

# INTRODUCTION TO CLOUD SYSTEMS

---

Lecture 7 – Dynamic scheduling –  
Kubernetes Telemetry Aware Scheduling

# Kubernetes TAS

Telemetry Aware Scheduling (TAS) - Automated Workload Optimization with Kubernetes (K8s\*)

# Kubernetes TAS

TAS is a K8s\* add-on that consumes platform metrics and makes intelligent scheduling decisions based on operator-defined policies. TAS can be used to direct workloads to specific nodes based on up-to-date platform telemetry during pre-runtime. For example, when the pod is initially started on the cluster and during runtime, the workload is already up and running on the platform.

TAS implements a K8s scheduler extender, meaning it modifies the decisions made by the core K8s system, rather than replacing them entirely. The standard resource checking, workload awareness, and other scheduling rules remain unchanged, and all standard pods are compatible with a cluster running TAS. Telemetry is collected using the K8s custom metrics Application Programming Interface (API), a standard metric interface in the orchestrator.

# Kubernetes TAS

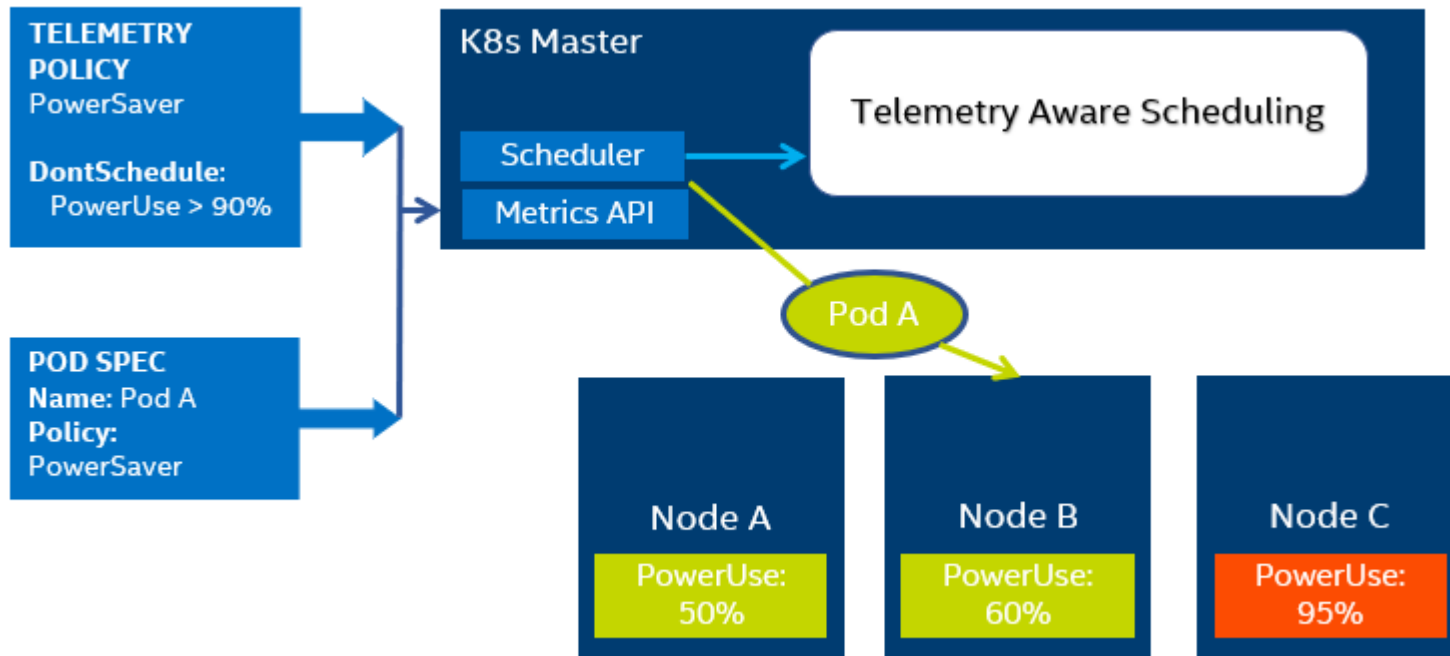
Placement decisions are modified by TAS in three ways:

