



# PLATFORM OBSERVABILITY - BUILDING THE FOUNDATION FOR 5G AND EDGE AUTOMATION

September 4<sup>th</sup> 2020

Presenters: Sunku Ranganath and Killian Muldoon

# Agenda

- Telemetry Capabilities in Intel's Network Infrastructure Platforms
- Telemetry Reports – Providing actionable insights
- Enabling Closed Loop Automation
- Use Case Demos:
  - Reporting and visualization: Platform telemetry reports
  - Automation and App metrics: Power and packet processing
  - Profiling trends and prevention: Platform resiliency


# The Evolution Toward Cloud Native



SAME GOALS AS WITH NFV – BUILDS ON YEARS OF INDUSTRY EXPERIENCE AND UTILIZES CLOUD NATIVE TECHNOLOGIES TO ACCELERATE ACHIEVING INDUSTRY GOALS

# Scale Efficiency with Data-Driven, Closed Loop Automation

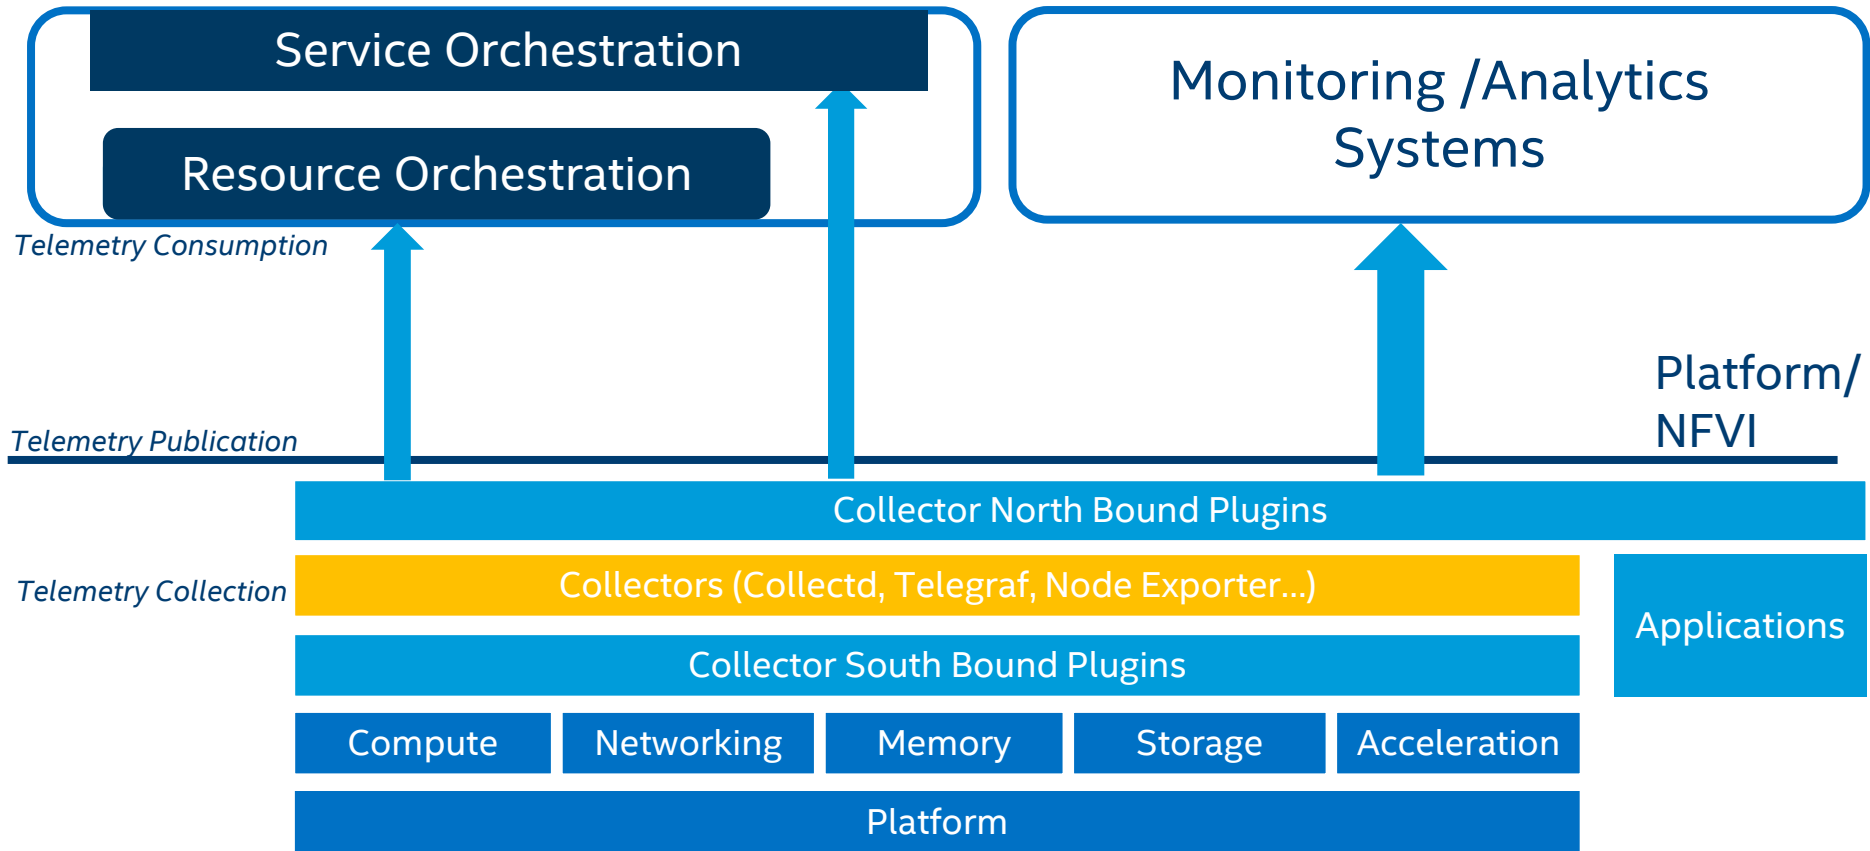
Telemetry capabilities in Intel platforms enable intelligent, closed loop solutions that are reactive, proactive and predictive, delivering new levels of efficiency for IT and network infrastructure.



