Mobile Edge Strategy Leveraging Containers

August 2019

Nick Barcet
Senior Director Technology,
OCTO/CPBU, Red Hat

Red Hat intel

These slides are created for the strict purpose of discussion only. Red Hat Inc. does not plan or commit to offer any functionality described in this material as part of its portfolio of products.

Red Hat Inc. does upstream open source software development work and will continue to participate in upstream activities and communities relative to technologies described in this presentation. This does not imply Red Hat will or have products associated with those software communities Nor is Red Hat committing to having such offerings in future.

Red Hat Inc. does not commit to or subscribe to any opinions, statements regarding third parties described in this presentation. Any disclosures associated with third party companies are a matter of conjecture and should not be taken as known facts or stated strategy of those third parties.

Acknowledgements: I would like to acknowledge the contributions of my colleagues at Red Hat.
I have used some of their material in this presentation to make my point.

INDUSTRY TRENDS

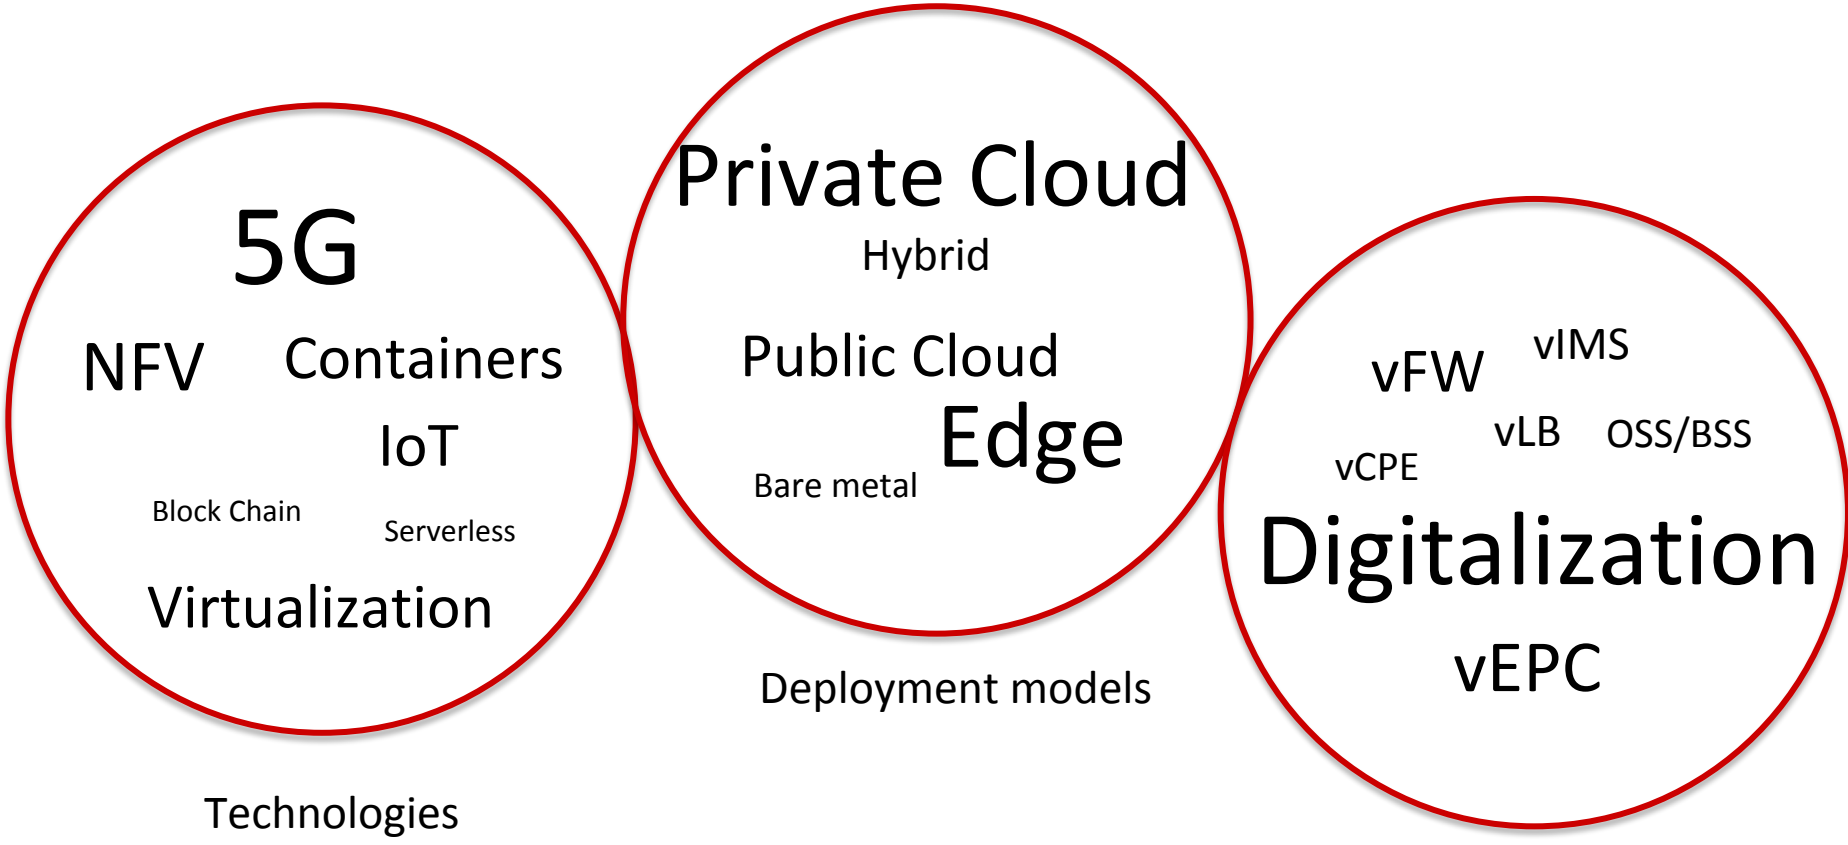

HEAVY READING SURVEYS

Figure 1: MEC Deployment Commitment

We will de

We are

Question:
Source: He

Figure 2:

