Innovating at the Edge Meetups

Edge Big Bang

- Francesc Guim, Chief of Edge Architecture, Intel

intel®

Notices and Disclaimers

- Intel technologies may require enabled hardware, software or service activation.
- No product or component can be absolutely secure.
- Your costs and results may vary.
-

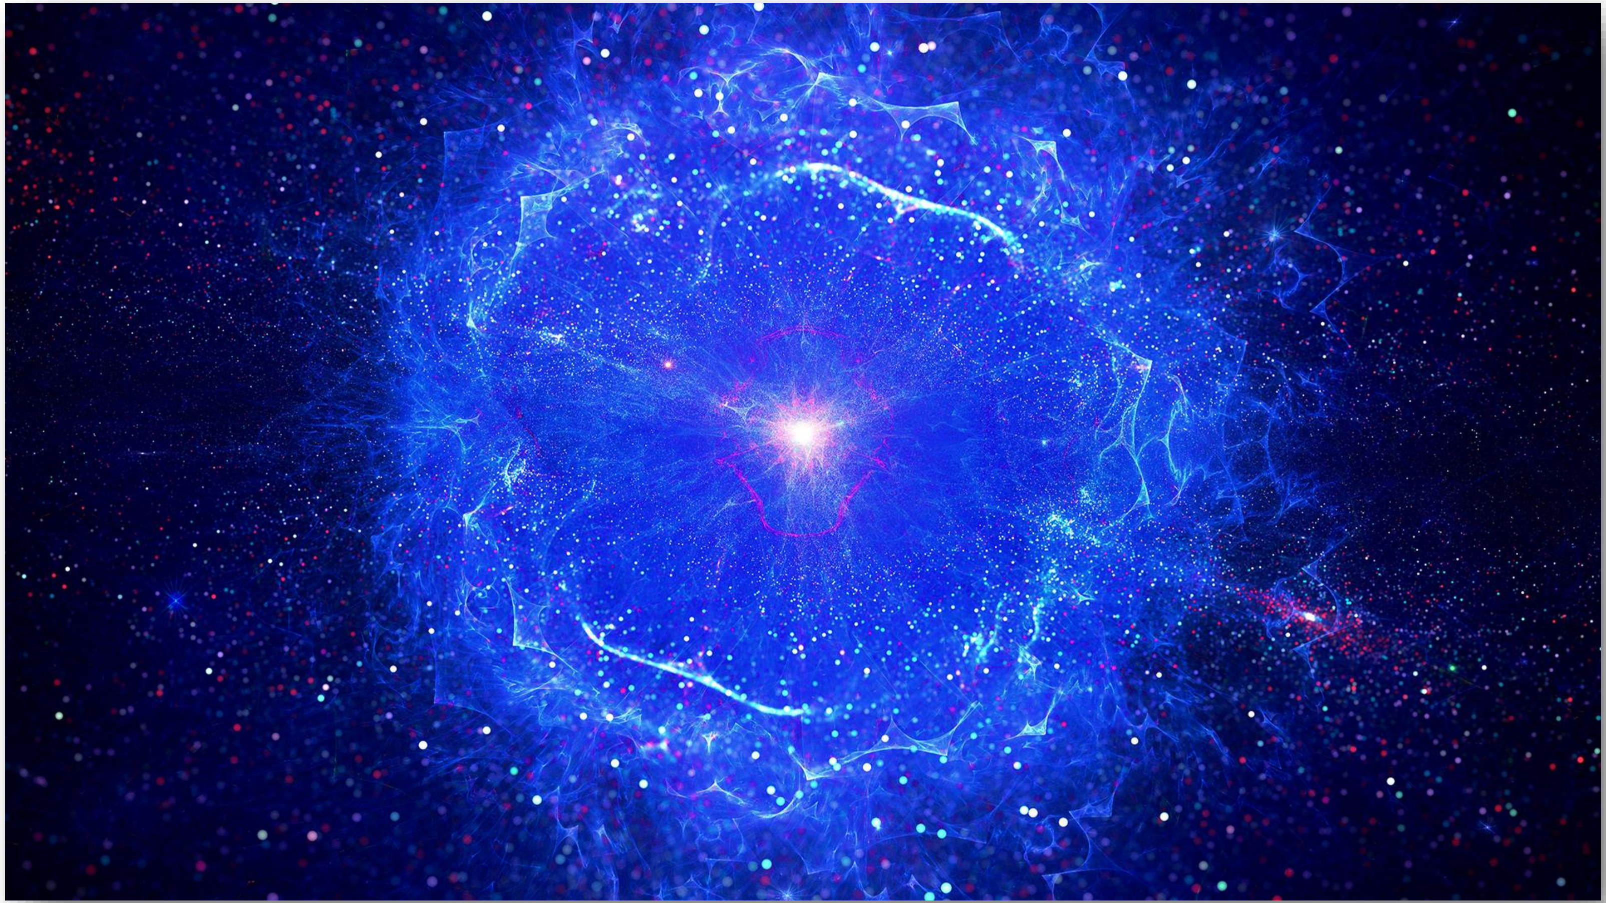

POCs
On-Premise or Controlled env.
10/100s Edges per deployment

Production deployments
Far Edge, Near Edge, Data Center Edge, Cloud
Up to 100K Edges per deployment
Challenges: Temperature, Energy, Reliability and security @ scale

2016 → 2023

Outline

Context & Problem statement

- Where we come from
- Current edge to cloud deployments
- Key areas to address

Foundation of Technologies

- Need of common frameworks
- Cloud native design
- Distributed cloud
- Adaptive Edge

Network current vectors

- Life cycle management
- Standards
- Security

Addressing the future

- Sustainability
- Design trade-offs
- What does zero emission mean?

Outline

Context & Problem Statement

- Where we come from
- Current edge to cloud deployments
- Key areas to address

Foundation of Technologies

- Need of common frameworks
- Cloud native design
- Distributed cloud
- Adaptive Edge

Network Current Vectors

- Life cycle management
- Standards
- Security

Addressing the Future

- Sustainability
- Design trade-offs
- What does zero emission mean?

