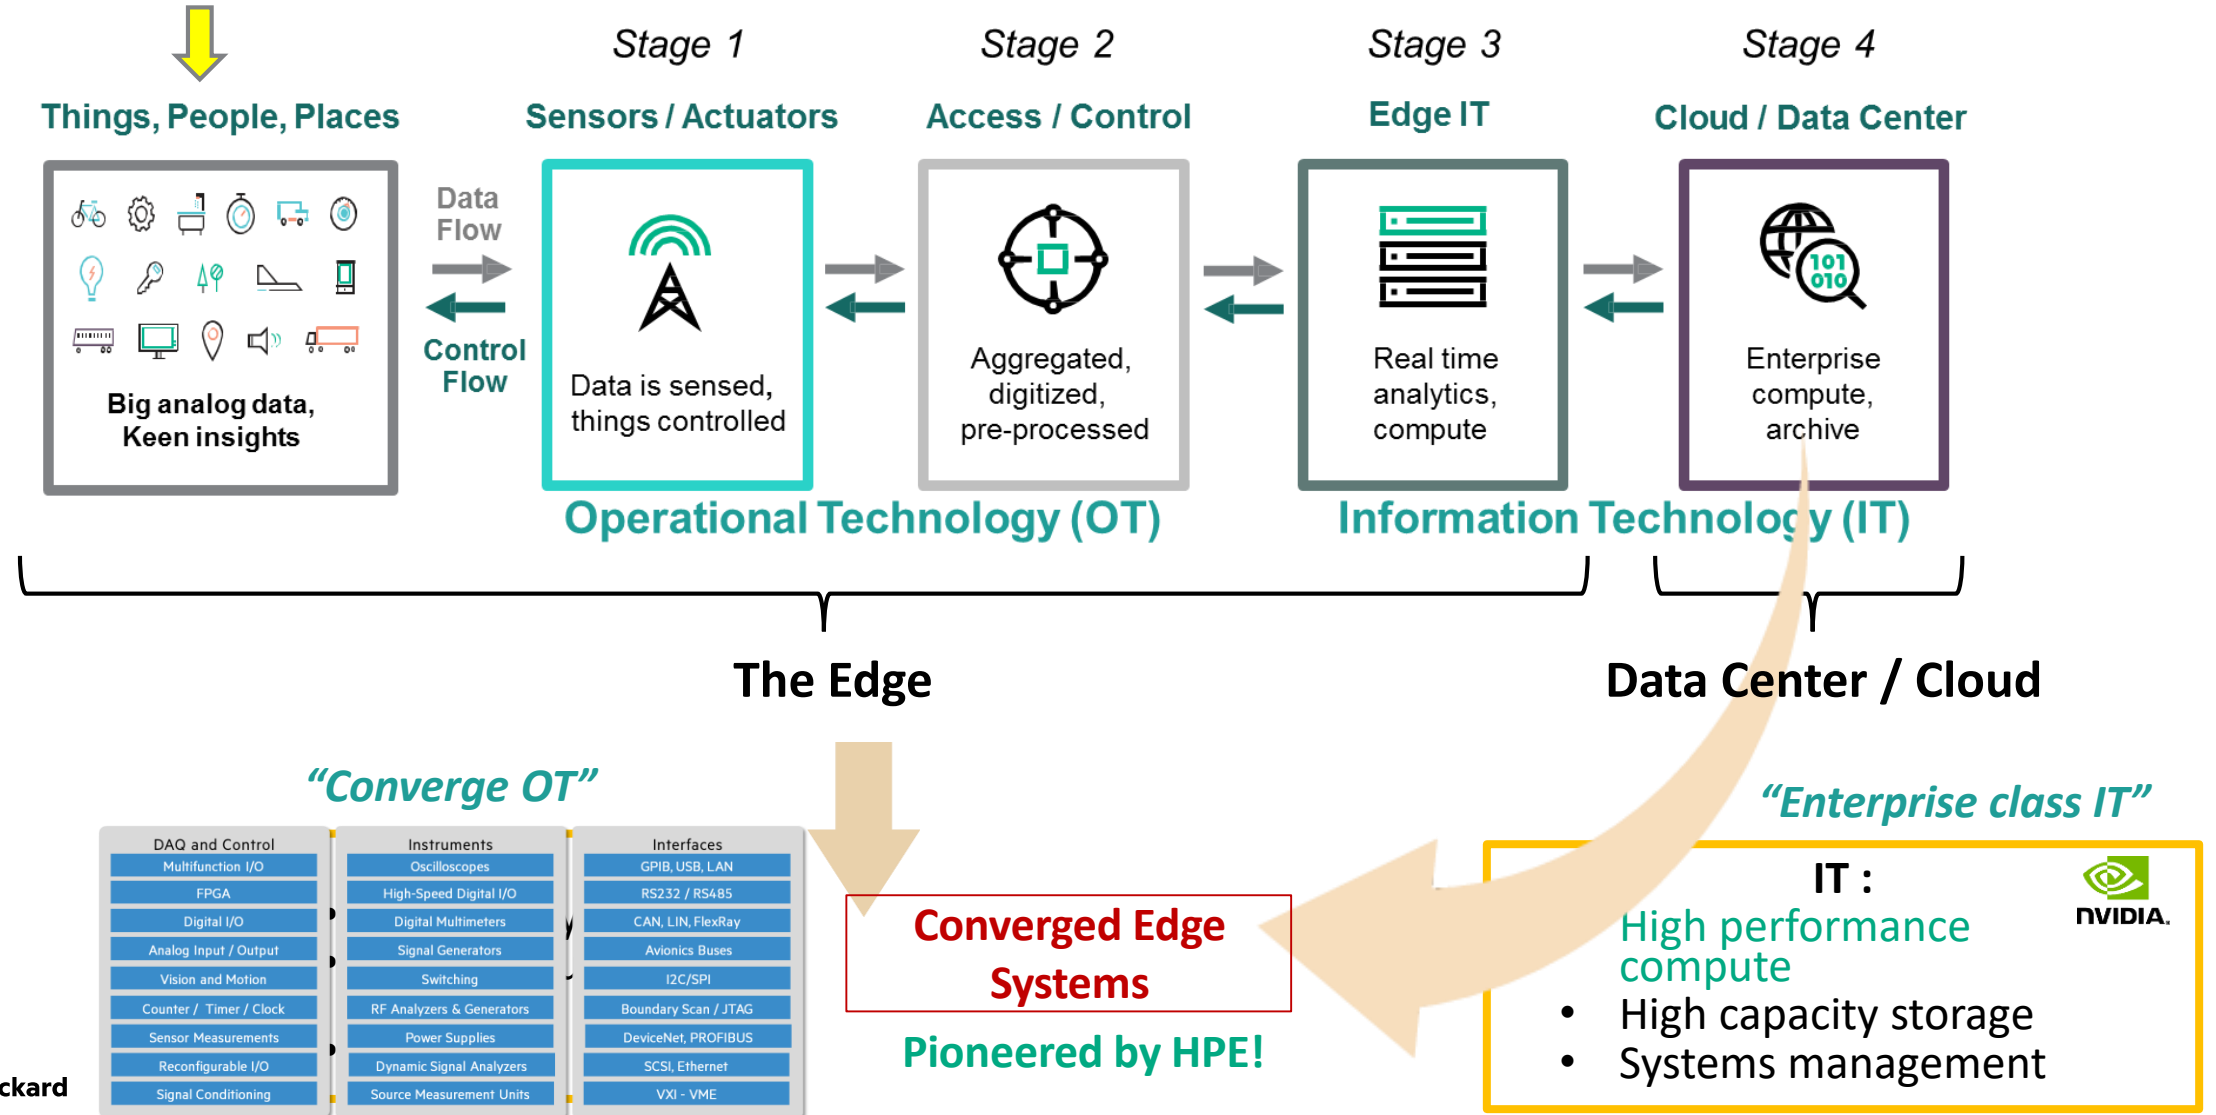


Challenges at the Operations Edge



Introducing the “Converged Edge System”

Emerging source of “Big Data”



HPE Edgeline – Unleashing AI on “Big Data” problems at the Edge

HPE Edgeline EL1000 v2



1 Blade per System
1 Xeon-E/D CPU : 2 Tesla GPUs

 **Tesla P4/T4**
NVIDIA



HPE Edgeline EL4000 v2



4 Blades in 1U System
1 Xeon-E/D CPU : 1 Tesla GPU

 **Tesla P4/T4**
NVIDIA



HPE Edgeline EL8000

Launching: 2H CY2019



222mm D x 220mm W x 431mm H
(17.5" D x 8.7" W x 16.9" H)

Rugged & Compact Enterprise-class Edge System
Perfect for the Operations Edge!