- Filtering
- Prioritization
- Descheduling

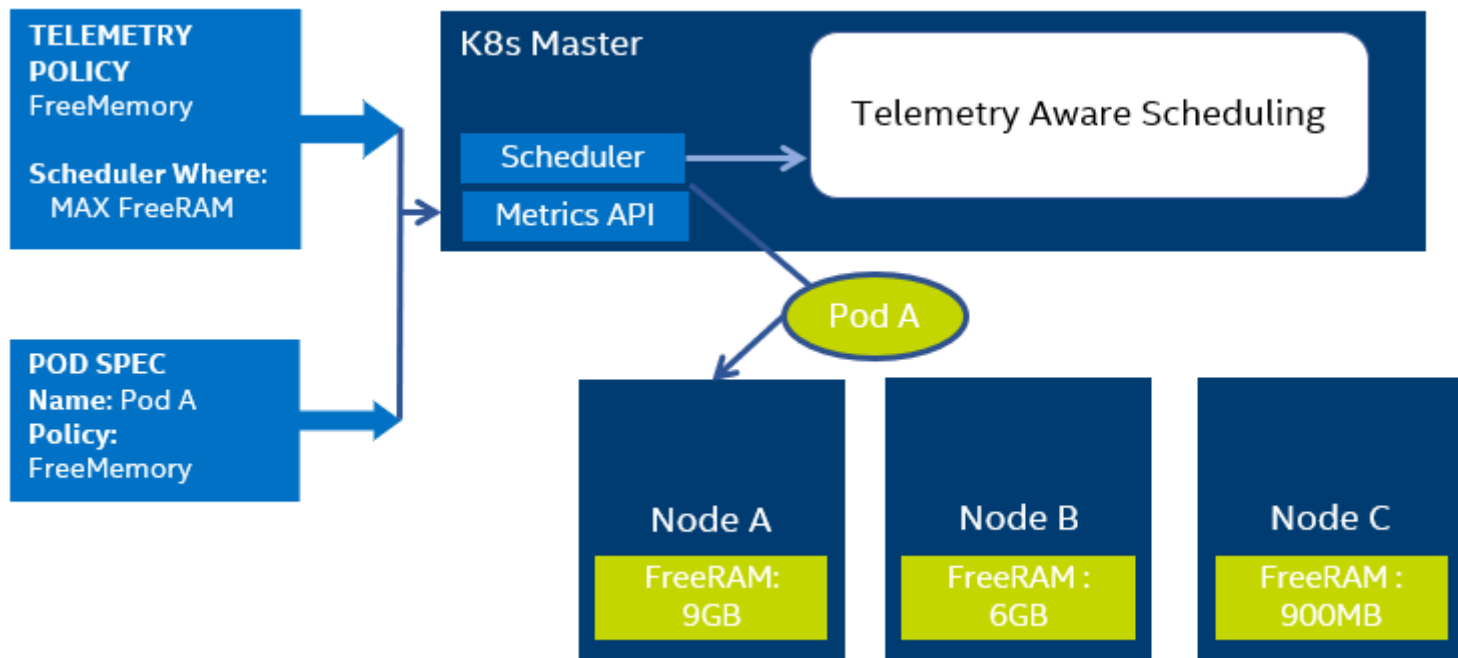
Each one can be customized as an individual telemetry scheduling policy, and all three can be combined into a single telemetry scheduling policy for holistic workload placement automation.

