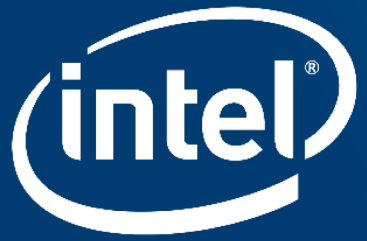

EDGE & E2E ORCHESTRATION & MACHINE LEARNING

Joe Butler , Sharon Ruane – Intel Labs Europe.

May 11, 2018.

Orchestrating apps (content) and network.

Application
And Content
Complexity

& demand for
network
performance.

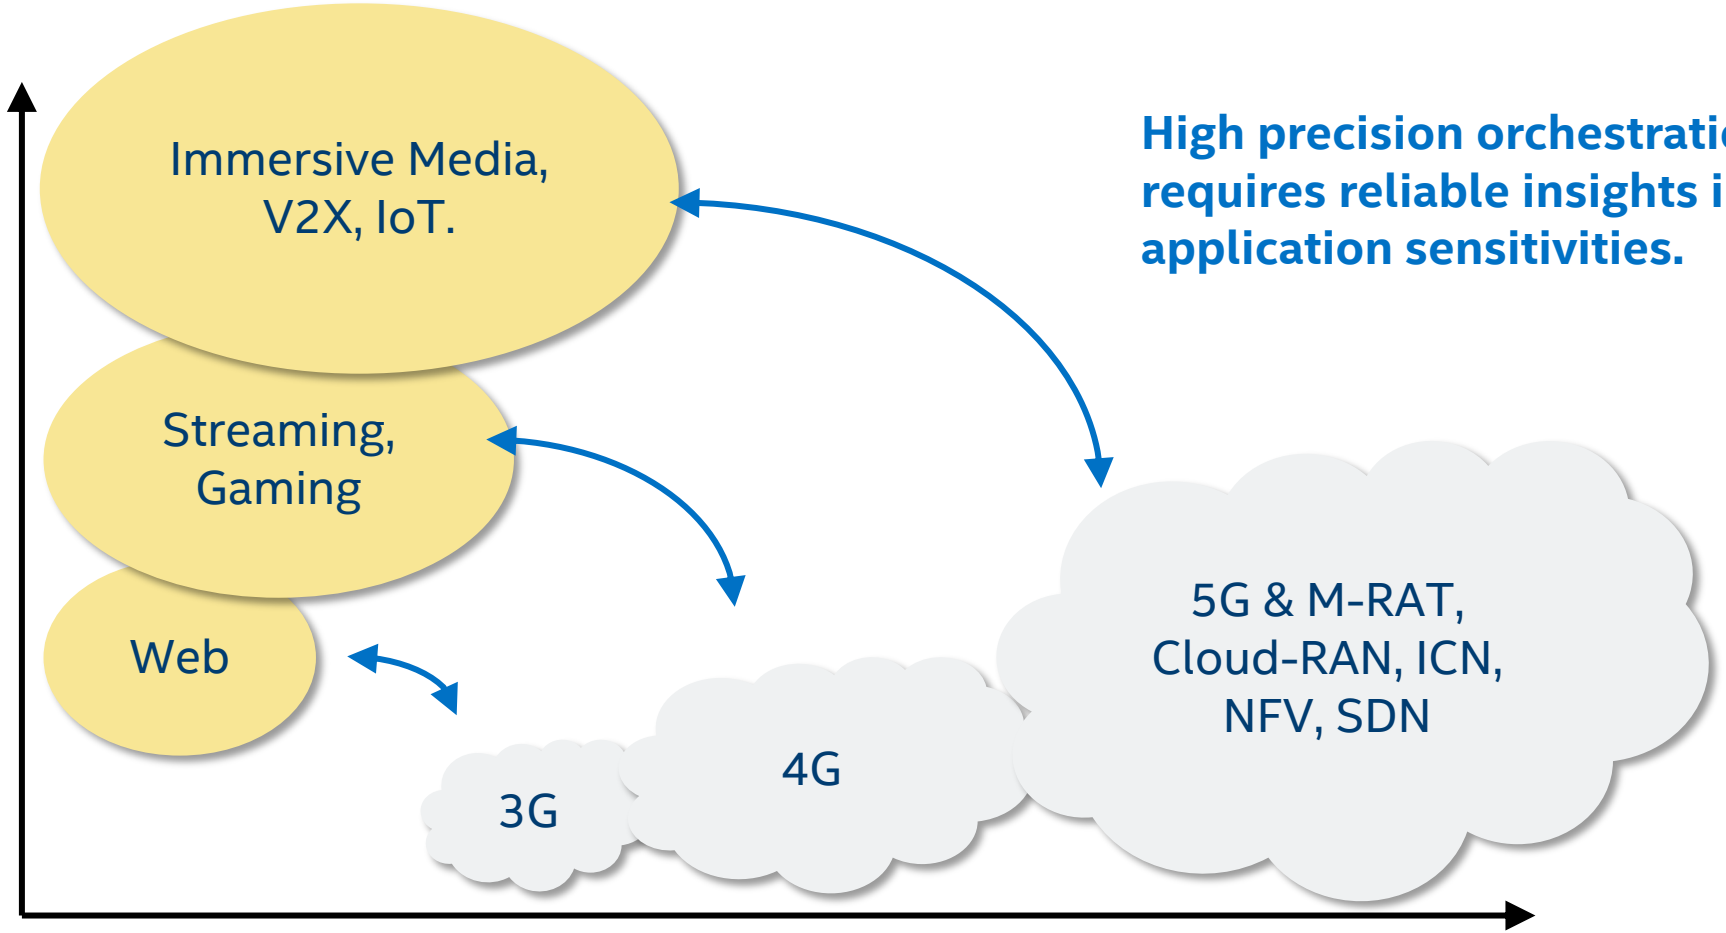

High precision orchestration requires reliable insights into application sensitivities.

Network diversity, complexity,
& need for efficiency, flexibility and automation.

Orchestration for E2E & Network Edge

- Stringent SLAs.
- Sensitive workloads.
- Complex chaining.
- Dense packing.
- Distribution.
- Heterogeneity.
- Scale.

→ Service Chain

Placement

↓ Infrastructure

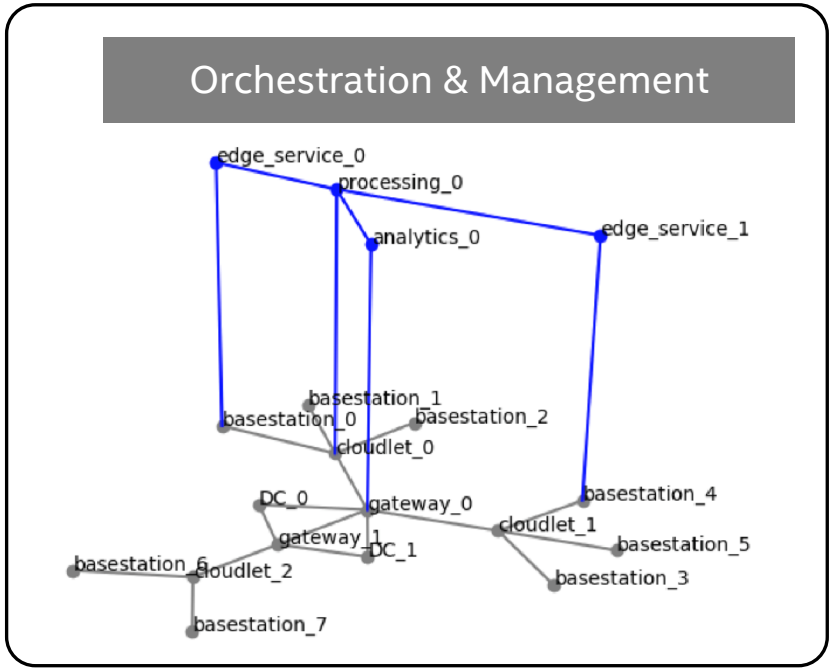

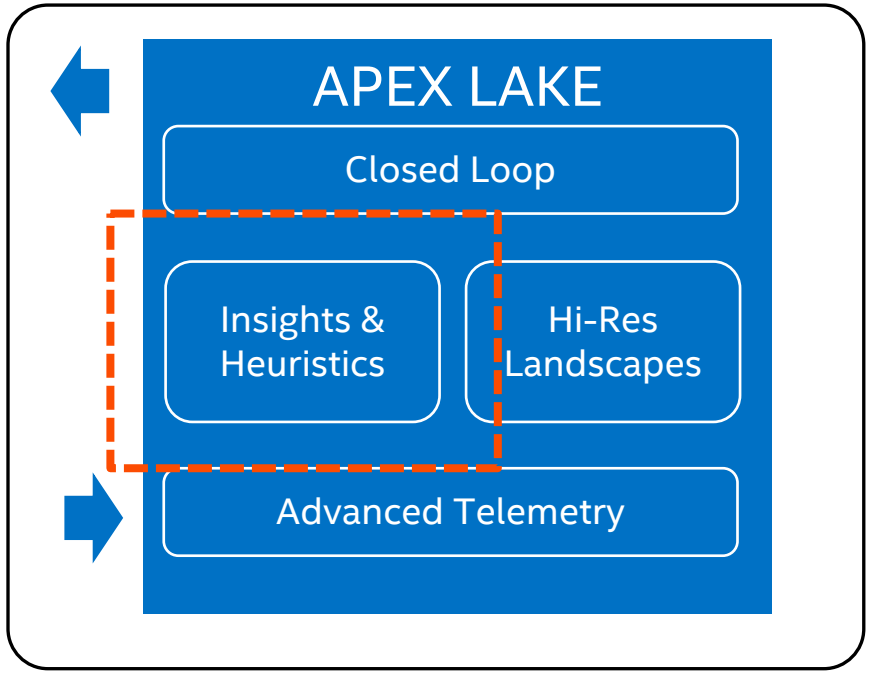

Apex Lake: Intel Labs Orchestration Research Platform.

Automating generation of usable insights for intelligent, high-precision orchestration.

Virtualised network & edge.

- **Fine grained view of resources and services.**
- **Resource constraints at edge.**
- **Heterogeneity.**
- **Densification.**
- **Dynamic consumption patterns.**

Virtualisation & Cloudification brings network as-a-service.

FP7 Mobile Cloud Networking (2013-2016) was a collaborative research project focused cloud-style delivery of Telco network services.

Mobile Cloud Networking

Cloud Computing disrupts IT consumption and operations

- On-Demand, Self-Service, Elasticity, Pay-as-you-Go, Remote Access.
- Infrastructure- Platform- Software-as-a-Service (IaaS/PaaS/SaaS) -> EPCaaS.

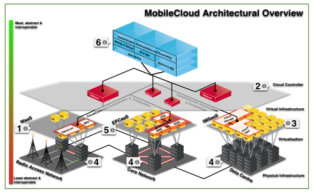

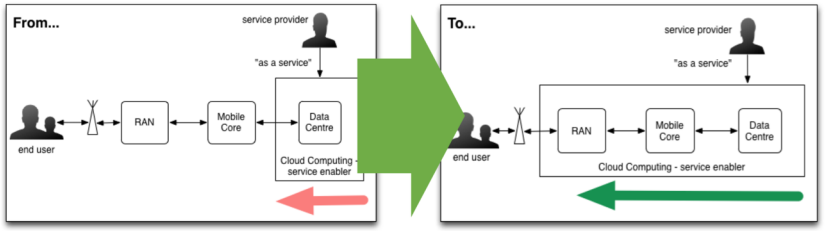

Small cell + cloud enablement extends on-demand to edge.

H2020 5G Essence is a collaborative research project focused on optimised deployment of 5G services across multi-tenant, cloud-enabled, Small Cell edge infrastructure..

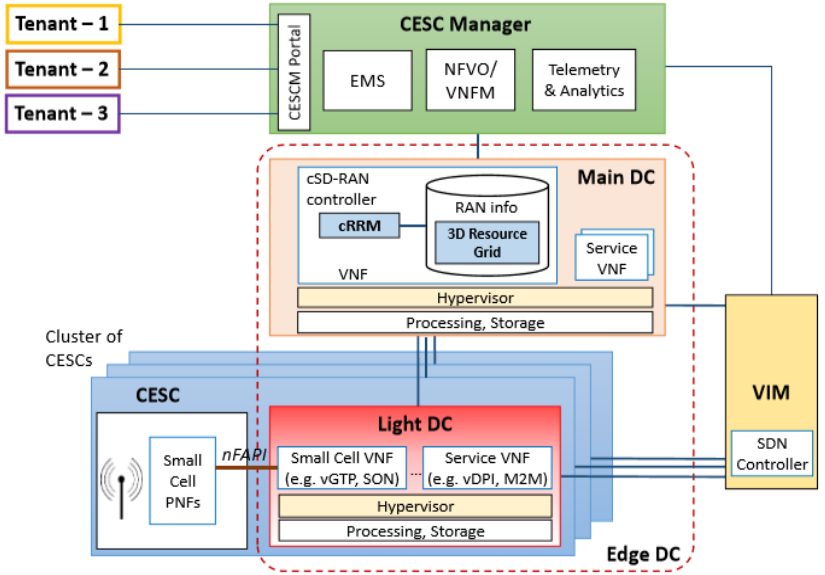

Service Function Chaining.