 **Tesla V100 (PCIe)**
NVIDIA



 **Tesla T4**
NVIDIA



Option A:

4 Blades in 5U System
1 Xeon-SP CPU : 1 Tesla GPU


Option B:

2 Blades in 5U System
1 Xeon-SP CPU : 1-4 Tesla GPUs

Option C:


Mix and Match

HPE Apollo 6500 Gen10 – Enterprise AI acceleration at the Core

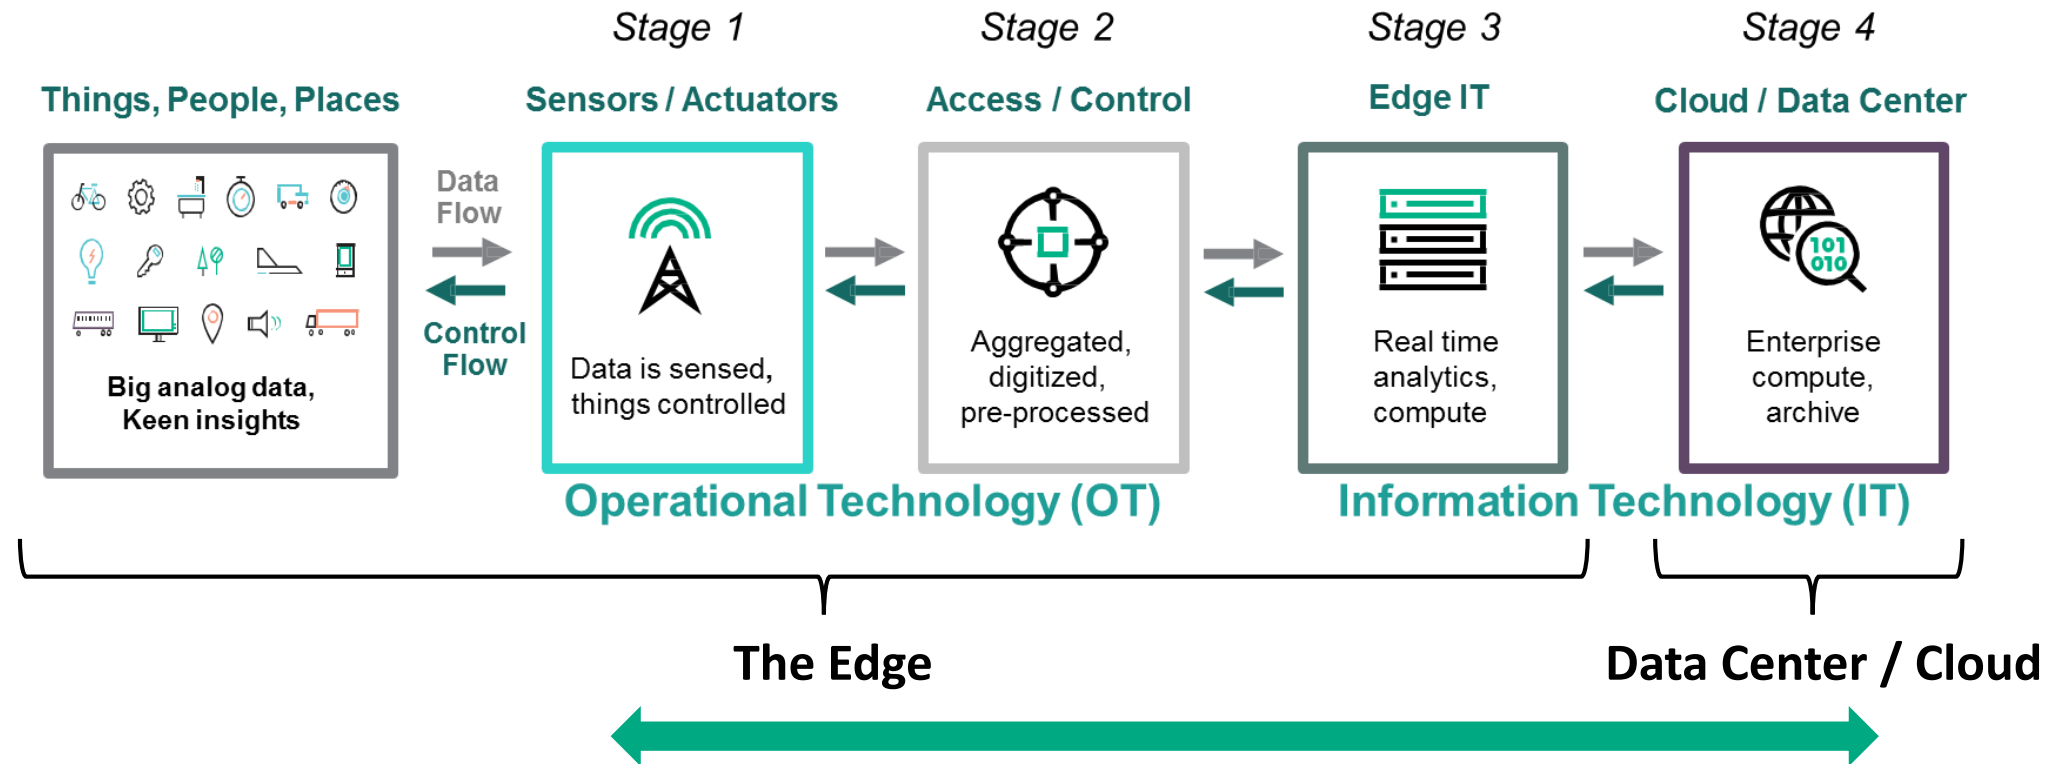


Tesla V100 (PCIe & SMX2)
Tesla T4
Quadro RTX 6000 & RTX 8000

8



Technologies that can scale from Edge to Cloud



Vision of Edge-to-Cloud


**Hewlett Packard
Enterprise**

- Identical Capabilities including Accelerators
- Identical Performance Levels
- Identical Software Stacks
- Management with a single-pane-of-glass
- Secure everywhere



**Hewlett Packard
Enterprise**



Real Life Example of AI from Edge to Cloud

Detecting defects in wafer manufacturing at Seagate Technology

VIDEO: Optimized manufacturing using AI analytics & HPE Converged Edge Systems
<https://www.youtube.com/watch?v=TPf0qC7itcA>

VIDEO: Seagate Transforms Manufacturing with Deep Learning from Edge to Cloud
https://youtu.be/WW_4z7qB7rs


From Wafers to Heads

Each 200 mm wafer contains 10,000's of sliders.

A slider is also called a recording transducer

We make 1 Billion transducers per year

Operations in Normandale, Springtown, and in Korat, Wuxi, Singapore



Data Generation and Analytics Constraints

**Millions of
Electron
Microscope
& Optical
images per
day**

**Real Time
Insights &
Automated
Controls
demand low
latency**

**High
bandwidth
costs
(10TB/day)
precludes
public cloud**

**ML Training
demands
large datasets
near GPU**

**Real time
sensor data
from 100's of
manufacturing
tools**



Seagate on Edge Project Athena

Edge Analytics in the Factory: Enable real time AI & ML

Challenge

Seagate factory's produce Billion+ recording transducers per year. ML/AI lab projects consistently show entitlement. How to scale low latency ML Inference?



High data volume millions of images & thousands of factory tool sensors



Long latencies and bandwidth cost to use Cloud analytics



Need for real time analytics

Solution

Apply Machine Learning to transducer images and sensor data to find anomaly's and defects early and reliably.



Seagate All-Flash for Training Data Storage



NVidia GPU and TensorFlow



Future: Real Time Inference at the Edge

Benefits

Reduced test time, increased efficiency and improve throughput.