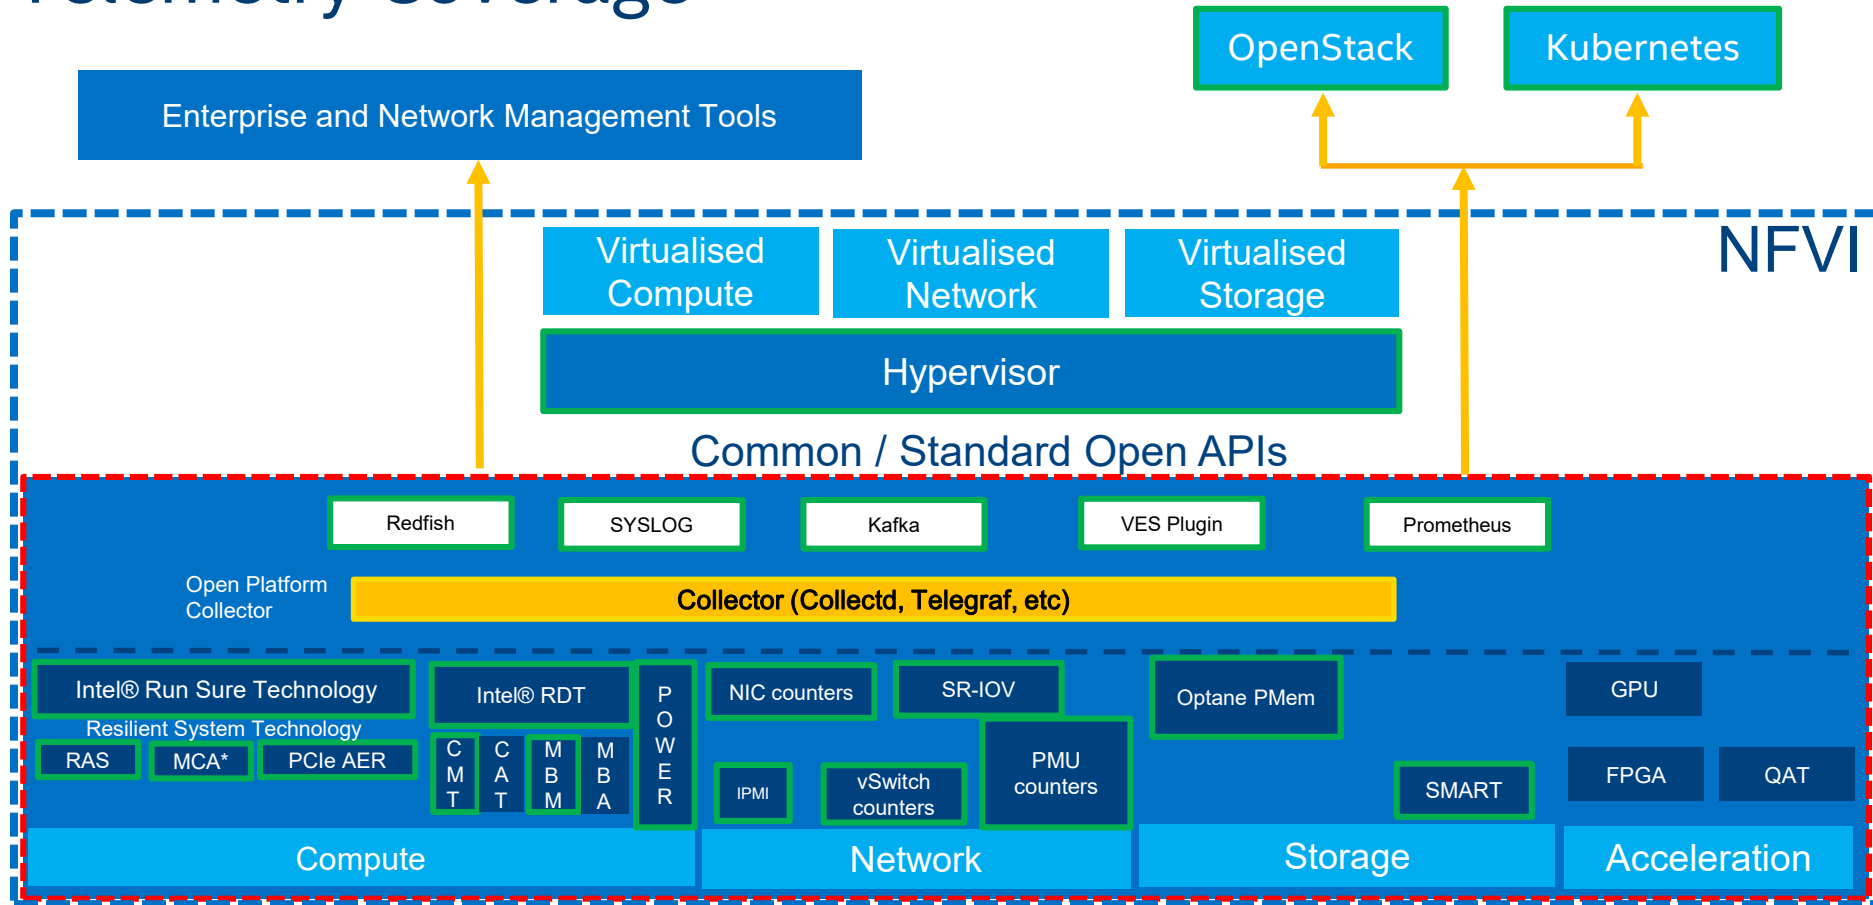
PLATFORM OBSERVABILITY & NETWORK INSIGHTS CRUCIAL FOR OPERATIONAL INTELLIGENCE AND AUTOMATION

# PLATFORM TELEMETRY

# Platform Telemetry Collection and Publication



# Telemetry Coverage



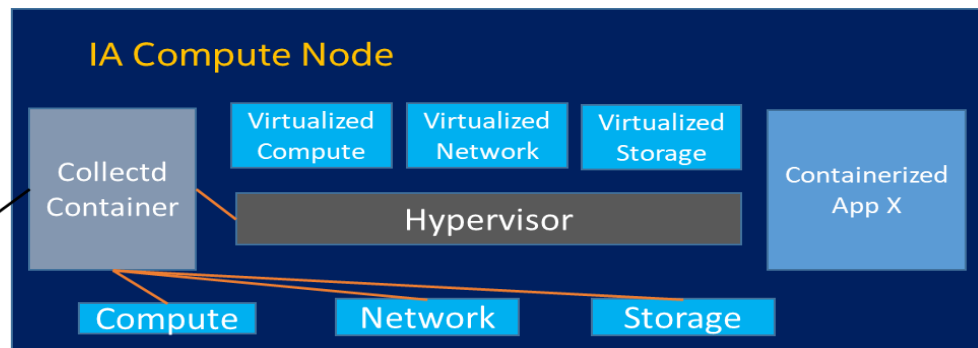
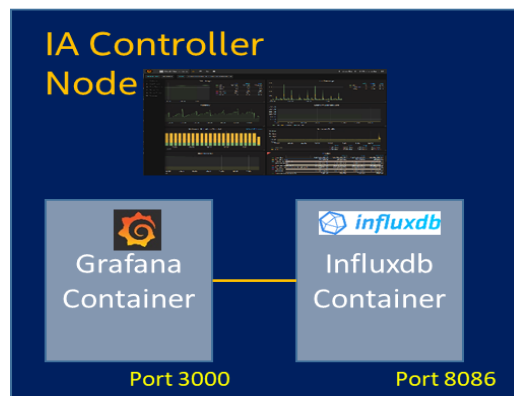
# OPNFV Barometer

## Barometer Strategy:

- Ensure platform metrics/events are accessible through open industry standard interfaces.
- Demonstrate platform & network technologies can be monitored, consumed and actioned in real time

## One Click Install:

- Easy install/configuration for customers
- One command to install Collectd/Influxdb/Grafana



More Info: <https://wiki.opnfv.org/display/fastpath/One+Click+Install+of+Barometer+Containers>