Challenge 1: Vertical View of Edge

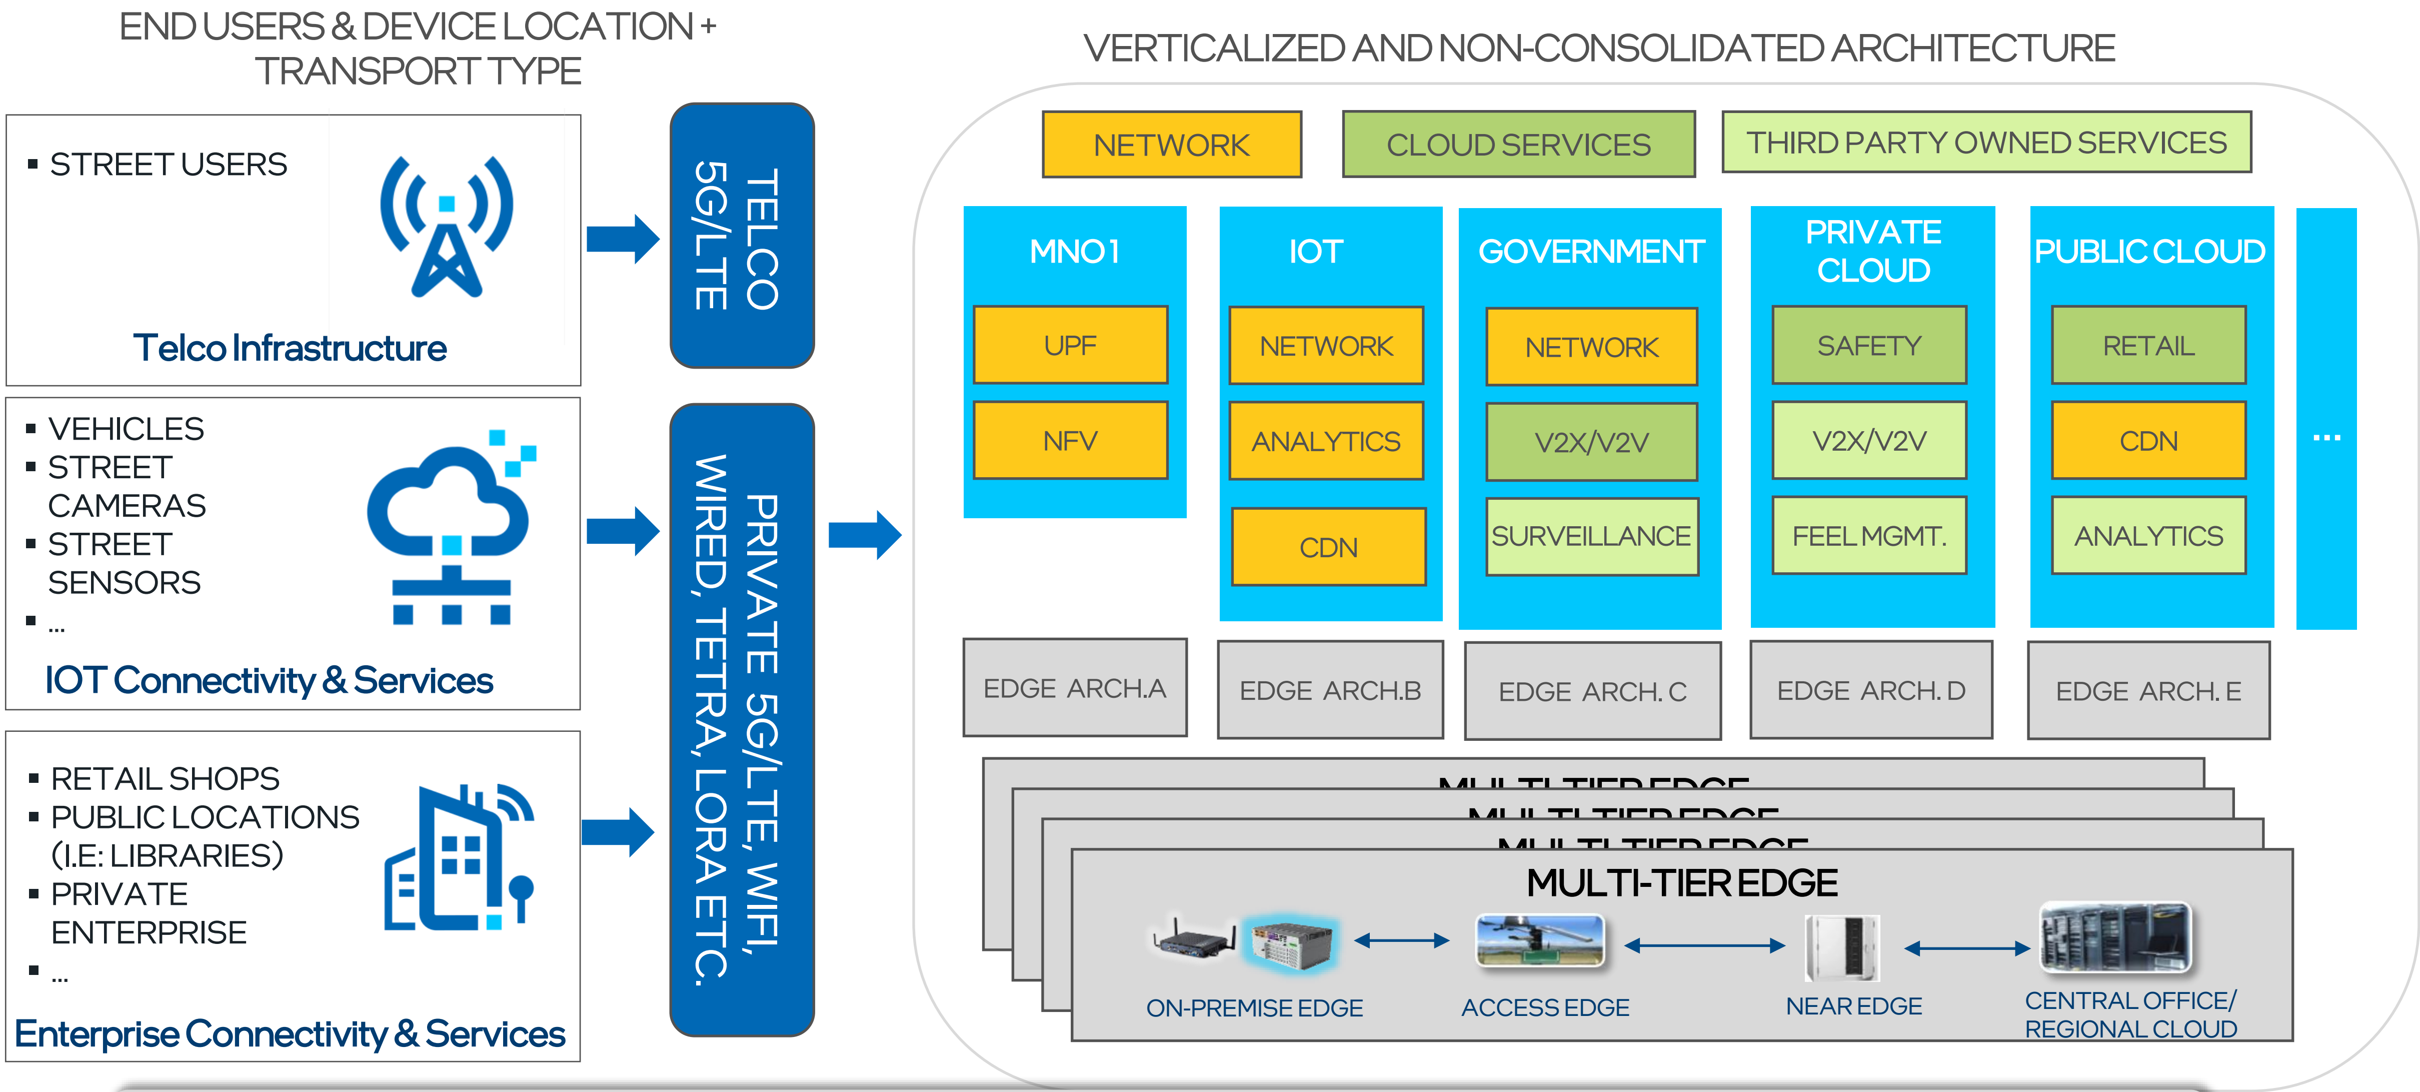

Challenges: CSPs/CoSPs have multiple edge verticals (IOT, Enterprise, Telco) with differing architectures.
How do we drive a scalable & converged architecture?

Challenge 2: Edge is Distributed, has Life and Requires Scales

	ON-PREMISE EDGE: INTELLIGENT SENSOR/GW	ON-PREMISE EDGE: INTELLIGENT EDGE	NETWORK EDGE	DC EDGE	CORE DC	PUBLIC CLOUD
DEPLOYMENT TIERS	ON-PREMISE EDGE INTELLIGENT SENSOR/GW	INTELLIGENT EDGE	NETWORK EDGE	DC EDGE	CORE DC	PUBLIC CLOUD
DEPLOYMENT LIMITATIONS	AVAILABLE POWER: [5,60W] THERMALS: -40 to +85C	AVAILABLE POWER: [100,3000W] THERMALS: -40 to +70C	AVAILABLE POWER: [400,3000W] THERMALS: -20 to +70C	AVAILABLE POWER: 9KW/Rack THERMALS: 20 to 40C	TYPICAL AVAILABLE POWER: GEO DEPENDENT 7KW – +20KW/RACK	
FORM FACTORS	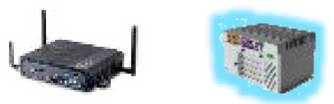 Ruggedized/ Custom FF	 Ruggedized / Custom FF	 Server/FF optimized for extreme conditions	 Short depth server Chassis	 Standard 1U/2U Chassis	 CSP Custom FF
PROCESSOR TYPE	intel ATOM – ENTRY ATOM intel CORE i7 – ENTRY CLIENT	intel XEON – XEON-D (SoC) intel XEON SILVER – ENTRY XEON	intel ATOM – ENTRY ATOM intel XEON – XEON-D (SoC)	intel XEON SILVER – ENTRY XEON	XEON FAMILY: intel XEON SILVER, intel XEON GOLD, intel XEON PLATINUM, intel XEON BRONZE	
ACCELERATION	intel MOVIDIUS – VPU intel CYCLONE 10 – FPGA LOW POWER	intel eASIC, intel STRATIX – INTEL® FPGA, GPU (<75W)			intel MOVIDIUS – VPU intel STRATIX – FPGA GPU (<150W)	
STORAGE/ MEMORY		intel OPTANE SSD – HIGH PERF SSD		intel OPTANE SSD – HIGH PERF SSD	intel OPTANE MEMORY – PERSISTENT MEMORY intel OPTANE SSD – HIGH PERF SSD	

Challenge 3: Amount of Technology Building Blocks

End-to-End Services: AI for Imaging | Content Distribution | VNFs | NLP | KVS | AR/VR

Layer	Building Blocks
Service Orchestration	K8s/K3 based Local/Multi-Edge Orchestrator; Hybrid-cloud / edge; Cloud Service Based Orchestration
SW Cloud Service Provides Solutions	AWS (outpost, sage maker ...); Azure (ARC, stack edge, HCI ..); Google (functions, ..)
VIM	KVM; OpenStack; ACRN; Containerization tech (VM or bare metal)
Infrastructure	Infrastructure Telemetry; SD-WAN
Network Infra	NIC; Smart NIC / IPU; Ethernet Switch; CXL switching; NZTP; iPXE
Frameworks	EdgeX; OpenVINO™; Intel® Smart Edge; Edge Insights for Industrial; Edge Control Stack
Managed Network Services	CUPS vEPC User Plane (PGW-U / SGW-U); CUPS vEPC Control Plane (MME / HSS / SGW-C / PGW-C); 5G FWA – Fixed Wireless Access; Full vEPC (MME / HSS / PGW / SGW); vBNG; FMC – AGF - FMIF; Ethernet / MPLS; vRAN
Access	LTE / 5GNR (ORAN/FlexRan); Fiber; WiFi / Industrial WiFi; LoRA; Ethernet; IIoT protocols
Form Factors	Smart Camera (VPU); No AC Cooling; AC/Liquid Cooling; Industrial Fanless; AI Appliance; Edge Servers; Edge DC Servers; Data center Servers
SW Platform	Security (IsecL); Quality of Service; Acceleration; Management; Telemetry
xPUs	GPU; AI Accelerators; FPGA(s); Low power CPU(Atom); Client CPU (i3, i5, i7); IPU; Xeon D; Xeon E; Xeon SP
Platform and xPU Features	Data protection; Workload Isolation; Edge Attestation; Edge Quality of; AMT; TSN, TCC

How Do We Architect Edge To Cloud?

- At scale
- With security
- With an optimized total cost of ownership
- With compute democratization
- Satisfying edge use cases requirements
- Being sustainable and scalable at the long run

Outline

Context & Problem Statement

- Where we come from
- Current edge to cloud deployments
- Key areas to address

Foundation of Technologies

- Need of common frameworks
- Cloud native design
- Distributed cloud
- Adaptive Edge

Network Current Vectors

- Life cycle management
- Standards
- Security

Addressing the Future

- Sustainability
- Design trade-offs
- What does zero emission mean?

Common Framework to Address Different Verticals

EDGE END TO END ARCHITECTURAL FOUNDATION: INTER-OPERABLE, OPEN & SELF MANAGED

END USERS & DEVICE LOCATION + TRANSPORT TYPE

- STREET USERS

Telco Edge

- VEHICLES
- STREET CAMERAS
- STREET SENSORS
- ...

IOT Edge

- RETAIL SHOPS
- PUBLIC LOCATIONS (I.E: LIBRARIES)
- PRIVATE ENTERPRISE
- ...

Enterprise Edge

INFRASTRUCTURE (5G/LTE, Wireless, Lora, Wired ...)

MULTI-SERVICE (NFV & non-NFV) COMMON SERVICE TAXONOMY

NFV

Flexible NFV

De-Centralized Ran

vEPC, UPF

Resources: CPUs, storage, **FPGAs , pNIC/sNIC**

Latency: To be analyzed

Internet of Things

IoT devices in many fields such as factory automation, process automation, smart grids, V2V

Resources: Communication, CPUs, storage, **Movidius, SPH**

Latency: Factory automation : 0.25ms to 10ms

Smart grids: 3-20ms / Process automation: 50-100ms

Autonomous

Assist in autonomous driving

Examples: overtaking systems, V2V comm, navigation

Resources: CPUs, storage, **Mobileye, SPH**

Latency: Ideally <20ms, up to 100 ms

AR/VR & Gaming

Process images (image recognition) from devices, wearables and annotate useful information

Examples: Google glass

Resources: CPUs, **FPGAs, ATS**

Latency: seamless - <20 ms , sensitive- <25ms Tolerable 50-100ms

Data caching & Storage GW

Cache data at the edge for faster loading at user end

Using Edge as main storage for the devices

Examples: cache popular videos in a region, Netflix

Resources: Storage, CPUs

Latency: Not latency bound

Video/Video Analytics

Live video analytics and video pre-processing, video transcoding

Examples: Traffic video analysis and alarm systems

Resources: CPUs, **FPGAs**, storage, **ATS, SPH for Inference?**

Latency: To be analyzed

FaaS

Perform web page related pre-processing at the edge and send page to user device

Examples: web page rendering, ad block, content evaluation

Resources: CPUs, **FPGAs**,

Latency: Not slower than current web page load times

Speech Recognition

Speech-to-text, User commands, Biometric Recognition

Resources: Communication, CPUs, **GNA-s, FPGA, SPH**

Latency: To be analyzed

Medical Applications

Assist medical appliances through connectivity and analysis

Examples: Tele surgery

Resources: CPUs, Communication, storage, **ATS, SPH**

Latency: Tele surgery <150ms without haptic, <10ms with haptic feedback

Enterprise

IMBD, Specific Enterprise WL (i.e: Linked In)

