Multi K8s Cluster App & Network Orchestration

Trend : Geo Distributed Computing trend with Edge-computing

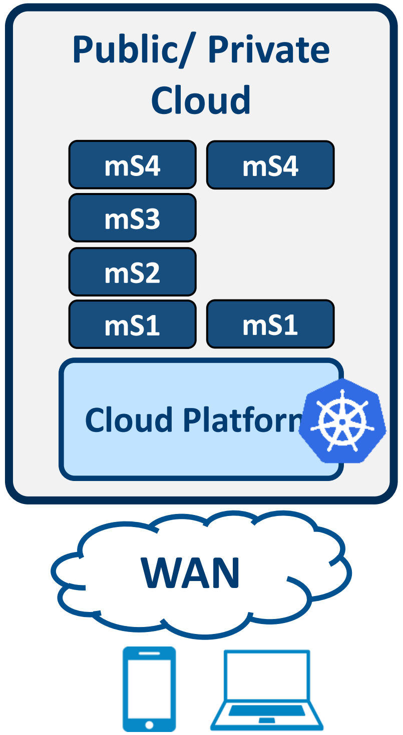

Drivers

- Latency/ Physics
- Bandwidth/ Economics
- Context/ Proximity
- Privacy/ Legal

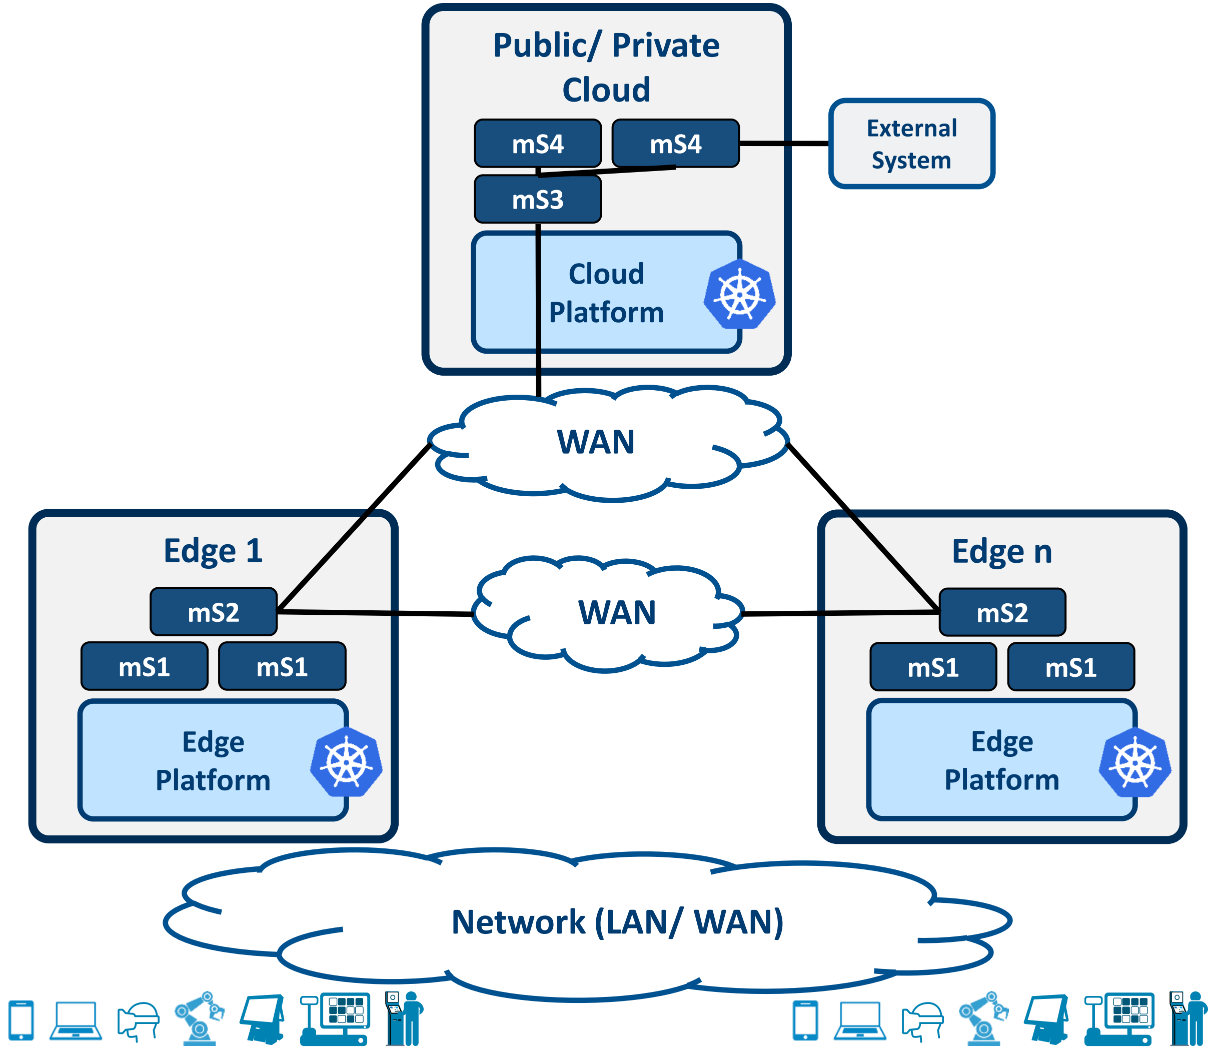

On Demand

Geo-Distributed

App Centric infra Config

Geo-Distributed Computing - few use cases

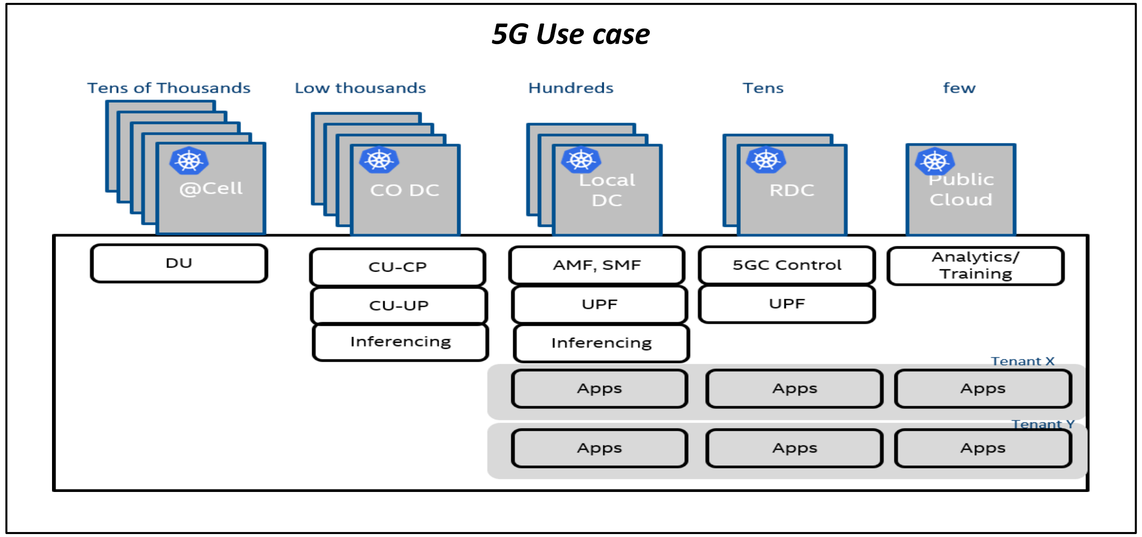

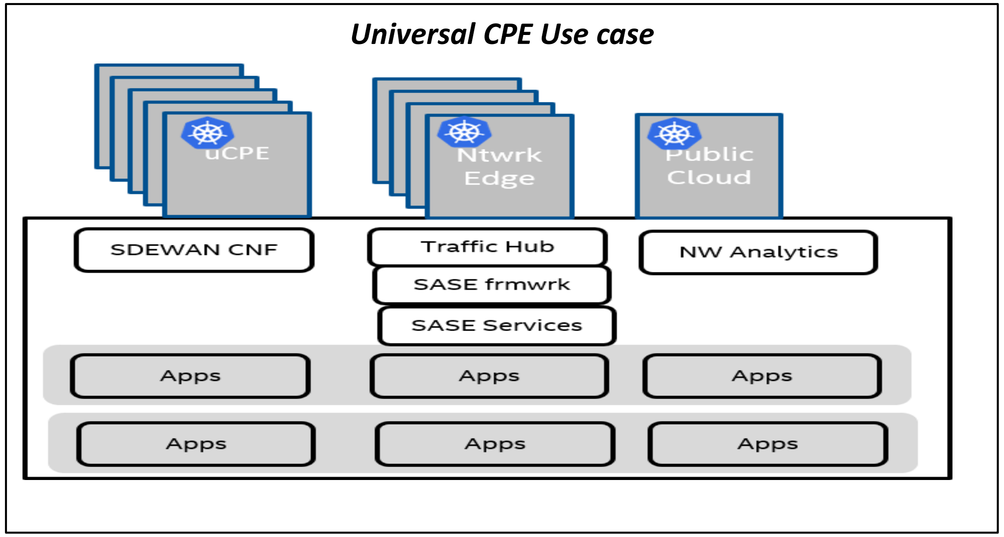

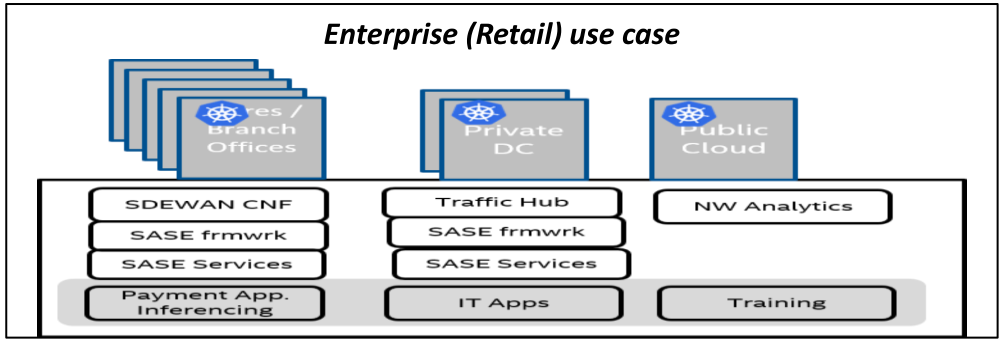

- *Large Number of sites*
- *Computing (Apps across sites) – MEC*
- *Multiple tenant applications along with operator CNFs.*
- *Workload types - VMs, VNFs, CNFs, CNAs and Functions (FaaS)*
- *Note: K8s is becoming choice of workload orchestrator in each cluster*

Multi Edge/Cloud computing scale is similar (or even higher) to Hyper-scalers' scale

Now Telcos, MSPs and Enterprises need @scale Orchestration and Automation solutions

Edge Computing – Similar to Cloud-computing, but with some special needs

Uniform Developer Experience across Clouds and Edges	*Support for all kinds of workloads (VM, Containers and Serverless functions)*
	Easy migration of workloads among Edges and Clouds
	Multi Cloud Uniform Networking (Overlay)
Resource Constraints (Power, Cost, Space)	*Converged Edge supporting IT, OT applications & Network functions*
	Optimized infrastructure software
	Accelerator usage (Hence, awareness without losing platform independence property
Edge requires high security assurance (No physical security in far edges)	*Platform attestation; Confidentiality*
	SASE framework for End clients
	Multi-tenancy isolation, Slice isolation
Ease-of-Use (@Scale requirements are higher than the clouds)	*Infrastructure Orchestration (K8s Cluster Life cycle management)*
	Multi Cluster Distributed Application Orchestration & LCM, Slicing, MEC Orchestration
5G based Edge	*5G dUPF, RAN Acceleration; Analytics (RIC, Non real time RIC)*
	Slicing – Performance & Security isolation; Per Slice SFC of security CNFs