# TELEMETRY ANALYSIS AND REPORTS

Intel's Network Infrastructure Platforms

# Scale Efficiency with Data-Driven, Closed Loop Automation

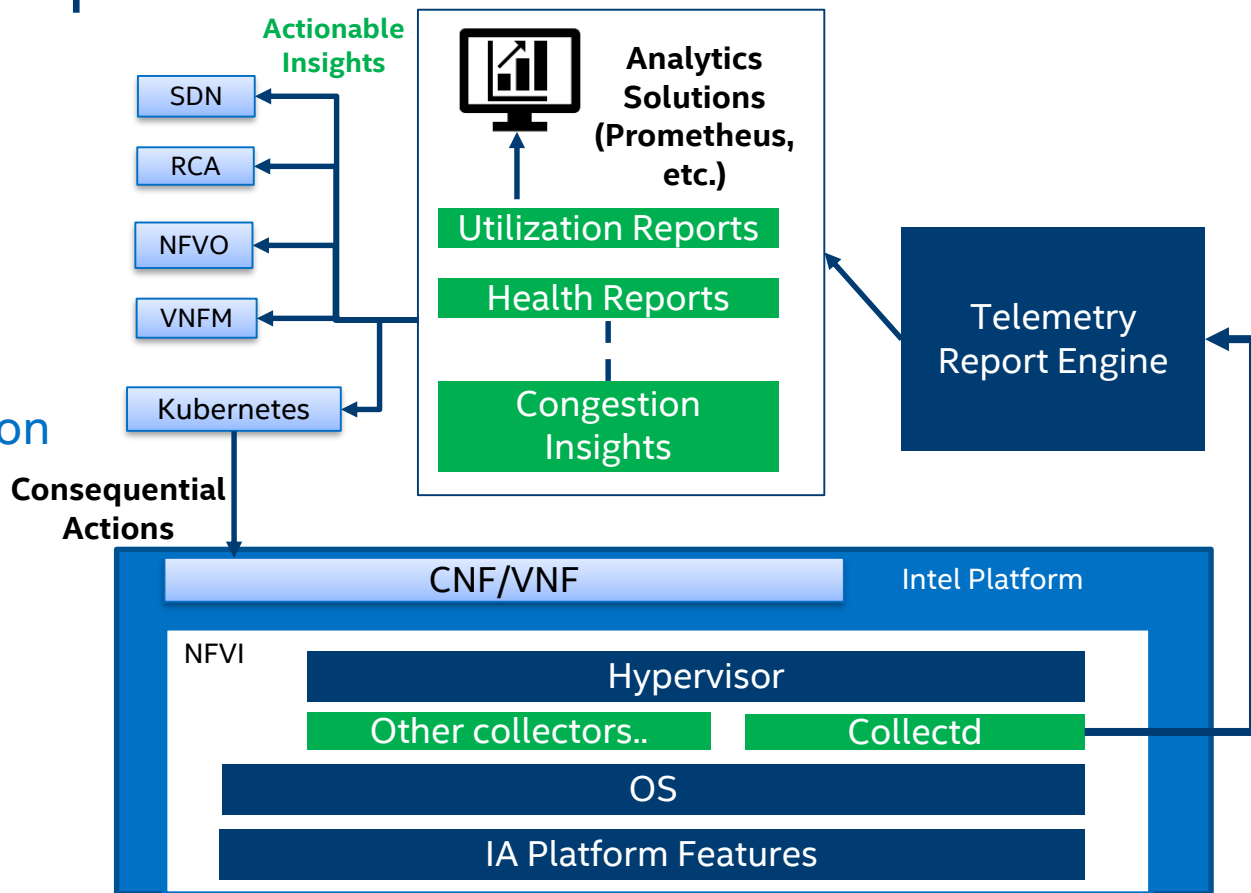
Telemetry capabilities in Intel platforms enable intelligent, closed loop solutions that are reactive, proactive and predictive, delivering new levels of efficiency for IT and network infrastructure.



PLATFORM OBSERVABILITY & NETWORK INSIGHTS CRUCIAL FOR OPERATIONAL INTELLIGENCE AND AUTOMATION

# Use of Telemetry Reports

- ✓ Provide the capability to easily decipher platform metrics
- ✓ Provide 'actionable data'/insights that management/orchestration systems can make decisions on
- ✓ Show the value of Intel platform in a monitored environment




# Telemetry Reports - Examples

Insight Type	Value
<b>Platform Health Insights</b>	<ul style="list-style-type: none"><li>- Used as inputs to drive corrective actions in MANO layers, SDN controller</li><li>- E.g. - failover and other service reliability related actions.</li></ul>
<b>Platform Utilization Insights</b>	<ul style="list-style-type: none"><li>- Used by management system/VIM to detect platforms trending to resource exhaustion</li><li>- Improved work-load placement decisions.</li></ul>
<b>Platform Configuration Check Insights</b>	<ul style="list-style-type: none"><li>- Used by management system/VIM detect misconfigured platforms</li></ul>
<b>Platform Congestion Insight</b>	<ul style="list-style-type: none"><li>- Used by management system/VIM detect platform capacity across compute, interfaces and virtual switching to determine overloaded platforms.</li></ul>

\*<https://builders.intel.com/docs/networkbuilders/SKU-343247-001US-queue-management-and-load-balancing-on-intel-architecture.pdf>