Reduced Bandwidth Costs



Faster test cycles




Increased Utilization

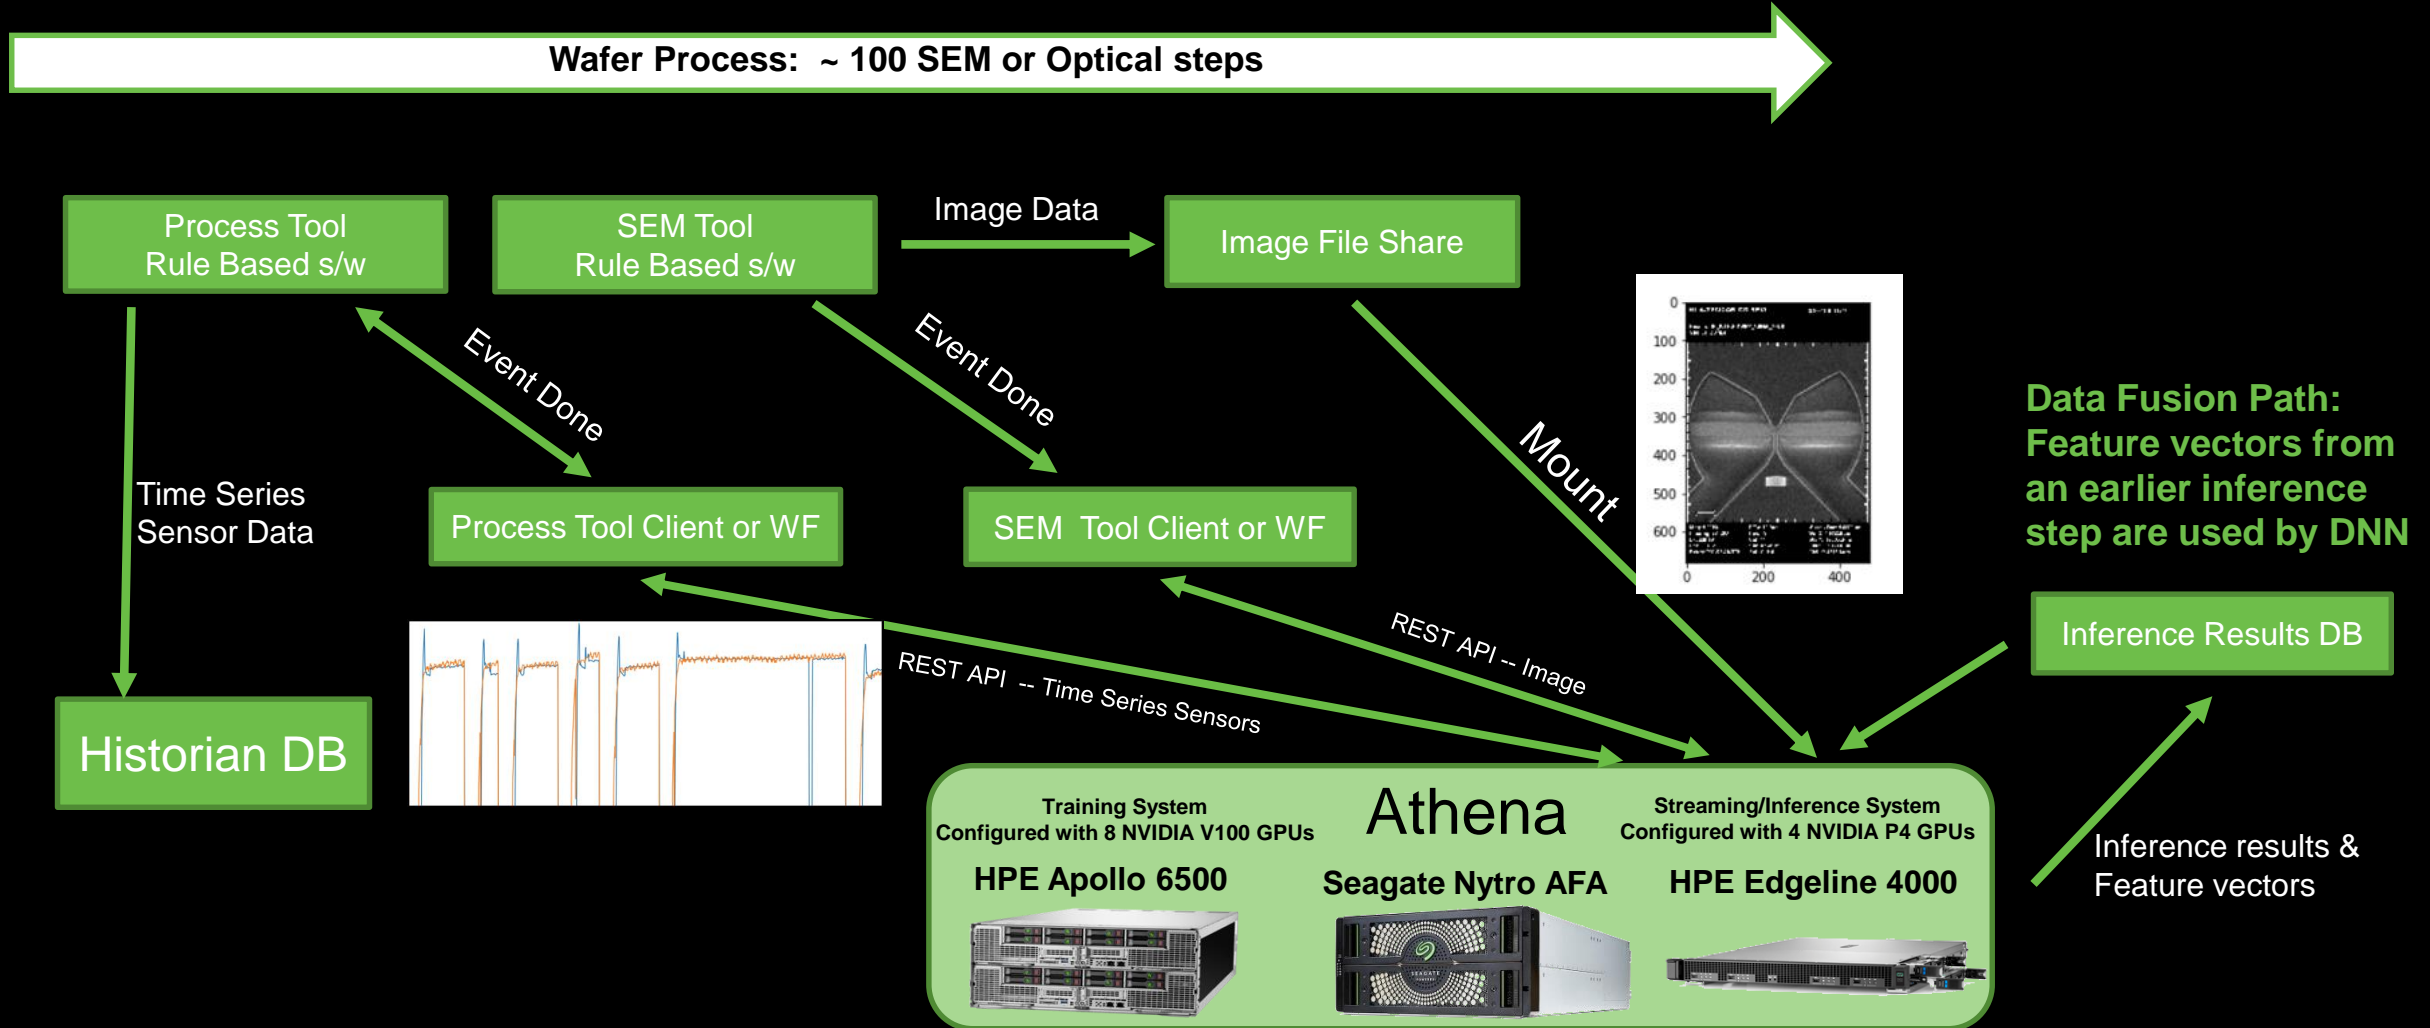
Athena Pilot Inference Dataflow

Wafer Process: ~ 100 SEM or Optical steps

Seagate's mastery of component technology, combined with vertically integrated operations, makes the company a leading provider of hard drives and storage solutions.