E2E Edge Stack

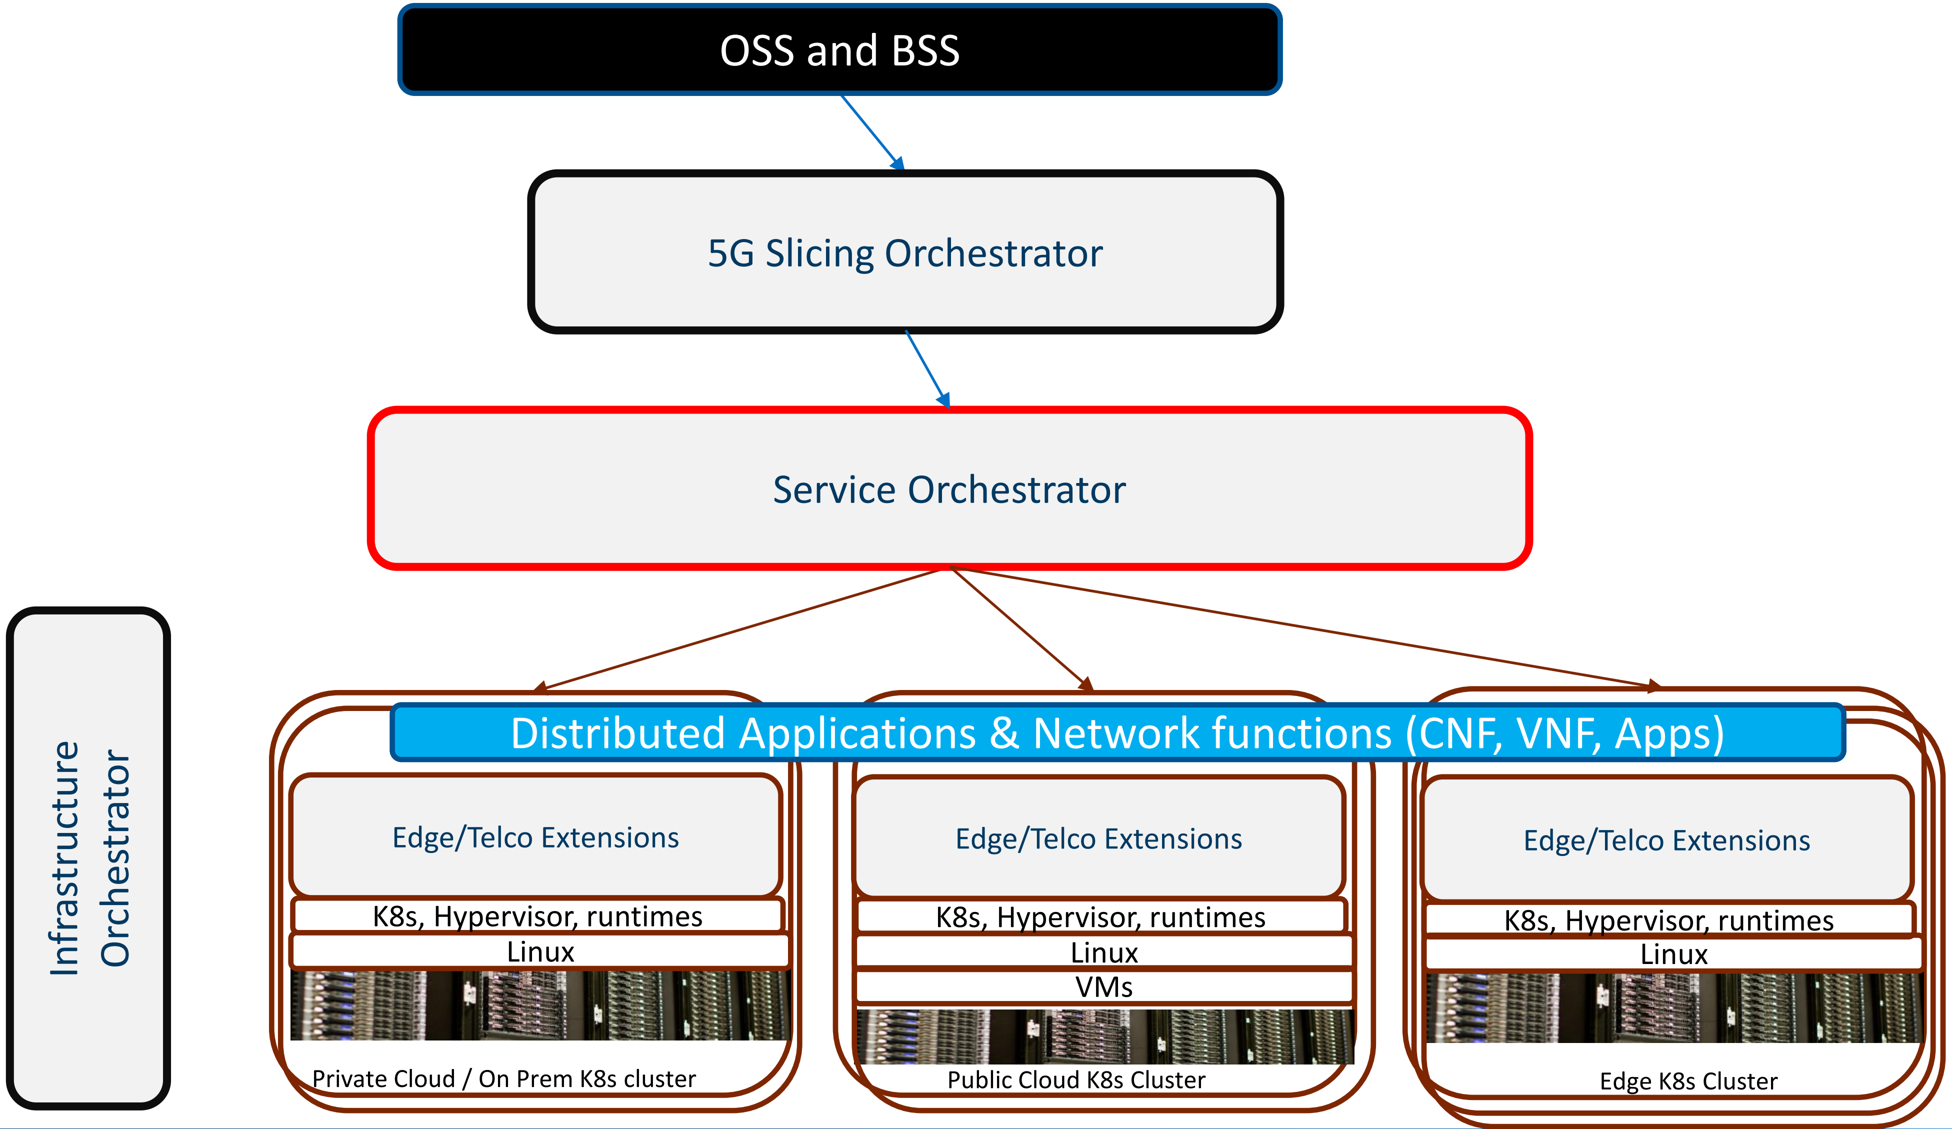
OSS and BSS
5G Slicing Orchestrator
Service Orchestrator
Infrastructure Orchestrator
Distributed Applications & Network functions (CNF, VNF, Apps)
Edge/Telco Extensions
K8s, Hypervisor, runtimes
Linux
Private Cloud / On Prem K8s cluster
Edge/Telco Extensions
K8s, Hypervisor, runtimes
Linux
VMs
Public Cloud K8s Cluster
Edge/Telco Extensions
K8s, Hypervisor, runtimes
Linux
Edge K8s Cluster

Service Orchestrator – Big Picture

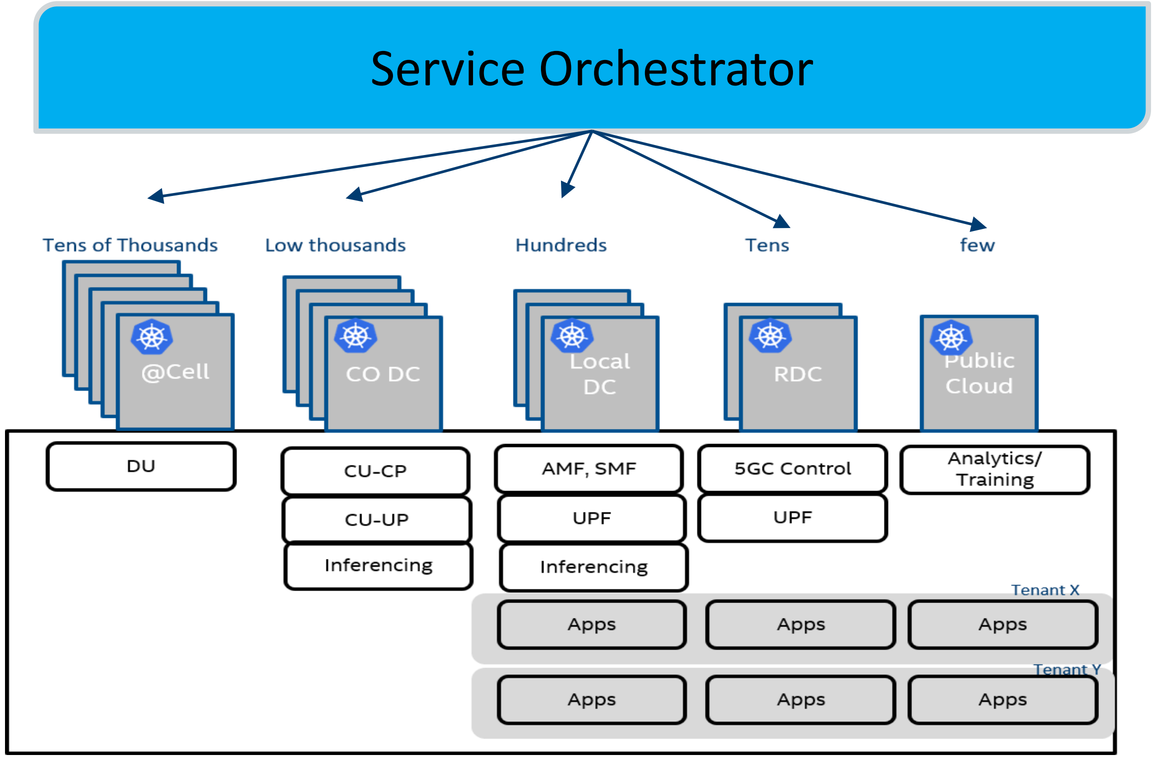

One Click deployment of complex applications & network services across multiple K8s clusters

Comprehensive Status monitoring of deployed complex applications

One Service Orchestrator for both CNF/CNA, VNF/VMs

Self Service Portal for multiple tenants

Comprehensive Analytics platform for Day2 operations

App Centric infrastructure configuration (Service Mesh, SDWAN, L2/L3 switches)

Needs/Requirements – Preparation

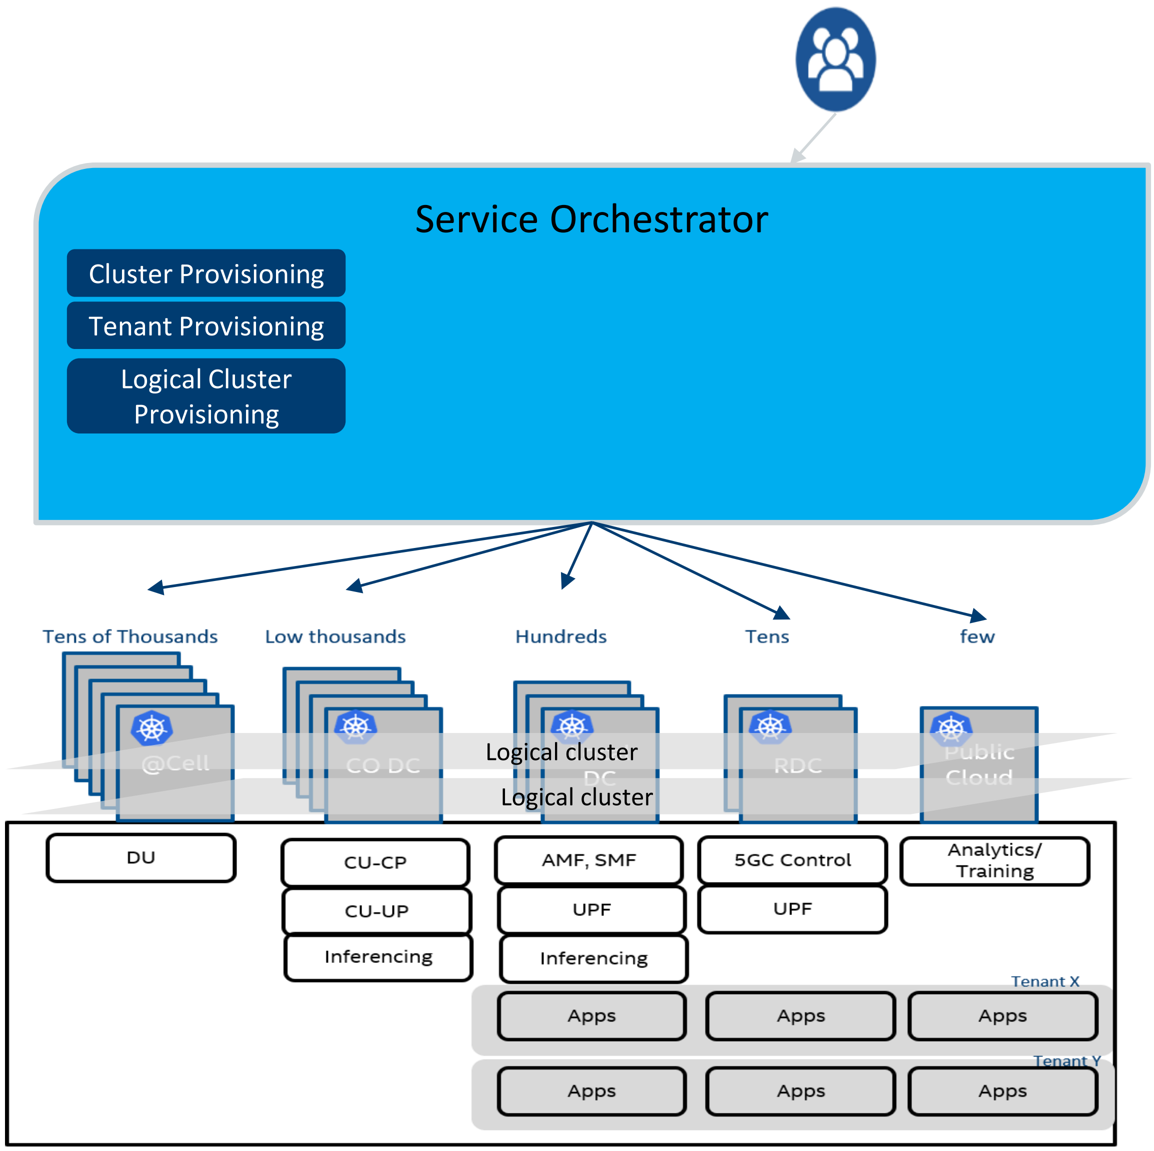

Registration of Clusters

Cluster labels
(Example: Cell tower Edge, CO Edge etc..)
Needed for identifying multiple clusters

Cluster specific configuration
(Few: ISTIO CA provisioning;
Virtual/Provider network preparation)

Tenant registration
Ability to use tenant specific OAUTH2 servers for authenticating tenant admins