Memory Availability

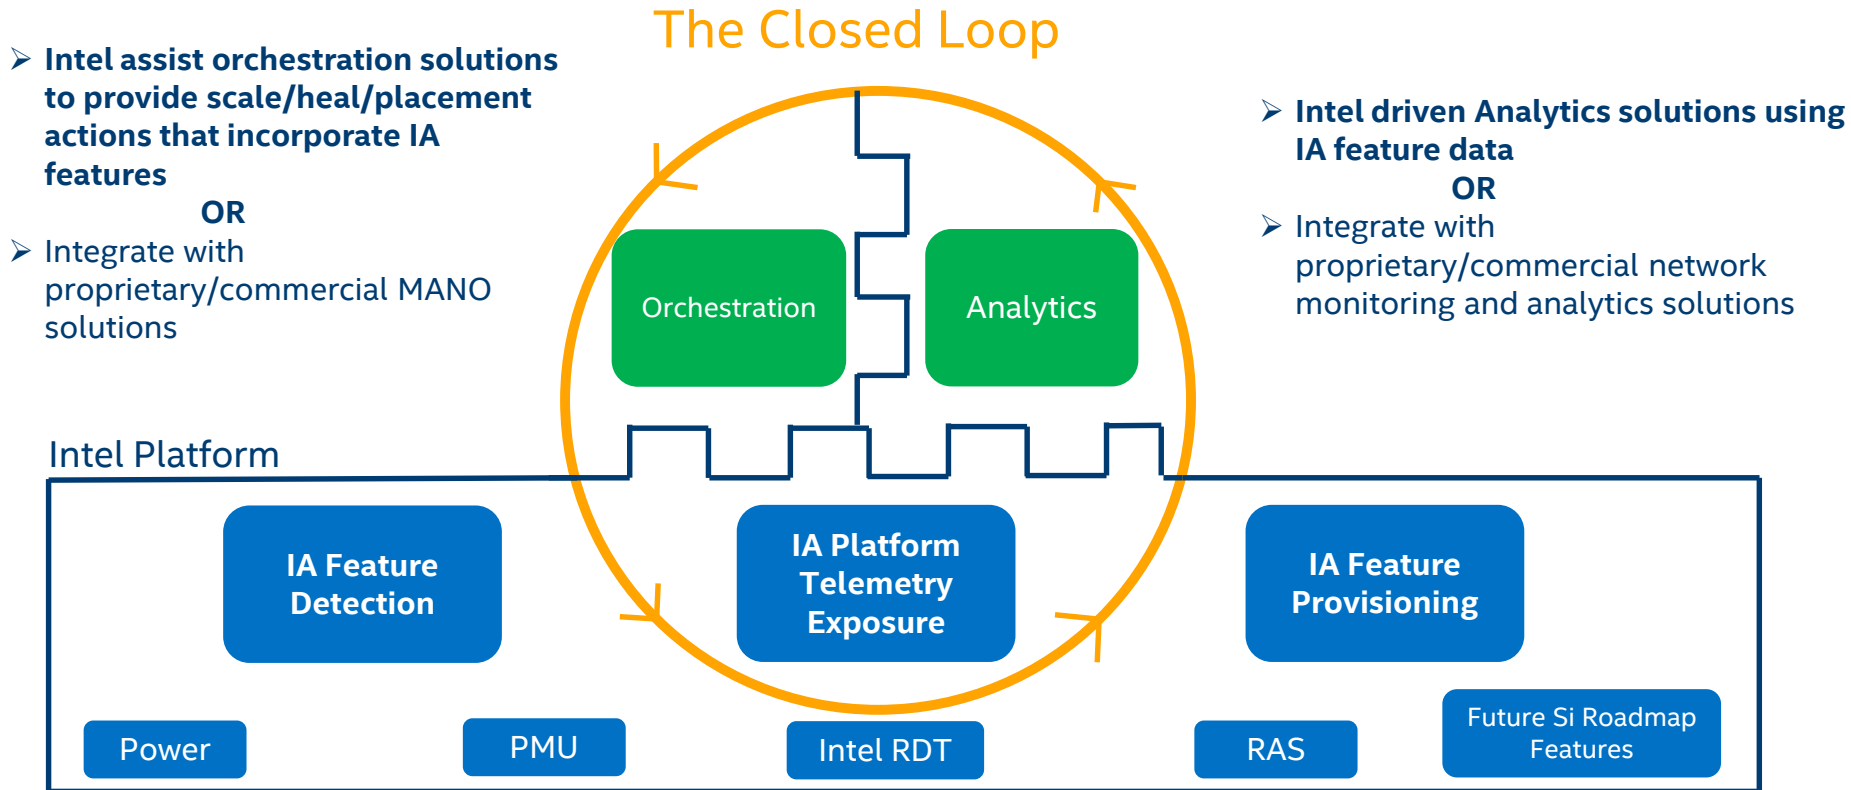


Memory Error Seconds



# ENABLING CLOSED LOOP AUTOMATION

# Platform Features Enable Closed Loop Automation



IA – Intel Architecture

# Use Case Overview

Category	Monitor	Action
Service Healing/ Platform Reliability	Platform Feature Reliability Counters	<ul style="list-style-type: none"><li>• Deploy workloads to the most reliable platforms</li><li>• Proactively move workloads, in planned maintenance windows, to avoid potential outage</li></ul>
	Platform shows high load	<ul style="list-style-type: none"><li>• Inform Backoff or scale out decision</li><li>• Move new workloads to platforms with sufficient capacity.</li></ul>
	Platform Feature Reliability Counters	<ul style="list-style-type: none"><li>• Use to inform horizontal Scale out as part of an overall resiliency scheme</li></ul>
Power Management/ Energy Optimisation	Application shows low load Or Other methods	<ul style="list-style-type: none"><li>• Scale Down Frequency(scale up also needed)</li></ul>
	Platform/Application	<ul style="list-style-type: none"><li>• Deploy new workloads to nodes benefitting from Higher Frequency cores</li></ul>
	Platform shows high load	<ul style="list-style-type: none"><li>• Backoff or scale out</li><li>• Deploy based on power consumption</li></ul>
Application QoS	Platform Feature Utilization insights	<ul style="list-style-type: none"><li>• Configure platform feature/resource per priority application or workload</li><li>• Move applications between different priority groups as part of NF rollout</li></ul>






# **TAKING ACTION: HORIZONTAL POD AUTOSCALER AND TELEMETRY AWARE SCHEDULING**

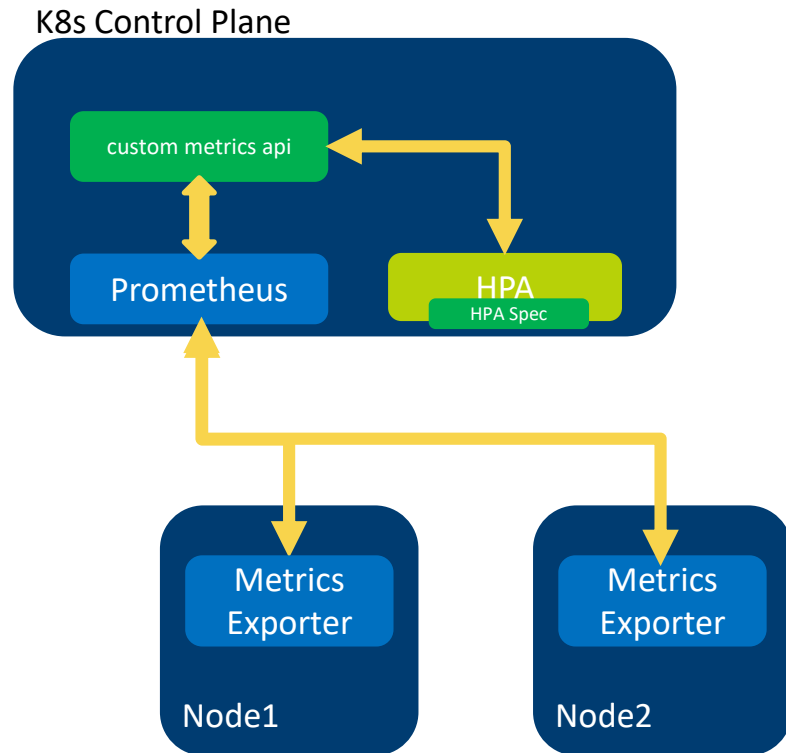
# Automated Action

Telemetry collection and analysis not only gives us the power to see what's going on in a system – but also the power to automate operations.



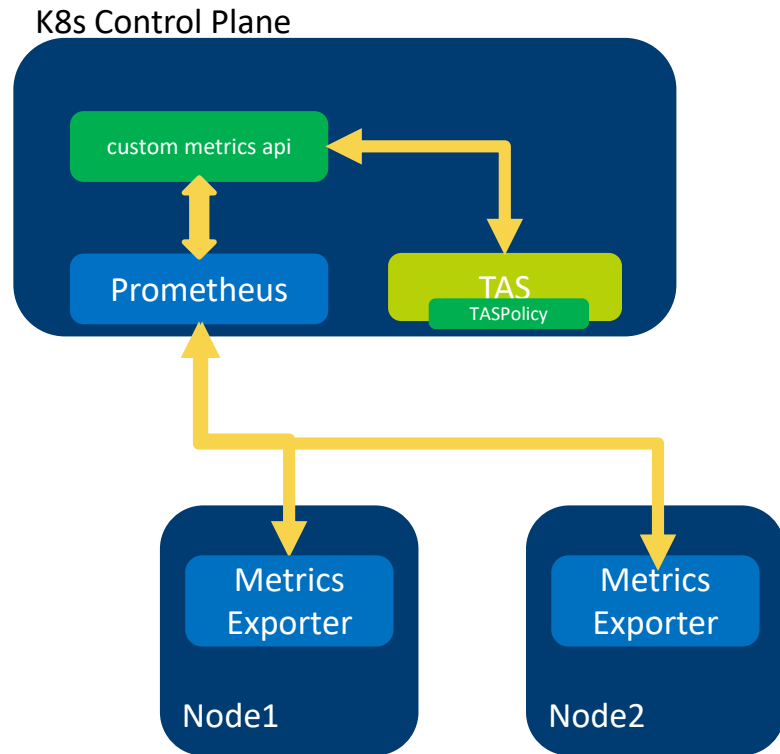
# Horizontal Pod Autoscaler

- Native Kubernetes tool to scale up and down workloads on custom metrics.
- Uses Prometheus and the Kubernetes Custom Metrics API for telemetry
- Allows scaling up or down within a range based on metrics changes
- `kubectl autoscale deployment ...`