We a
compute
computin

We are co
of clou
access to
current a

We ar
refe
netwo
perfor

Question:
network t
Source: H

Figure 4: MEC Use Case Drivers

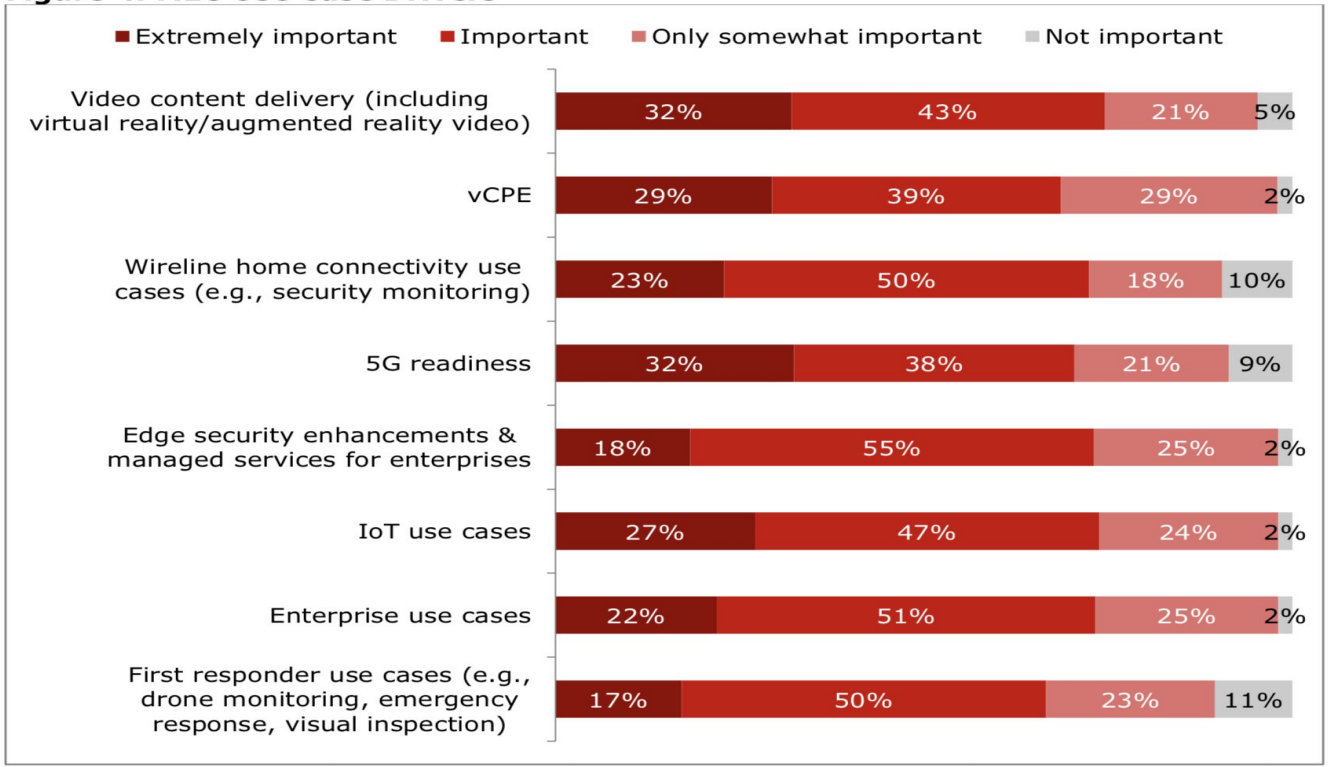

Question: Please rate the importance of the following use cases in driving your company's MEC deployment strategy. (N=102-107)
Source: Heavy Reading: Intel Custom Survey Q417

Red Hat intel

5G

5G

Ultra-broadband (mmWave):
Offering higher bitrates and supporting extreme traffic densities

Ultra Low Latency (Sub 6Ghz):
Mission critical specialized services and immersive virtual reality

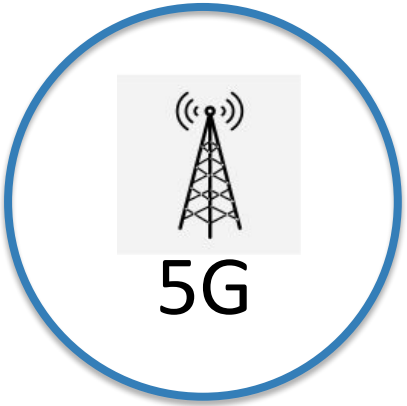

Ultra-narrowband (sub 1GHz):
Efficient sensing and control; massive densities of low traffic devices and bearers

- Enabler Technology
- Seen as a leveling the playing field between incumbents and new entrants
- Resolves last mile access challenges – provides instant access consumers and Enterprise customers (No wires)
- New services possible through enhanced use cases and capabilities
- Promise of massive scale
- Better user experiences via Edge Compute
- Huge Investments for large scale deployments
- Still 18-24 months away from mainstream/mass scale

Radio Access

Option 1: Non Stand Alone
LTE Evolution and New Radio Access

OR

Option 2: Standalone
New Radio Access Only

LTE Evolution will be backwards compatible with existing LTE deployments but the New 5G RAT (Radio Access Technology) will not be

Additional Spectrum

- cmWave
- mmWave

Network Architecture

- Hetnet (Heterogenous Network)
- Core Edge Network Architecture
- Virtualization