Tenant level isolation via RBAC rules

Logical Cluster provisioning across multiple selected clusters

Logical Cluster user and permission provisioning

Needs/Requirements – Application deployment design

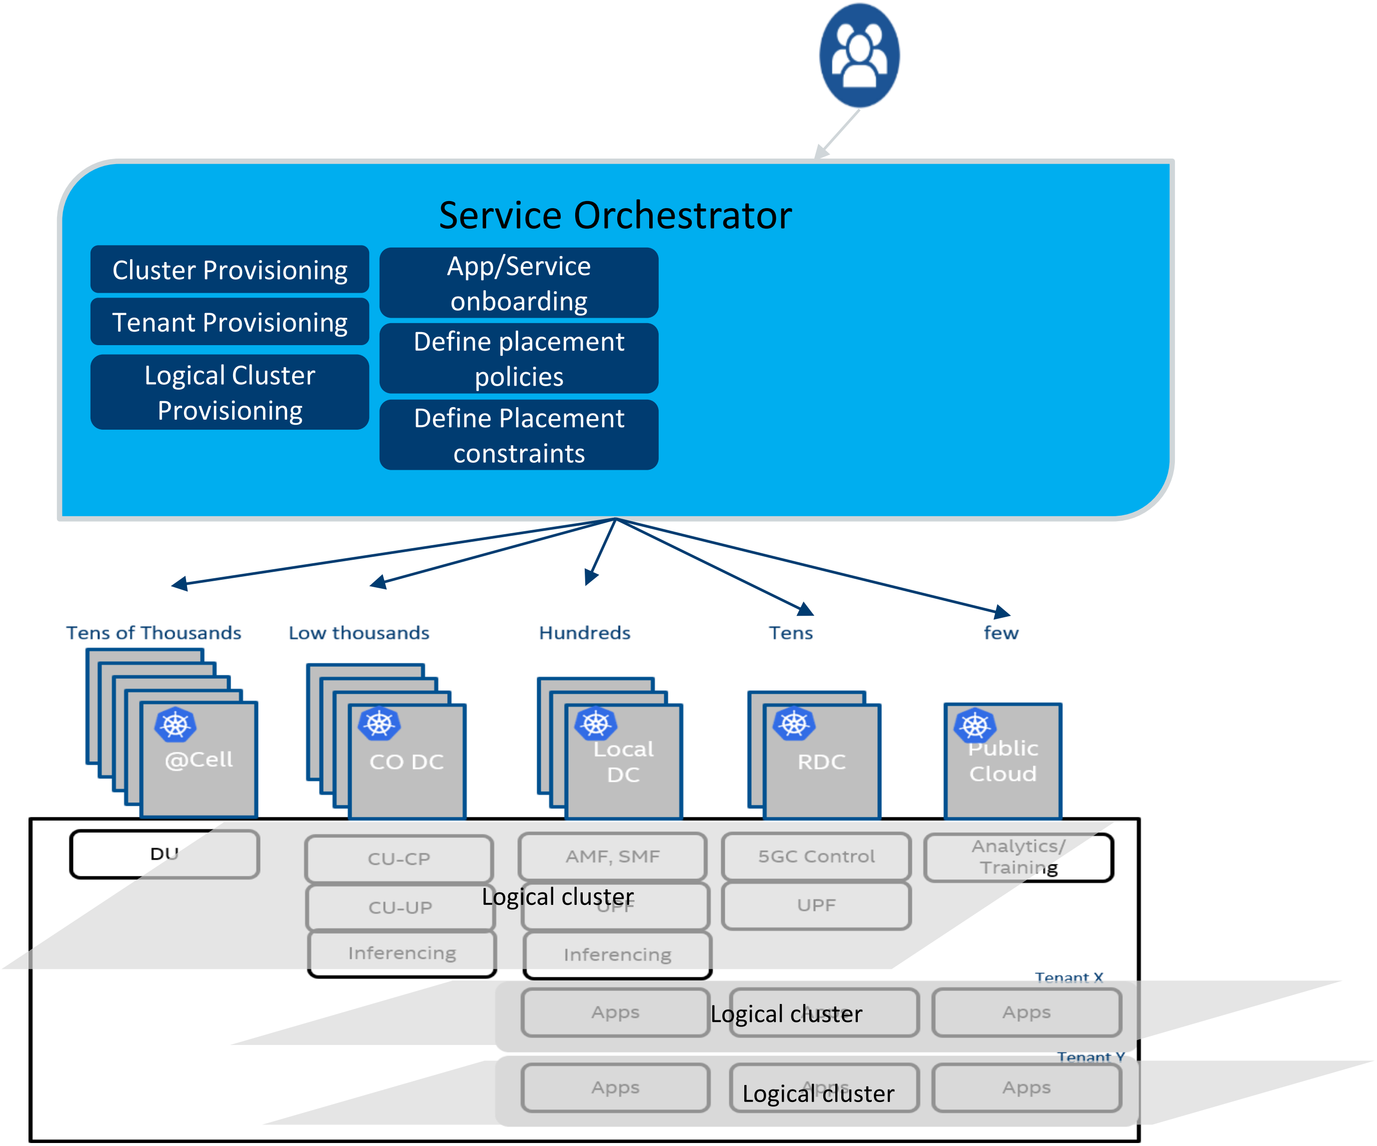

App Onboarding
(Complex Apps & Network Services)

Multiple deployment profiles to ensure same APP can be instantiated multiple times

Placement policies to replicate and distribute workloads across clusters

Placement constraints :
Affinity and Anti-Affinity;
Platform capabilities;
Latency; Cost

Requirements – Workload Customization & Connectivity management

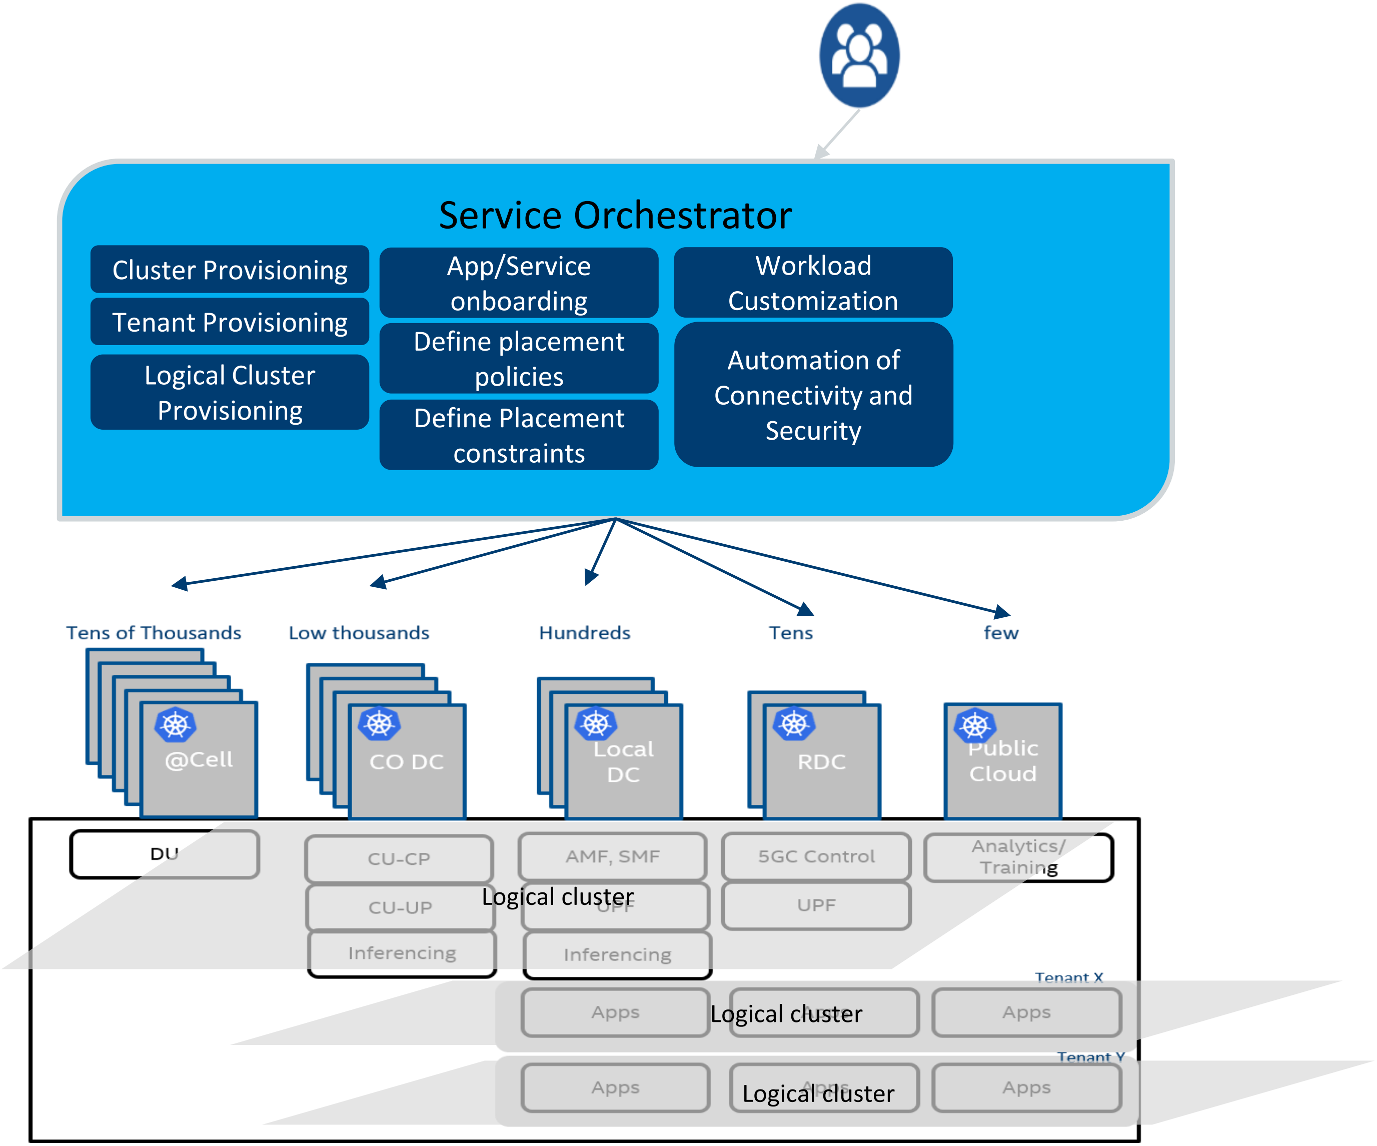

Service Orchestrator

Cluster Provisioning

Tenant Provisioning

Logical Cluster Provisioning

App/Service onboarding

Define placement policies

Define Placement constraints

Workload Customization

Automation of Connectivity and Security

No changes to helm charts/K8s description of applications

Each deployment may have its own customization

Connectivity intent provisioning

- *Enabling inter-micro service communication within or across clusters*
- *Enabling communication to external entitles*
  - *With/Without Mutual TLS*
- *Multi Cluster DNS management*

Dynamic provisioning with LCM of Applications

Extensible framework to add new capability controllers

Requirements – Operations

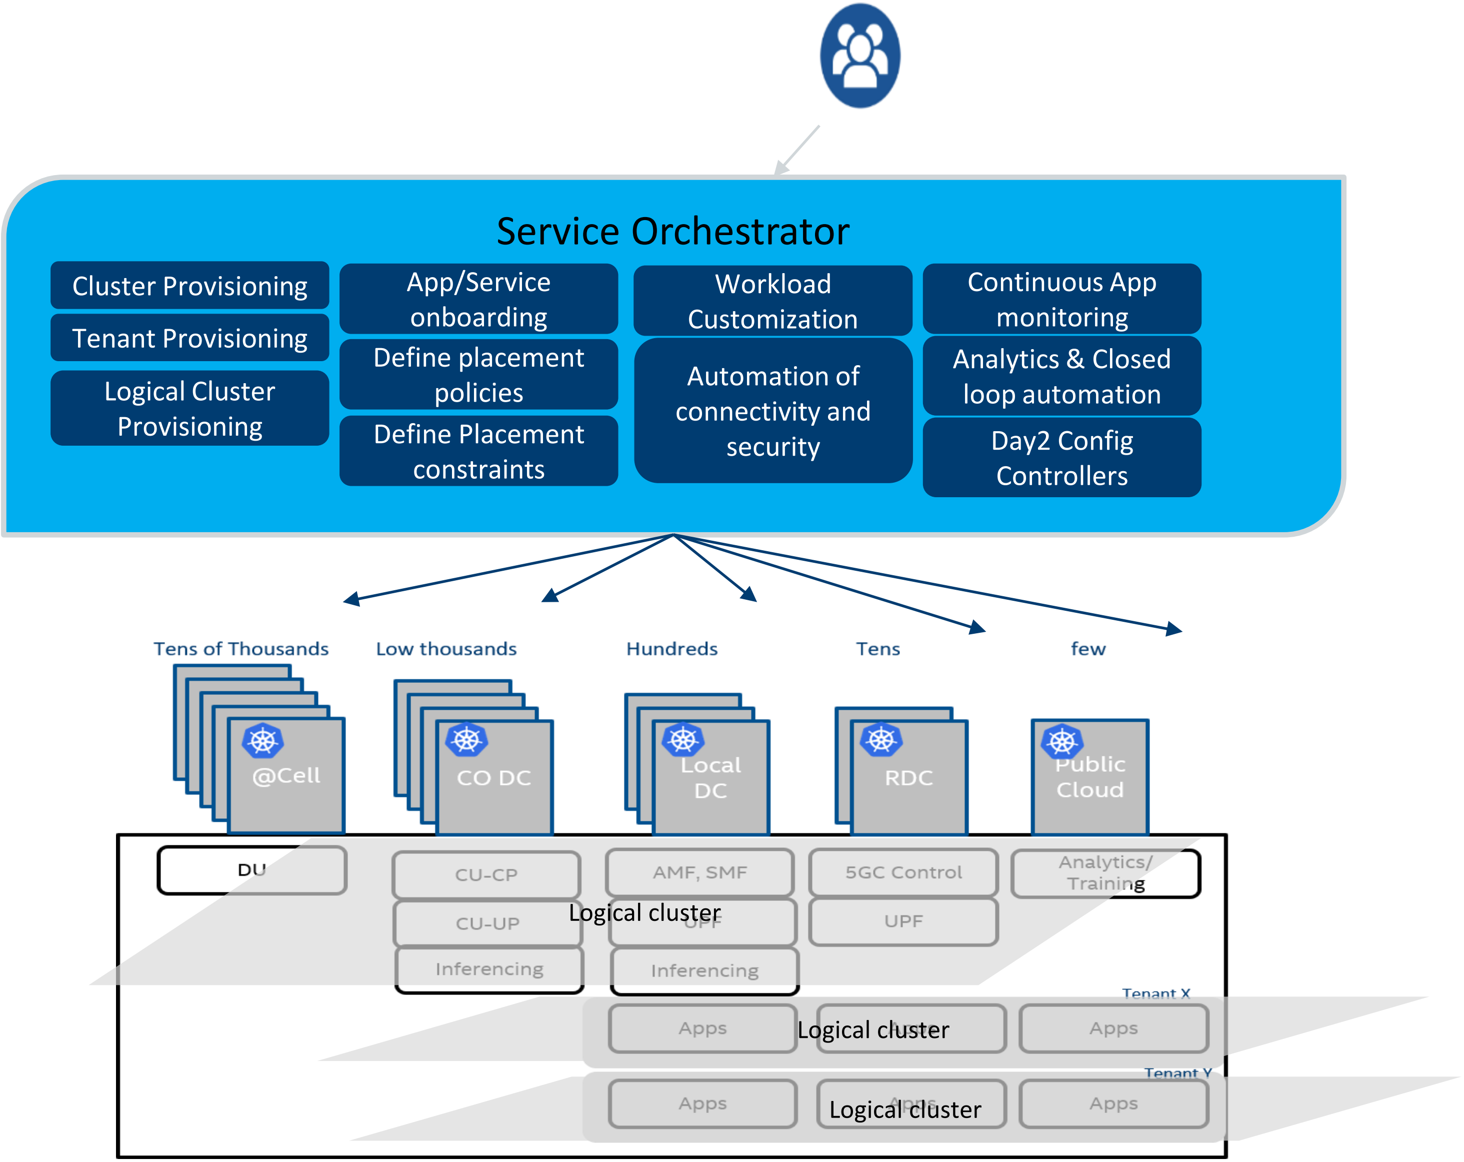

Service Orchestrator

Cluster Provisioning

Tenant Provisioning

Logical Cluster Provisioning

App/Service onboarding

Define placement policies

Define Placement constraints

Workload Customization

Automation of connectivity and security

Continuous App monitoring

Analytics & Closed loop automation

Day2 Config Controllers

Continuous monitoring of complex Application (Across clusters, apps and micro-services)