# Telemetry Aware Scheduling

- Intelligent Kubernetes workload placement using platform telemetry
- Blocking and prioritizing of placement based on telemetry signals from nodes
- Can combine with analytics to generate multi-factor scheduling hints
- Integrates with native Kubernetes scheduler
- <https://github.com/intel/telemetry-aware-scheduling>





# **POWER DRIVEN SCHEDULING AND SCALING WITH CPU TELEMETRY IN KUBERNETES**

# The Challenge

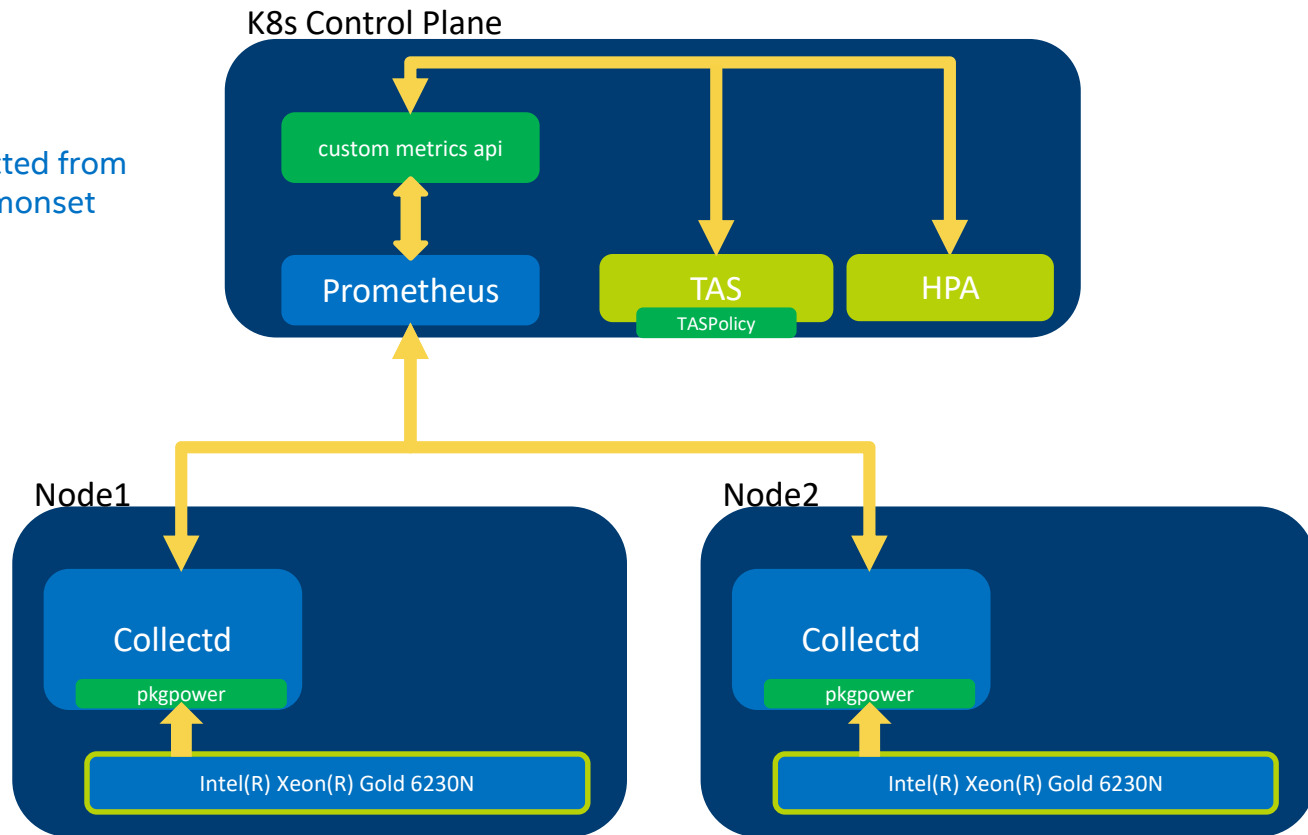
- Understanding the utilization of a platform is key to understanding current load and potential overload conditions, giving the network operator an opportunity to monitor and react to meet Service Level Agreements (SLAs)
- Service providers can identify platform load for workloads and automate actions to resolve incidents meeting service level agreements
- Insights based on power metrics leads to intelligent automated orchestrated decisions
- **This demo implements automation of workload deployments and scaling out based on the power (watts) consumed by the CPU**

# Demo Key Components

- Use platform power telemetry to provide insights into platform load, and intelligently orchestrate workloads based on the reported platform load.
- Horizontal Pod Autoscaler (HPA) is used to scale out workloads when platform overload is detected.
- Telemetry Aware Scheduling (TAS) uses platform metrics to intelligently schedule workloads created by HPA.
- The telemetry agent Collectd reports the platform metrics such as power consumption

# Data Collection

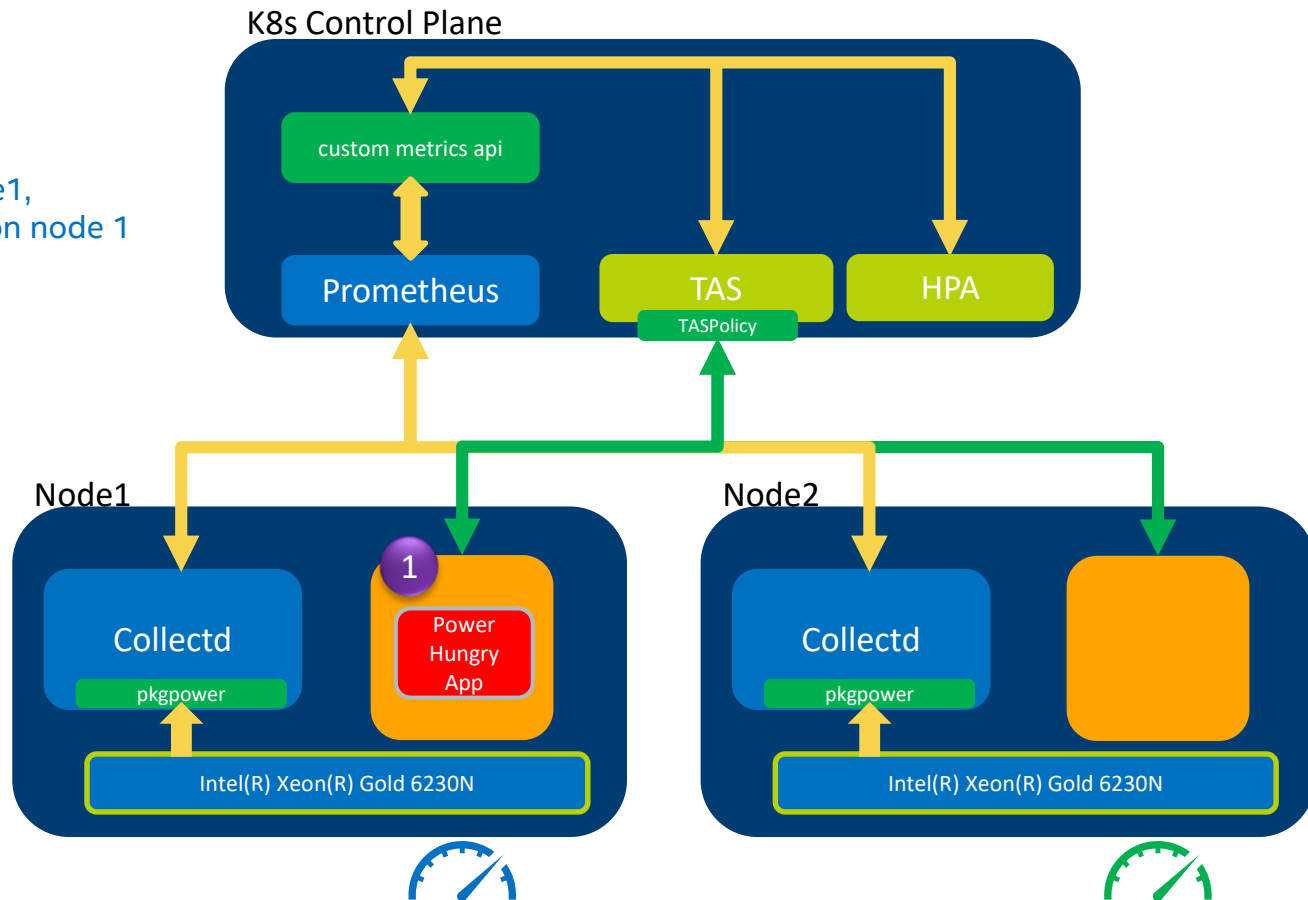
Platform power telemetry collected from each node using a Collectd daemonset