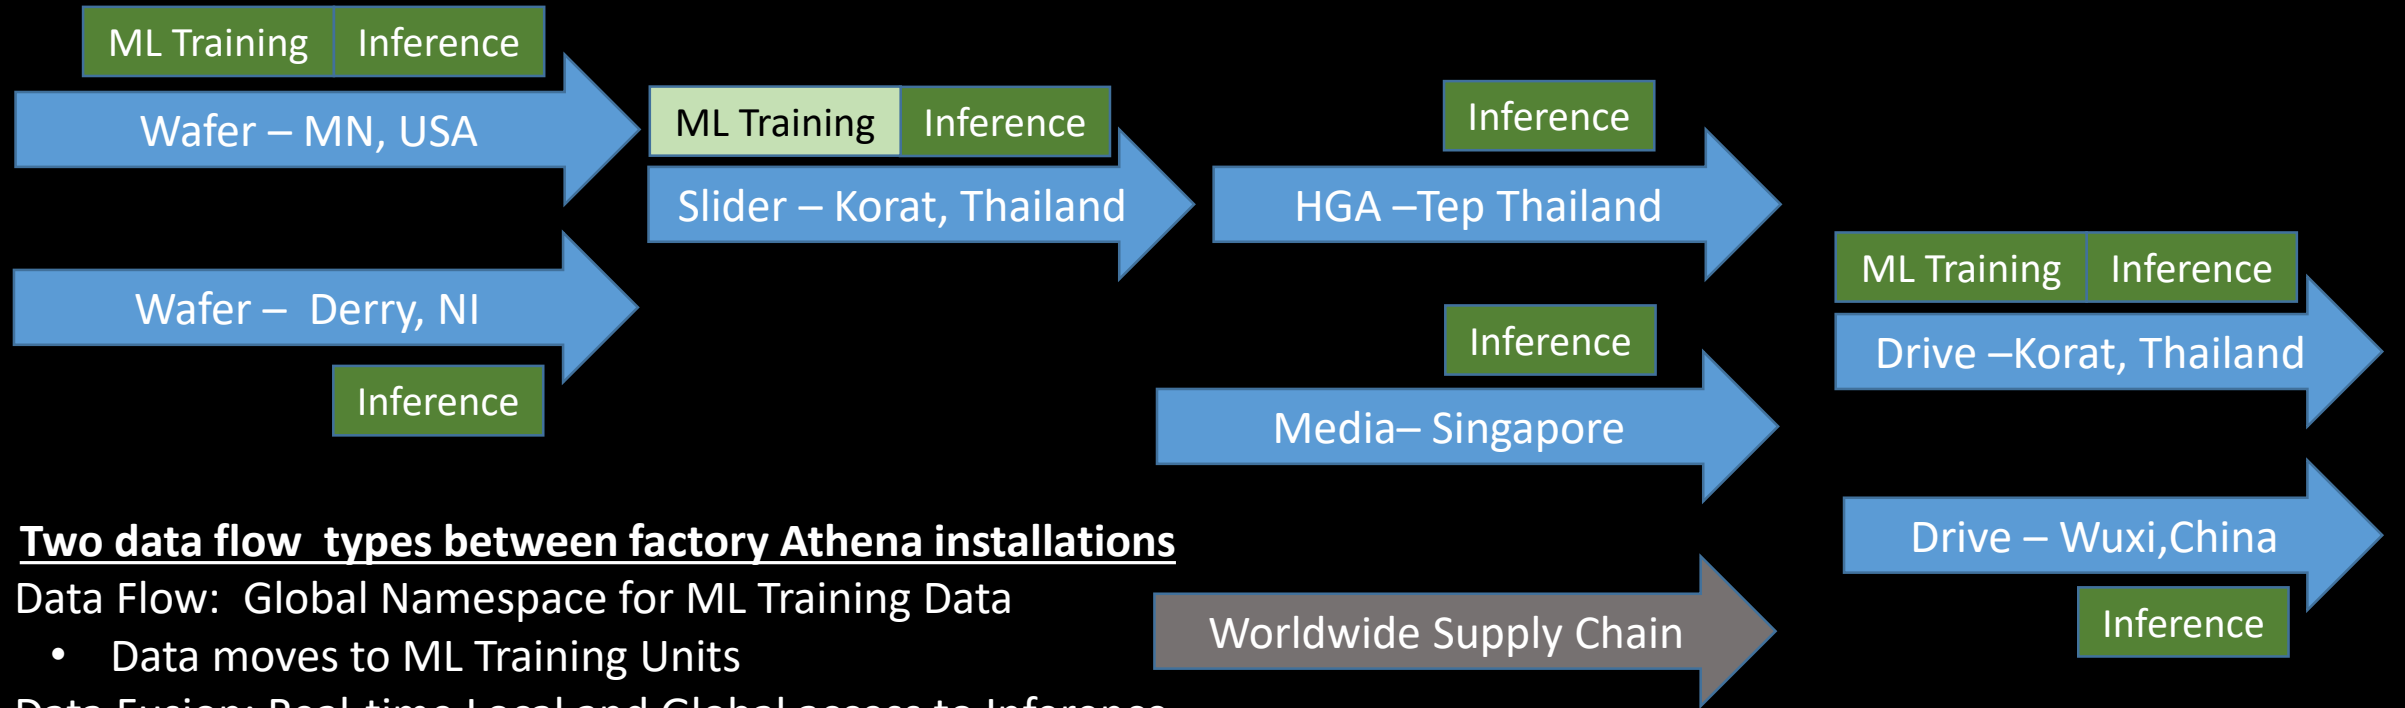


Athena Pilot Inference Dataflow



Seagate on Edge Project Athena

Machine Learning Data Flow through the Vertical Supply Chain



Two data flow types between factory Athena installations

- Data Flow: Global Namespace for ML Training Data
 - Data moves to ML Training Units
- Data Fusion: Real-time Local and Global access to Inference results at successive Inference points
- Trained Model Files flow from ML Training to Inference

Two Athena version: with and without ML Training Hardware (HPE Apollo 6500)

Seagate on Edge Project Athena

Edge Analytics enables new capabilities – The outlook for 2020

Expert Assistant

SME trains AI/ML Agent to help do work that augments human abilities



Interactive ML tools will allow Machines to learn from SME's



SME Assistant will find "interesting" relationships and anomaly's in huge datasets.

Machine assisted AI/ML model building

Data Fusion/Digital Twins

Image, Time-Series Sensor, etc. data from across the vertical supply chain is integrated for ML/AI



Pervasive AI/ML
Athena Real-time
Edge Inference



"Fuse" Results
from Local Edge
Inference



Global storage
of inference
results

System Optimization

ML/AI Applied to Supply-Chain
"Fused Data"