Comprehensive report on the application status

Analytics framework

- *Metric collection across clusters & apps*
- *Long term central store (Time Series)*
- *Training framework*
- *Closed loop policy management*

Day 2 Configuration

- *Configuration of apps/network-functions that are already deployed.*
- *Various types of configurations (CR based, RESTful based or Netconf/yang based)*

EMCO – Edge Multi Cluster Orchestrator (Opensource)

EMCO is an implementation of Service Orchestrator
Addressing majority of requirements; Extensible architecture allows new automation requirements

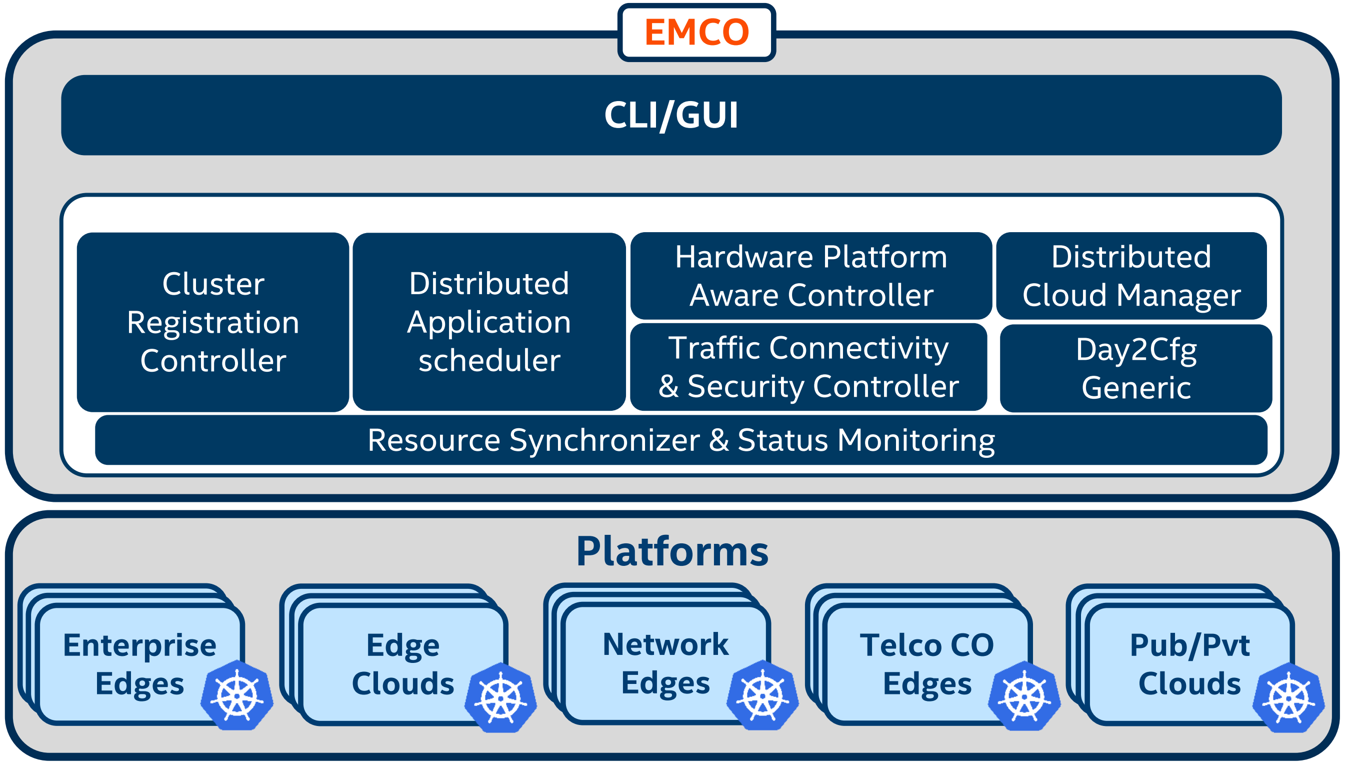
EMCO
CLI/GUI
Cluster Registration Controller
Distributed Application scheduler
Hardware Platform Aware Controller
Distributed Cloud Manager
Traffic Connectivity & Security Controller
Day2Cfg Generic
Resource Synchronizer & Status Monitoring
Platforms
Enterprise Edges
Edge Clouds
Network Edges
Telco CO Edges
Pub/Pvt Clouds

- **Cluster Registration Controller** registers clusters by cluster owners
- **Distributed Application Scheduler** provides simplified, and extensible placement; tenant mgmt; LCM implementation
- **Hardware Platform Aware Controller** enables scheduling with auto-discovery of platform features/ capabilities; Others: Cost, Power Savings, Latency aware… (WIP)
- **Distributed Cloud Manager** presents a single logical cloud from multiple edges
- **Traffic Connectivity controller** auto-configure service mesh (ISTIO) and security policy (NAT, firewall), DNS and SLB entities of edges - WIP
- **Day2 generic configuration** configures Day2 configuration of any app/network function via templates & configs - WIP
- **Resource Synchronizer & Monitoring** synchronizes resources across multiple edge/cloud platforms and then monitors the status of deployed resources

Plan for 2021 (subject to change)

21.03/21.06 release

Quality improvements and fixing some technical debt

- Robustness of RSYNC (retries, restart of RSYNC resumes the synchronization)
- Increase unit test coverage
- Automation of use cases

Features

- Helmv3 templating support
- Platform capabilities-based selection (Placement)
- Automation of CoreDNS Servers for inter cluster microservice communication
- Upgrades & Updates (when new helm chart is released; when new cluster is added; when existing cluster is removed from placement; when intents are changed)
- Network Policy Automation

For 21.09 and beyond

Continue to increase quality levels mainly with respect to scalability and security (including E-W security)

Features:

- Traffic Controller (Automation of SD-WAN, ISTIO resources)
- Distributed Cloud enhancements (Namespace Quotas, Labels and as many that can be configured on per namespace basis).
- Application Configuration support via CRs
- Generic CRD Controller to JSON based RESTful API
- SFC Automation
- Few Closed loop action related actions (Increase replica count; Restart etc..)
- External DNS Automation support
- Helm hooks support
- Dependency logic support across multiple apps in a composite app
- Referential integrity fixes.

Plan for 2022 (subject to change)

Continue to increase quality levels mainly with respect to scalability and security

Features:

- Capacity aware Placement
- Kustomize packaging support
- KUDO.IO operator packaging support
- Security and Performance Isolation automation (as required by network slicing)
- ISTIO CA Certificate Automation
- Dynamic Slicing Orchestration
- Latency Aware placement (Based 5GFF)
- MEC Support
  - Automation of UPF via SMF/PCF for traffic redirection
  - SNAT/DNAT/FW/IPSEC sidecar for UPF that help in traffic redirection
- Changing MongoDB to more open DB
- Network Slicing Orchestration support
- Support for ARC and Anthos for resource synchronization
- Support for GitOps at EMCO API level.

EMCO Integrations

EMCO Integrations

Part of these commercial solutions
Part of AMCOP solution from Aarna: https://www.aarnanetworks.com/amcop
Commercial Support
Calsoft SI: https://www.calsoftinc.com/news/calsoft-announces-commercial-support-for-akraino-led-icn-integrated-cloud-native-blueprint/

Few Blueprints in LFE/Akraino use EMCO for Multi Cluster Orchestration
free5GC deployments using EMCO by Aarna networks

ONAP uses EMCOv1 to Onboard/design network services and deploy on K8s clusters

ONAP Slicing Orchestrator uses EMCOv1 for Day 2 LCM

With Azure and Anthos (To be done)

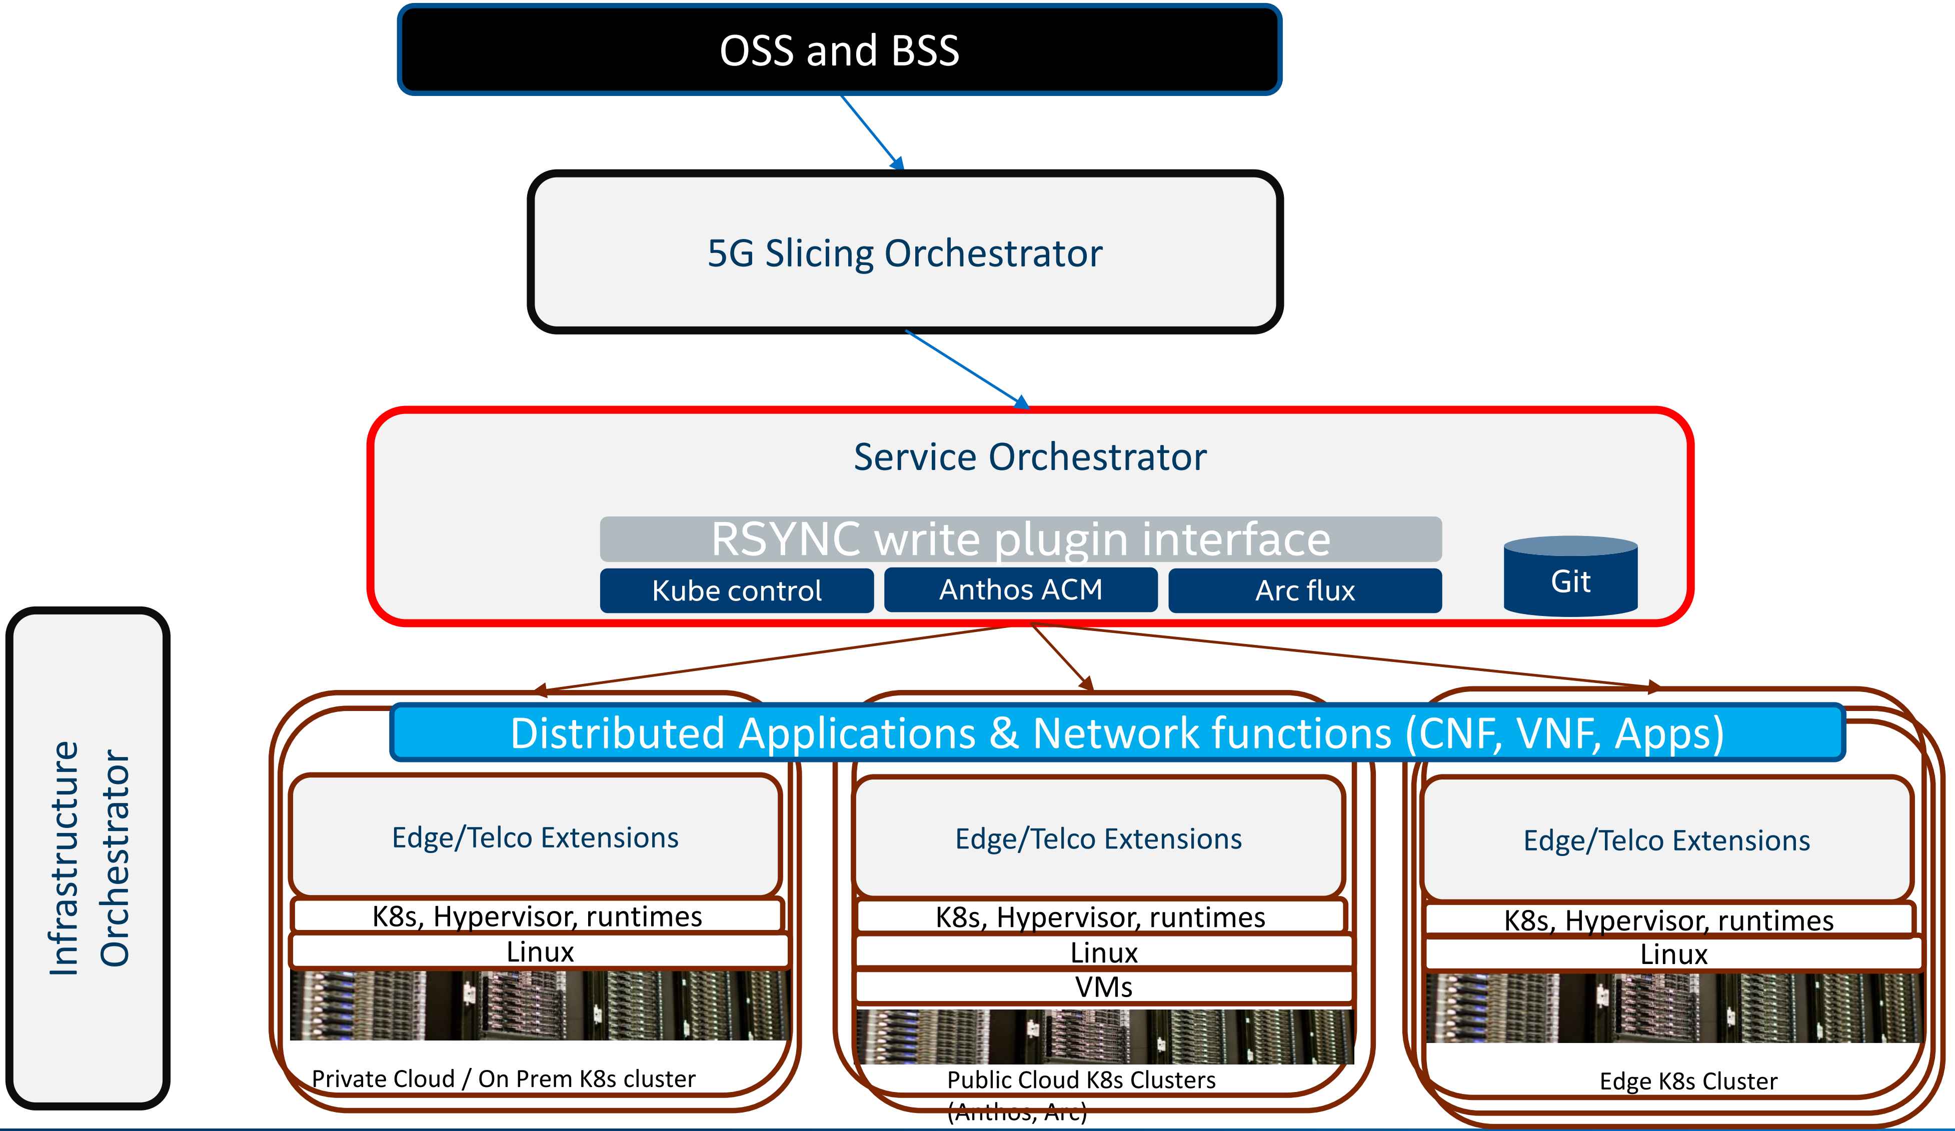

OSS and BSS
5G Slicing Orchestrator
Service Orchestrator
RSYNC write plugin interface
Kube control
Anthos ACM
Arc flux
Git
Infrastructure Orchestrator
Distributed Applications & Network functions (CNF, VNF, Apps)
Edge/Telco Extensions
K8s, Hypervisor, runtimes
Linux
Private Cloud / On Prem K8s cluster
Edge/Telco Extensions
K8s, Hypervisor, runtimes
Linux
VMs
Public Cloud K8s Clusters (Anthos, Arc)
Edge/Telco Extensions
K8s, Hypervisor, runtimes
Linux
Edge K8s Cluster