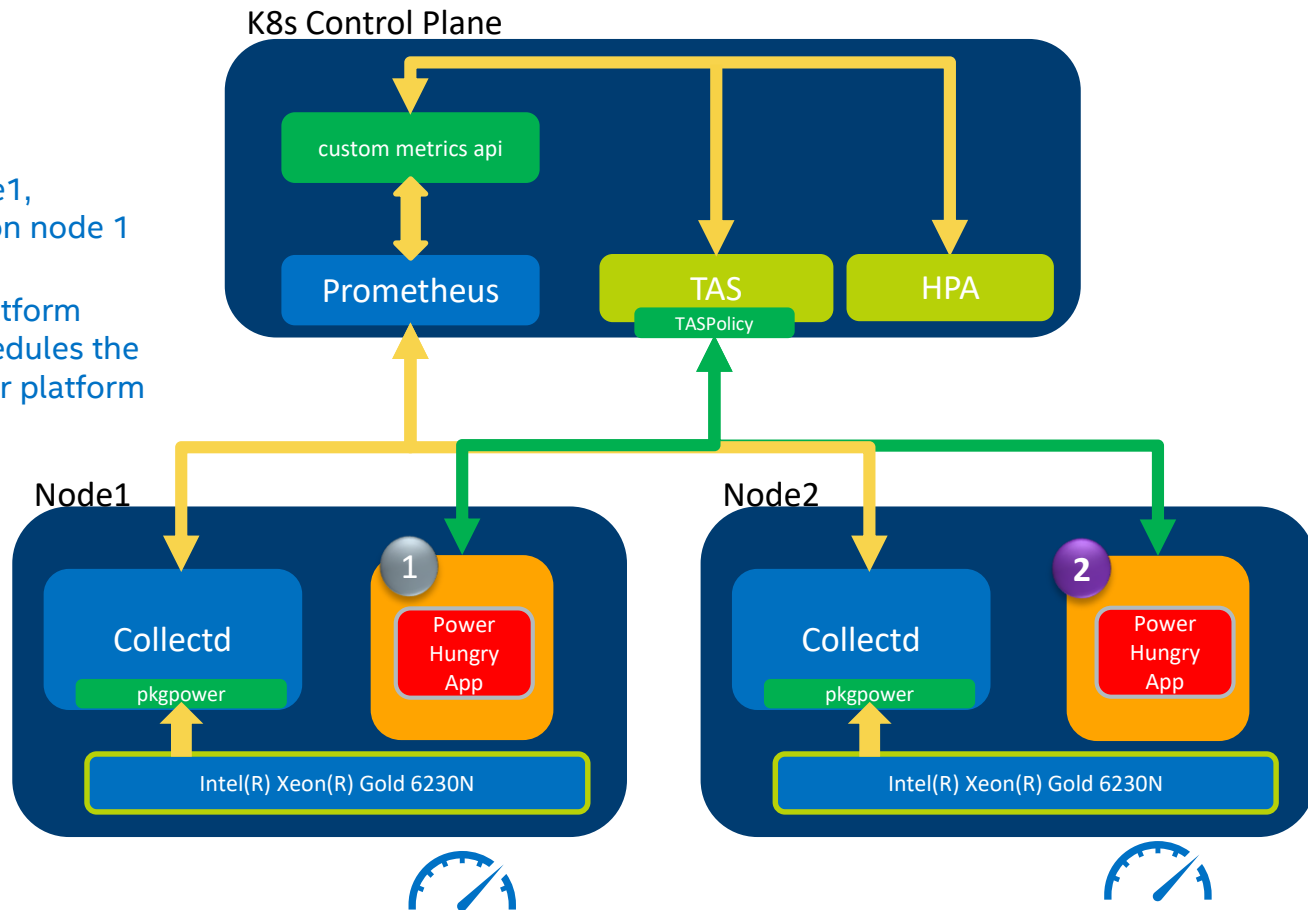
# Demo Scenario

- 1 First pod is deployed to node1, increasing the power usage on node 1



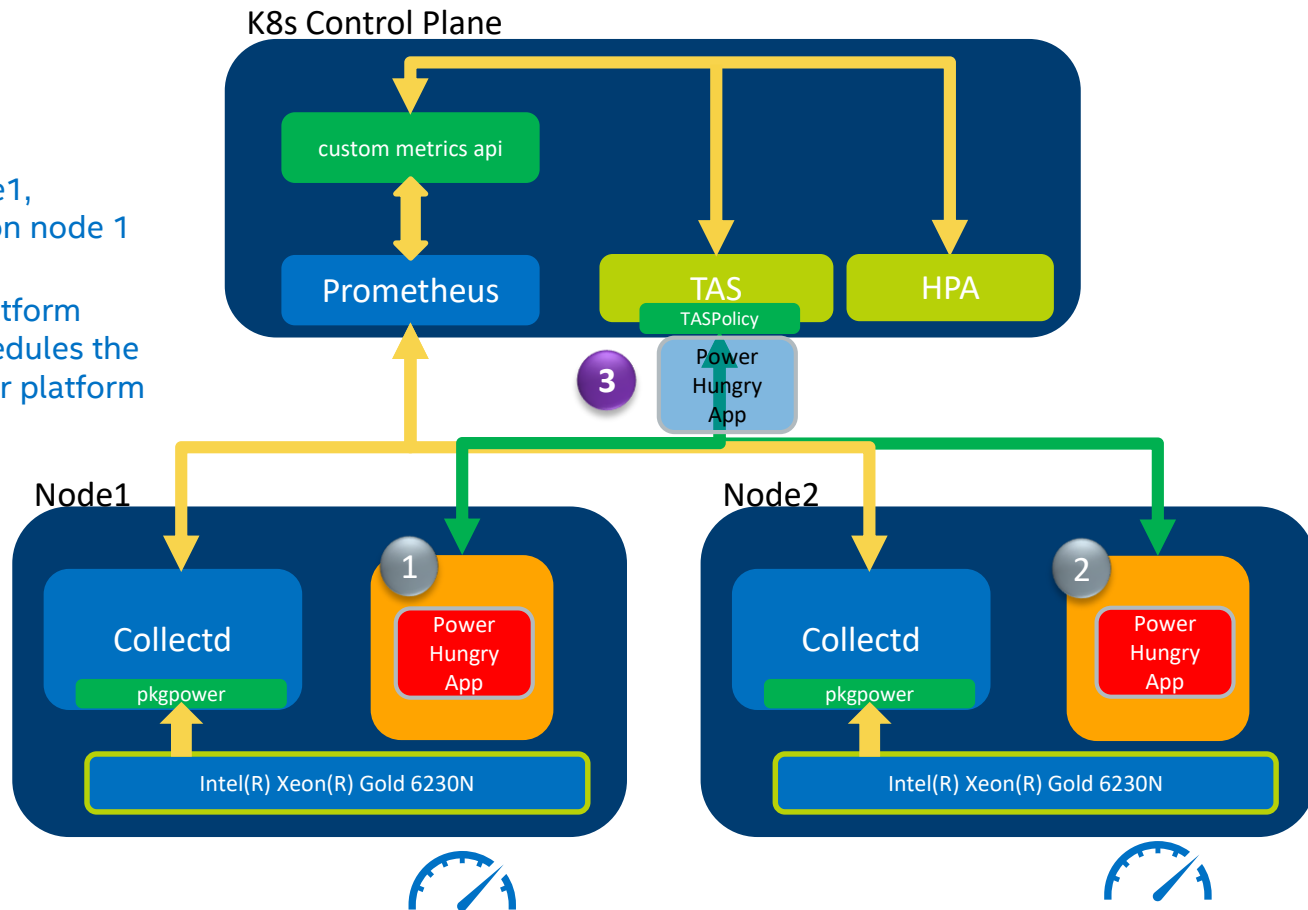
# Demo Scenario

- 1 First pod is deployed to node1, increasing the power usage on node 1
- 2 HPA is triggered from the platform overload on node1. TAS schedules the new pod on node2 with lower platform load



# Demo Scenario

- 1 First pod is deployed to node1, increasing the power usage on node 1
- 2 HPA is triggered from the platform overload on node1. TAS schedules the new pod on node2 with lower platform load
- 3 HPA is triggered again from node2 platform overload. TAS blocks scheduling of the new pod as no node meets power requirements as per policy. Pod will remain in a "pending" state until a suitable node becomes available.




**DEMO VIDEO**

- 1 First pod is deployed to node1, increasing the power usage on node 1
- 2 HPA is triggered from the platform overload on node1. TAS schedules the new pod on node2 with lower platform load
- 3 HPA is triggered again from node2 platform overload. TAS blocks scheduling of the new pod as no node meets power requirements as per policy. Pod will remain in a "pending" state until a suitable node becomes available.