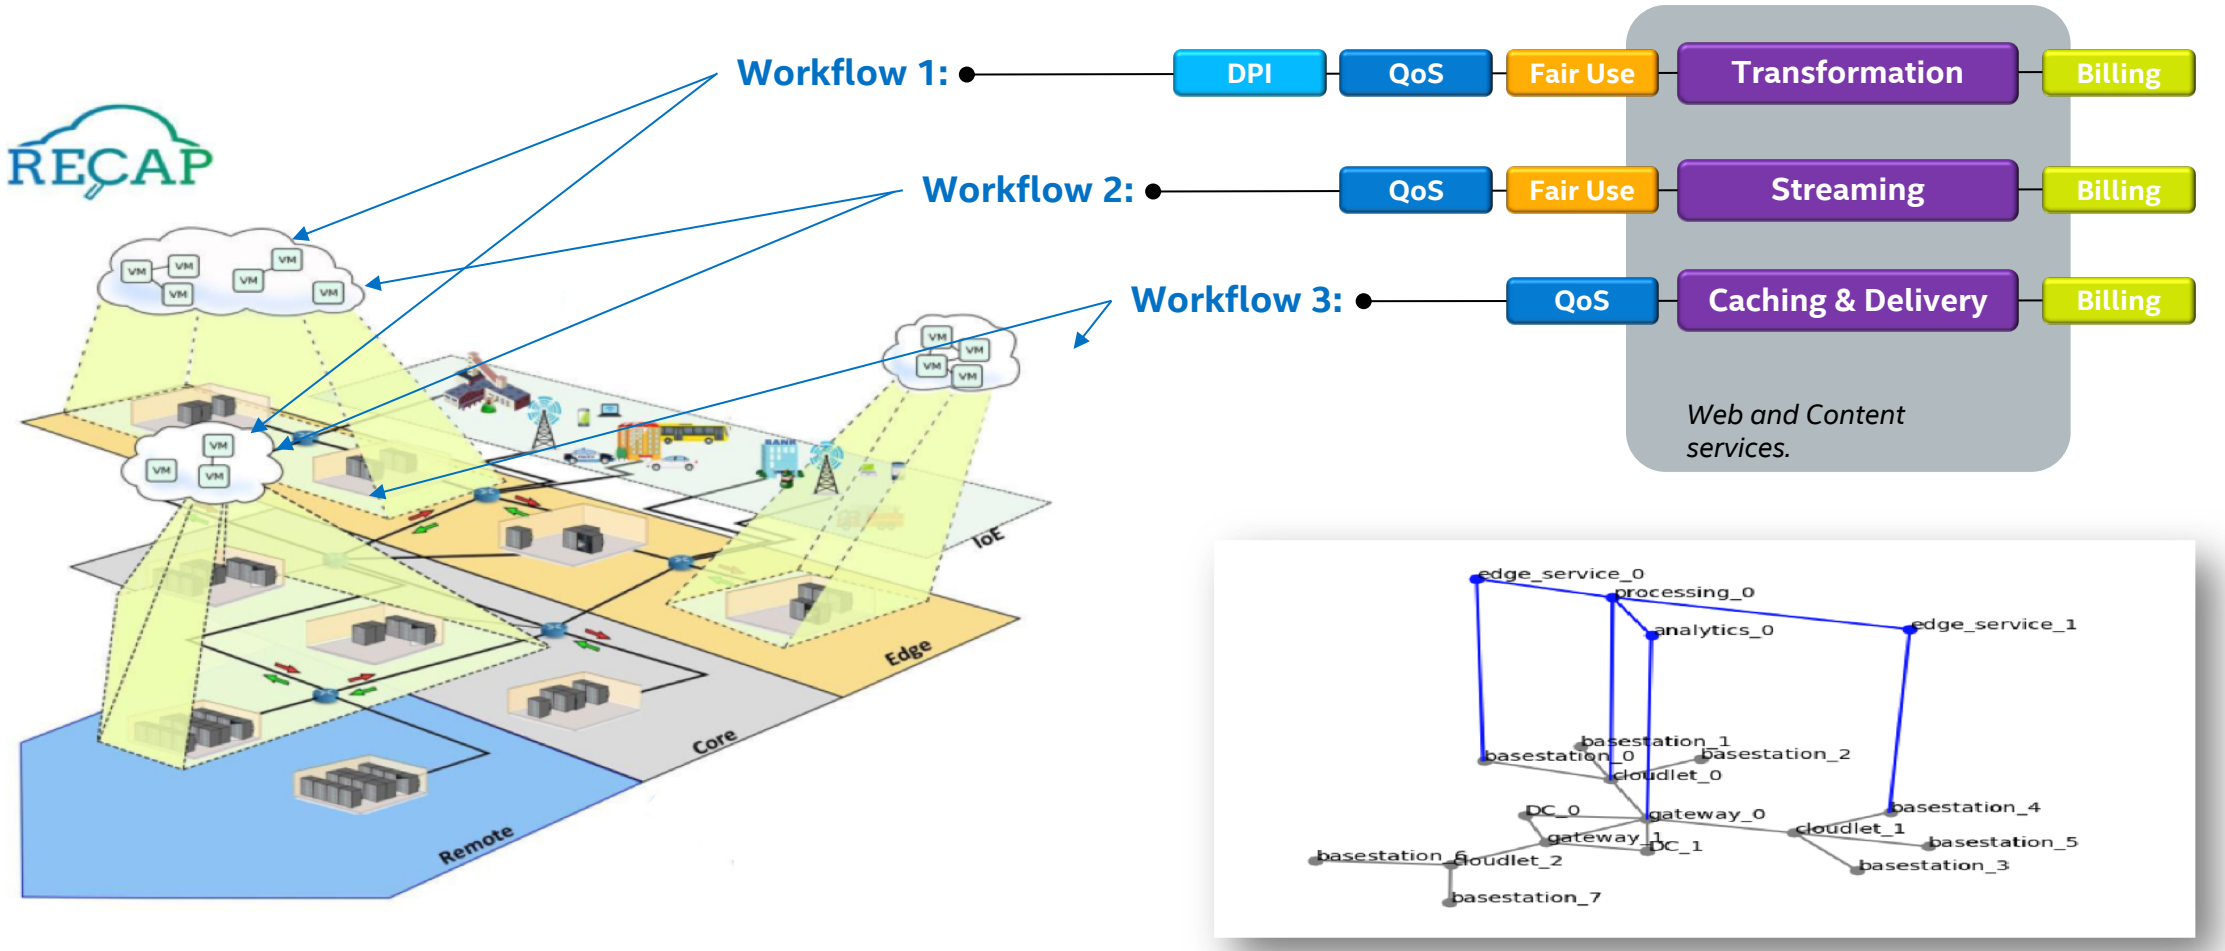

H2020 RECAP is a collaborative research project focused on capacity planning of heterogeneous Cloud to the Edge via infrastructure optimization, modelling, simulation and automated self-adaption

Service chain components and distributed resource placement options as a directed acyclic graph..

Toward Network Functions -aaS.

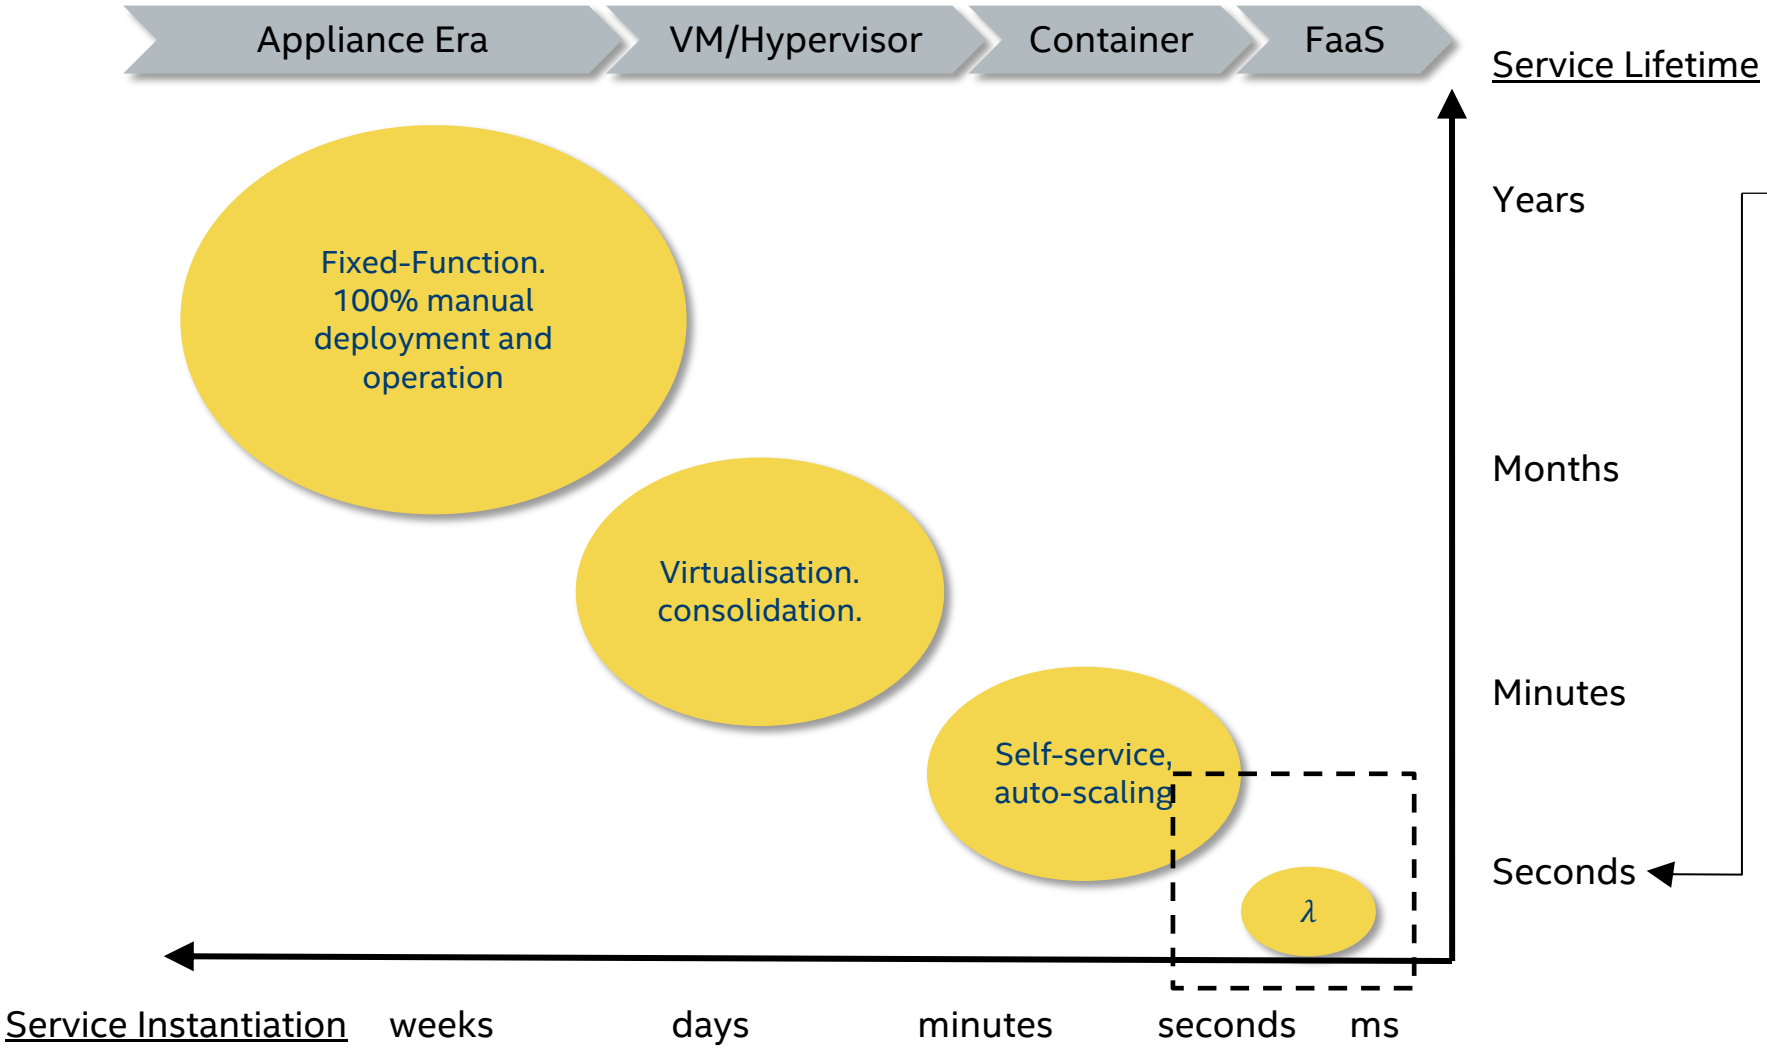

Next stage of progress.