Slicing

Network Slicing Phase 1 – Starting Point for providers

(Note: NF placement in various Edges is just one example)

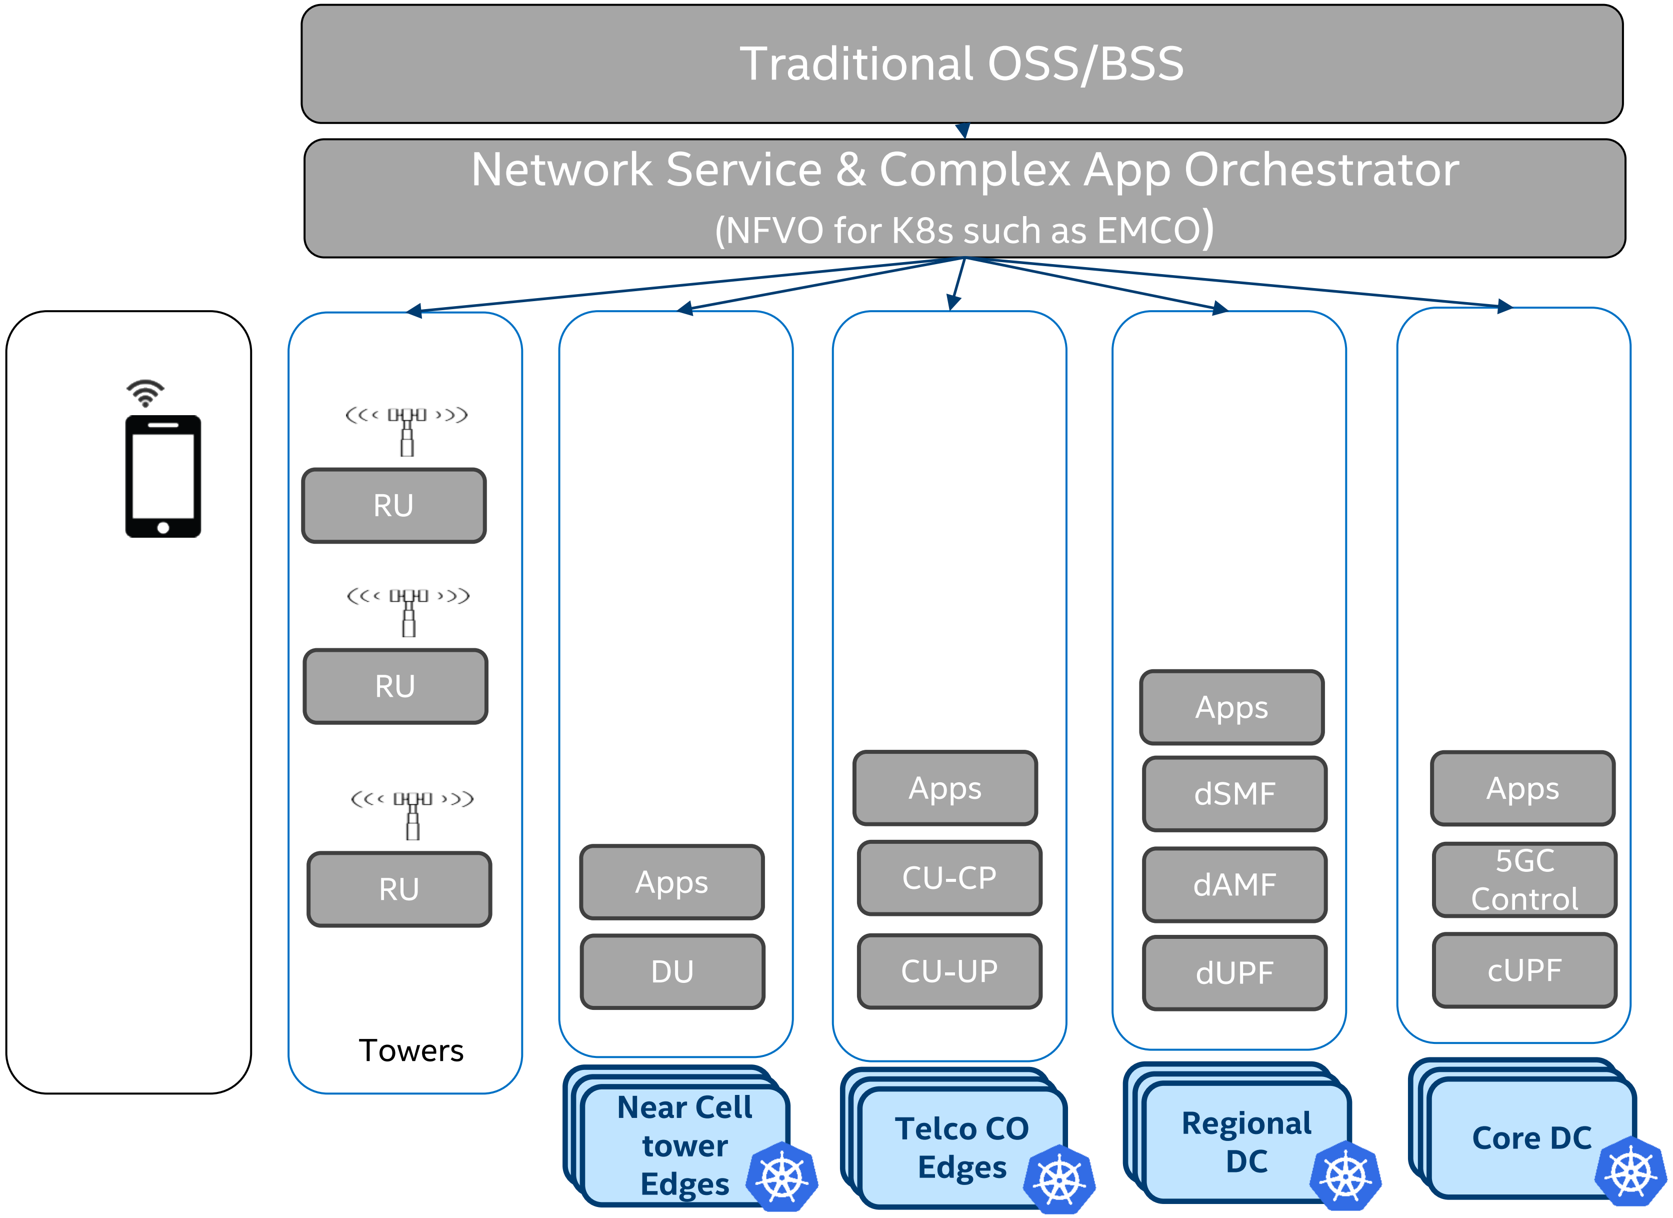

Traditional OSS/BSS

Network Service & Complex App Orchestrator
(NFVO for K8s such as EMCO)

RU
RU
RU
Towers

Apps
DU
Near Cell tower Edges

Apps
CU-CP
CU-UP
Telco CO Edges

Apps
dSMF
dAMF
dUPF
Regional DC

Apps
5GC Control
cUPF
Core DC

All kinds of traffic flows are considered equal. Only QoS treatment is taken care based on the traffic priority. Essentially, no slicing (or considered as one slice).

Works good for generic traffic such as Voice calls, Consumer broadband and SMS etc..).

Only for eMBB by majority of operators

Notes:
At each location, there could be multiple instances of a given NF. 3GPP Selection procedures will take care of selecting the right NF during PDU session establishment.
'd' in dUPF, dSMF and dAMF indicating 'Distributed'. Normally, they are expected at central place. But, for scalability reasons, thinking is to deploy at multiple places as these are mostly commonly used Control Plance CNFs

Network Slicing Phase 1.5 – Enterprise Application namespaces

(aka Multi Access Edge computing)

Enterprise Apps and security/non-5G network functions

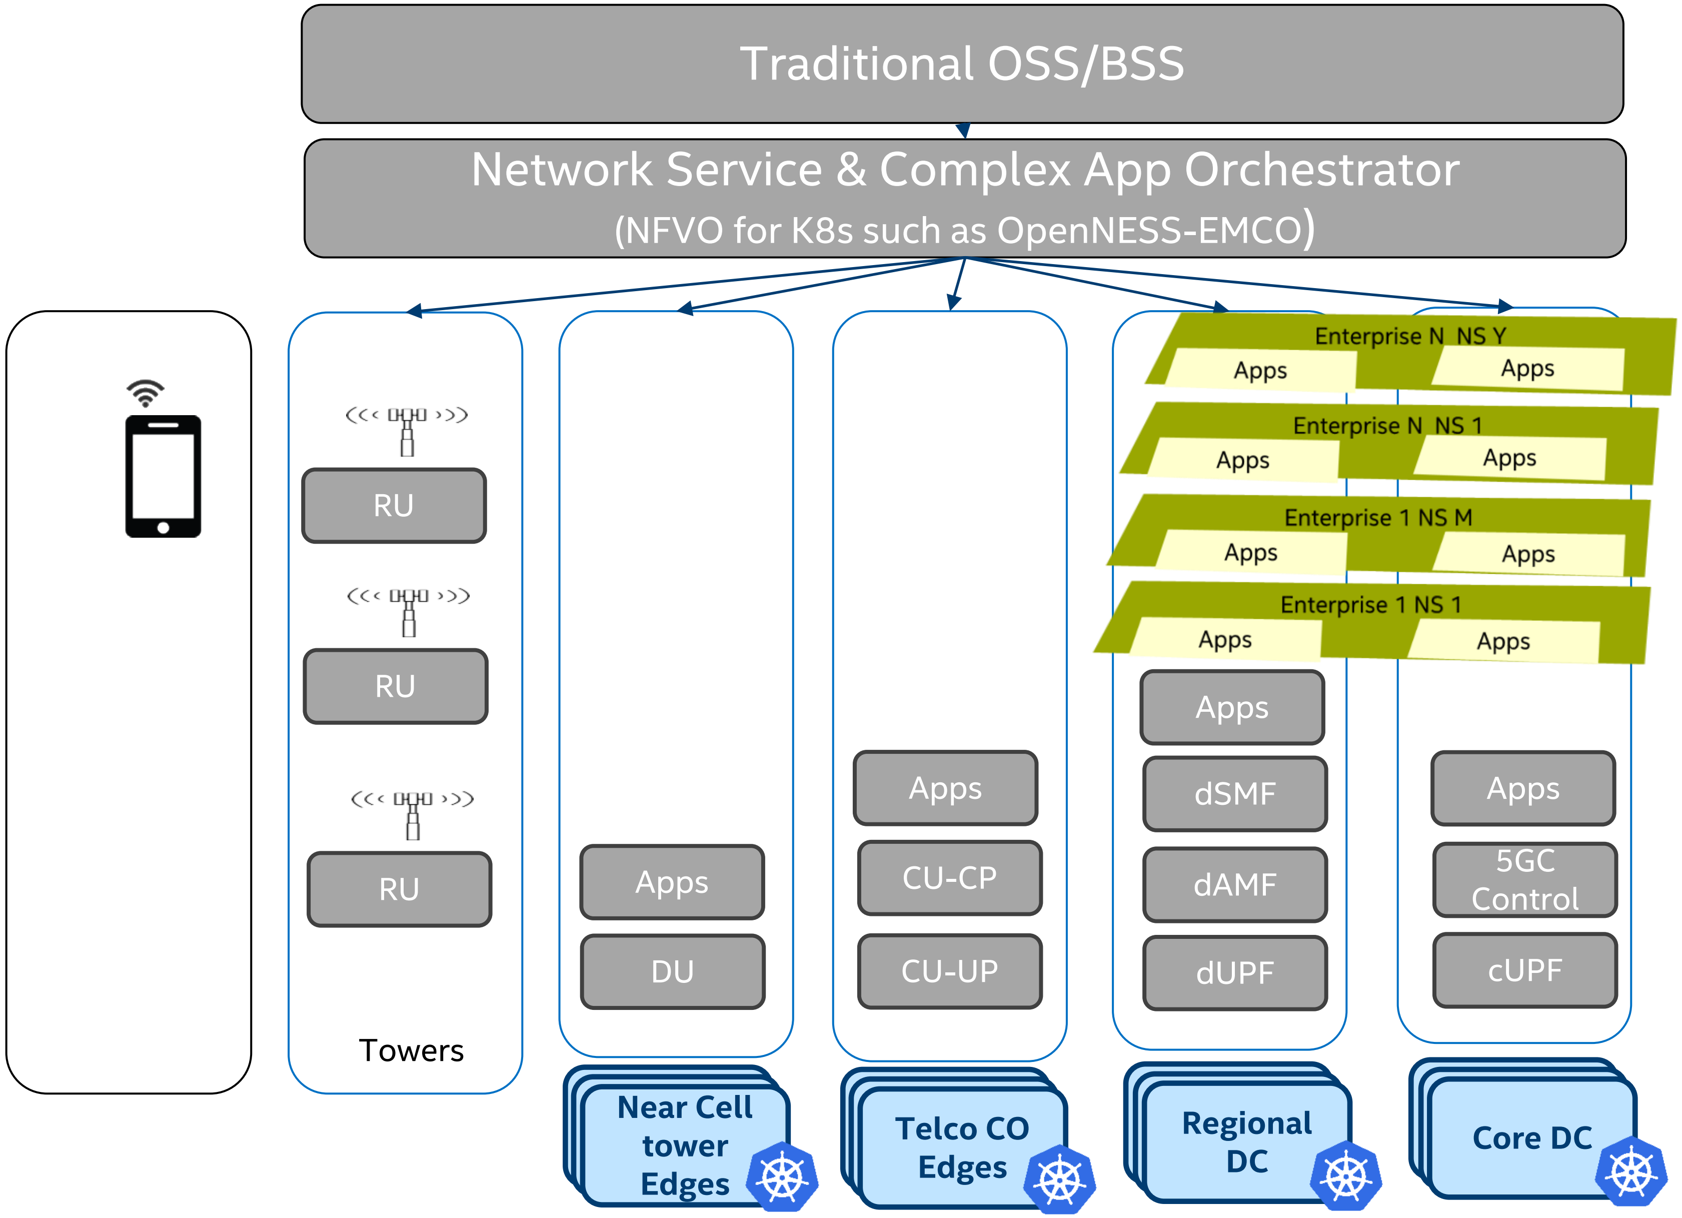

Operators allow Enterprises to deploy their applications near UPFfor distributed computing.