# 5G App Metrics

5G workloads benefit from app specific as well as platform telemetry - Autoscaling will particularly be helped by using app telemetry that exposes specific usage measures



# 5G app specific insights used for autoscaling

- User Plane Function – a core feature of the 5G Service Based Architecture.
- Autoscaling here based on packets received using Kubernetes Horizontal Pod Autoscaler
- As traffic reaches the threshold new workloads are created and new traffic directed to them




# PROFILING TRENDS IN PLATFORM RESILIENCY



# Profiling trends and prevention: Platform Resiliency

Platform telemetry feeds into analysis and automated action to bring proactive decision making to orchestration.




# Platform Resiliency Demo

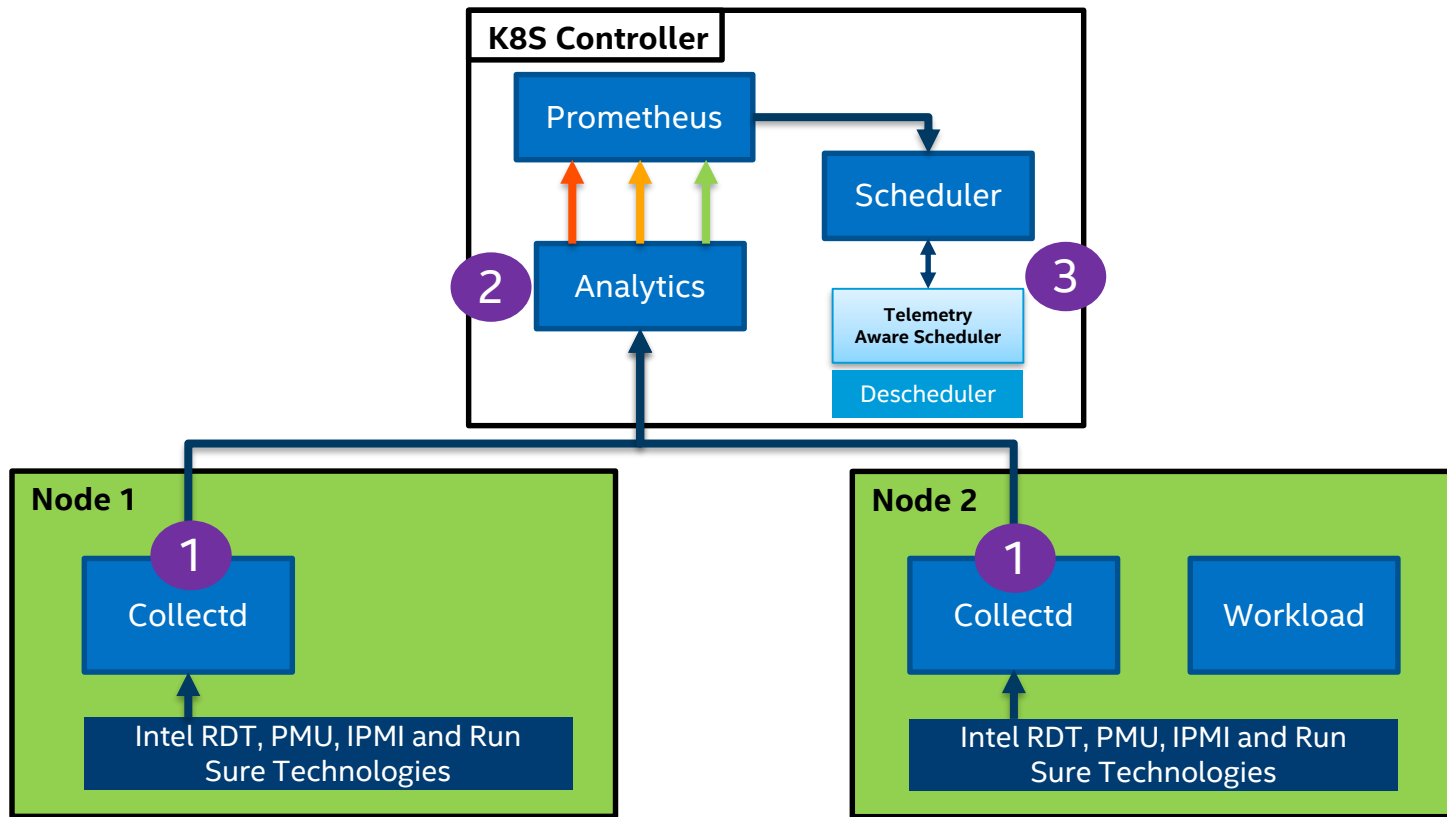
A “Host Health Indicator” is determined from multiple Platform Telemetry metrics

Host Health Indicator triggers Intelligent Scheduling decisions using Kubernetes Telemetry Aware Scheduler

Remediation actions taken at the VIM layer that have Service impacts will be indicated