- Micro-services with millisecond instantiation times, sub-second billing.
- Tight packing.
- Fully Automated.

Emerging edge use cases.

Industry

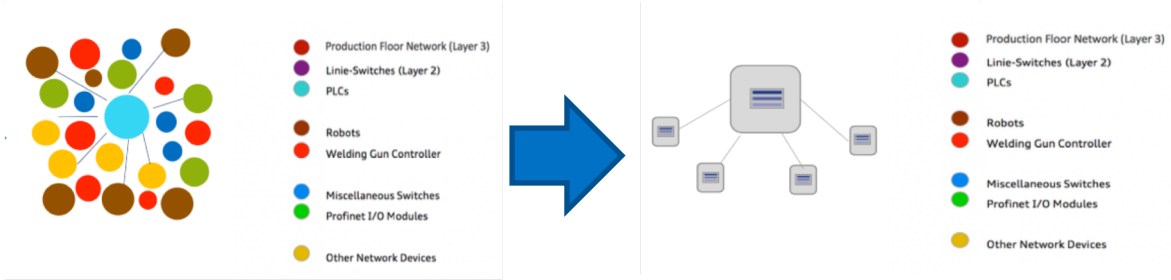

Manufacturing Automation Cell today

Manufacturing Automation Cell of the future

Transport / V2X

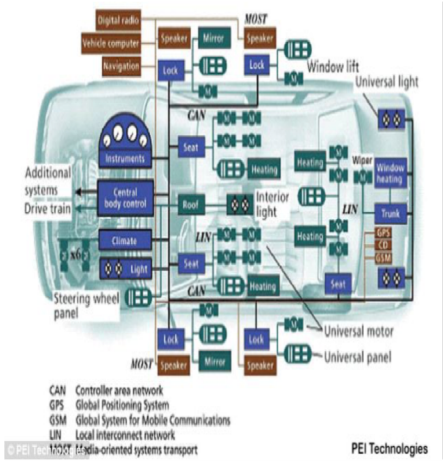

Control contained within the car

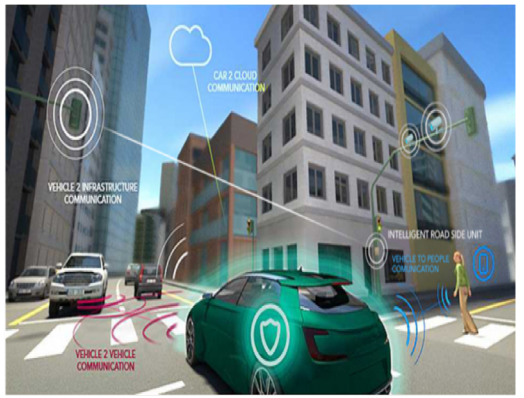

Control augmented with the environment and infrastructure

-> dynamic resource topologies.
-> hard constraints on data processing.
-> additional constraints: data timeliness and provenance, security and privacy.

Resource Allocation: Prescription.

Descriptor metadata

Goals:
- Service Performance
- Manageability.

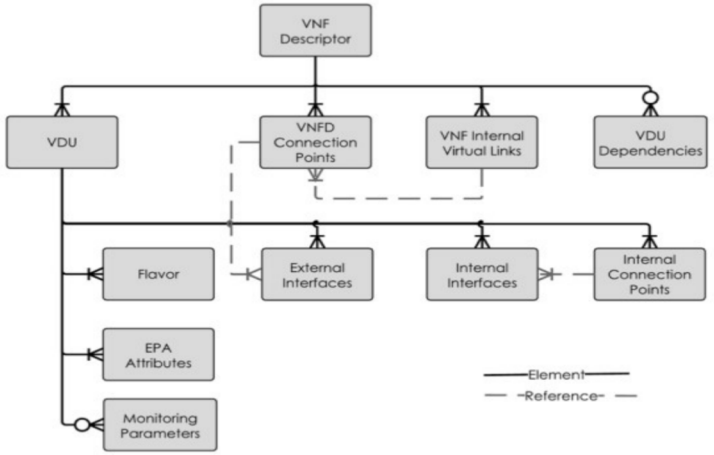

VNFD high level object model, source.
https://osm.etsi.org/wikipub/images/2/26/OSM_R2_Information_Model.pdf

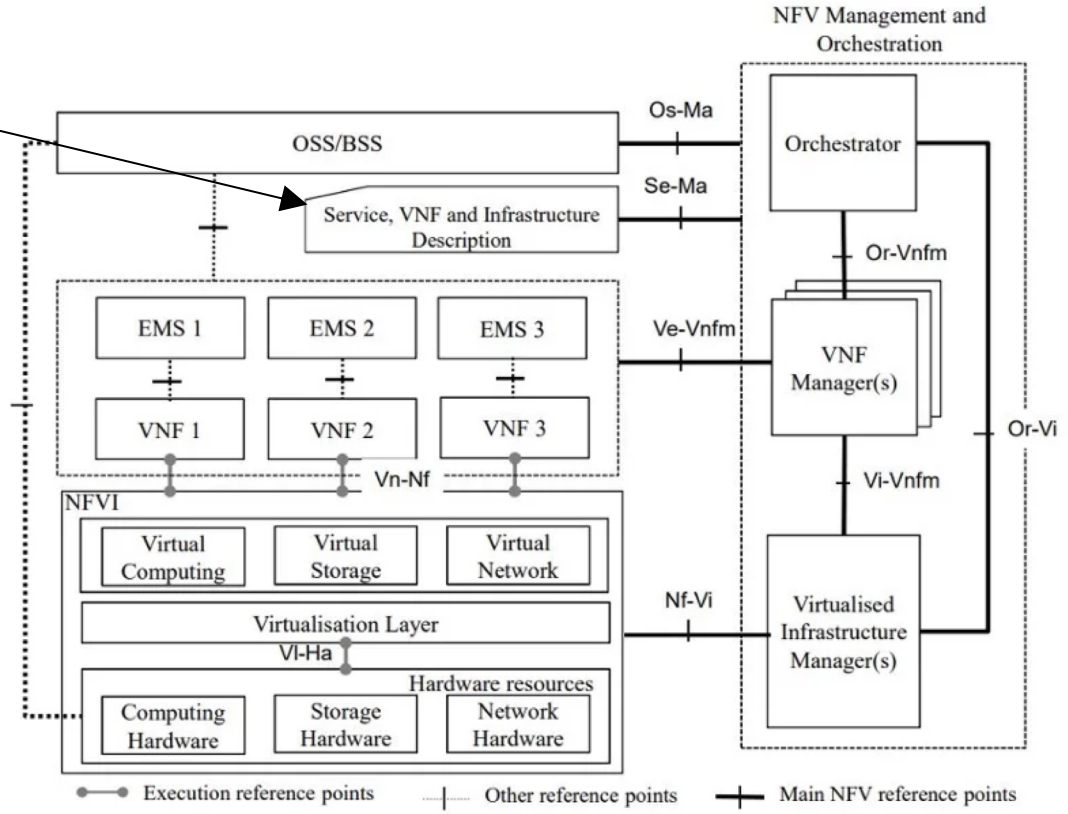

NFV reference architecture diagram.
Source: ETSI specifications documents circa 2013..

Resource Allocation: Learning.

Monitoring + Analytics + Automation

Goals:
- Tight Packing,
- Platform Awareness,
- Resource Affinities,
- KPI Mapping / SLA Compliance.

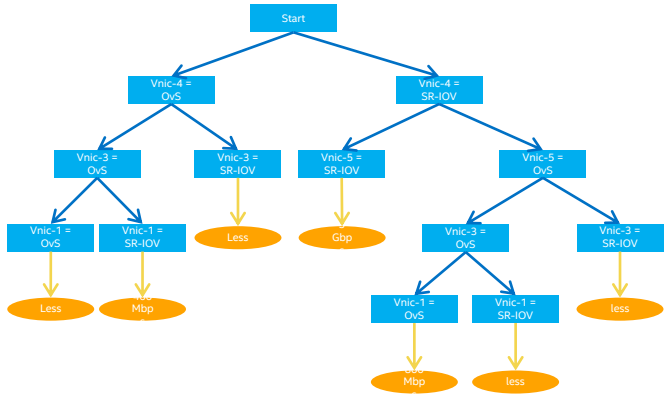

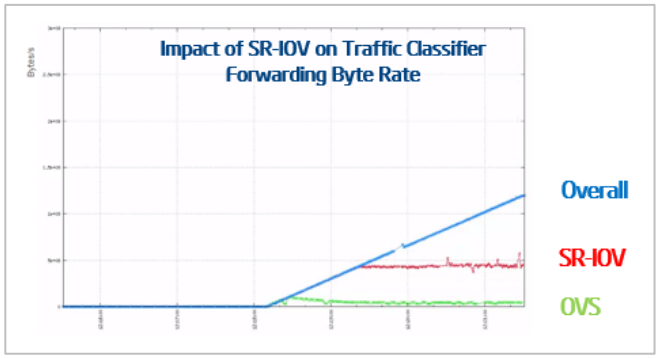

Decision Tree automatically guides Orchestrator to select SR-IOV enabled NIC based on service characterization..

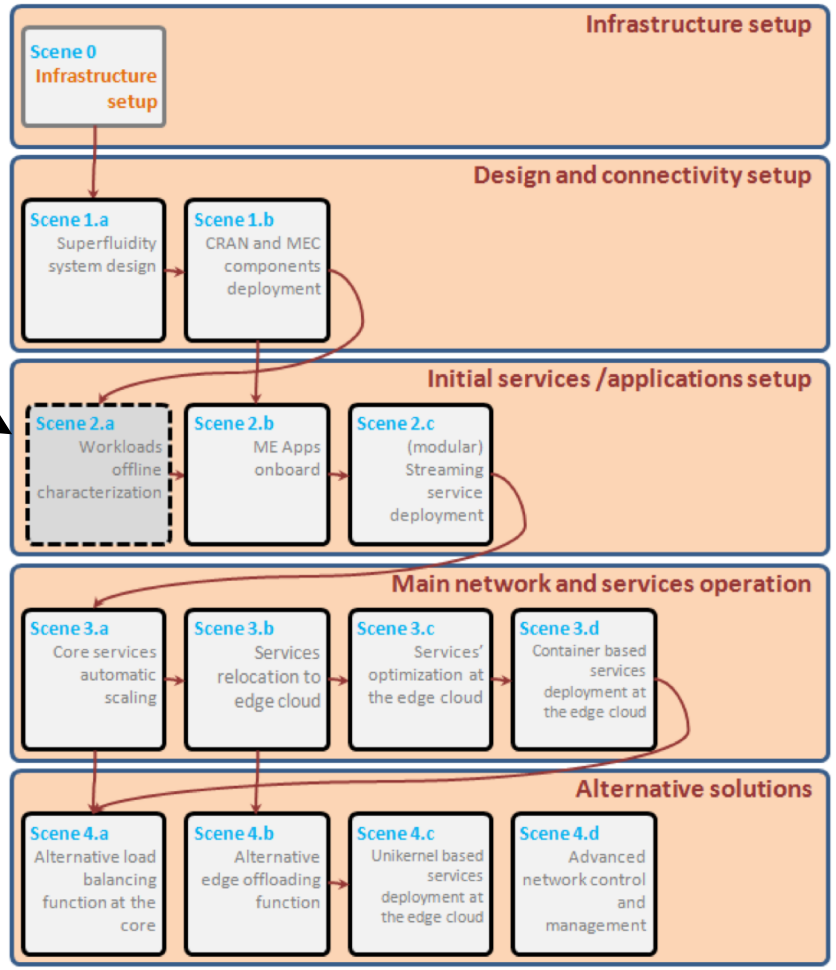

H2020 Superfluidity is a collaborative research project targeting a super-fluid, cloud-native, converged edge system.