Resources: CPUs, Communication, storage, **FPGA**

Latency: <10 ms

MULTI-TIER COMMON HW EDGE TAXONOMY

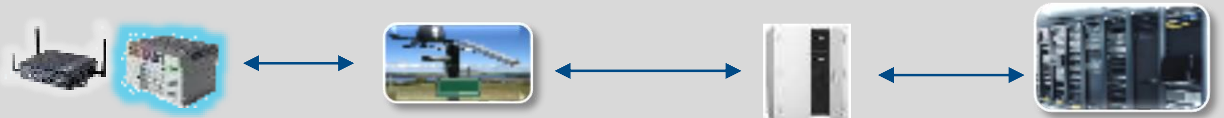

Use Case and End Customer Driven Architecture

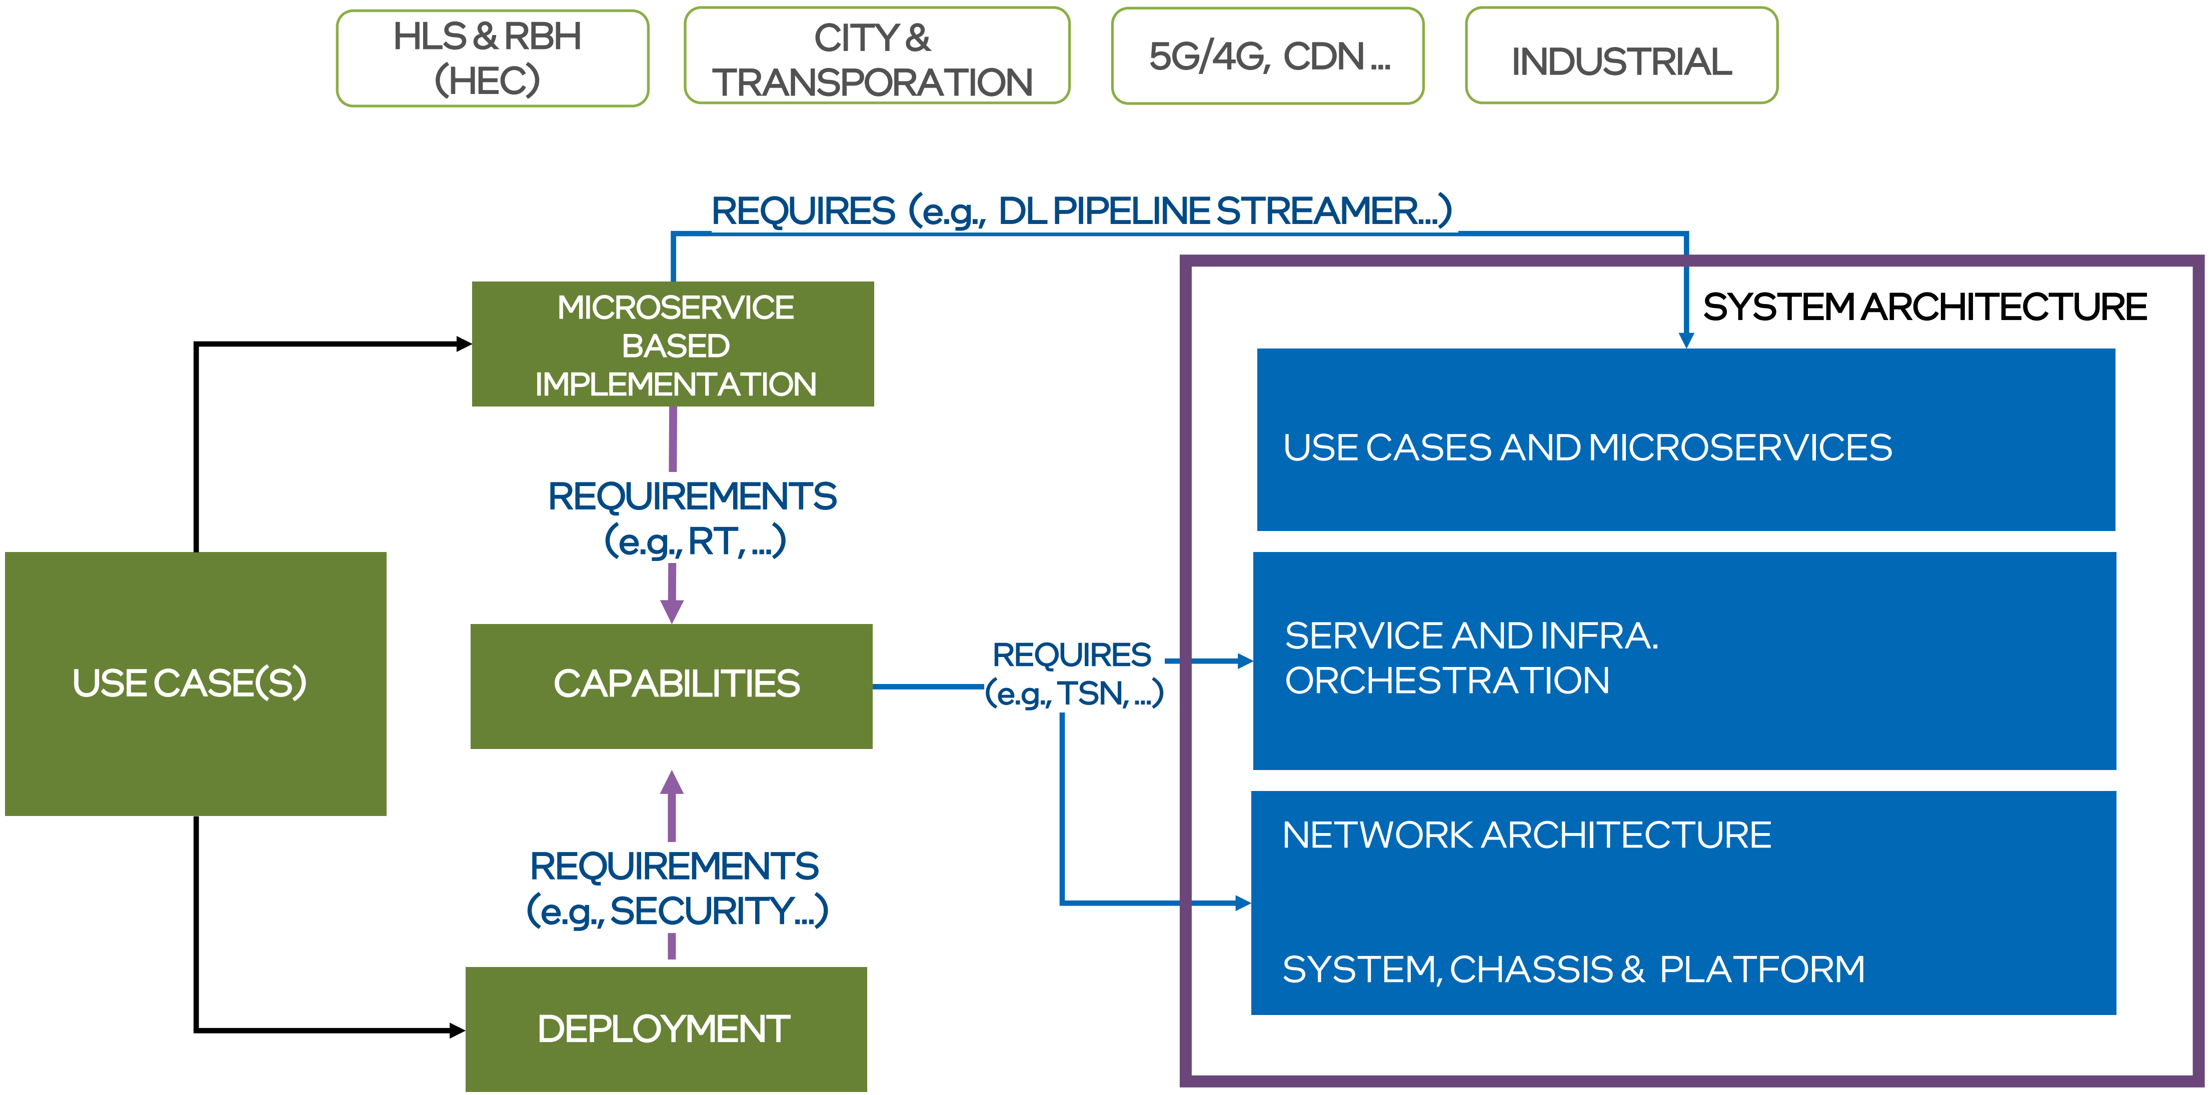

Tiers	On-Premise Edge		Network Edge	DC Edge	Public Cloud
	Intelligent Sensor/GW	Intelligent Edge			
Network Latencies (Wire Round trip)	< 1ms	< 1ms	1-5ms	5ms + 1-2 ms (every 100kms)	5ms + 5 ms (every 400kms)
Deployment Requirements	Compute Available Power: < 50 W Form Factor: Small Box Thermals: NEBs Mgmt.: Remote	Compute Available Power: ~10KW Form Factor: Rack(s) Thermals: NEBS or Standard DC Mgmt.: Remote	Compute Available Power: <600 W Form Factor: Pizza box Thermals: NEBS Mgmt.: Remote	Compute Available Power: 9KW/rack / 1KW sqm Form Factor: Rack(s) Thermals: NEBS or Standard DC Mgmt.: Remote	Standard Data Center (DC)

Where, What & Why

Intelligent Sensor/GW

Use Case	KPI
Intelligent Transportation	Data Privacy, Backhaul Traffic Savings, Reliability
V2V	Latency
Retail	Same as Int. Transp.
Video Analytics	Same as Int. Transp.

Intelligent Edge

Use Case	KPI
AR/VR	Latency. Backhaul Traffic Savings, Network scalability
Retail	Data Privacy, Backhaul Traffic Savings, Reliability
RT Streaming Healthcare	Same as AR/VR Access to services

Network Edge

Use Case	KPI
Intelligent Transportation	Data Privacy, Backhaul Traffic Savings, Reliability throughput Latency
V2V	Latency
Video Analytics	Same as Int. Transp.
Drone/IoT	Same as Int. Transp.
Rural	Access to services

DC Edge

Use Case	KPI
Intelligent Transportation	Data Privacy
Video Analytics	Same as Int. Transp.
Drone/IoT	Same as Int. Transp.
Healthcare	Access to services
CDN & Storage	Backhaul Traffic Savings Throughtput
Faas	Latency
AR/VR/MR	Latency

Public Cloud

Use Case	KPI
CDN	Backhaul traffic savings, Throughput
Storage GW	Same as CDN

Requirements to Distributed System Architecture

- Use case decomposition
- Intel HW Architecture Positioning w.r.t (1)
- Platform Architecture
- Edge Appliances (Rack, PODs, Platform etc..)

(1) DECOMPOSE USE CASE IN MICROSERVICES AND UNDERSTAND DEPLOYMENT REQUIREMENTS/OPTIONS

NETWORK EDGE / STREET CABINET

DATA CENTER EDGE

CLOUD

(a) Media → (b) Composition → (c) Video Analytics → (d) Biz. Logic → (e) Data Base