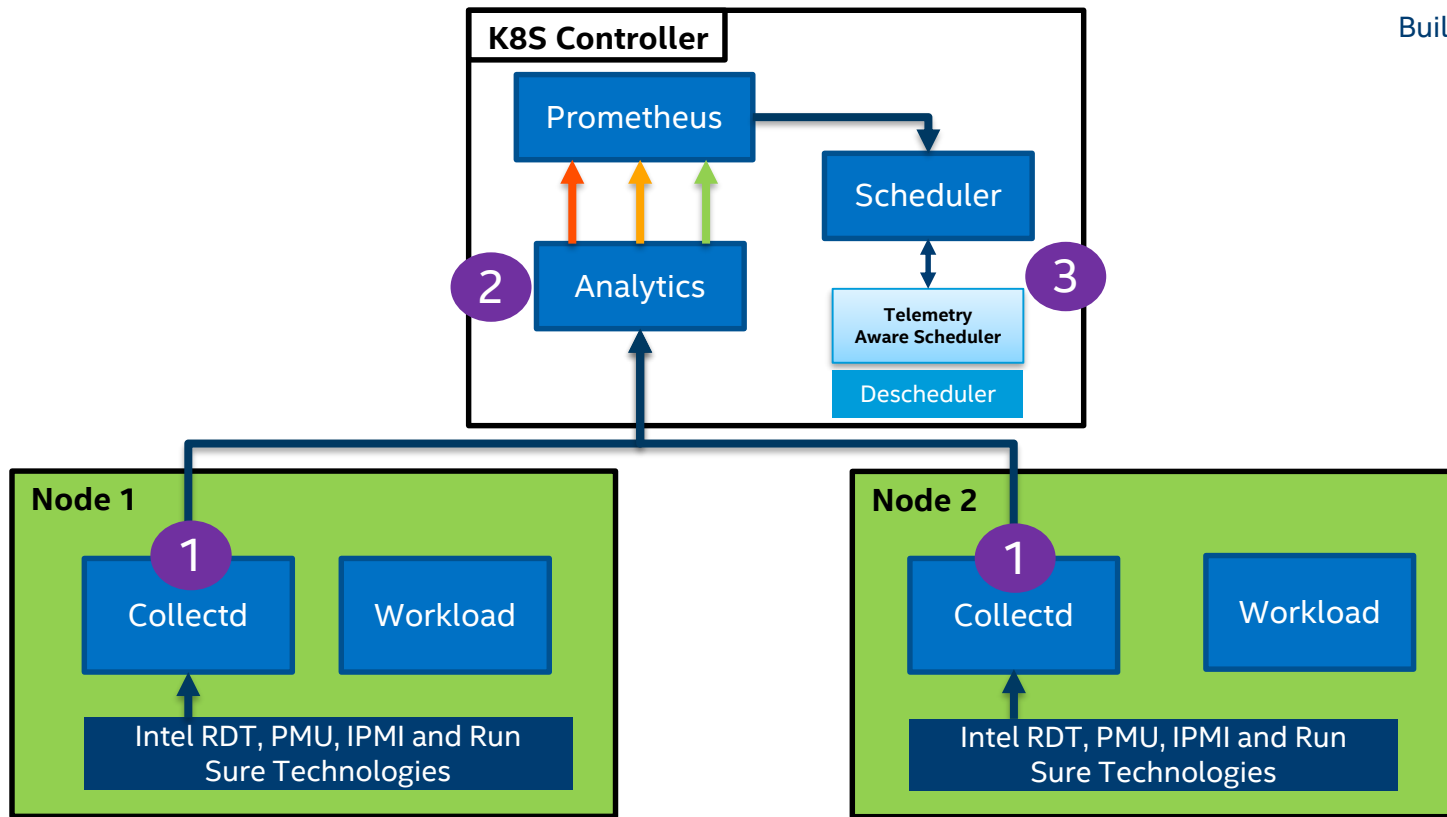


# Platform Resiliency Prototype – Minor alert Scenario



# Platform Resiliency Prototype – Critical Scenario


Build slide



**DEMO VIDEO**


# Platform Resiliency Prototype – Critical Scenario

Build slide



# Scale Efficiency with Data-Driven, Closed Loop Automation

Telemetry capabilities in Intel platforms enable intelligent, closed loop solutions that are reactive, proactive and predictive, delivering new levels of efficiency for IT and network infrastructure.



# Summary

Telemetry available from Intel servers spans a vast number of functional domains - utilization, power consumption, faults and performance and more.

Platform Telemetry, Application Telemetry, Analytics and Automated Action will underpin observability and automation pipelines for the cloud native network functions, including 5G and beyond

For further information, demo's, white papers etc:

<https://networkbuilders.intel.com/intel-technologies/network-transformation-exp-kits>

**Feel free to reach out for further information:**

[Sunku.Ranganath@intel.com/Killian.Muldoon@intel.com](mailto:Sunku.Ranganath@intel.com/Killian.Muldoon@intel.com)



**Q&A**

# FURTHER INFORMATION

# Where to Get Started?

More info on demos presented:

- <https://networkbuilders.intel.com/intel-technologies/network-transformation-exp-kits>
- <https://networkbuilders.intel.com/intel-technologies/container-experience-kits>

# OPNFV Barometer & Closed Loop Automation Working Group Links

Barometer Home: <https://wiki.opnfv.org/display/fastpath/Barometer+Home>

Metrics/Events through Barometer (not on Collectd site):

<https://wiki.opnfv.org/display/fastpath/Collectd+Metrics+and+Events#CollectdMetricsandEvents-Metrics>

Barometer “One-click” install:

<https://wiki.opnfv.org/display/fastpath/One+Click+Install+of+Barometer+Containers>

OPNFV CLA WG:

<https://wiki.opnfv.org/pages/viewpage.action?spaceKey=CLAWG&title=Closed+Loop+Automation+Working+Group+Home>

# Further Demo Resources

Visit our "**Network Transformation**" page at <https://networkbuilders.intel.com/network-technologies/network-transformation-exp-kits>

Under Automation section you will find:

Power Savings demo: <https://networkbuilders.intel.com/closed-loop-platform-automation-power-savings-demo>

Host Health/Platform Resiliency white paper: <https://builders.intel.com/docs/networkbuilders/closed-loop-platform-automation-service-healing-and-platform-resilience.pdf>

Host Health/Platform Resiliency demo video: <https://networkbuilders.intel.com/closed-loop-automation-telemetry-aware-scheduler-for-service-healing-and-platform-resilience-demo>

Some background information on the work we are doing around managing resources (cache and memory bandwidth in this case) to provide optimum VNF performance:

<https://builders.intel.com/docs/networkbuilders/intel-platform-service-assurance-platform-policy-enabling-resource-management-white-paper.pdf>

# Further Resources

Learn more from these helpful sites:

<https://networkbuilders.intel.com/network-technologies/serviceassurance>

<https://wiki.opnfv.org/display/fastpath/Barometer+Home>

<https://wiki.openstack.org/wiki/Telemetry>

<https://01.org/openstack/blogs/2015/openstack-enhanced-platform-awareness-white-paper>

<https://github.com/intel/telemetry-aware-scheduling>

<https://github.com/intel/telemetry-aware-scheduling/docs/power>

# Collectd 101 Materials

- Collectd 101
  - <https://wiki.opnfv.org/display/fastpath/Collectd+101>
- Write simple read plugin
  - <https://wiki.opnfv.org/display/fastpath/Collectd+how+to+implement+a+simple+plugin>