Hardware / Device

- Massive MIMO
- Dual Connectivity
- Device to Device
- Flexible Duplex

Spectrum

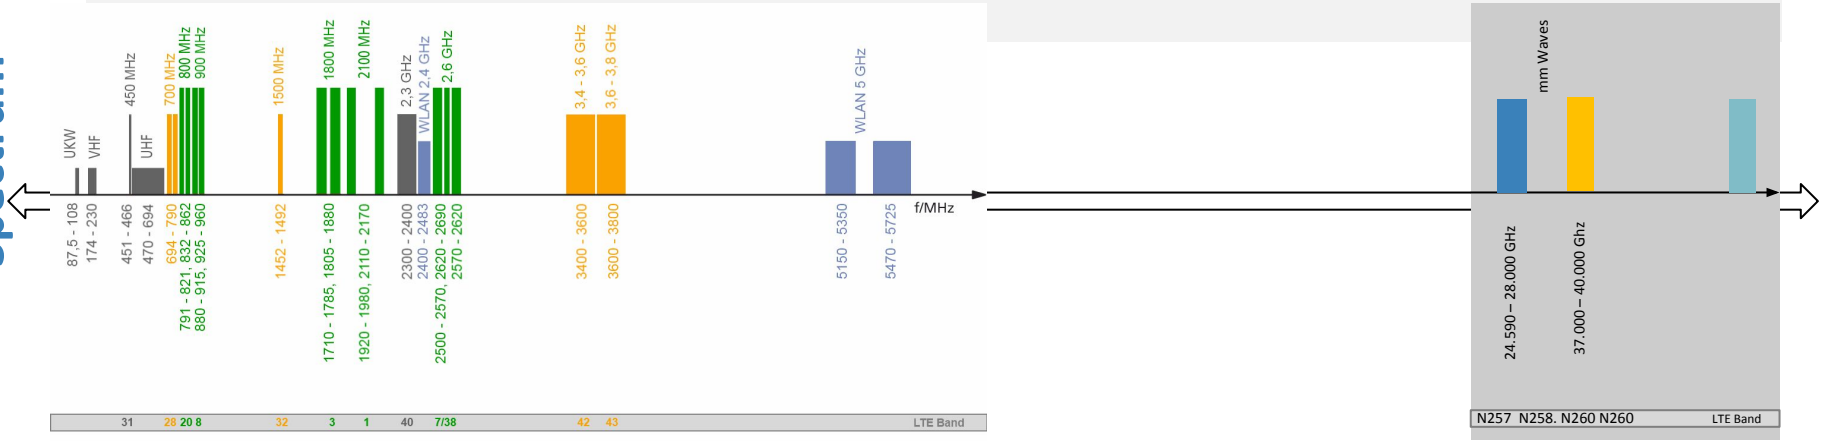

Red Hat intel