Reduce cost by
replacing "defense in
depth" with systemic
understanding

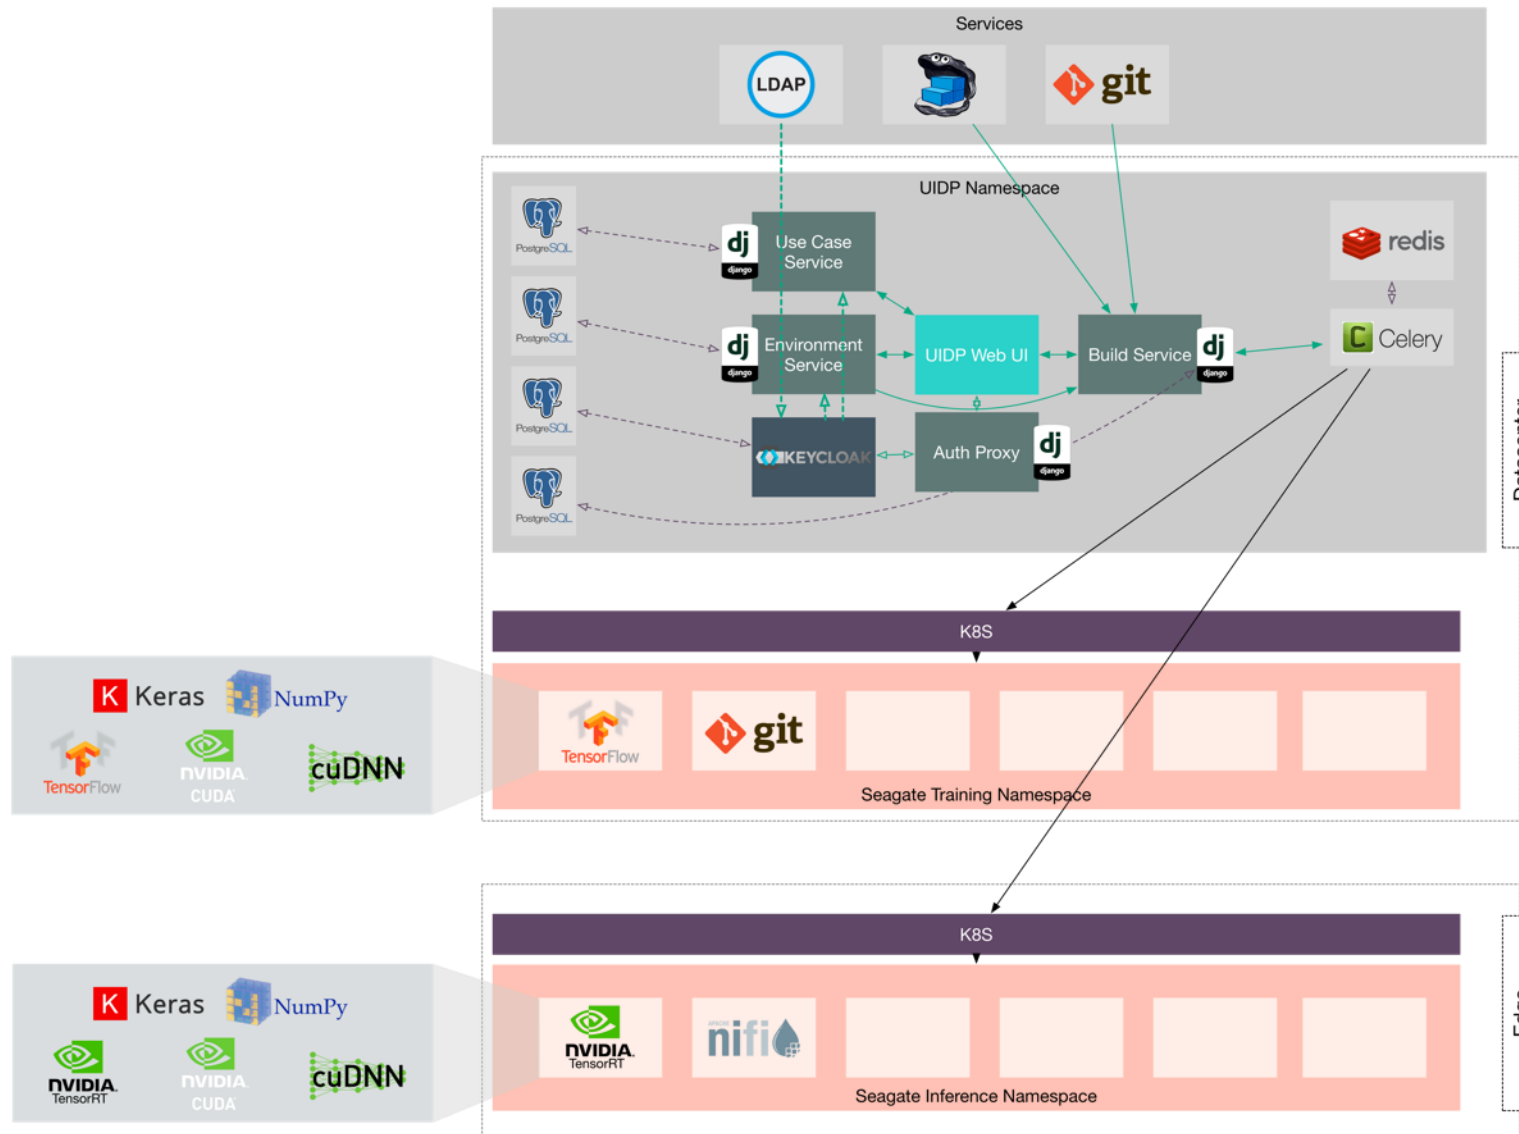


Improve
performance and
quality



Recognize
autonomous car or
robotics analogy

Example Use Case Architecture





**Hewlett Packard
Enterprise**



Thank You