Already happening to some extent – Operators partnering with Hyperscalers and others to provide managed services to Enterprises (Example: Verizon and AWS wavelength).

Belief is that operators would themselves provide application NS on the same infrastructure as the UPF for better resource usage and latency (Need good security and performance isolation solutions though)

Note: These are typical use cases – Enterprise application and there is no special need for ultra low latency.

EMCO provides this functionality to great extent (Traffic Steering is a roadmap feature)

Network Slicing Phase 2 – Multiple generic slices

Operator software - NFs, Apps, Mgmt

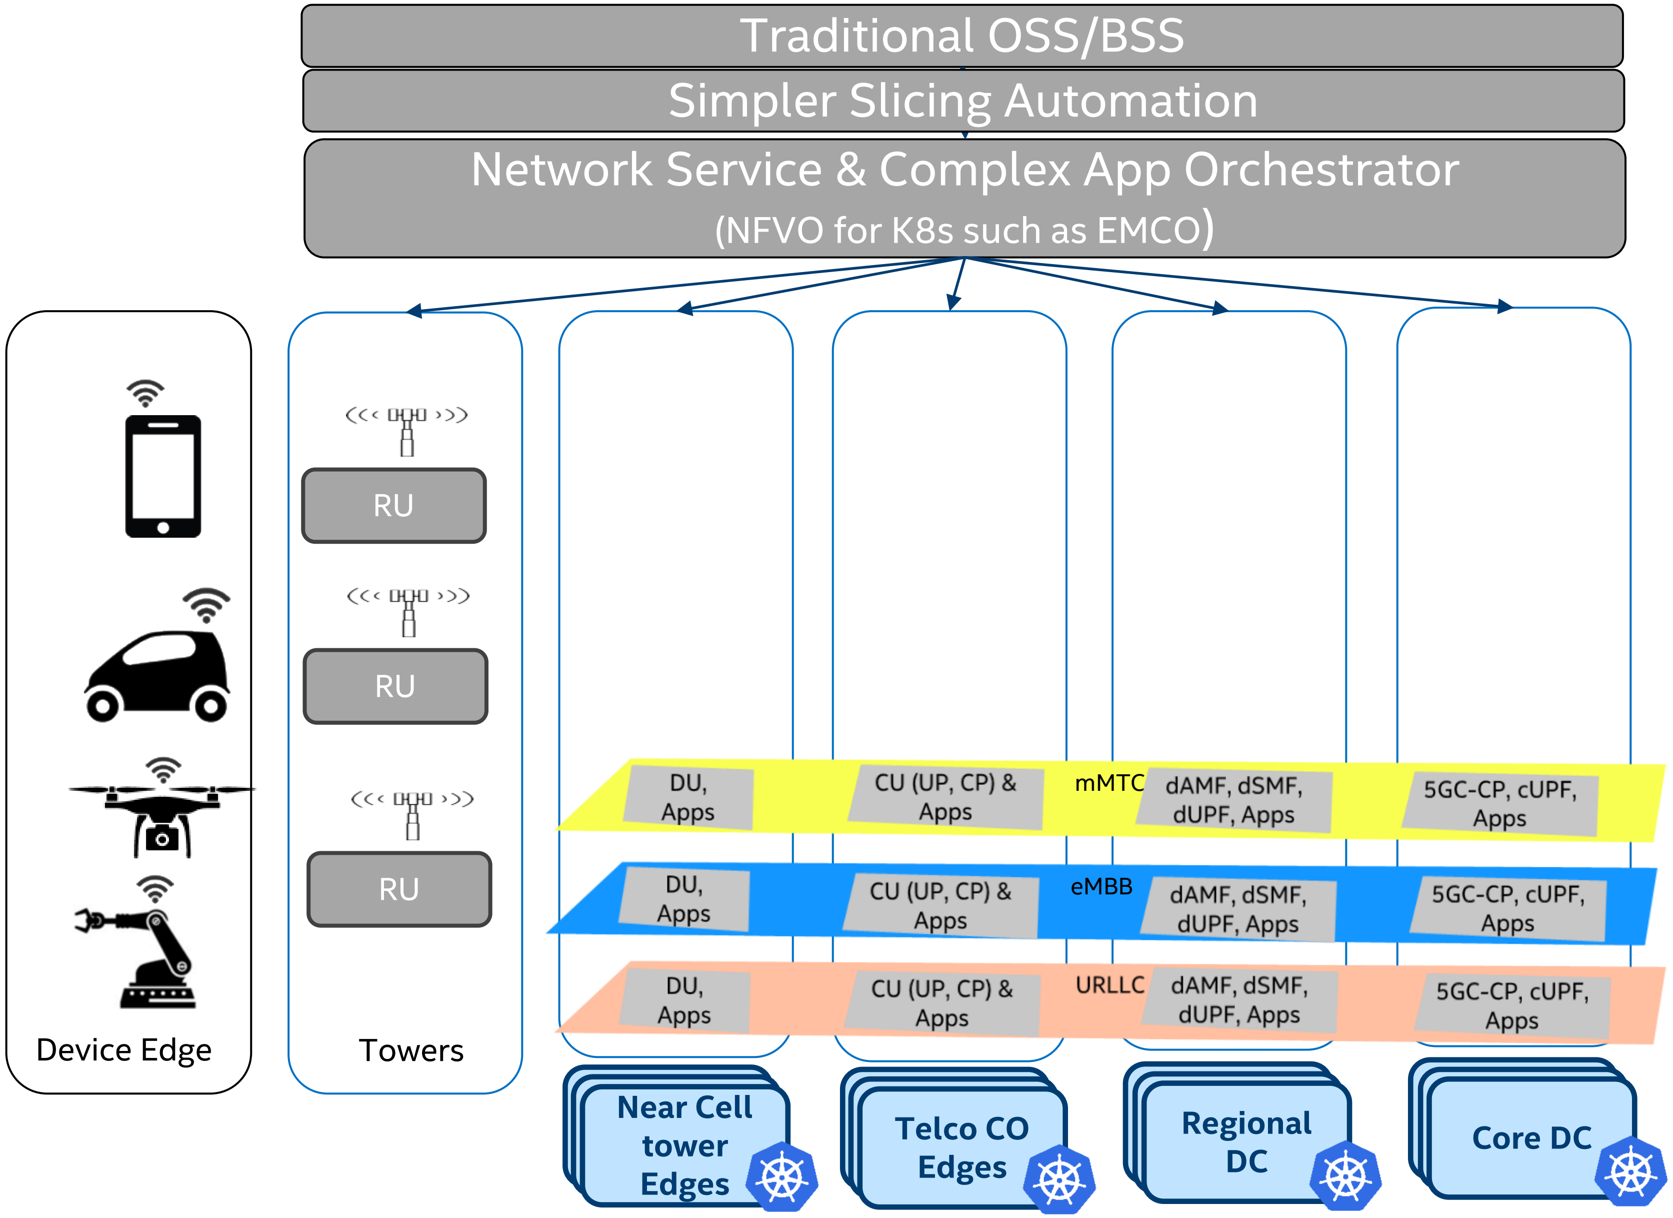

Hypothesis:

- Support for 3 or few generic slices for different use cases.
- No NF Sharing
- Common infrastructure
- NFs as CNFs.
- Simpler Slicing Orchestration as there is no dynamic slice creation/deletion.
- All slices are provider controlled.

Why?

- Simpler deployment
- Dedicated NFs to slices don't need to have any special knowledge in NFs in regard to isolation as it leaves the isolation to infrastructure.
- Easy to visualize and control resources as they share common infrastructure

Dynamic Network Slicing

What is Dynamic Slicing : A way to Create/Delete/Modify Slices on demand basis; Allocate resources from shared pool of resources; Adjust resources dynamically to meet SLA (Service Assurance)

Need for dynamic slices:

Enterprises/ Organizations that have highest security requirements

Data leaking concerns with No Sharing of 5GS NFs and associated services

Enterprises/ Organizations that have highest performance determinism

Processing isolation concerns resulting from shared NFs

Enterprises/ Organizations that don't trust operator technologies of other countries