Machine Learning in context.

Goal – automating precise and efficient placement & adjustment, driven by service objectives.

- *What are the usable key metrics, attributes, and expressions of constraints and objectives ?*

- *How can we automatically discover these in context ?*

Approach:

- *Full-stack, adaptive telemetry,*
- *Background and foreground analytics loops,*
- *Workload fingerprinting,*
- *Automation.*

Techniques:

- *Utility Theory,*
- *Cost Functions,*
- *Machine Learning,*
- *Evolutionary Algorithms,*
- *Hybrid Algorithms.*

Metrics:

- *Service performance,*
- *Resource allocation rightsizing,*
- *Resource utilization,*
- *Accuracy of predictions,*
- *Confidence of maintaining SLA.*

Example: dynamic placement.

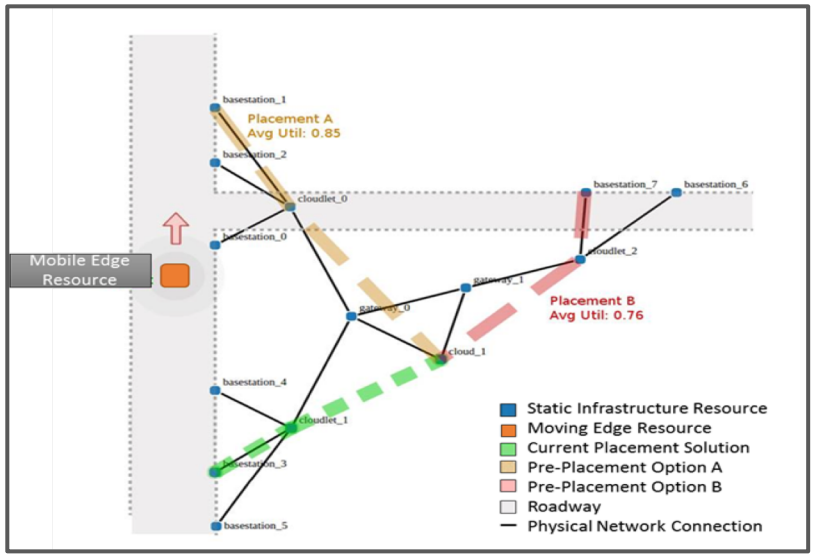

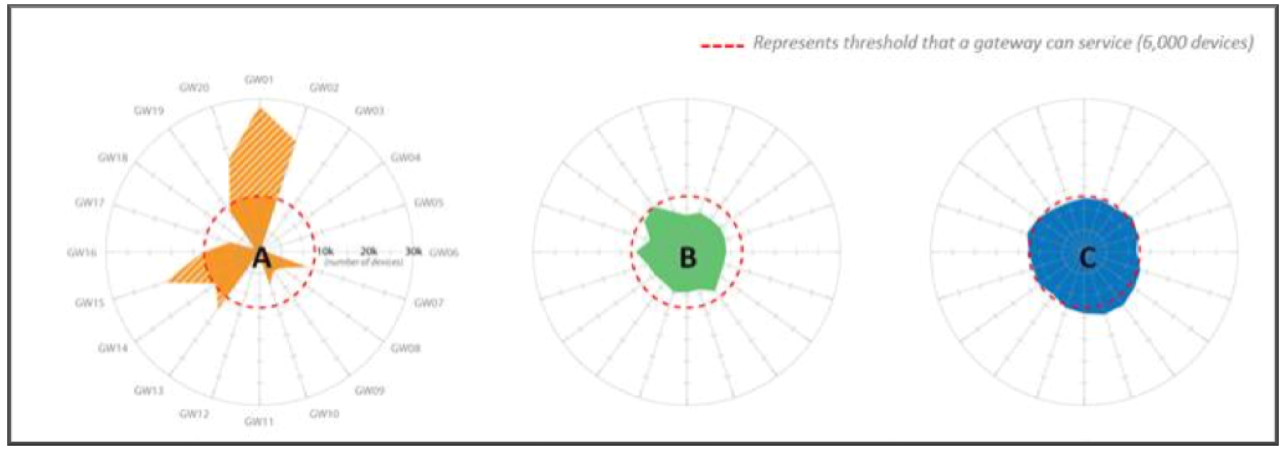

Test bed emulation of mobile urban environment.

20 gateways with overlapping. coverage, >100k mobile endpoints.

Nearest fit / current utilisation as baseline.

Hybrid multi-attribute utility function + evolutionary algorithm formulates capacity and endpoint trajectories.

A: Baseline

Max scale16k endpoints,

6/20 gateways over-subscribed,

14ms to compute placement.

B: 100k endpoints

Headroom and utilization balanced,

7ms compute placement.

C: 140k endpoints

Gateway saturation,

7ms compute placement.

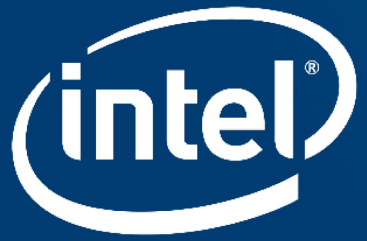

AUTOMATED PERFORMANCE OPTIMIZATION VIA MACHINE LEARNING

Sharon Ruane, Joe Butler

PREDICTION OF HETEROGENEOUS WORKLOAD BEHAVIOR:

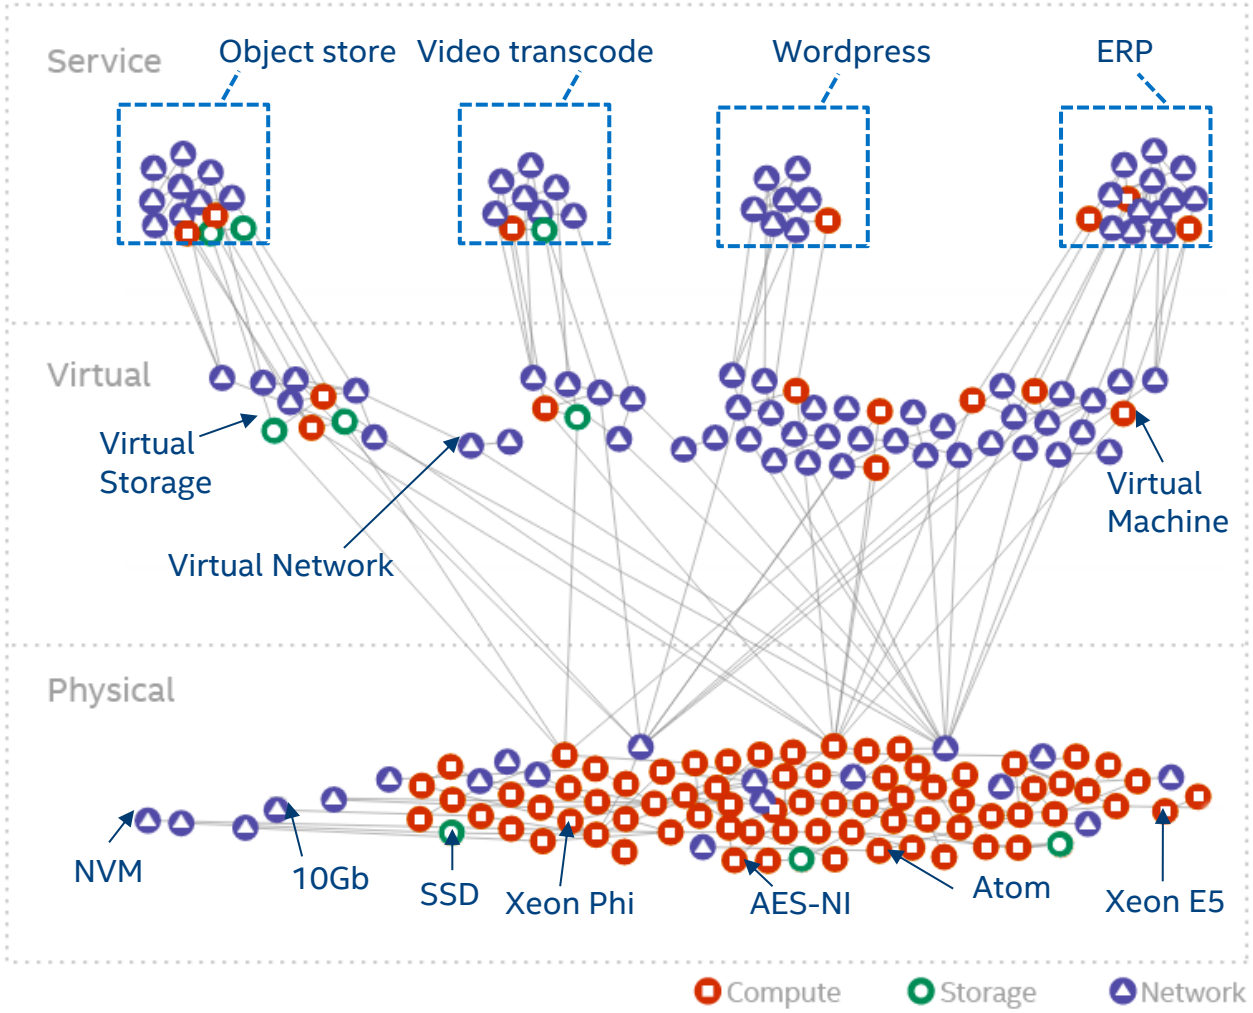

PREDICTION OF HETEROGENEOUS WORKLOAD BEHAVIOR:

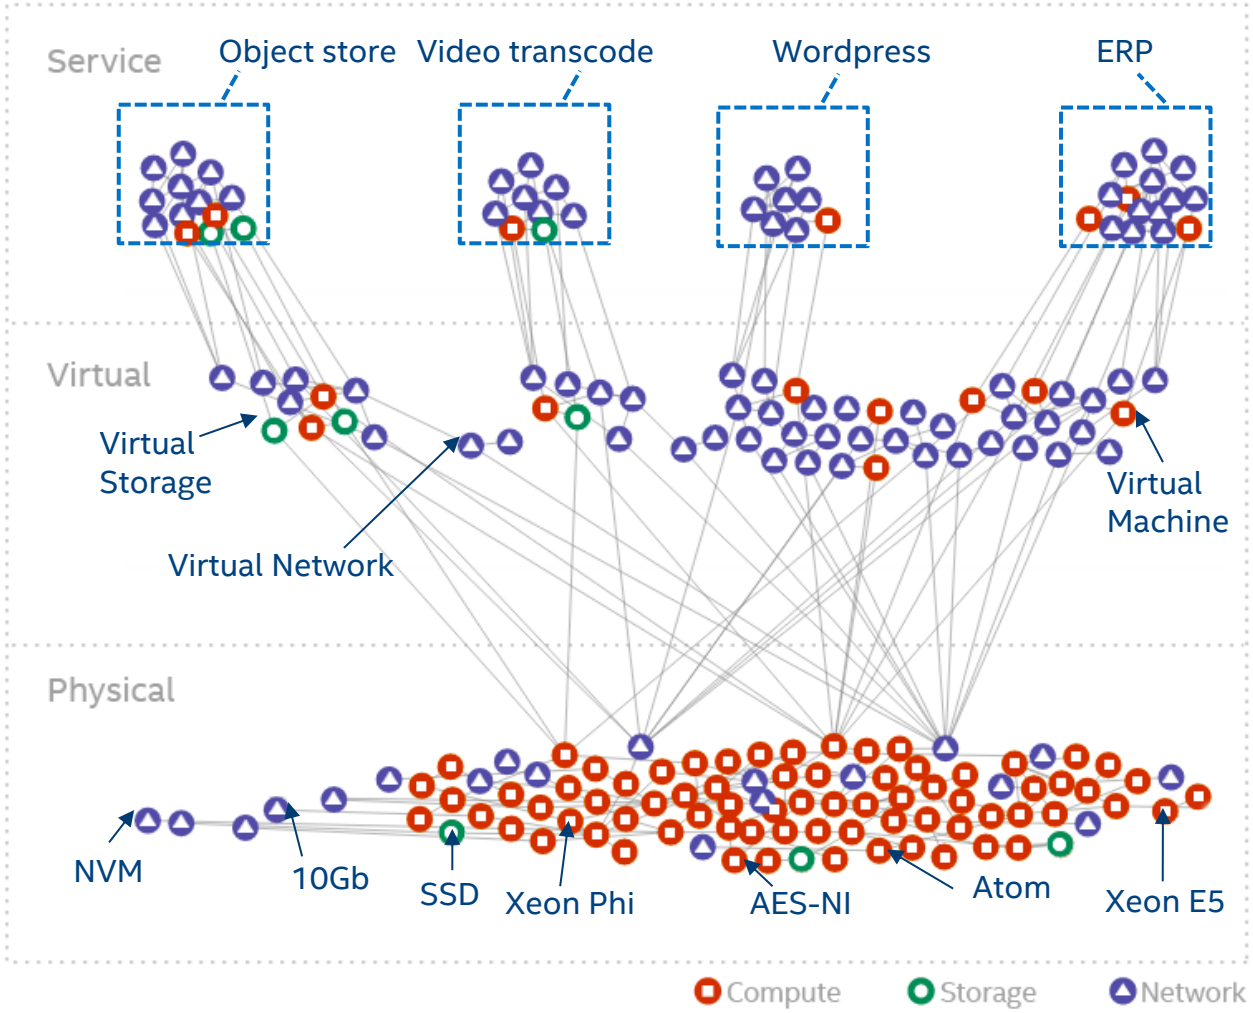

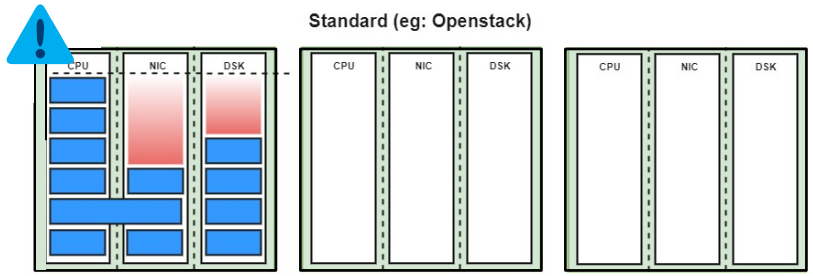

PREDICTION OF HETEROGENEOUS WORKLOAD BEHAVIOR:

Research Questions:

Can we predict the behavior of an incoming workload if it's placed on a resource which is already in use?
Using this, can we pack the workloads to maximize utilization of all resources, while avoiding overload?

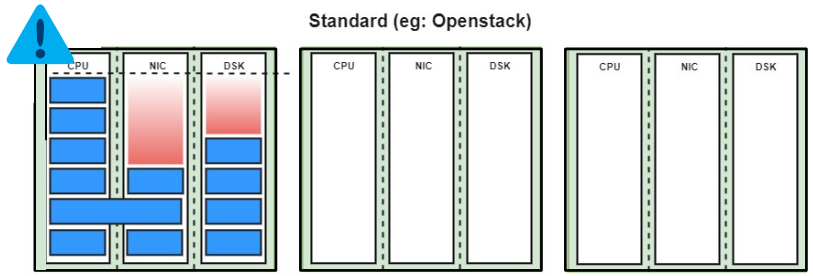

PREDICTION OF HETEROGENEOUS WORKLOAD BEHAVIOR:

Research Questions:

Can we predict the behavior of an incoming workload if it's placed on a resource which is already in use?
Using this, can we pack the workloads to maximize utilization of all resources, while avoiding overload?

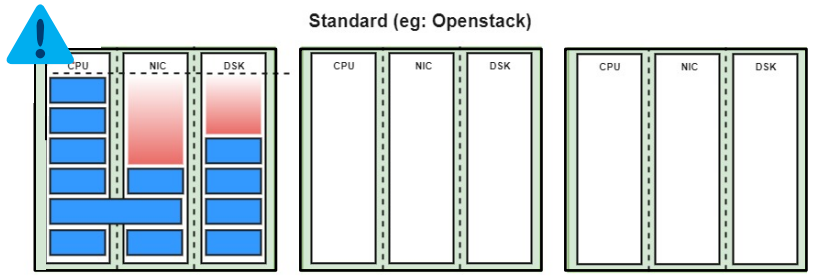

Experimentation:

CPU Utilization per workload instance

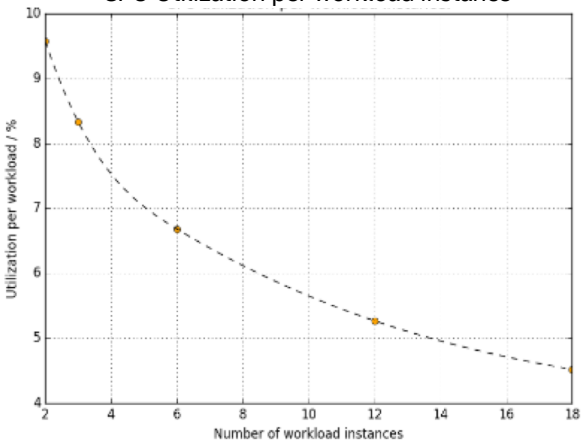

Subsystem Interference

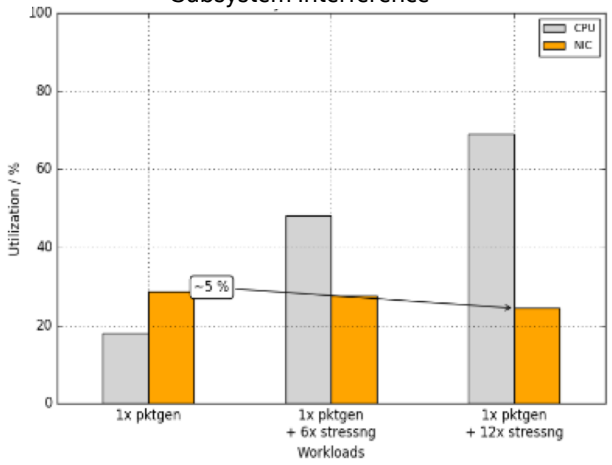

Different systems exhibit varied saturation patterns

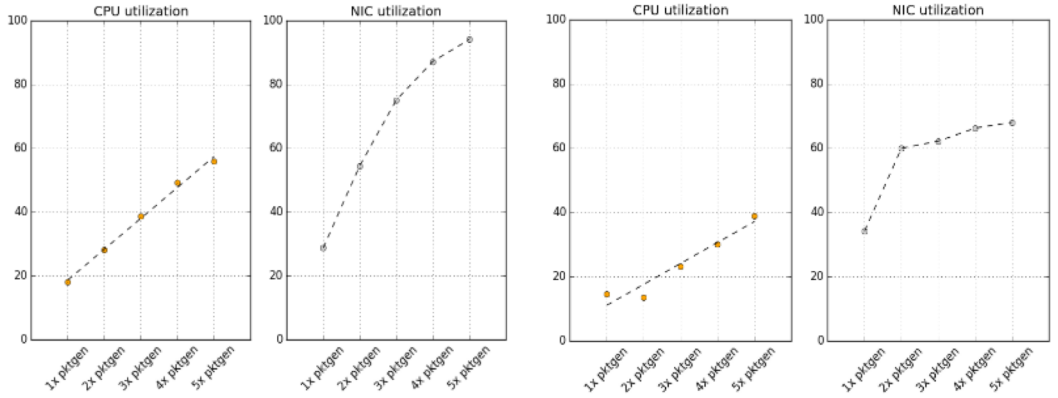

PREDICTION OF HETEROGENEOUS WORKLOAD BEHAVIOR:

Research Questions:

Can we predict the behavior of an incoming workload if it's placed on a resource which is already in use?
Using this, can we pack the workloads to maximize utilization of all resources, while avoiding overload?

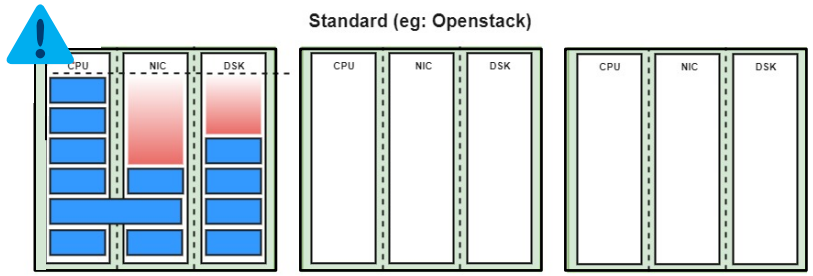

Experimentation:

	Stress-ng	Stress	Memcach	Pkg-gen	iozone
Max	24	12	10	3	3

PREDICTION OF HETEROGENEOUS WORKLOAD BEHAVIOR:

Research Questions:

Can we predict the behavior of an incoming workload if it's placed on a resource which is already in use?
Using this, can we pack the workloads to maximize utilization of all resources, while avoiding overload?

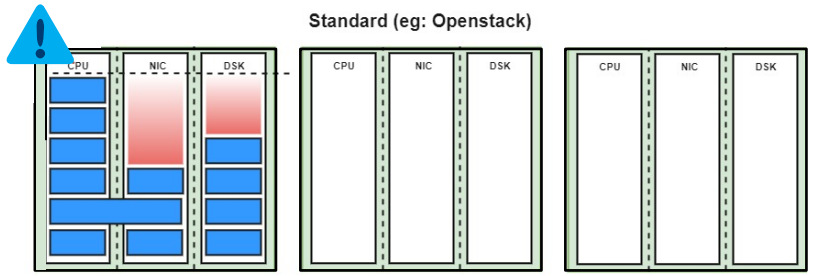

Experimentation:

	Stress-ng	Stress	Memcach	Pkg-gen	iozone
Max	24	12	10	3	3

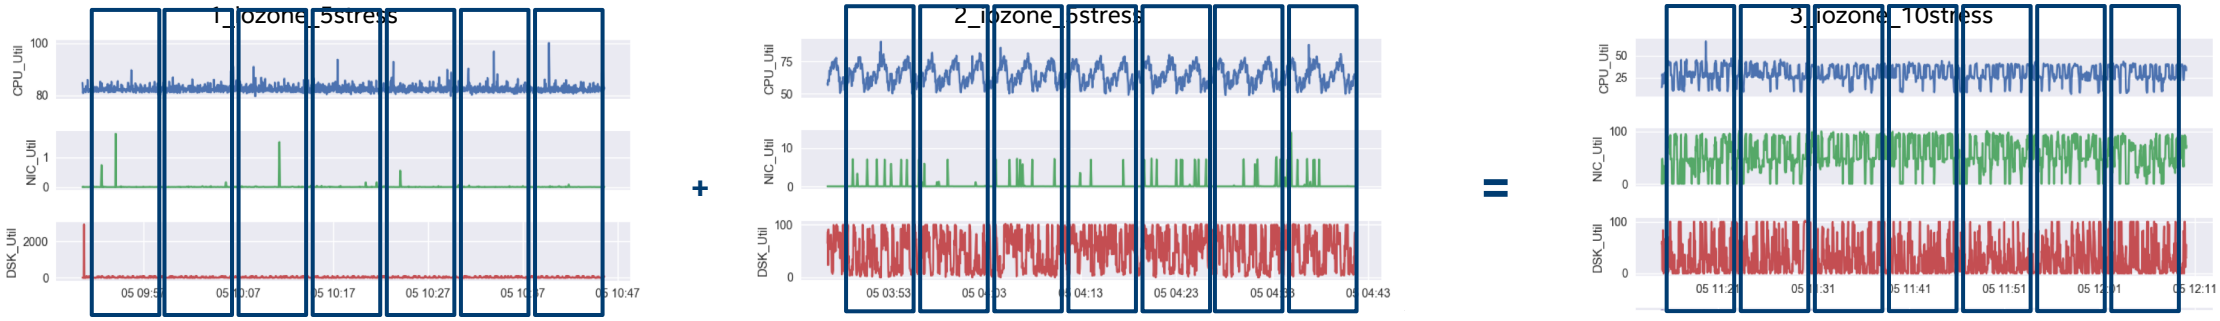

PREDICTION OF HETEROGENEOUS WORKLOAD BEHAVIOR:

Research Questions:

Can we predict the behavior of an incoming workload if it's placed on a resource which is already in use?
Using this, can we pack the workloads to maximize utilization of all resources, while avoiding overload?