(2) UNDERSTAND HOW MICROSERVICES CAN ALLOW ADOPTION AND STICKINES

(b) Composition → (c) DL Pipeline Streamer Server → (d) Biz. Logic

ABSTRACTION

intel MOVIDIUS

intel CORE i3

intel CORE i5

intel CORE i7

intel ATOM

intel XEON

DIFFERENT KPIS

DIFFERENT TRADE-OFFS

(3) TRANSLATE (1) AND (2) INTO SW/HW SYSTEM ARCHITECTURE

Tiny

Immersion cooling Technologies

- Orchestration
- Network
- Edge Appliances

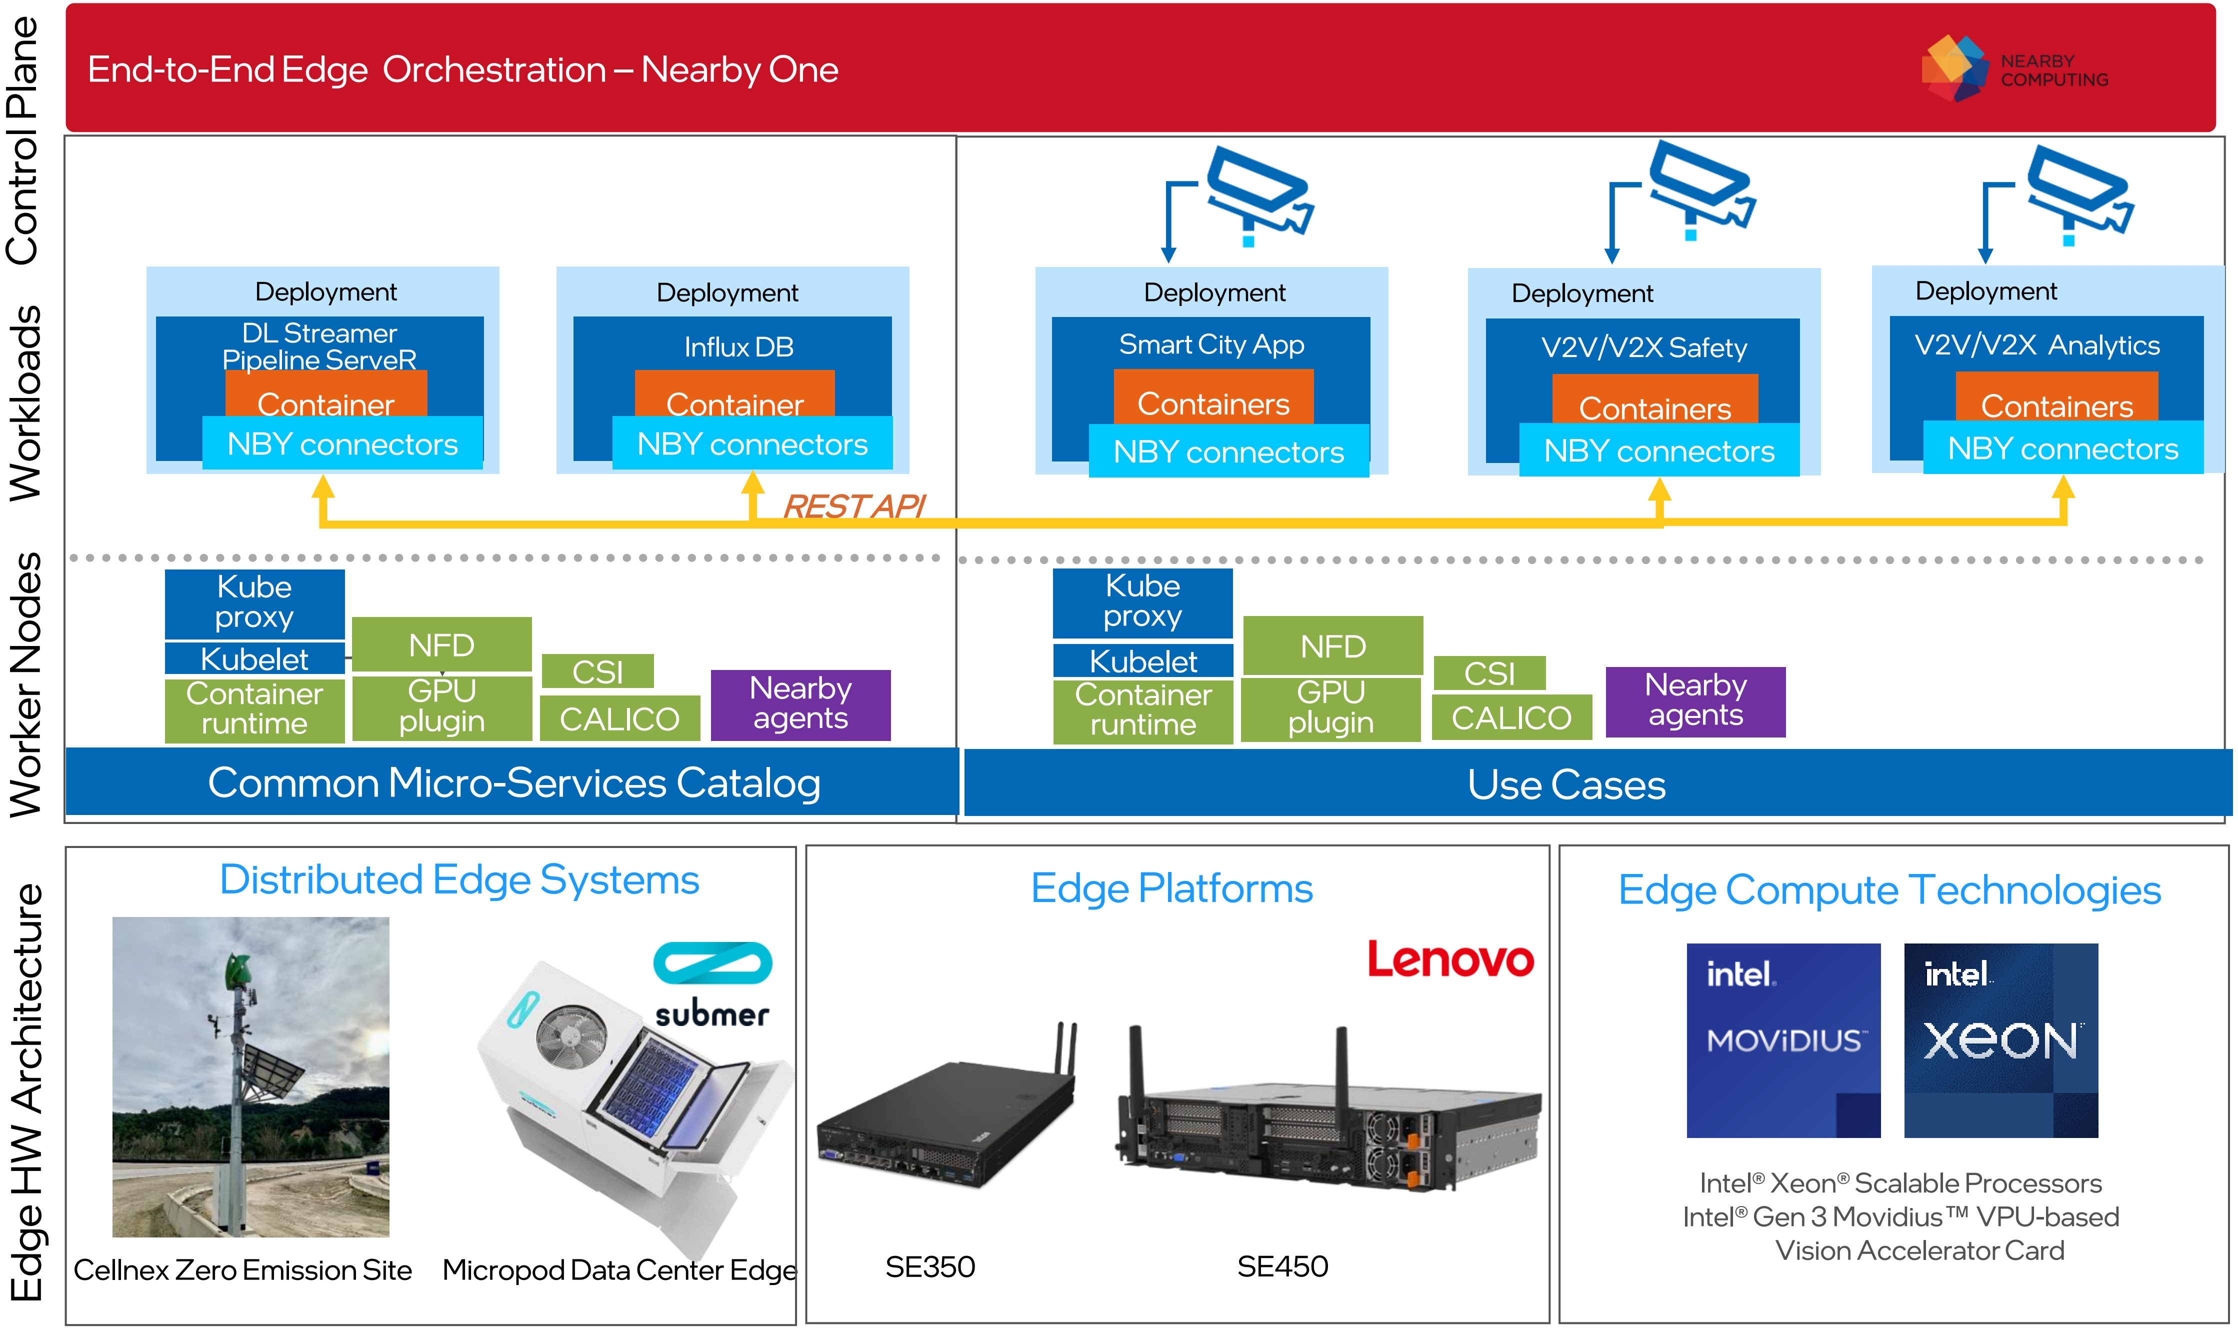

Control Plane
End-to-End Edge Orchestration – Nearby One
NEARBY COMPUTING
Workloads
Deployment
DL Streamer Pipeline ServeR
Container
NBY connectors
Deployment
Influx DB
Container
NBY connectors
Deployment
Smart City App
Containers
NBY connectors
Deployment
V2V/V2X Safety
Containers
NBY connectors
Deployment
V2V/V2X Analytics
Containers
NBY connectors
REST API
Worker Nodes
Kube proxy
Kubelet
Container runtime
NFD
GPU plugin
CSI
CALICO
Nearby agents
Kube proxy
Kubelet
Container runtime
NFD
GPU plugin
CSI
CALICO
Nearby agents
Common Micro-Services Catalog
Use Cases
Edge HW Architecture
Distributed Edge Systems
submer
Cellnex Zero Emission Site
Micropod Data Center Edge
Edge Platforms
Lenovo
SE350
SE450
Edge Compute Technologies
intel.
MOVIDIUS
intel.
XEON
Intel® Xeon® Scalable Processors
Intel® Gen 3 Movidius™ VPU-based Vision Accelerator Card

Outline

Context & Problem Statement

- Where we come from
- Current edge to cloud deployments
- Key areas to address

Foundation of Technologies

- Need of common frameworks
- Cloud native design
- Distributed cloud
- Adaptive Edge

Network Current Vectors

- Life cycle management
- Standards
- Security

Addressing the Future

- Sustainability
- Design trade-offs
- What does zero emission mean?

Network Matters for E2E QoS.. And Matters a Lot

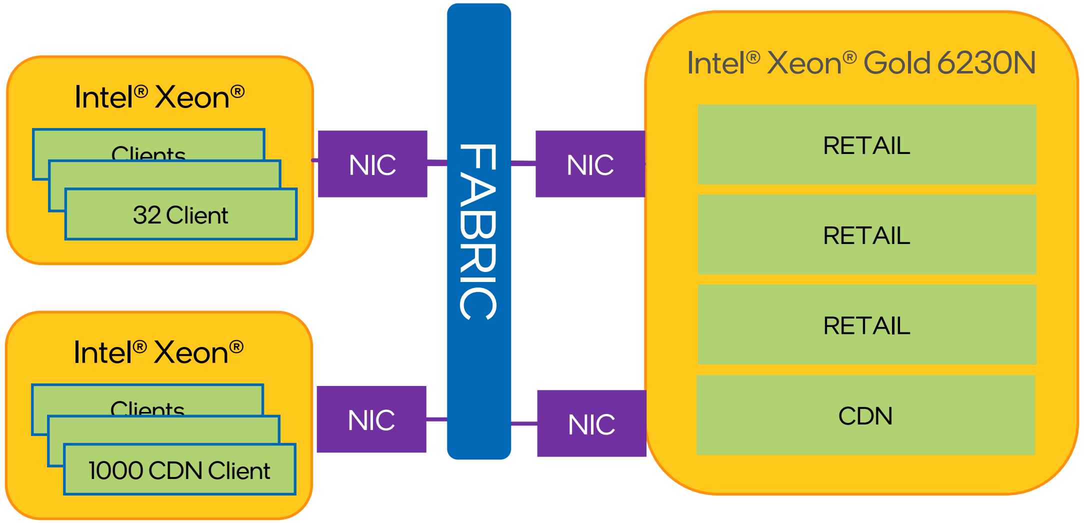

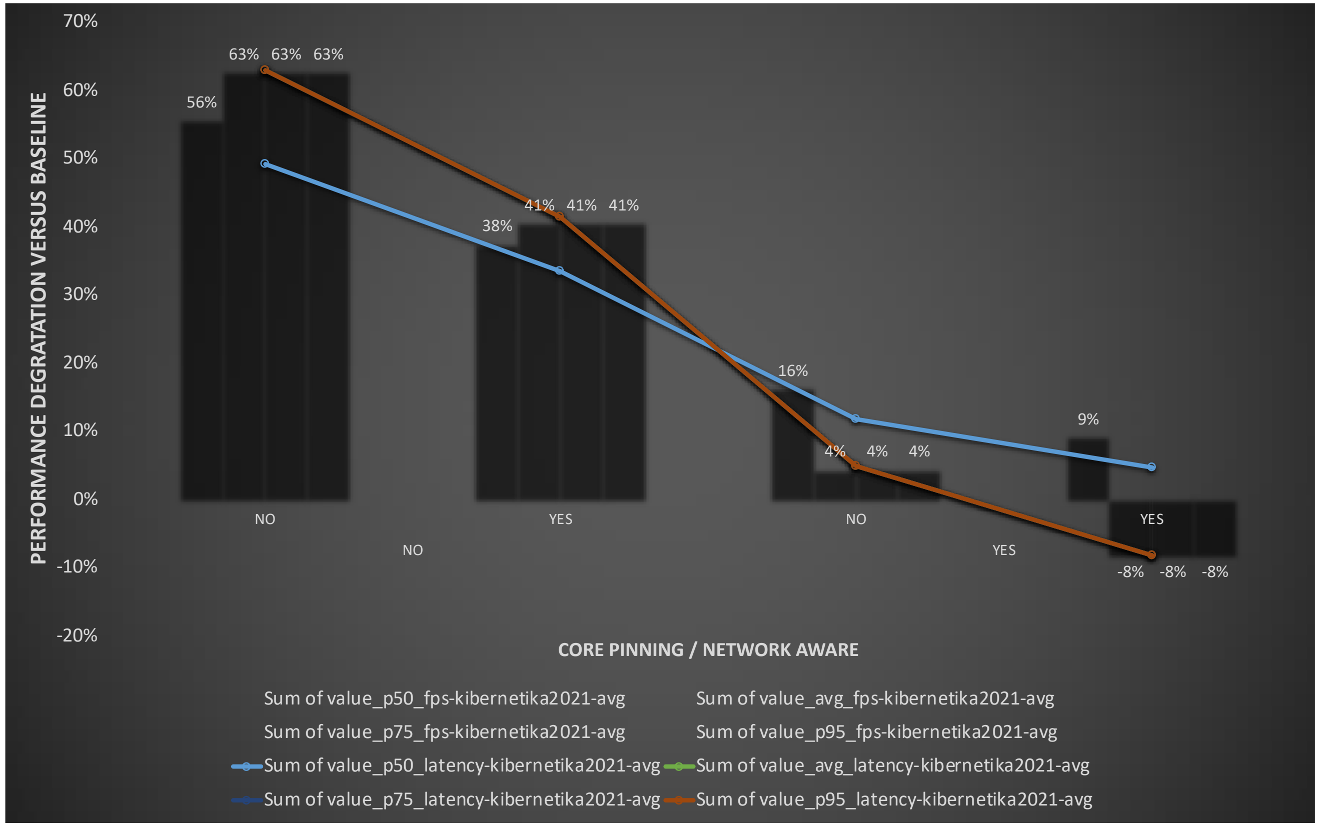

CONDUCTED FOUR TYPE OF EXPERIMENTS:

- COMPUTE RESOURCE SELECTION (CORE PINNING)
  a) YES = Workloads pinned into specific cores
  b) NO = No core pinning
- NETWORK AWARE (INTERFACE SELECTION)
  a) YES = Different interface for CDN and KIBERNETICA
  b) NO = Same interface for both

SHOWED RESULTS:

- Degradation of for each deployment with respect to baseline (Kibernetica alone)
- Application KPIS={fps, latency} – Statistics={avg, 50/75/95 percentile)

NOTES:

- COMPUTE AWARE REDUCE DEGRADATION
- NETWORK AWARE REDUCE DEGRADATION
- DISCLAIMER: we need to keep looking at the results
- DISCLAIMER 2: the fact that network + compute aware allocation has better performance that baseline can be something that we need more runs to normalize results

Application Life Cycle Management Matters a Lot

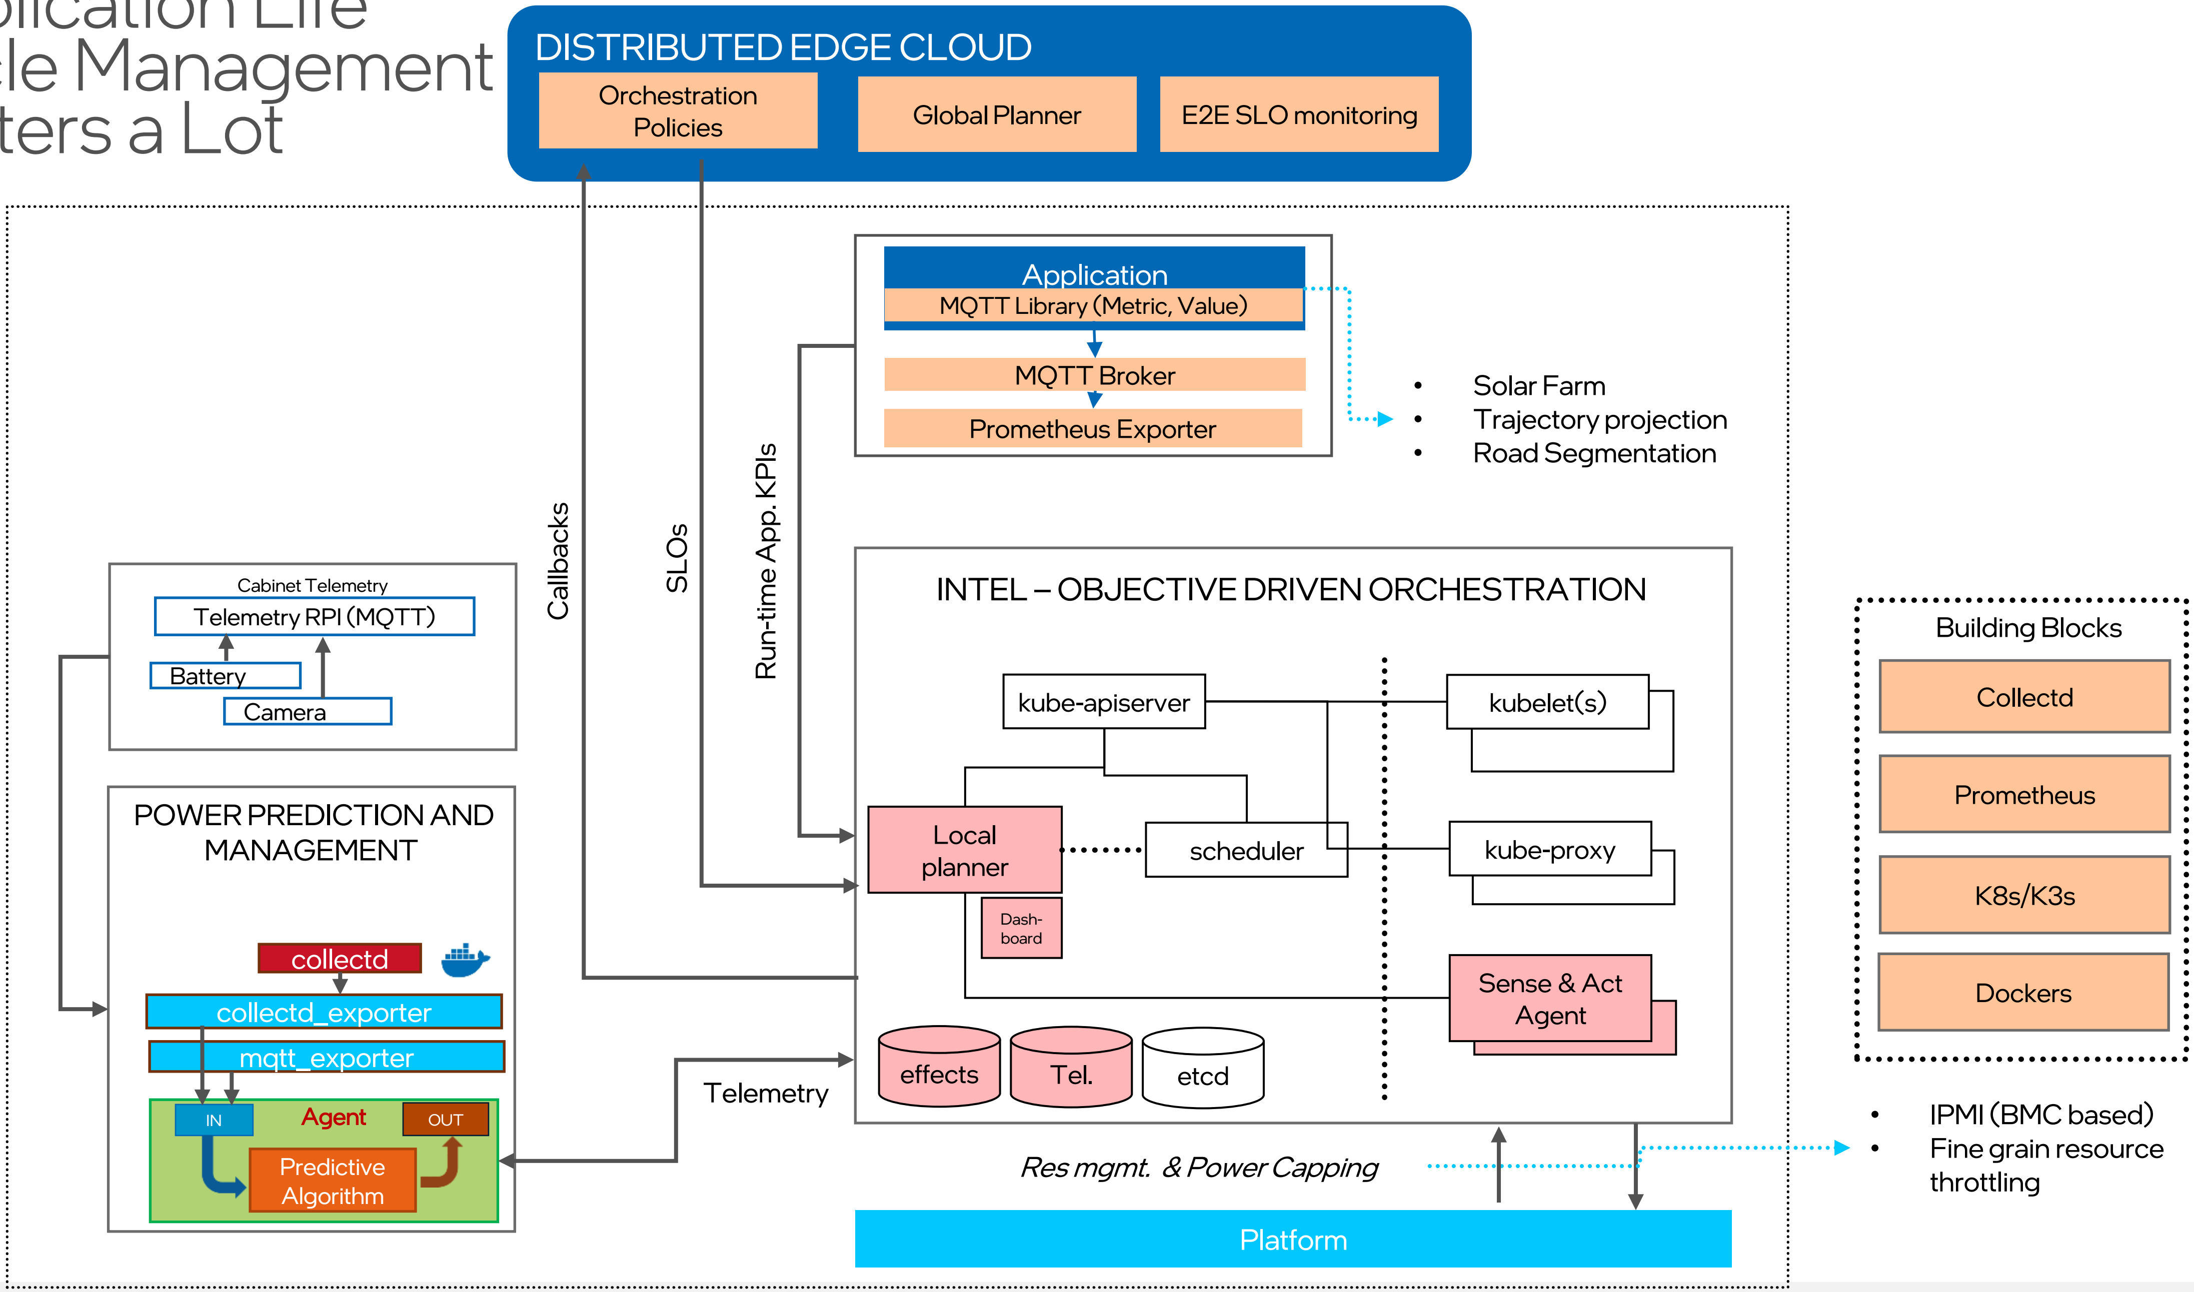

Application and Service Can Adapt Sometimes

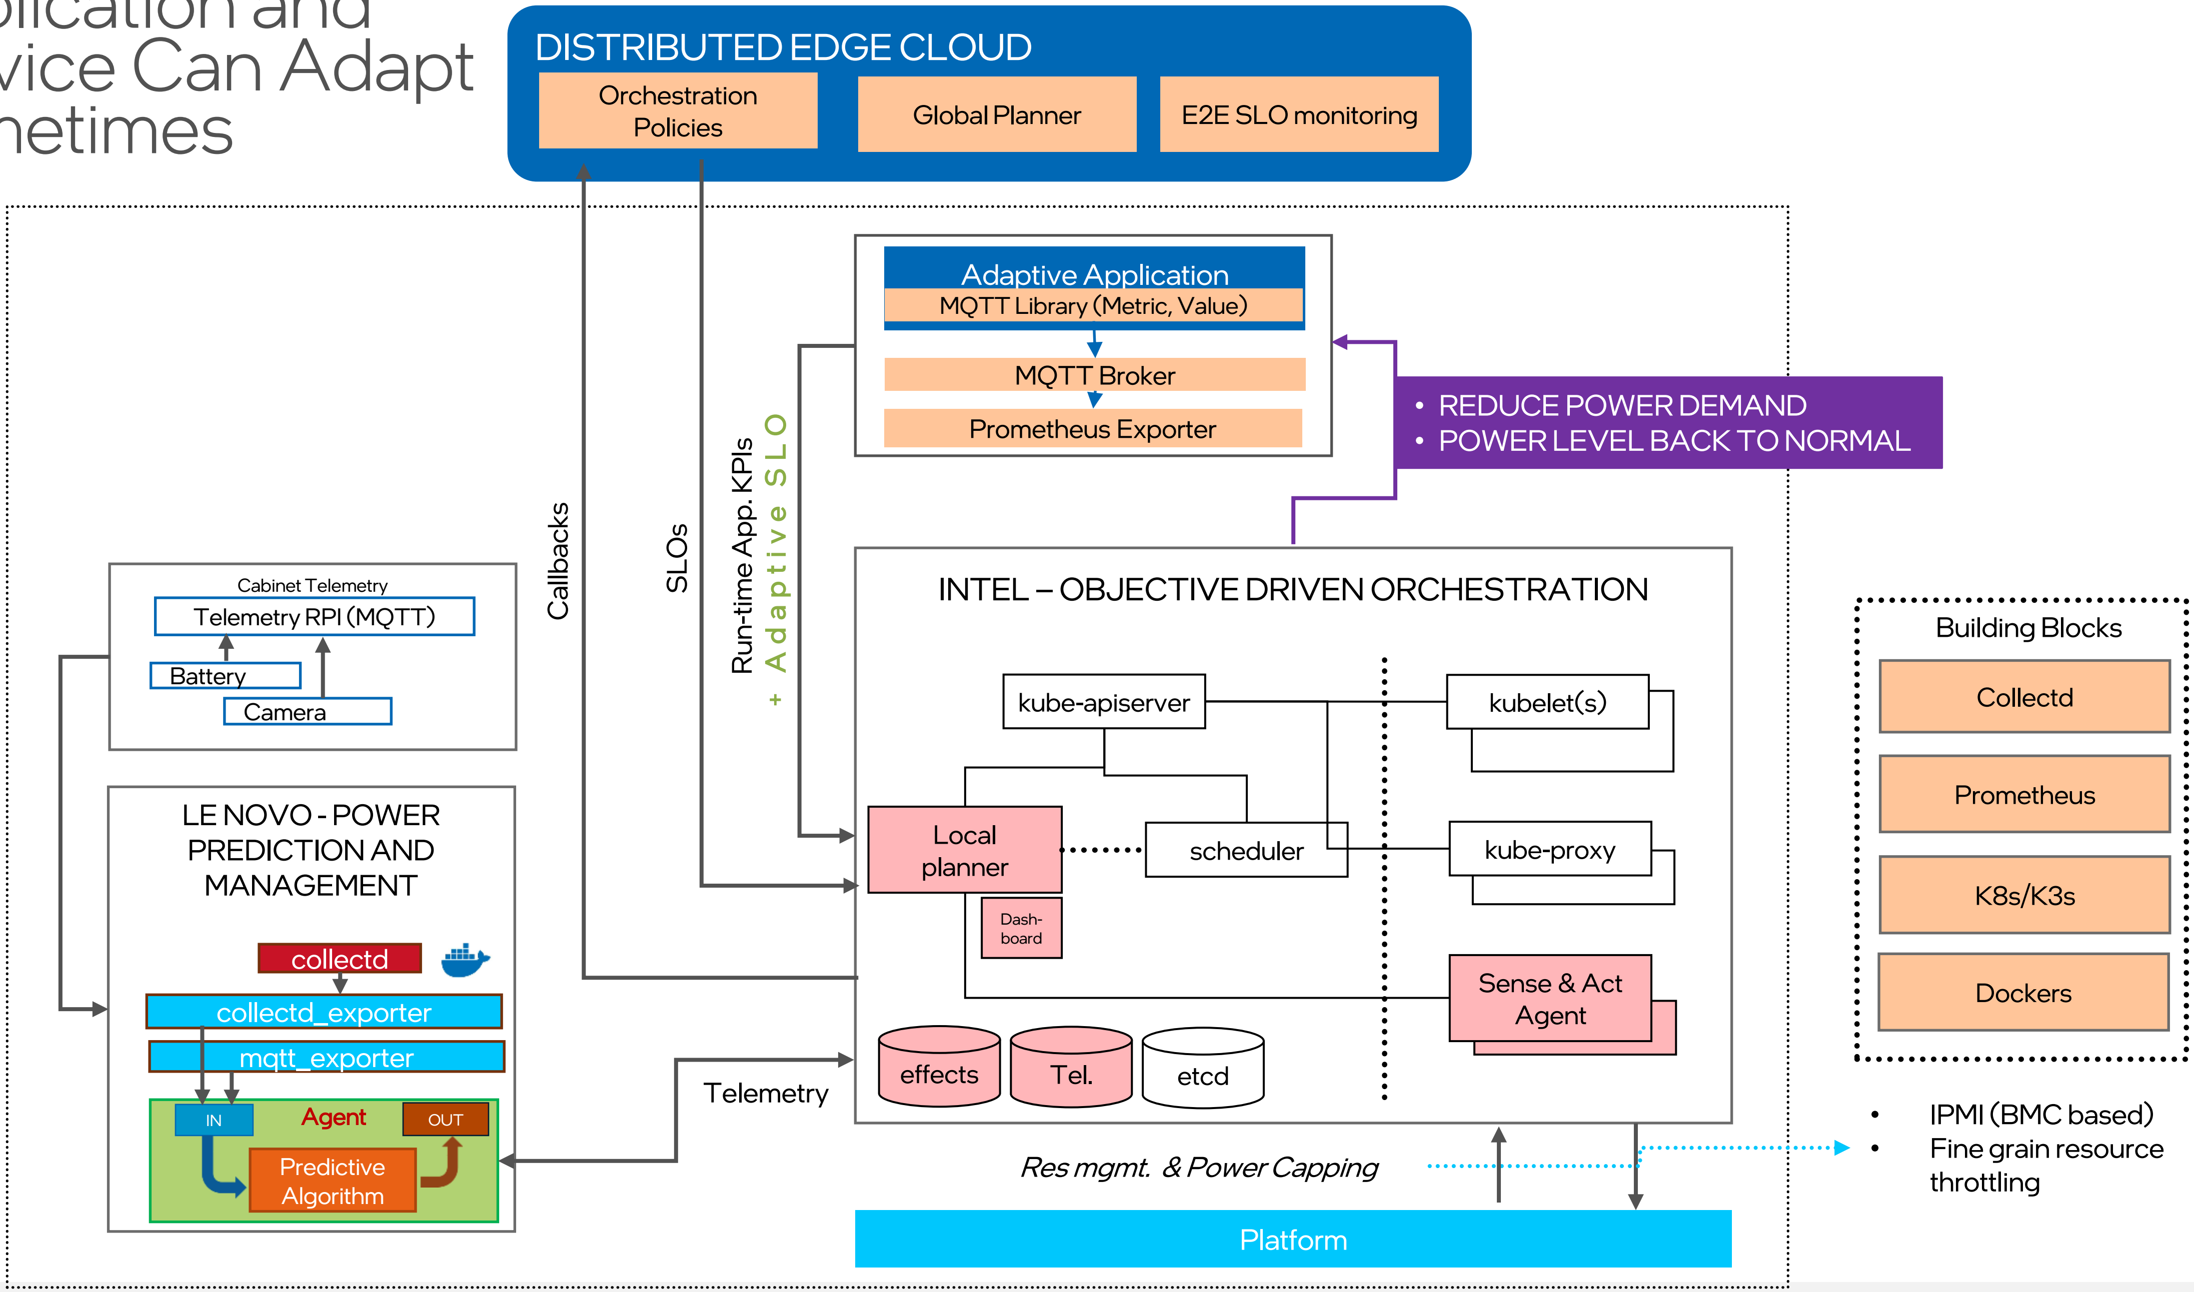

- IPMI (BMC based)
- Fine grain resource throttling

But … We Still Have Some Important Challenges

- Network end to end observability
- Network monitoring at service or application level
- Establish end-to-end QoS hooks between Edges and UE
- Application metrics vs system metrics in multitenant deployments
- And ….

We need standards and common APIS
We need network to as part of the service life cycle mgmt.
Re CAMARA PROJECT, OPG ...