Bring their own 5GS and associated software

Organizations provide managed services to Enterprises

Network Slicing Phase 3 – Dynamic network slices

Operator software - NFs, Apps, Mgmt

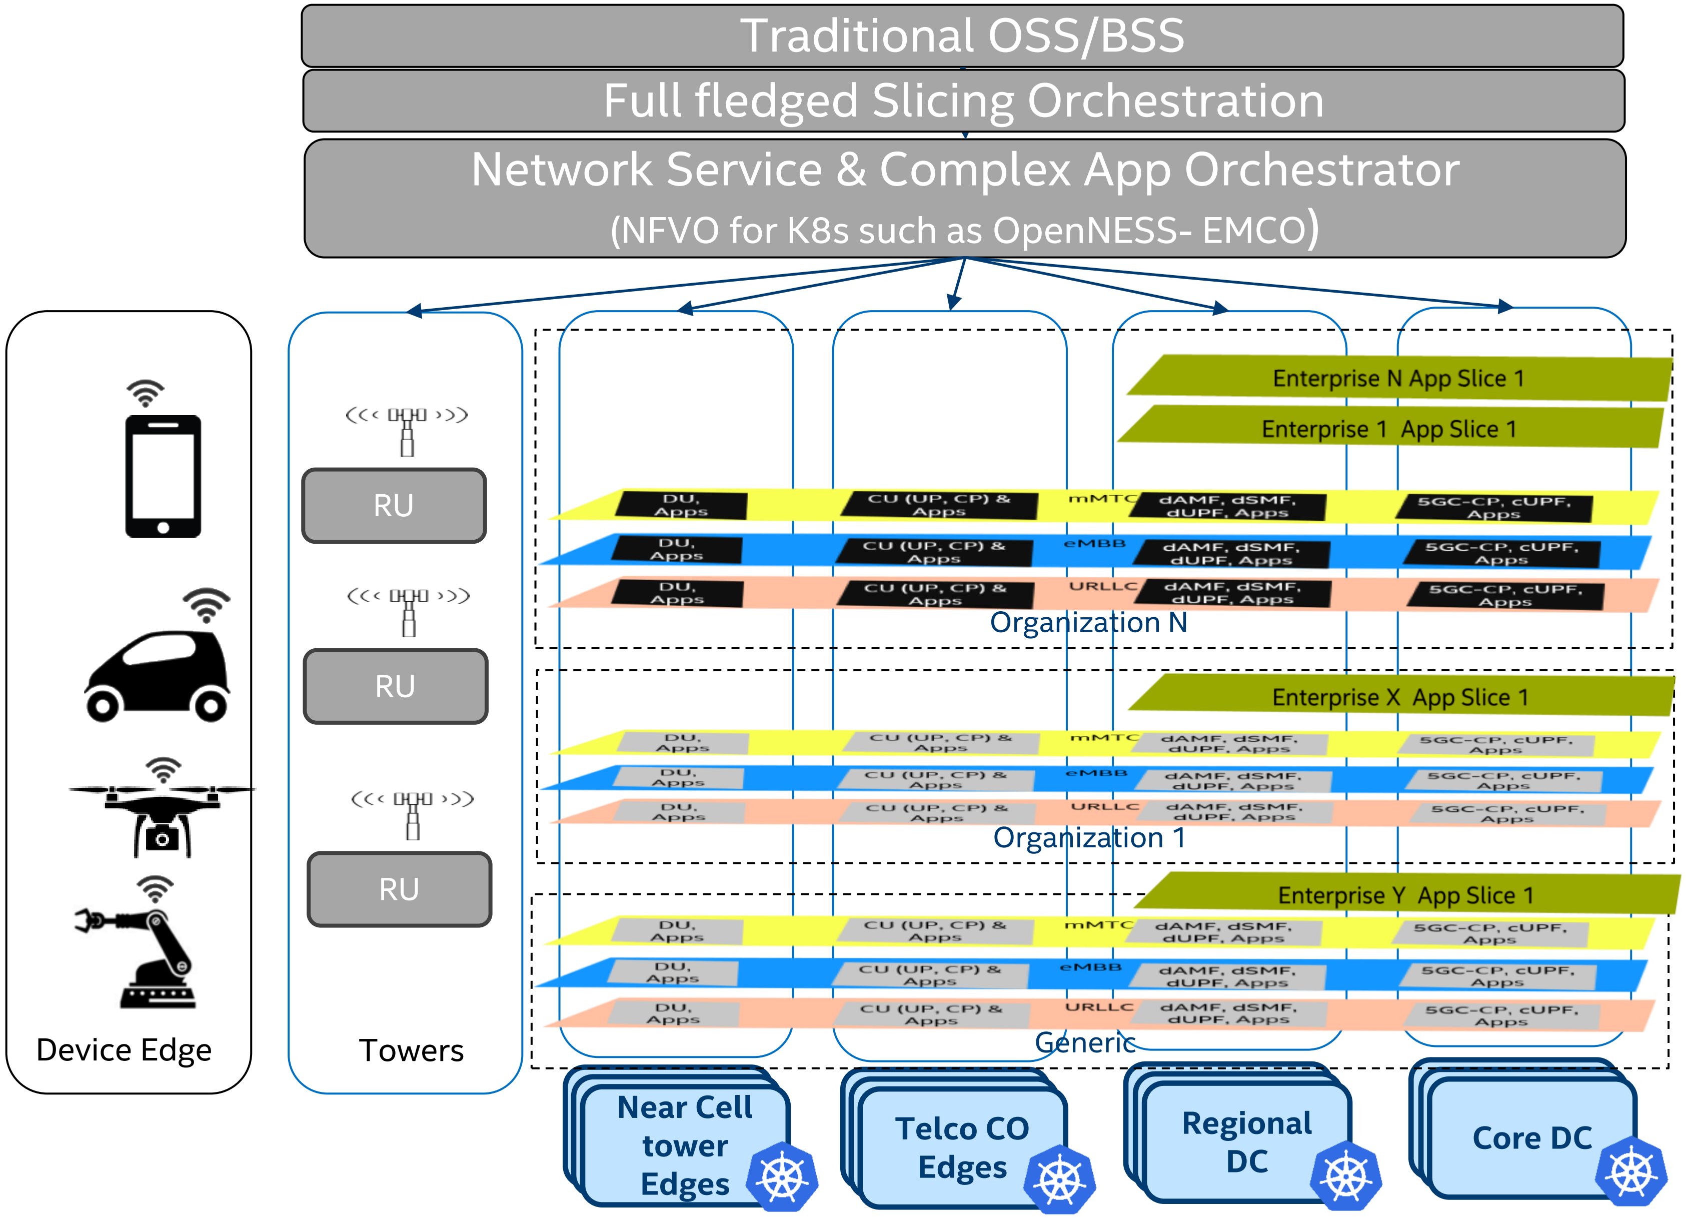

- Operator with generic slices.
- Operator provides network slices to its customers (org1, OrgN in this example).
- OrgN provides application slices to its customers (Enterprise 1 and Enterprise N). Org1 provides app slices to its customers (In this example, Enterprise X). Operator itself has its own customers for application slices (Enterprise Y).
- OrgN gets its own 5GS function too, whereas Org 1 lets the operator bring the 5GS functions.

Dynamic Slicing – Big picture Mapping

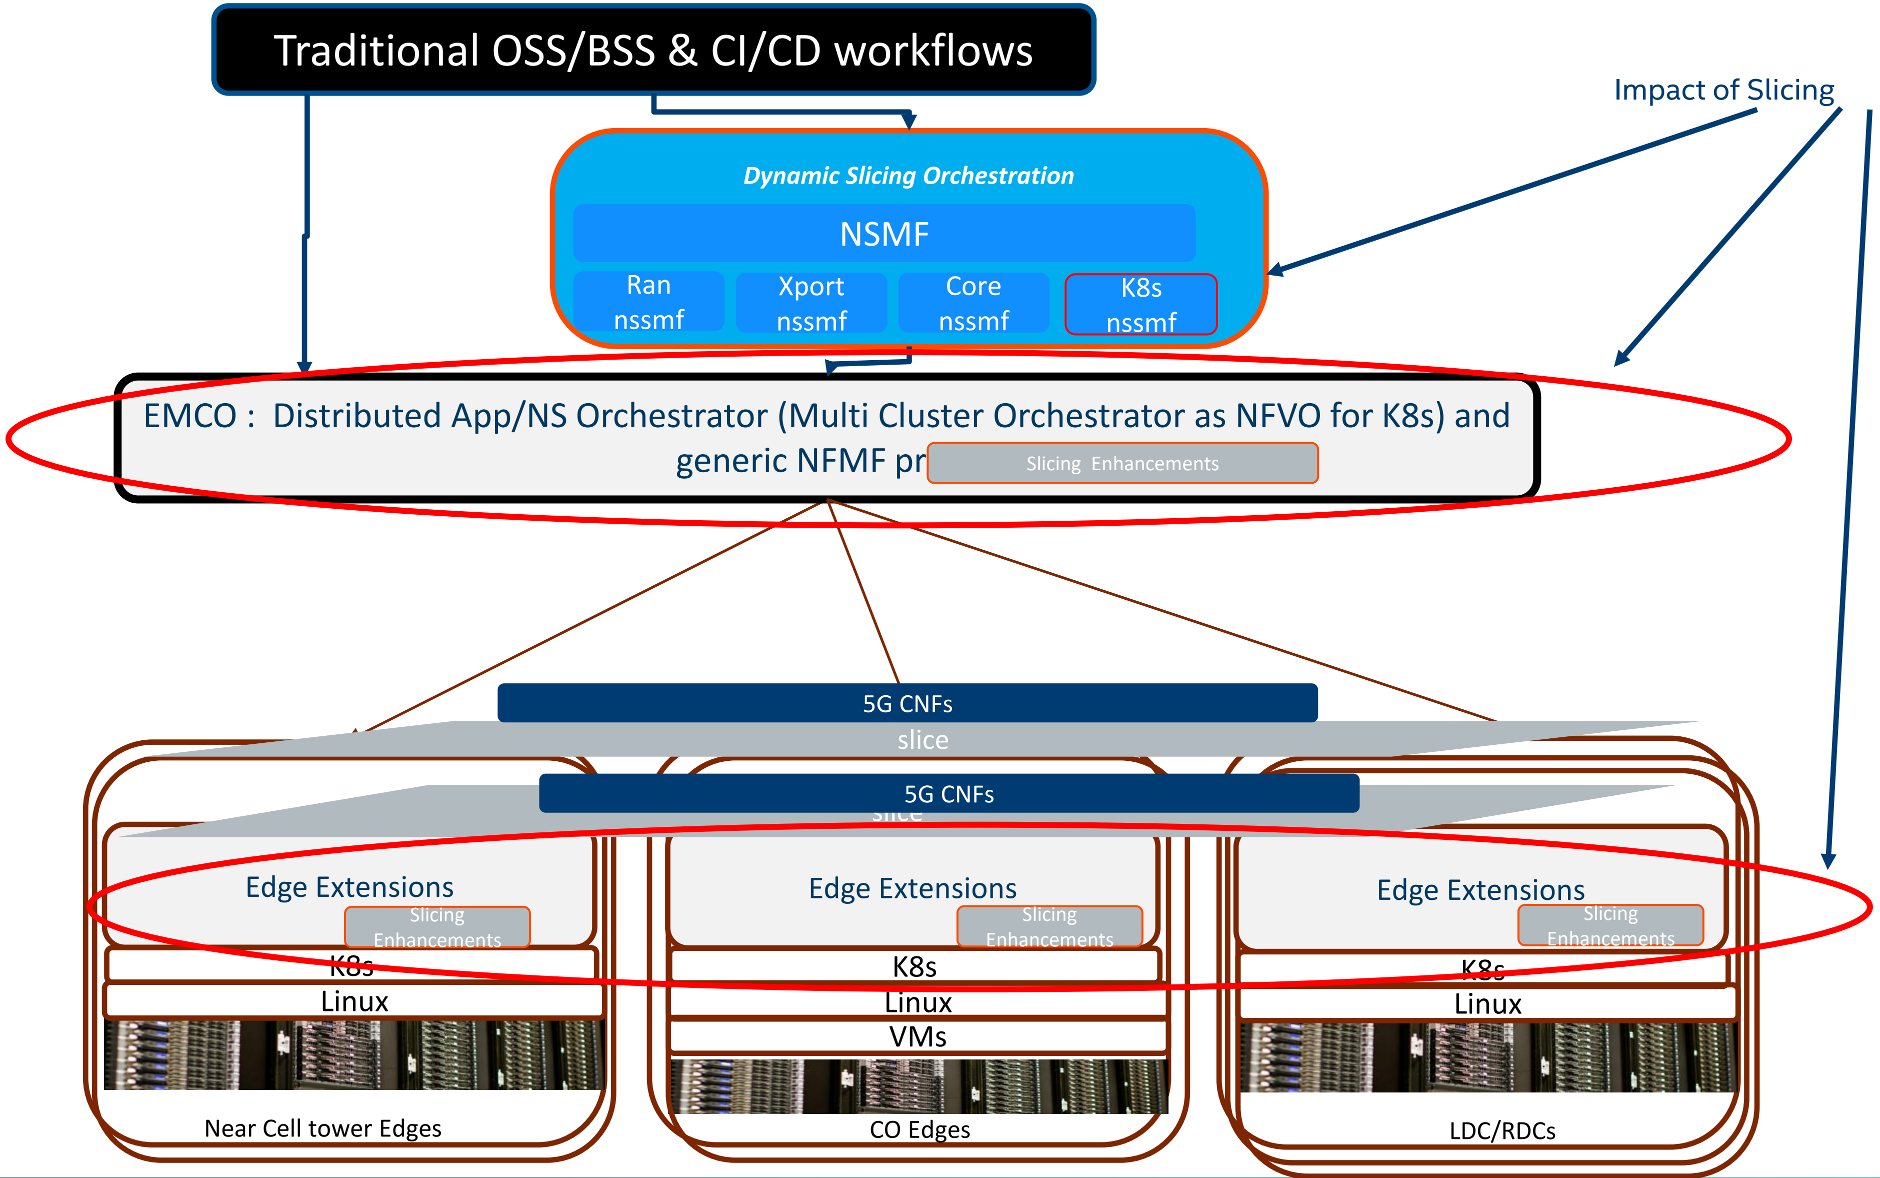

Slicing

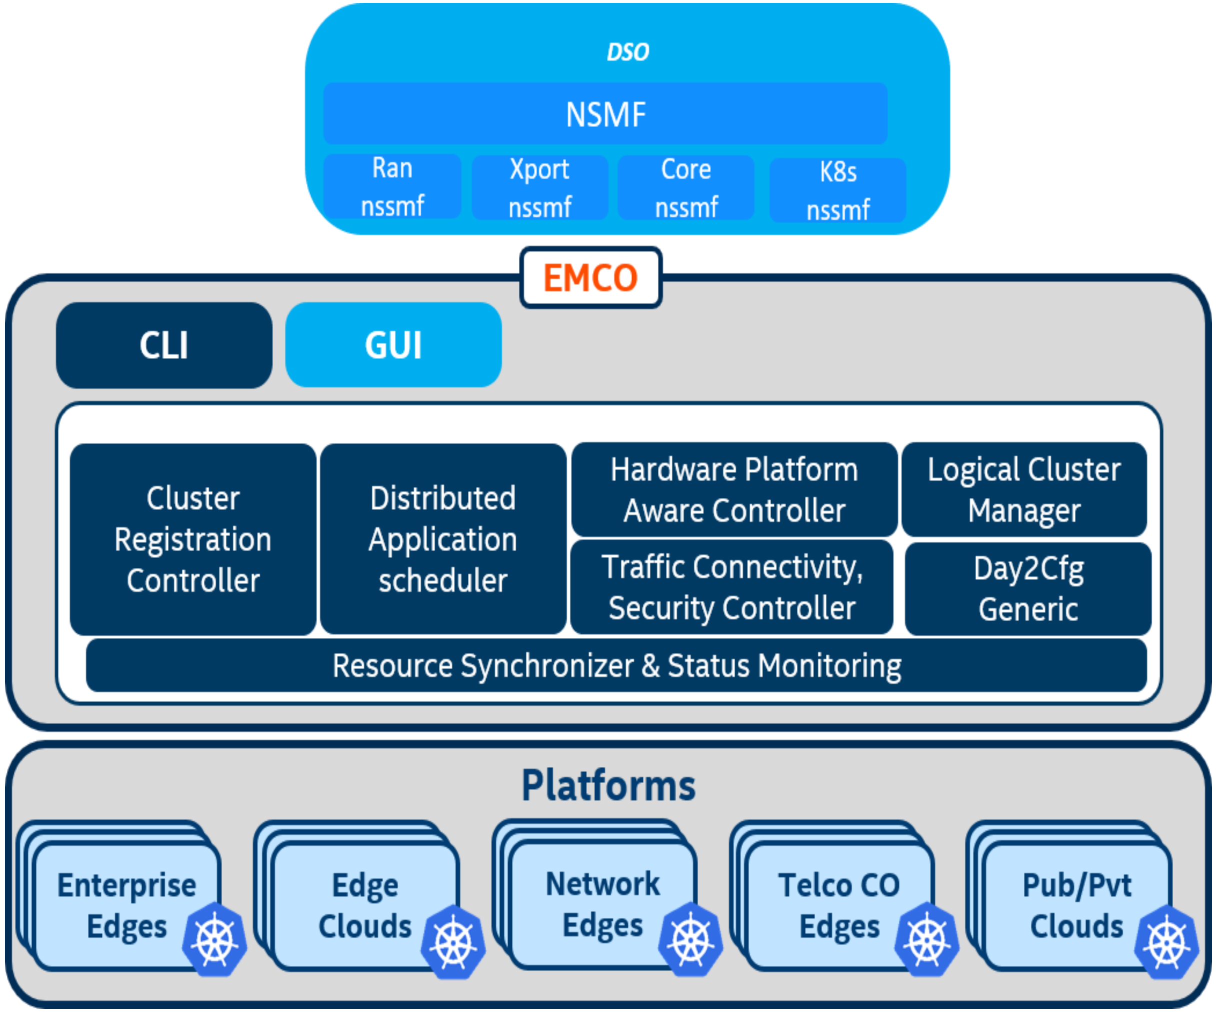

DSO

NSMF

Ran nssmf

Xport nssmf

Core nssmf

K8s nssmf

EMCO

CLI

GUI

Cluster Registration Controller

Distributed Application scheduler

Hardware Platform Aware Controller

Logical Cluster Manager

Traffic Connectivity, Security Controller

Day2Cfg Generic

Resource Synchronizer & Status Monitoring

Platforms

Enterprise Edges

Edge Clouds

Network Edges

Telco CO Edges

Pub/Pvt Clouds

Slicing expectations

Security Isolation from other slices. (Due to vulnerabilities of shared kernel)

Security Isolation from other entities stealing code and data from running CNFs/Apps

Performance isolation of slices from noisy neighbors of other slices

Performance isolation and determinism from unwanted traffic

Role of Domain Orchestrators (e.g EMCO)

Automation of CNF deployments (and customization) & LCM across multiple Edge locations

Selection of Edge locations for CNF and Application deployments

Automation of connectivity components of Edges to enable inter CNF communication across Edge clusters

Additional needs due to slicing in EMCO

Automation of security primitives for security isolation (Kata, TEE)

Automation of performance primitives for performance isolation (HQoS; RDT; Resource Reservations & Quotas;

Fixing gaps/technical-debt HPA (Hardware Platform Awareness)

Fixing gaps/technical-debt is automating the connectivity & Day 2 configuration for closed loops

New initiatives (Proposal)

Translation of KPI of SLA to IA HW resources

Reference NSSMFs and NSMF implementation

🟧 proposed

Thank You!!