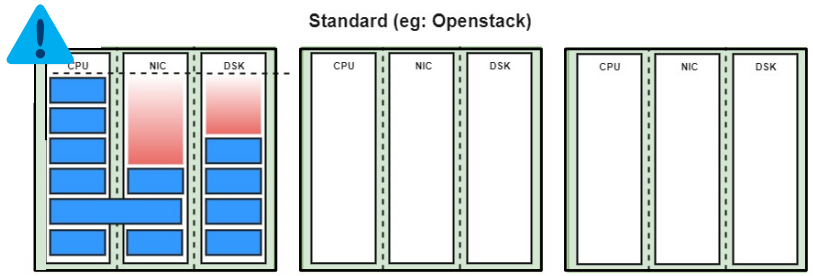

Experimentation:

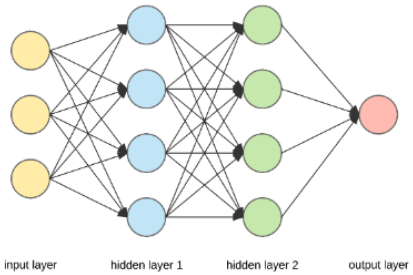

Model
99% CPU
99% NIC
93% DSK

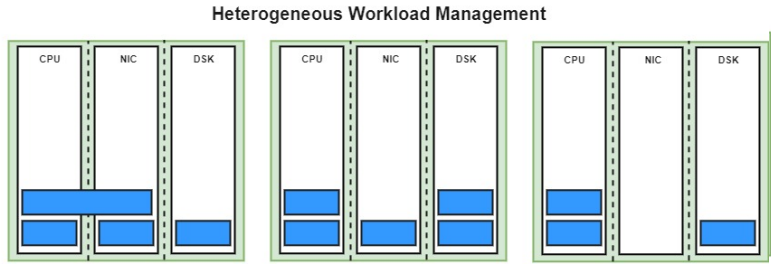

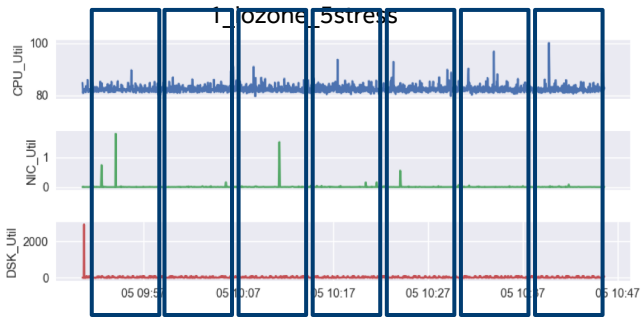

+

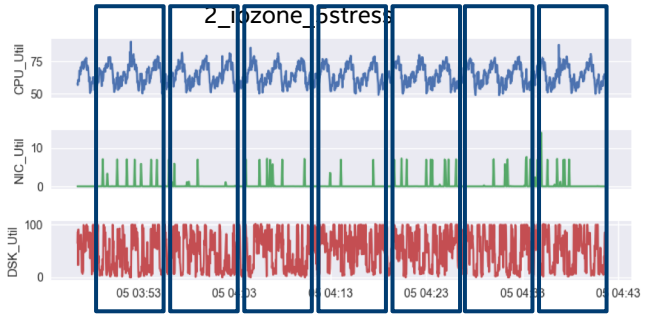

=

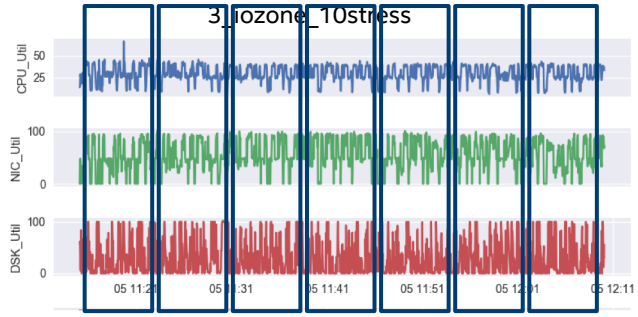

PREDICTION OF HETEROGENEOUS WORKLOAD BEHAVIOR:

Research Questions:

Can we predict the behavior of an incoming workload if it's placed on a resource which is already in use?
Using this, can we pack the workloads to maximize utilization of all resources, while avoiding overload?

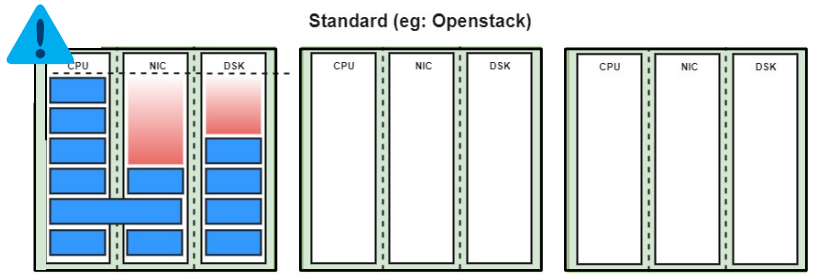

Experimentation:

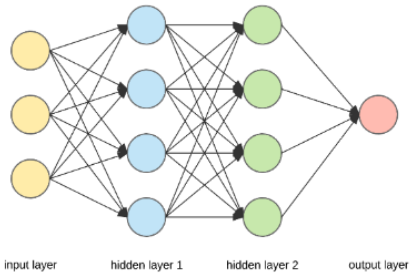

Model
99% CPU
99% NIC
93% DSK

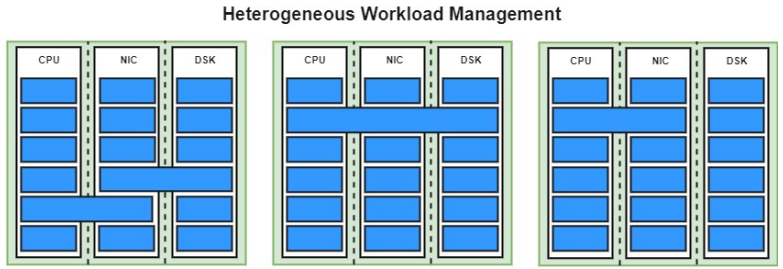

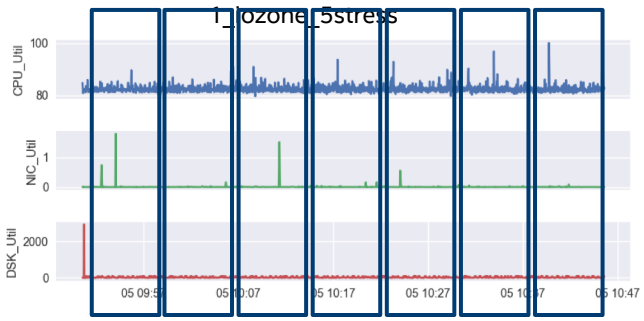

+

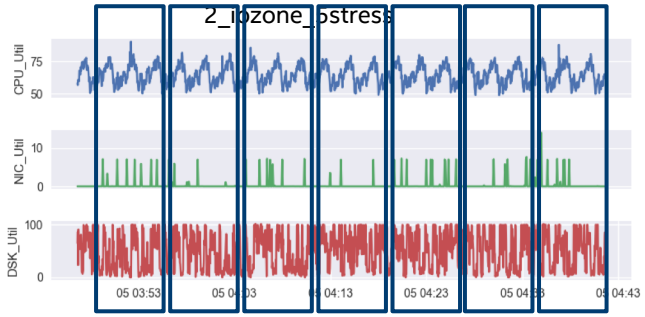

=

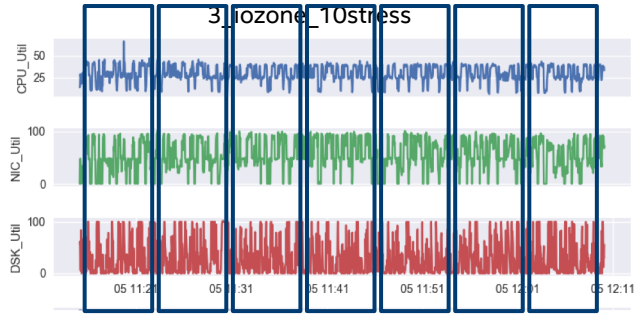

INTEGRATED ML APPROACH TO ORCHESTRATION

KPI identification

```
net_interfaces_network_utilization
net_interfaces_receive_bytes
net_interfaces_transmit_packets
net_receive_bytes
proc_stat_meminfo_slab
proc_stat_meminfo_sunreclaim
net_interfaces_receive_packets
net_tx_mb
net_transmit_bytes
net_tx
```

```
proc_stat_meminfo_active_file
proc_stat_meminfo_sunreclaim
proc_stat_meminfo_cached
proc_stat_meminfo_active
proc_stat_meminfo_active_anon
proc_stat_meminfo_anonpages
proc_stat_meminfo_memavailable
proc_stat_meminfo_anonhugepages
proc_stat_meminfo_sreclaimable
proc_stat_meminfo_free
```