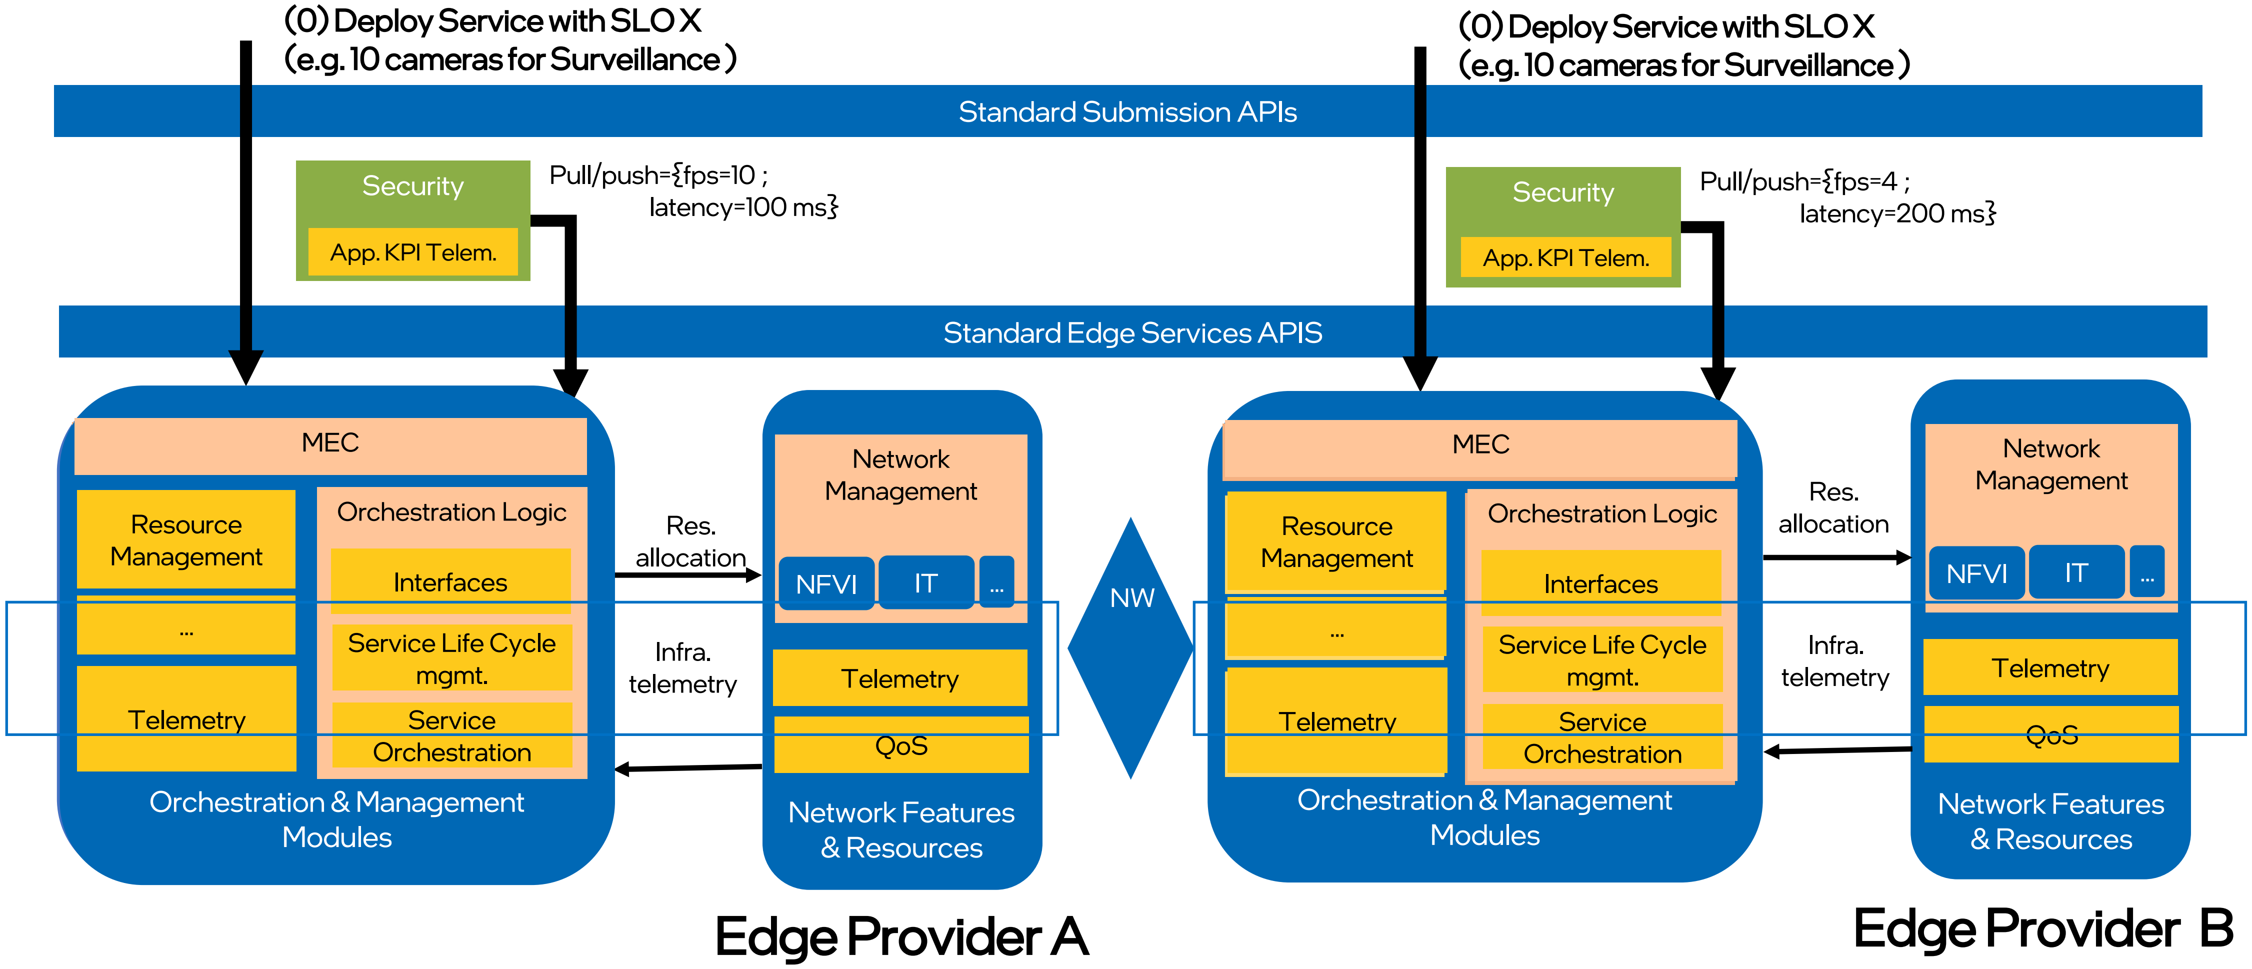

… and we need it secure

CHALLENGES

- (A) Expansion of Attack Surface
- (B) Sophistication and Evolving Intent for Hacking
- (C) Fragmentation of Products, Vendors, Solutions
- (D) Multi-Tier Security Dependencies
- (E) Zero Trust

REQUIREMENTS

- (1) Data Sovereignty & Privacy
- (2) Physical Attacks
- (3) Distributed Workloads
- (4) Heterogeneity
- (5) Data & IP Provenance

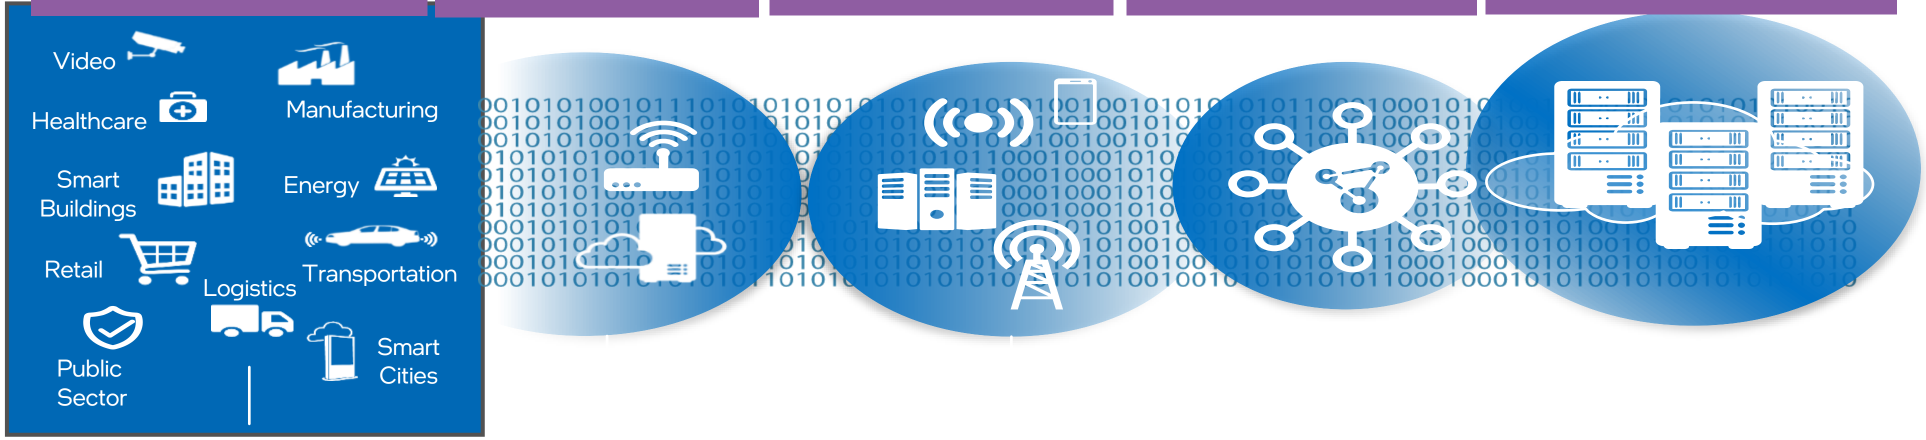

Devices / Things | Edge Compute Node | Network Hub or regional Data Center | Core Network | Cloud Data Center

ISVs (1-5/A-E) – 1/5 Critical

- Health Care (1,2,5/A,C):
- V2V/V2X (1,2,5/A-D)
- Enterprise – Private 5G (1,3,4/C-D)
- Industry 2.0 (1-5/A,C,D)
- Smart Cities (1-5/A,C,D,E)
- CoSPs + CSP Edges (1-5/A-E)

Outline

Context & Problem Statement

- Where we come from
- Current edge to cloud deployments
- Key areas to address

Foundation of Technologies

- Need of common frameworks
- Cloud native design
- Distributed cloud
- Adaptive Edge

Network Current Vectors

- Life cycle management
- Standards
- Security

Addressing the Future

- Sustainability
- Design trade-offs
- What does zero emission mean?

What Sustainable Distributed Edge May Require ...