EMCO & Security

A. Securing EMCO itself

B. Automation of Security at the Edges

EMCO & Security

A. Securing EMCO itself

B. Automation of Security at the Edges

Securing 5G - Attack Surfaces

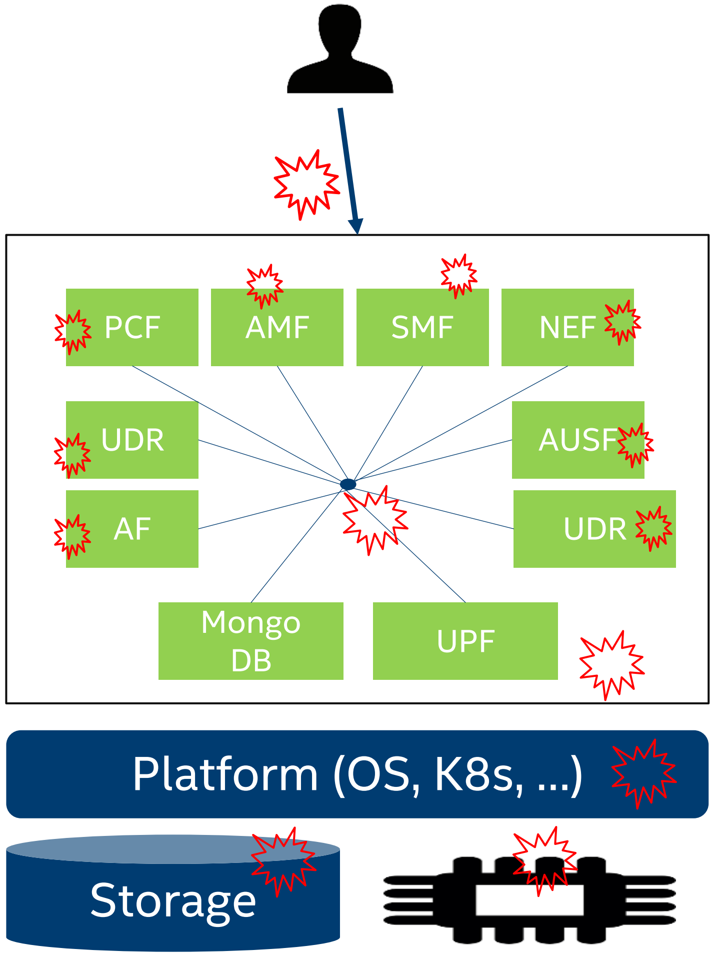

Attack surfaces include

- **Data at Rest** : Secrets, password, private keys, authentication credentials are stored in databases/storage. Any access (stealing or otherwise) can expose these to attackers.
- **Data in motion**: Inter micro-service communication and communication shall be secured to ensure that attackers does not get hold of data in clear. Data origin assurance
- **Data in memory**: Scraping of memory is one possibility for attackers to get hold of secret information.
- **Vulnerabilities & configuration mistakes** : Exploitation of any vulnerability and injecting any malware.
- **Insufficient authentication and authorization**: Can lead to access of data by unprivileged users
- **Platform:** Tampering of platform to get hold of confidential information.

Securing 5G

ISTIO/Envoy Security via ingress proxy:
JWT/Certificate authentication; TLS; OAUTH2, Request authorization; WAF, IDS/IPS
→ Protects from attacks from Internet

ISTIO/Envoy Security via sidecars (Zero trust security)
Inter Micro-service communication security (Mutual TLS, Authentication and Authorization, WAF, IDS/IPS) -> protects from other micro-services (avoid any lateral attacks) and protects from stealing information on the wire

Secure boot and Platform attestation:
To ensure that platform is booted with right software and configuration

Storage data security:
Encrypted file system (example: dm-crypt) with symmetric key secured in outside vault

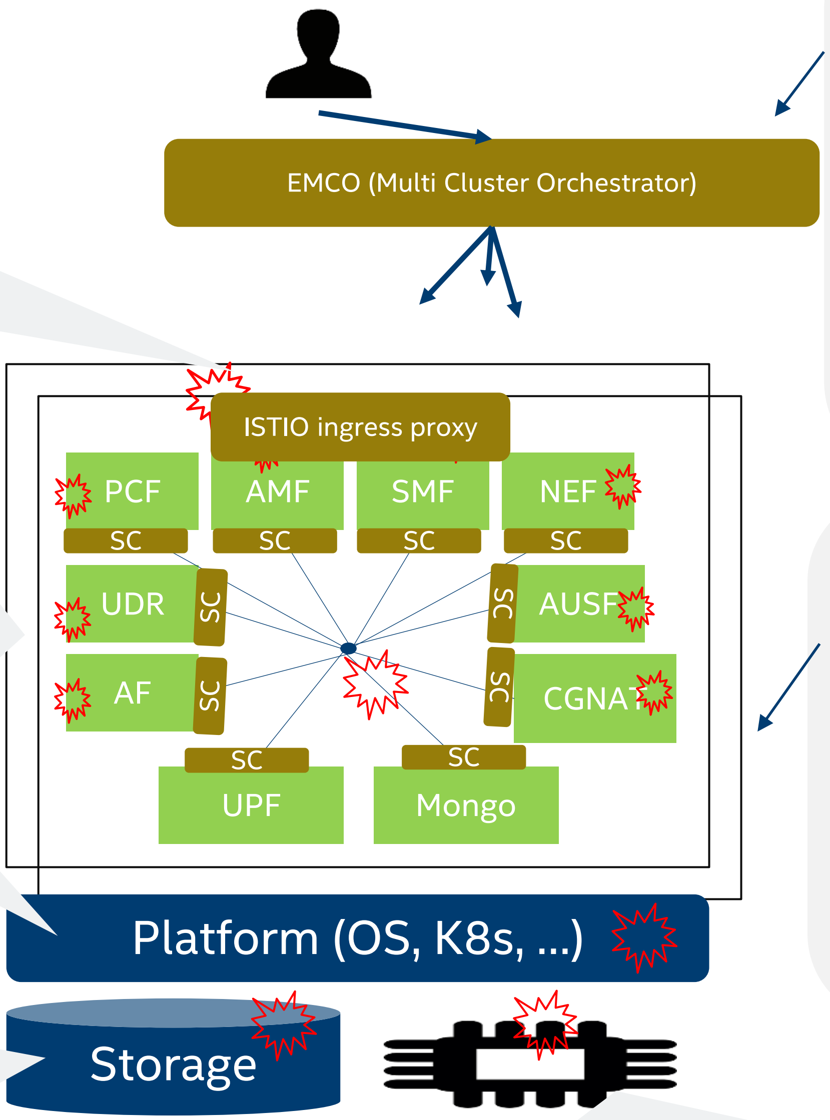

Confidential Computing:
Ensure that critical micro-services are run in TEE

EMCO automates security
EMCO places workloads on trusted cluster
EMCO Slice placement on set of trusted clusters
EMCO monitors security posture from other entities such as iSECL and uses this to place workloads /Slices (TBD)

EMCO community is working on AI analytics for security from K8s and Service mesh metrics/traces/logs and close the loops by automating security policies.

K8s secret protection; PKCS11 support for cert-manager, ISTIO, Envoy; iSecL for platform attestation;

Security Automation of Apps, Network Functions & Slices

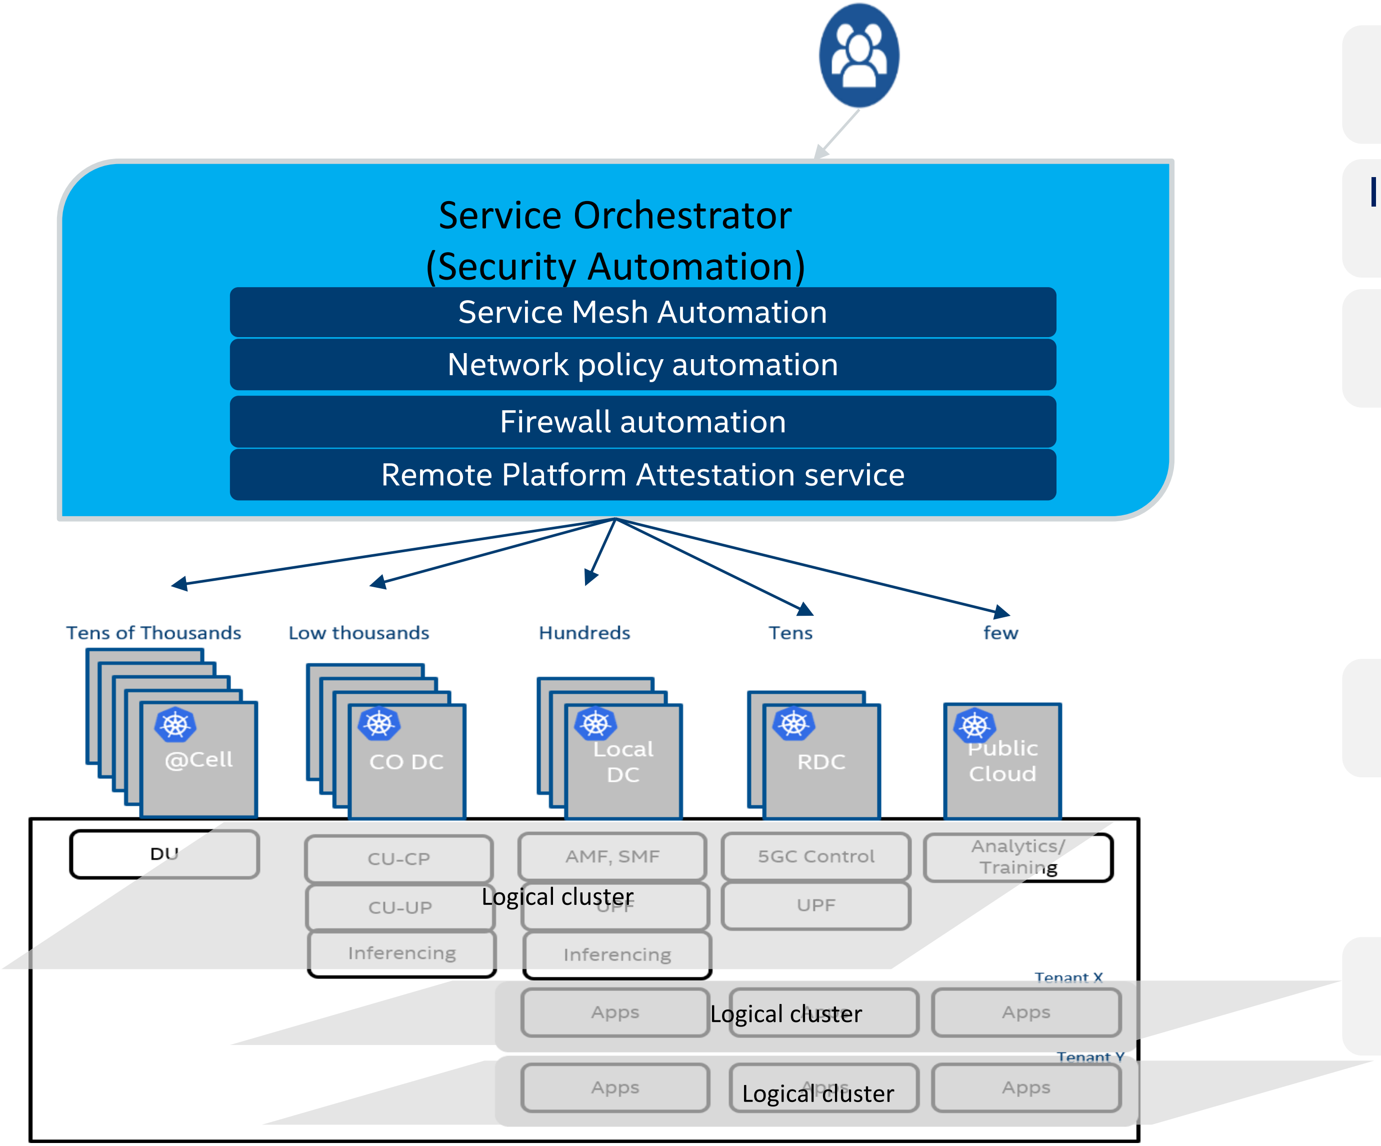

Service Orchestrator
(Security Automation)

Service Mesh Automation

Network policy automation

Firewall automation

Remote Platform Attestation service

Tens of Thousands

Low thousands

Hundreds

Tens

few

@Cell

CO DC

Local DC

RDC

Public Cloud

DU

CU-CP

CU-UP

Inferencing

AMF, SMF

UPF

Logical cluster

Inferencing

5GC Control

UPF

Analytics/ Training

Tenant X

Apps

Logical cluster

Apps

Tenant Y

Apps

Logical cluster

Apps

K8s clusters are being brought up with service meshes

Increasingly Telcos are using service mesh for security for control CNFs

K8s clusters are front ended by traditional security functions as CNFs

Need for security automation; Smart processing for Anti-DDOS

EMCO automates security configuration with LCM of Network services & Applications

EMCO Vision

Be a comprehensive geo-distributed Cloud native application orchestrator

Be a Multi-Party and Multi-Cloud Orchestrator

Be an orchestrator for Network services and Enterprise applications

Be an orchestrator for convergence of Network services and Enterprise applications

Be an orchestrator for Distributed Clouds with Edge-computing

What is not in the scope of EMCO?

EMCO does not expose ETSI and Tmforum APIs.

EMCO does not deploy workloads in non-K8s enviornments

EMCO CNF/App configuration is limited to K8s CR based apps/CNFs. It does not support NetConf, CLI and other mechanisms as of now.

EMCO does not include Analytics stack

Few distributions of thirdparty service orchestrators leveraging EMCO combine other projects such as ONAP CDS,   ONAP DCAE to address brownfield deployments.