A telemetry scheduling policy, which is defined as a K8s custom resource, contains a set of rules like those described below. Each policy may serve more than one type of workload, and more than one policy can be active in a cluster at any time. Each scheduling decision made by TAS is based on a single policy

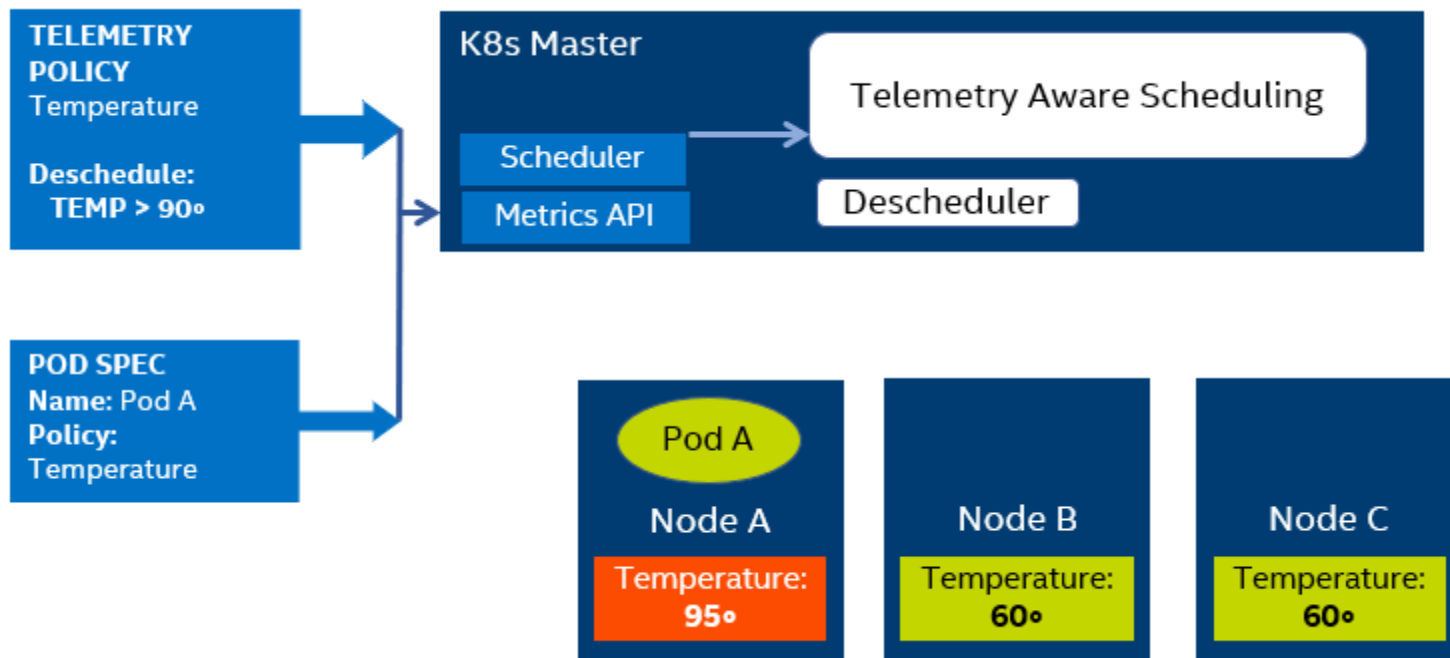
# Kubernetes TAS



# Kubernetes TAS



# Kubernetes TAS



# Kubernetes TAS

TAS allows arbitrary, user defined rules to be put in place in order to impact scheduling in a K8s cluster. Leveraging the K8s descheduler, it can evict workloads that are breaking some rules in order to have it replaced on a more suitable node.

In a modern cloud computing cluster, there is a torrent of data that only certain subject matter experts know how to interpret and act upon. In scheduling workloads, operators know on which worker nodes a workload may perform better based on up to date utilization metrics. Likewise, certain telemetry values, or combinations of values, can be recognized as signs that a node has some serious problem that is interfering with workload operation.



# Kubernetes TAS

The TAS policy system allows these insights to influence the scheduling and lifecycle placement process— turning implicit personal knowledge into formal, actionable information. As in the above use cases, the modified scheduling recommendations automate actions that would otherwise be done by operations technicians and increase resource utilization across the whole cluster.

TAS is an open source community project and the software is available at:  
<https://github.com/intel/telemetry-aware-scheduling>

Deployment is done directly on K8s\*, using Helm\* charts provided at the same repository.