Efficient resource allocation

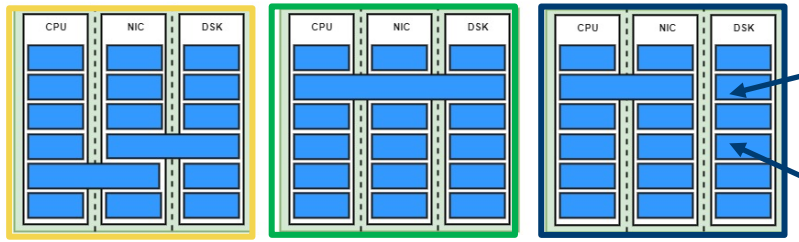

Energy management

Selected cores active
Rest powered down

Optimal placement of chained workloads

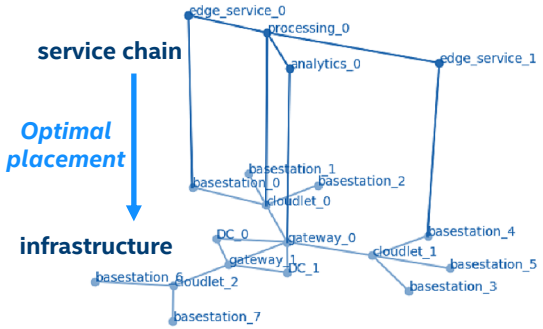

SLA assurance

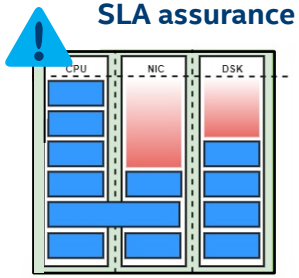

INTEGRATED ML APPROACH TO ORCHESTRATION

INTEGRATED ML APPROACH TO ORCHESTRATION

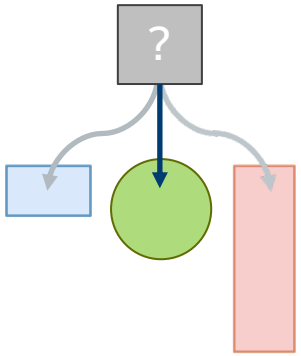

INTEGRATED ML APPROACH TO ORCHESTRATION

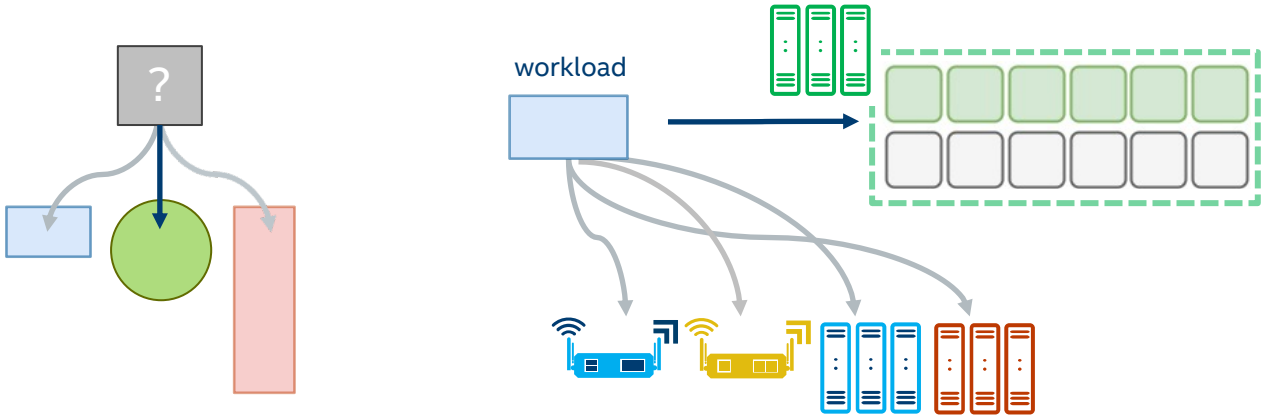

INTEGRATED ML APPROACH TO ORCHESTRATION

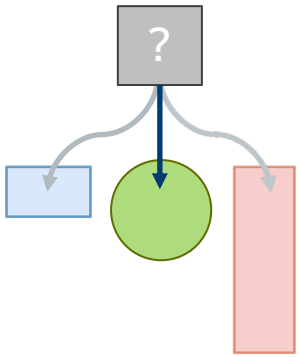

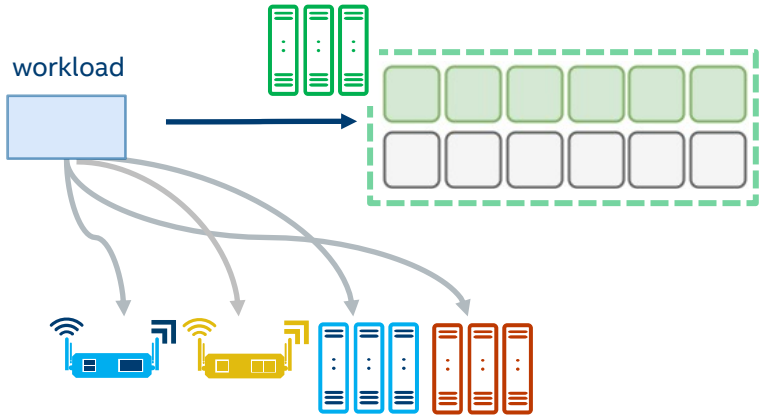

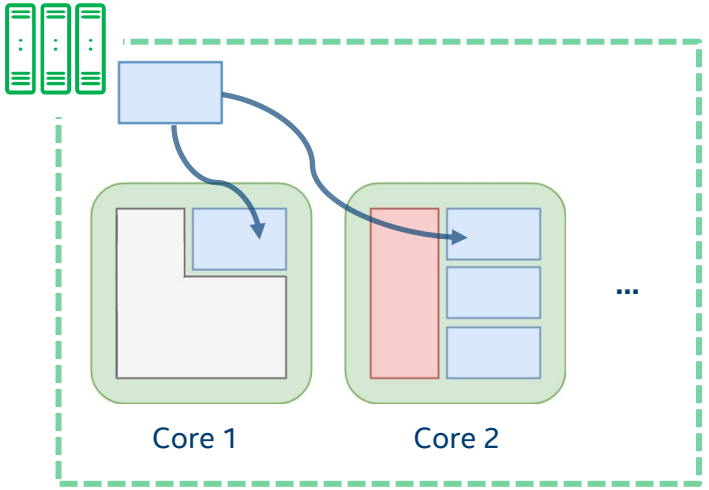

INTEGRATED ML APPROACH TO ORCHESTRATION

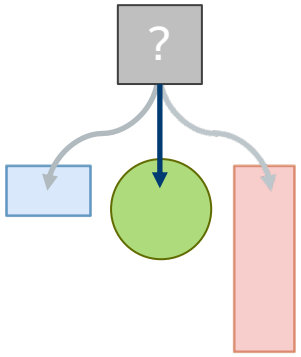

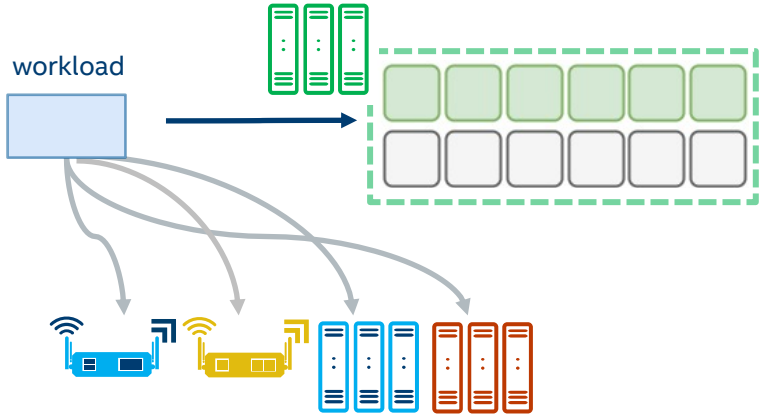

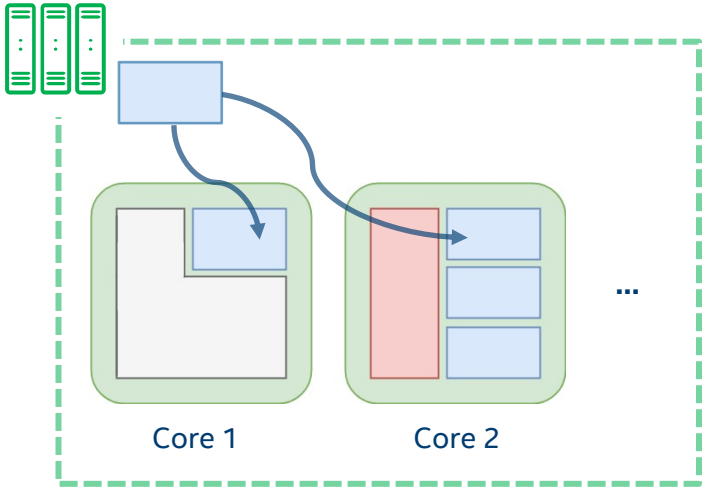

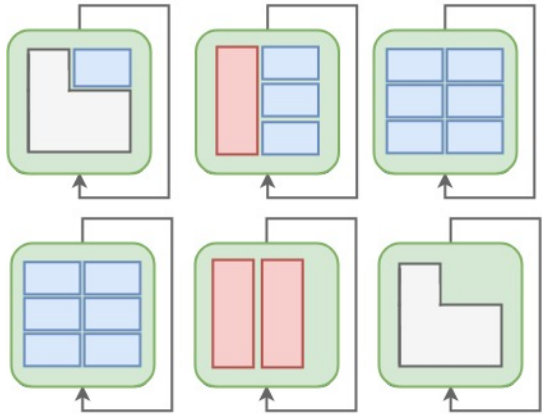