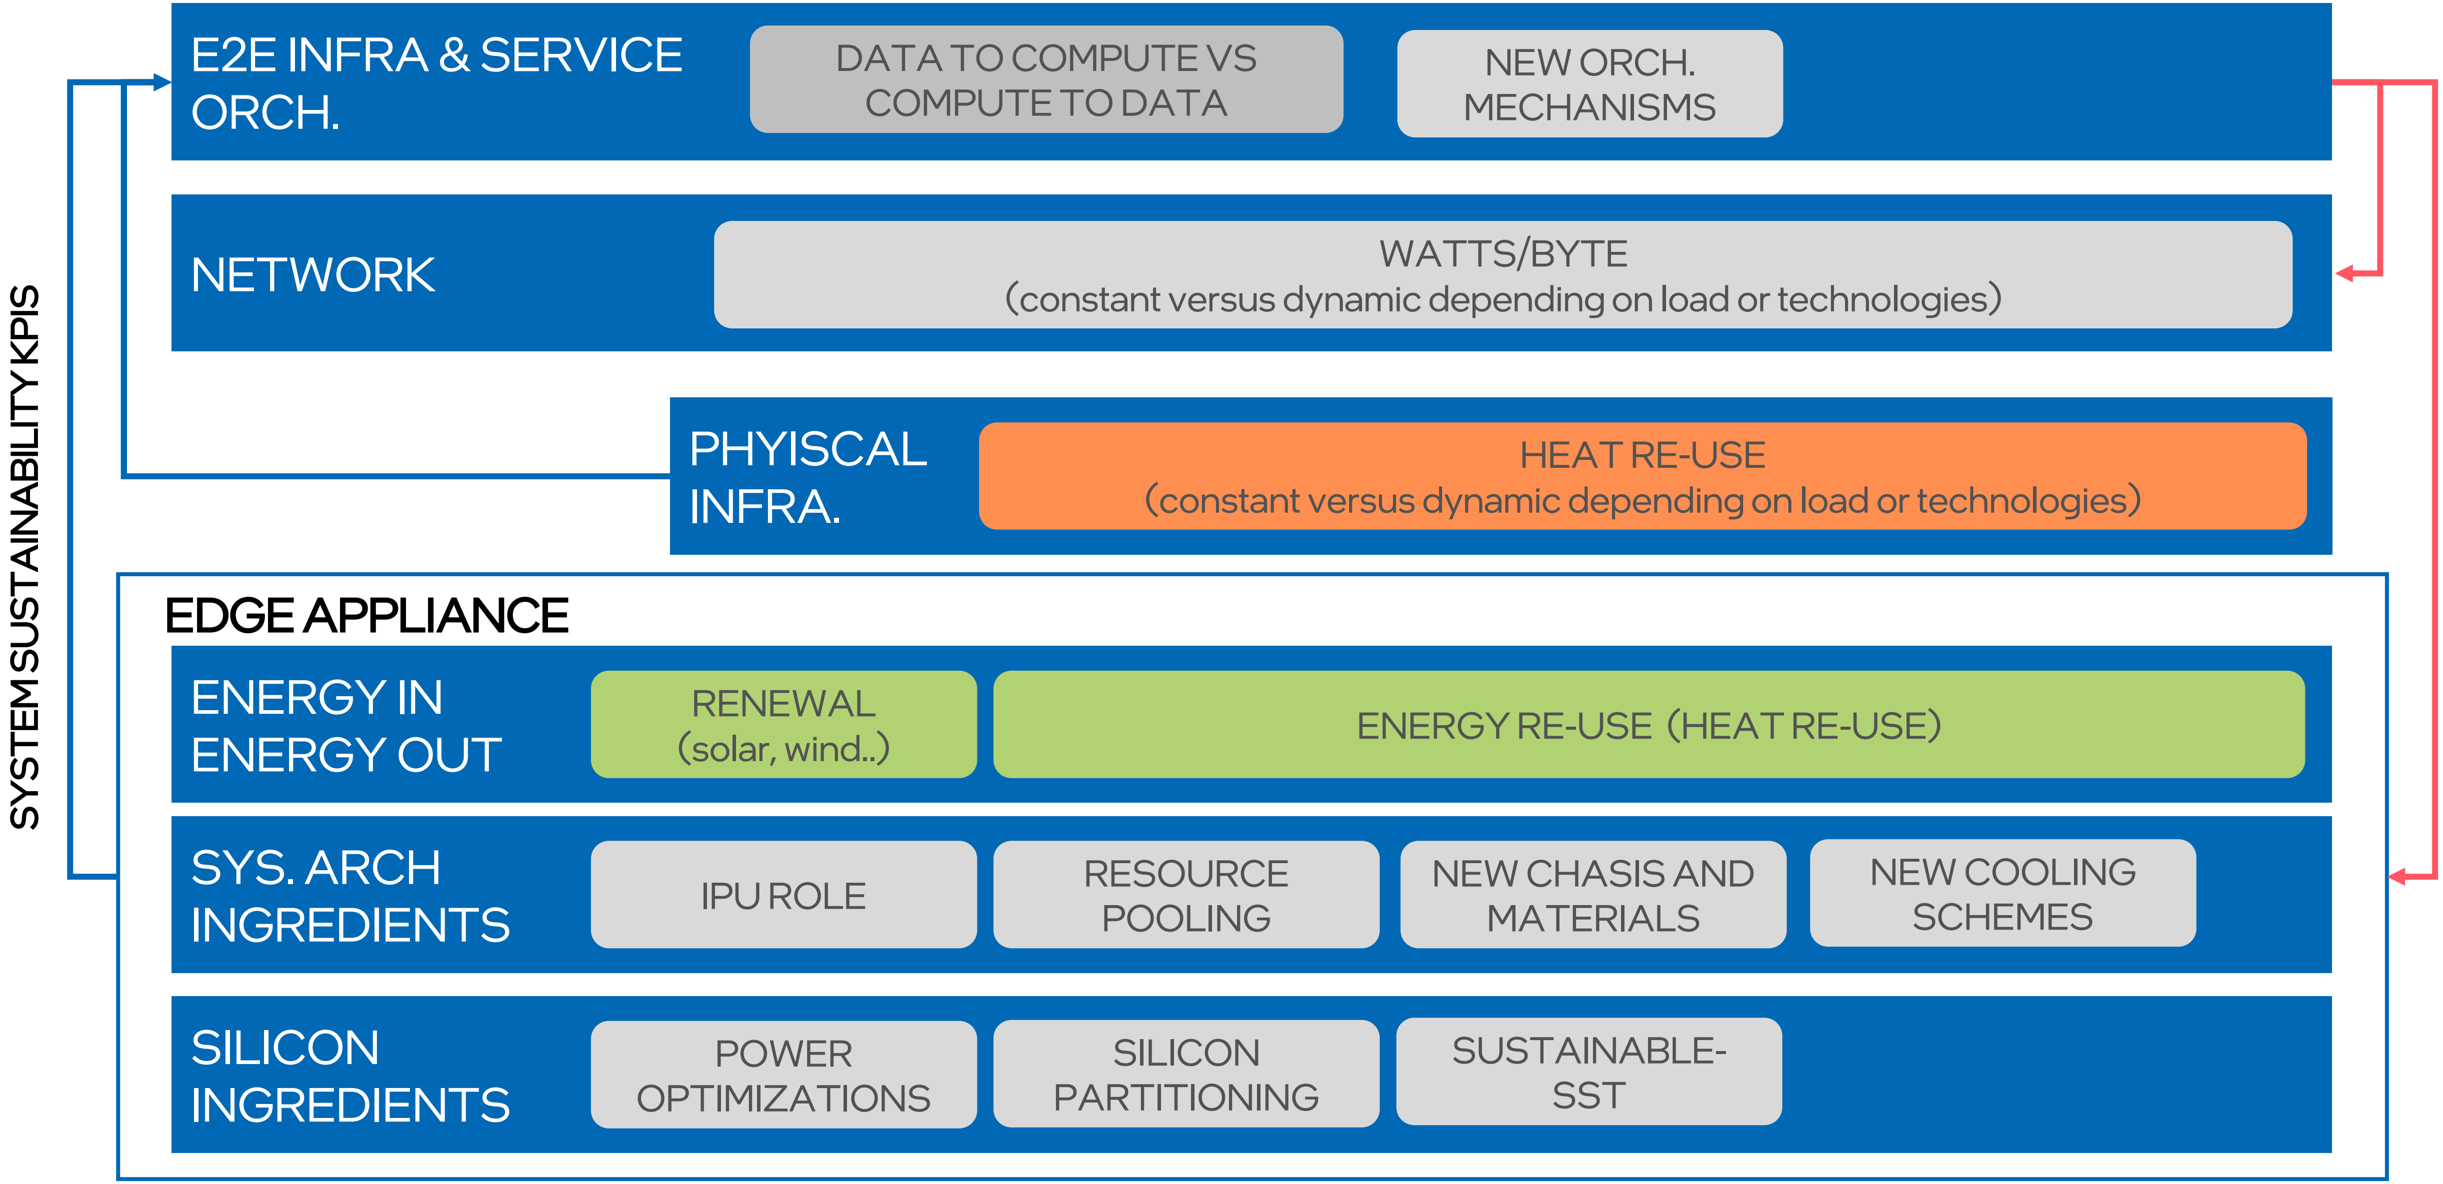

We Need to Adapt Distributed Edge System Architectures

Energy mgmt.,
Observability,
Networking,
Connectivity ..

Chassis architecture,
Cabinet design,
Telemetry, Security ..

System Trade-offs

- CORE VS XEON
- HEAT REUSE
- CORE VS XEON VS VPU VS GPU
- PERF VS POWER
- POWER VS RELIABILITY
- IMMERSION LIQUID VERSUS FAN
- PERF / $$
- CAPEX VERSUS OPEX
- SALABILITY
- RESOURCE DISAGGREGATION
- ACCELERATION VS GP
- …

Lenovo cellnex submer NEARBY COMPUTING

Edge Emersion Cooling Edge
Appliance (FF, mgmt., security,
connectivity, power, ..)..
Platform architecture

…

We Need Application Awareness

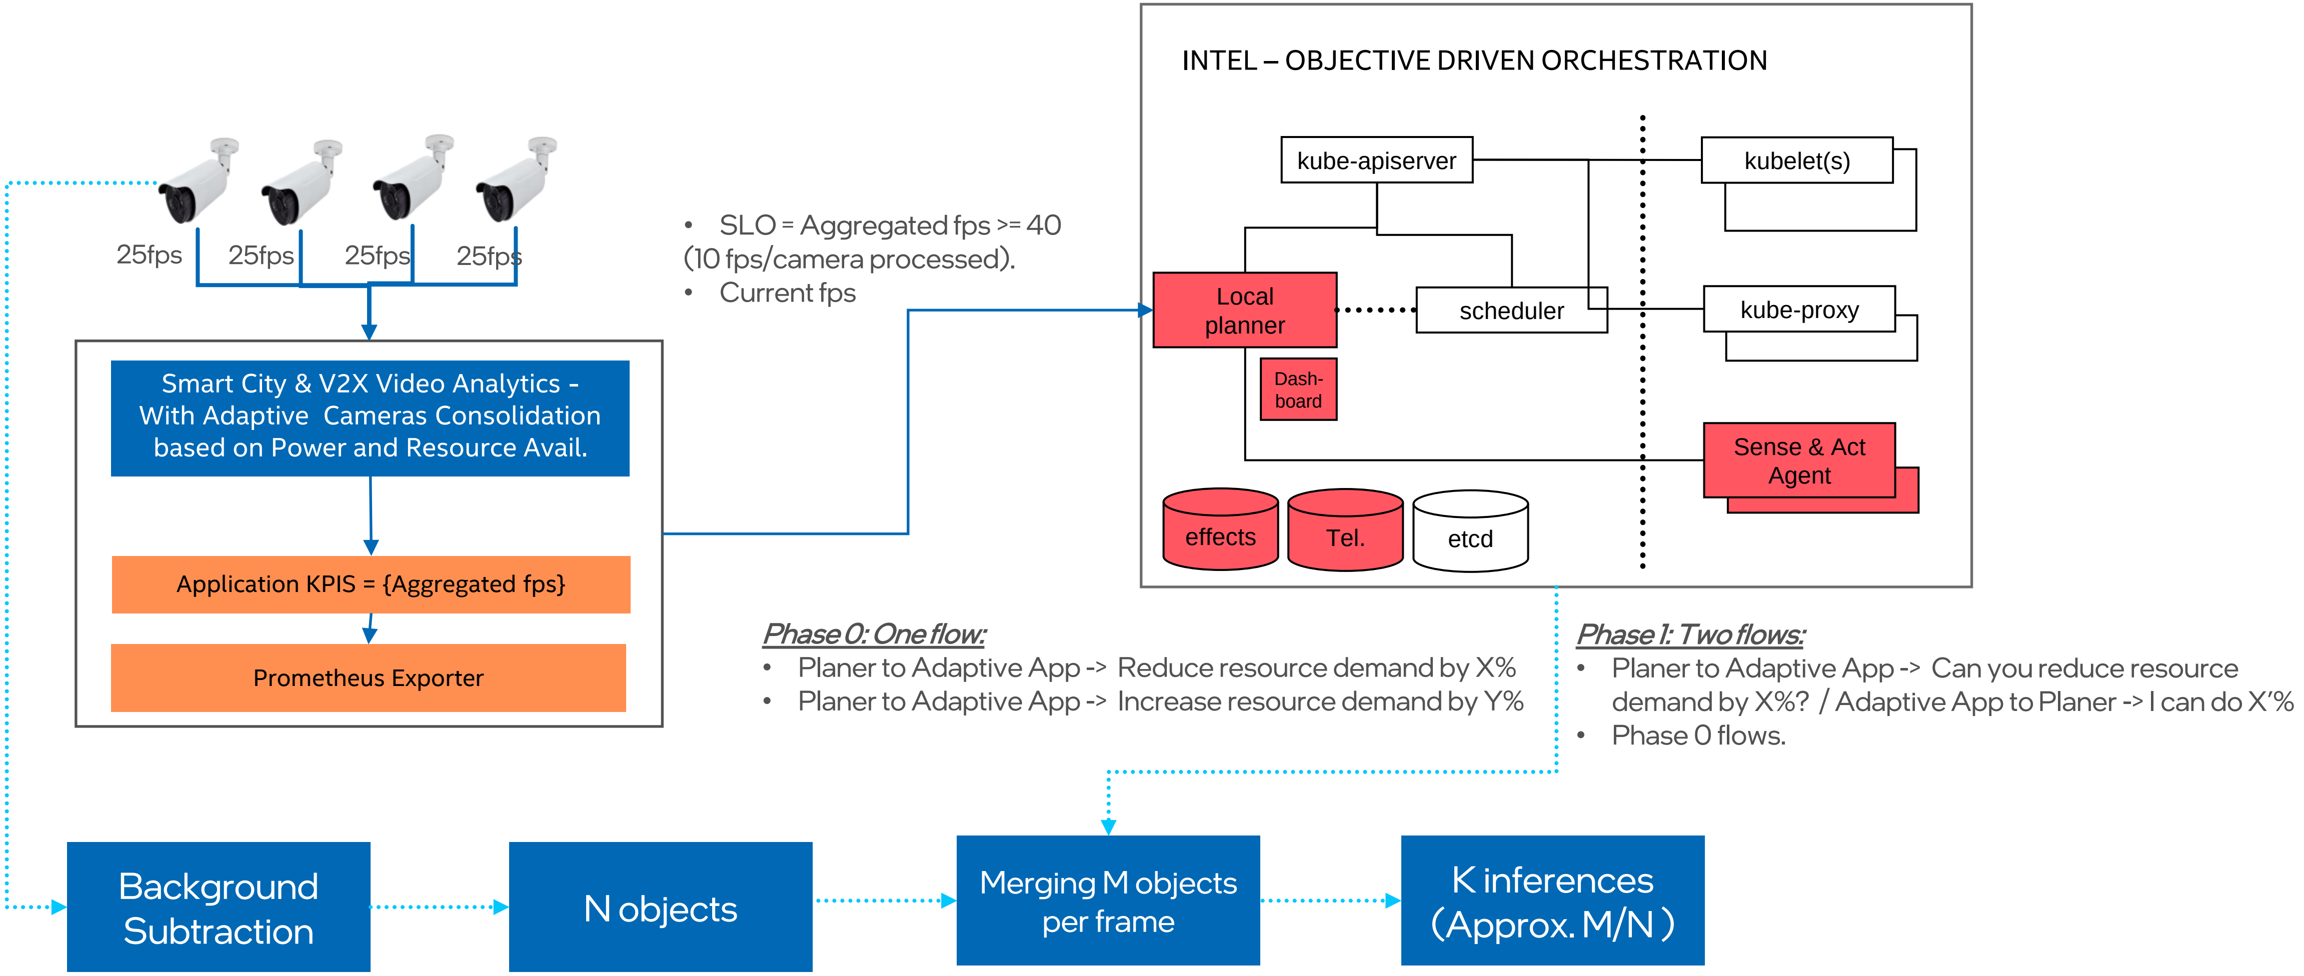

M = 1 -> Highest accuracy Highest Power Consumption
M > 1 -> The bigger M the lower accuracy and the Lower Power Consumption

Call For Action

Call For Action

1. Common services framework for Edge to Cloud for application development:

"Change the paradigm from "local/enterprise/private" or "specific cloud" to anywhere (Edge to Cloud)"

2. Event driven functions and Open APIs for interoperability

3. End-End Sustainability for Edge-Cloud
   - Awareness at infrastructure (HW/SW) and application level
   - Reporting, Telemetry, and feedback loop to achieve sustainability goals
   - Using sustainability as a metric for deployment of distributed apps from Edge to Cloud

References

References

Conference and White Papers

- Composable Architectures for a Sustainable Edge
  - https://networkbuilders.intel.com/solutionslibrary/composable-architectures-for-a-sustainable-edge
- Autonomous Lifecycle Management for Resource-Efficient Workload Orchestration for Green Edge Computing
  - https://ieeexplore.ieee.org/document/9612603
- Convergence of Edge Services & Edge Infrastructure
  - https://ieeexplore.ieee.org/document/9665021
  - https://dblp.org/pid/08/2452.html

Questions?

Francesc Guim, Solutions Architect
francesc.guim@intel.com

intel®