INTEGRATED ML APPROACH TO ORCHESTRATION

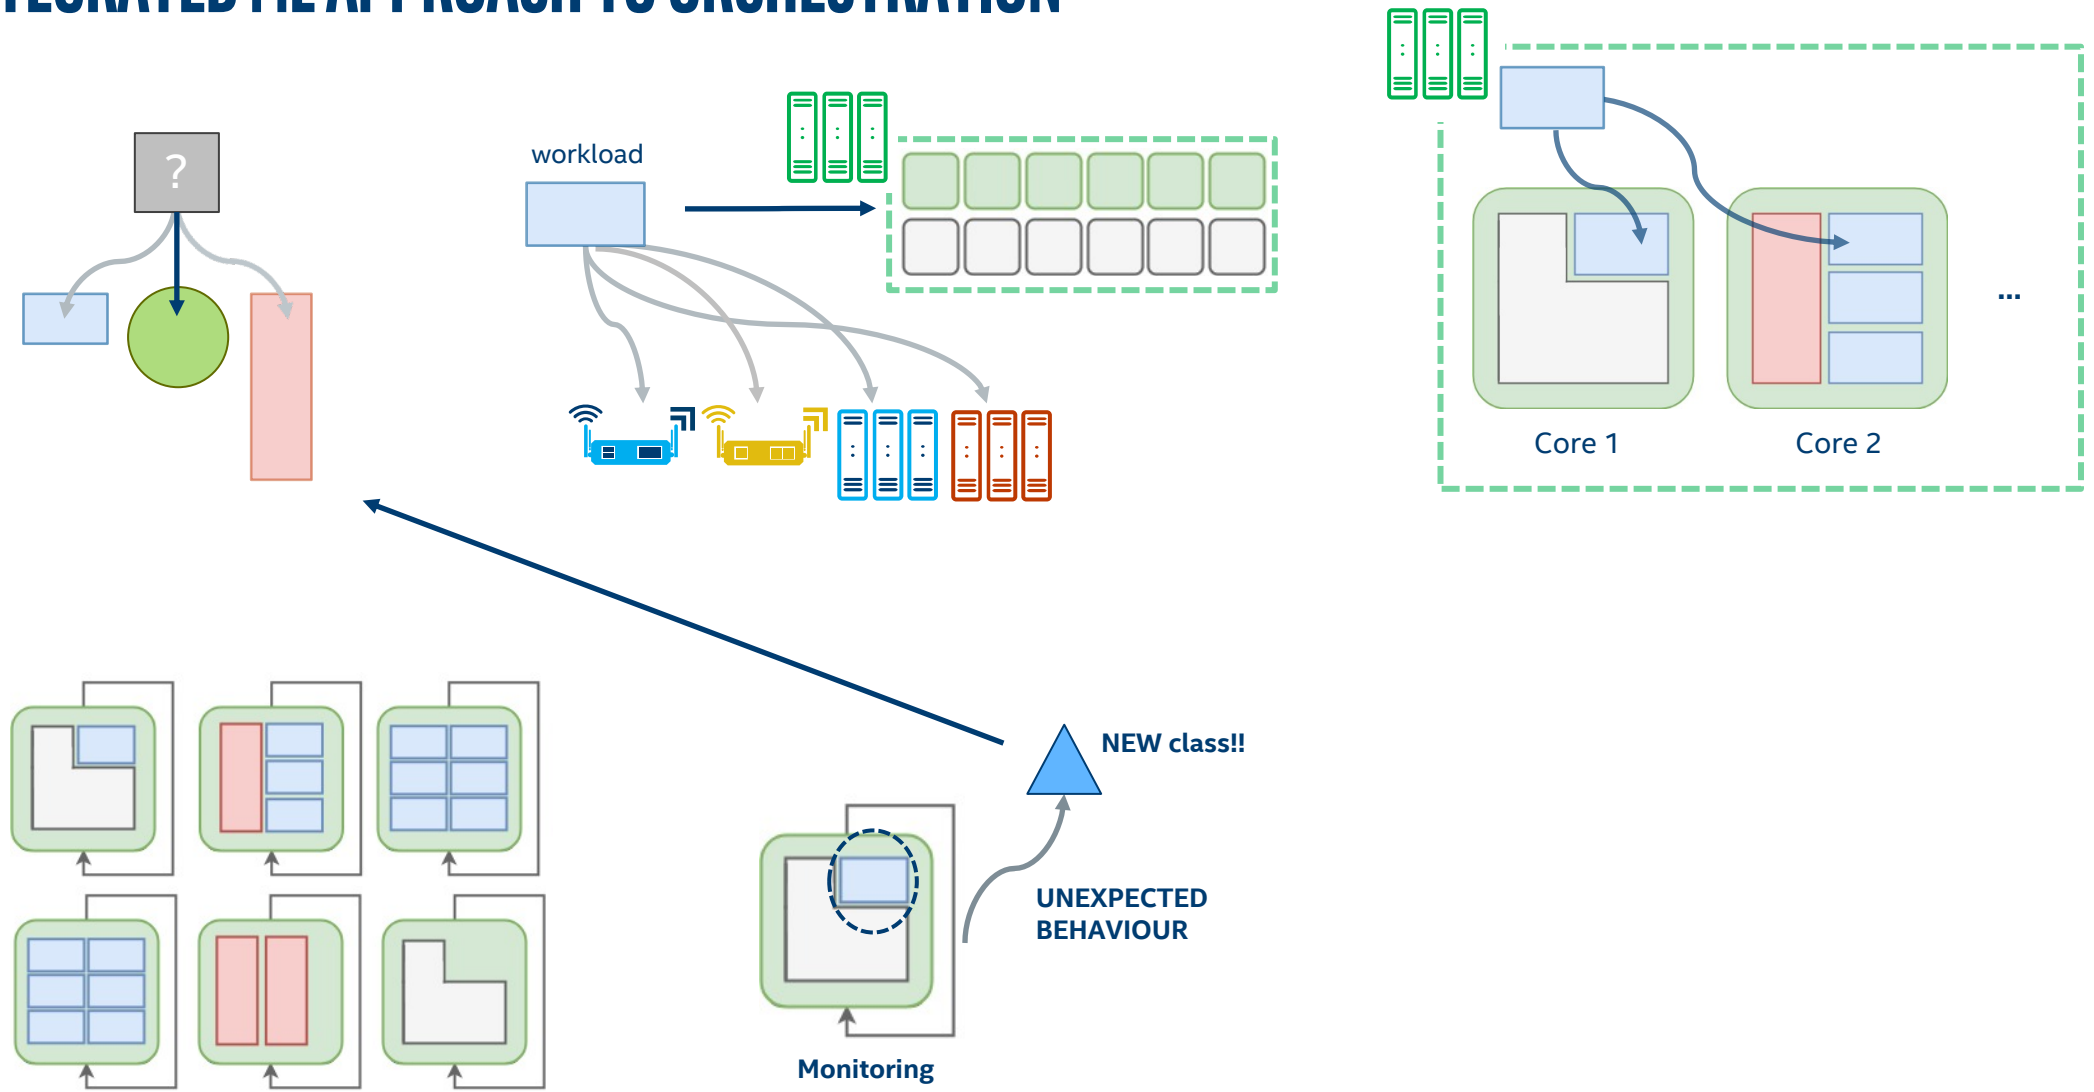

INTEGRATED ML APPROACH TO ORCHESTRATION

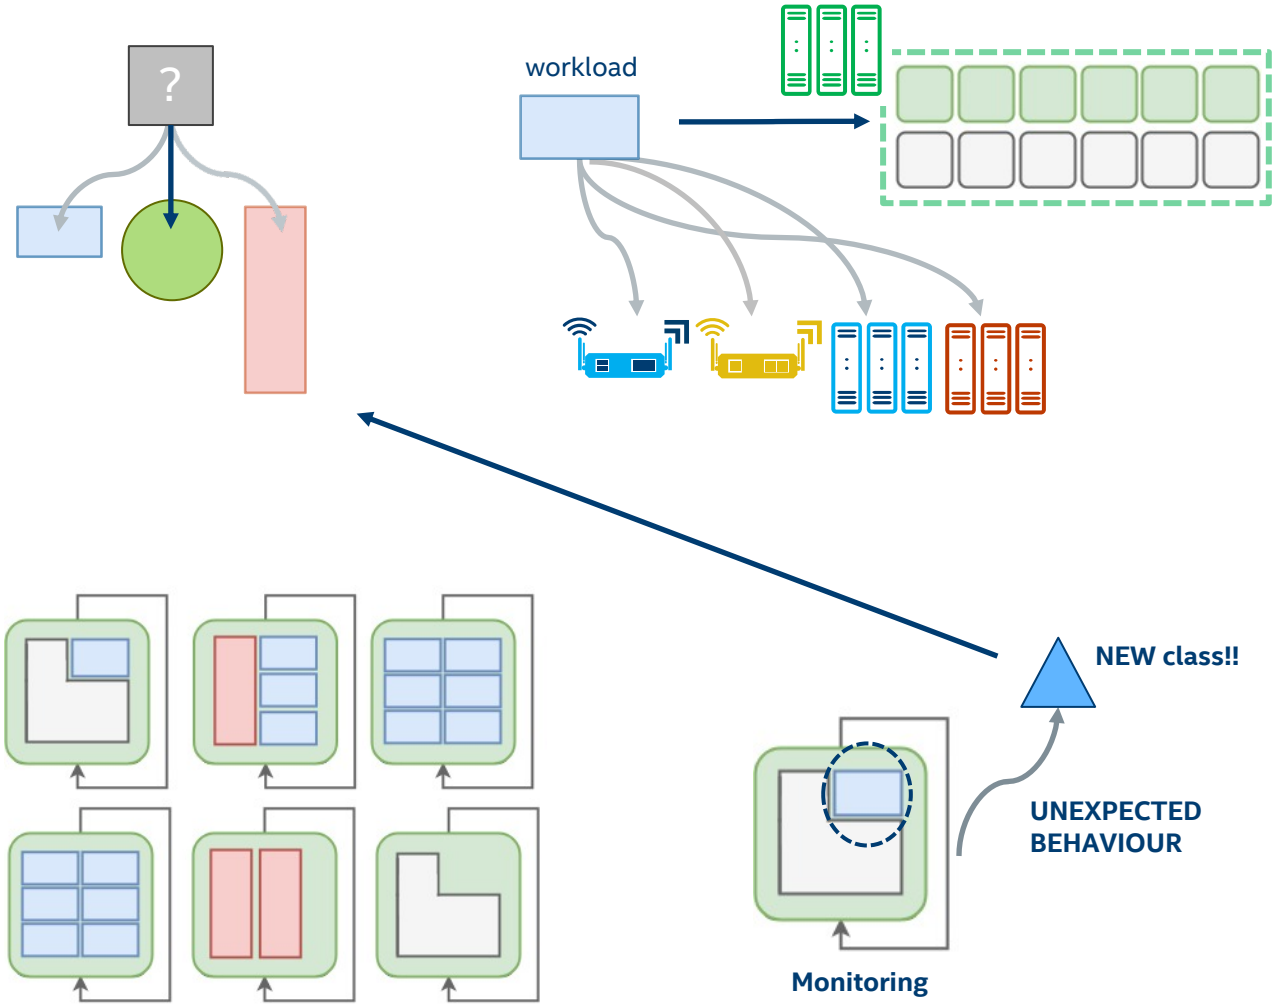

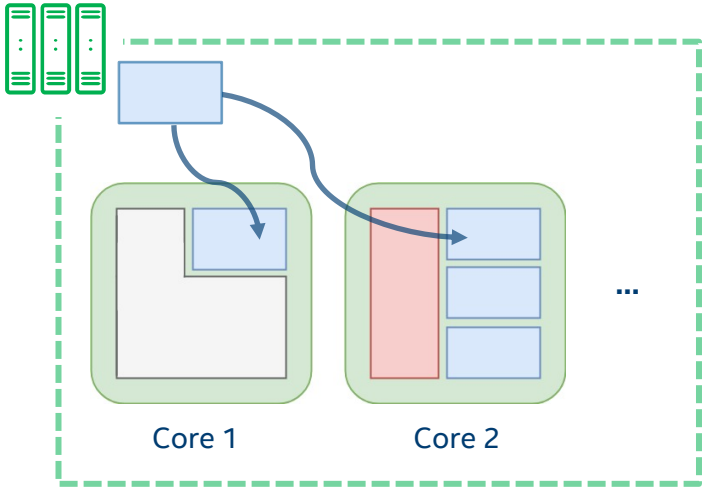

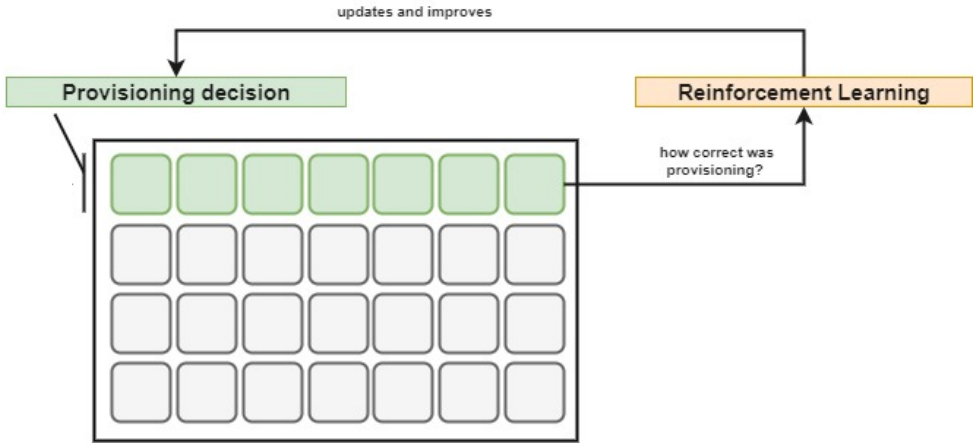

INTEGRATED ML APPROACH TO ORCHESTRATION

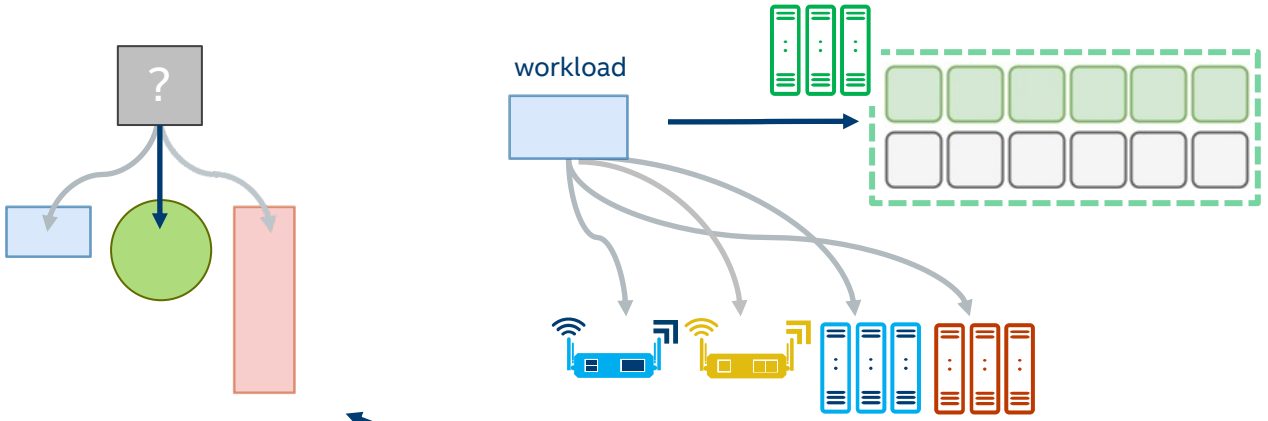

Efficient resource allocation
Power minimization

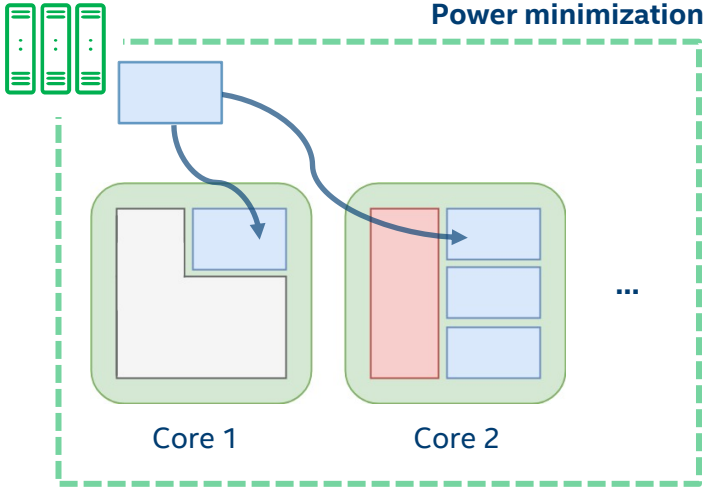

SLA awareness
Continuous learning of shared behavior

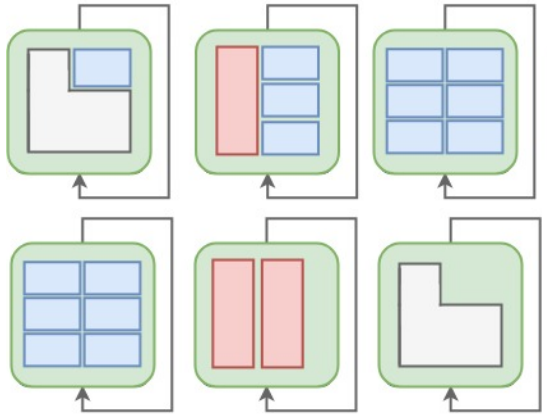

NEW class!!

UNEXPECTED BEHAVIOUR

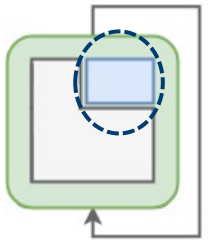

Continuous learning of new workloads
Ability to deal with changes in environment

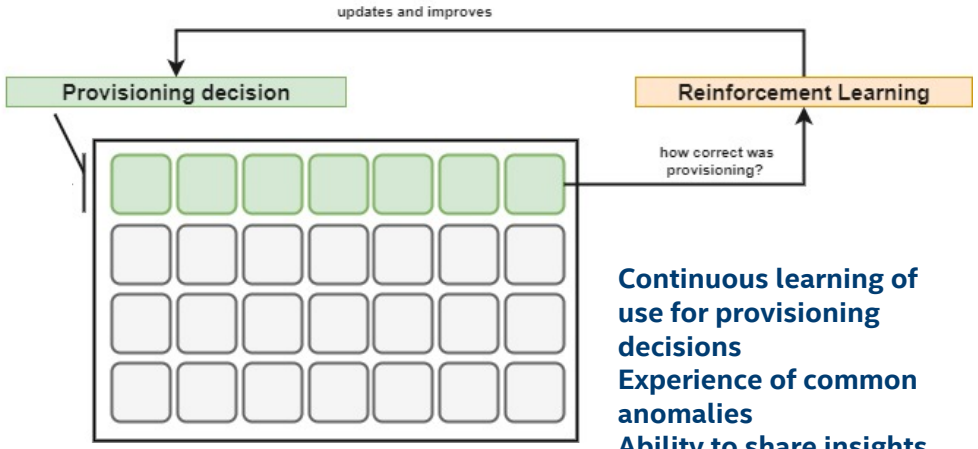

Continuous learning of use for provisioning decisions
Experience of common anomalies
Ability to share insights with other machines...


intel
experience
what's inside™