RAN EVOLUTION- LTE/4G
Traditional RAN to Centralized RAN to Virtualized RAN

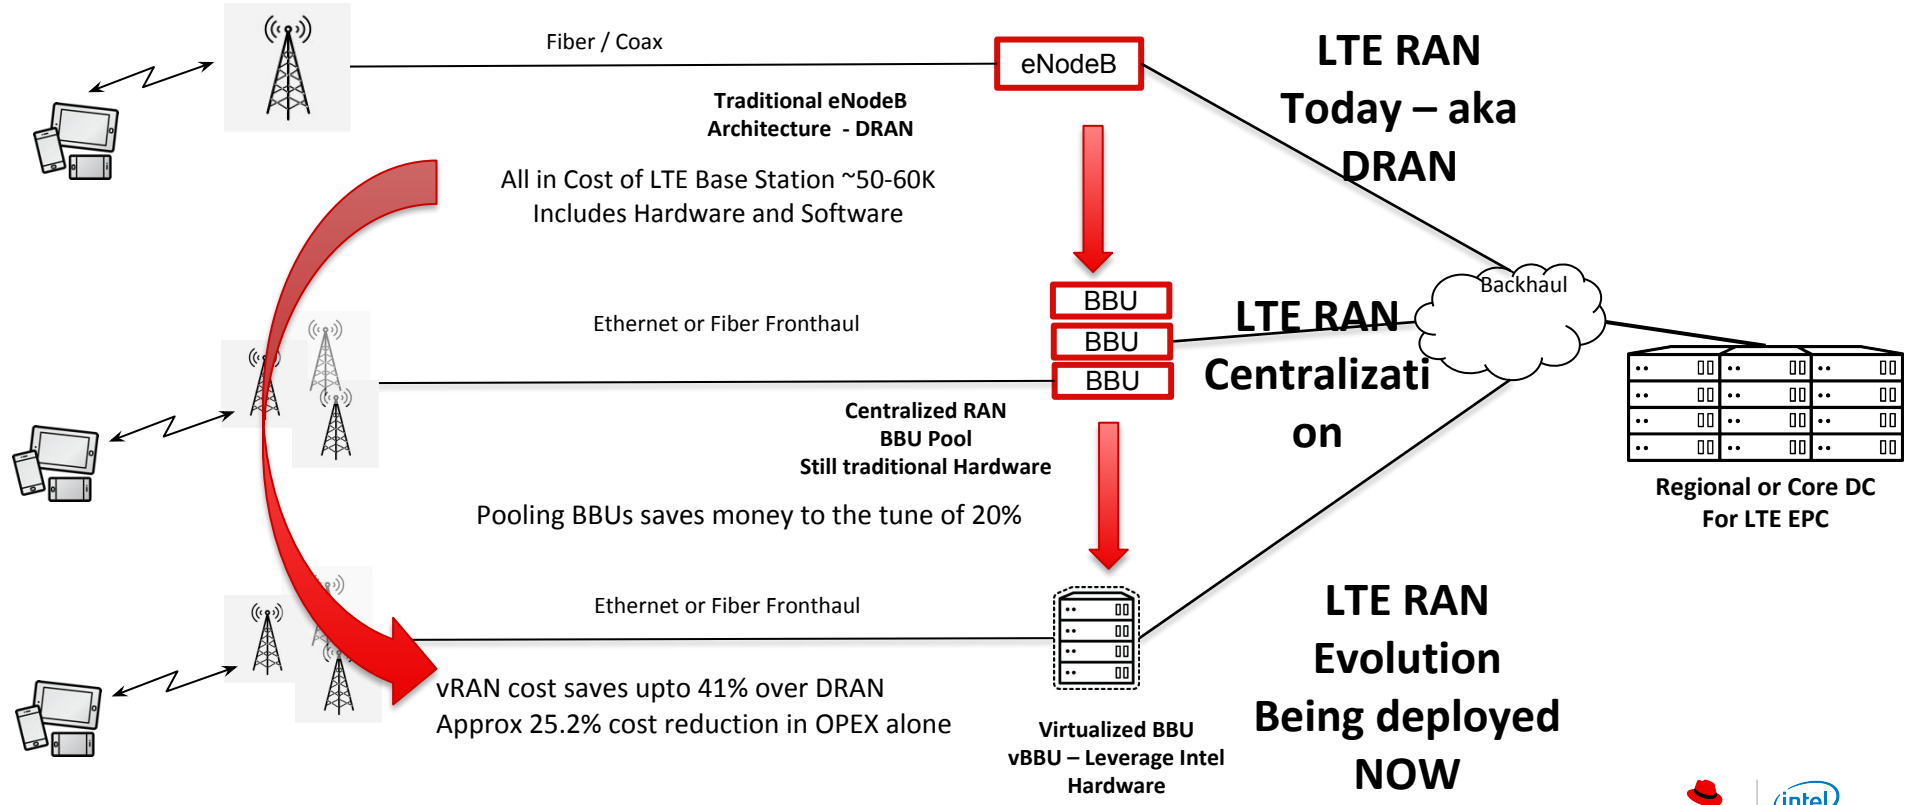

LTE RAN Evolution and 5G RAN …Cont'd

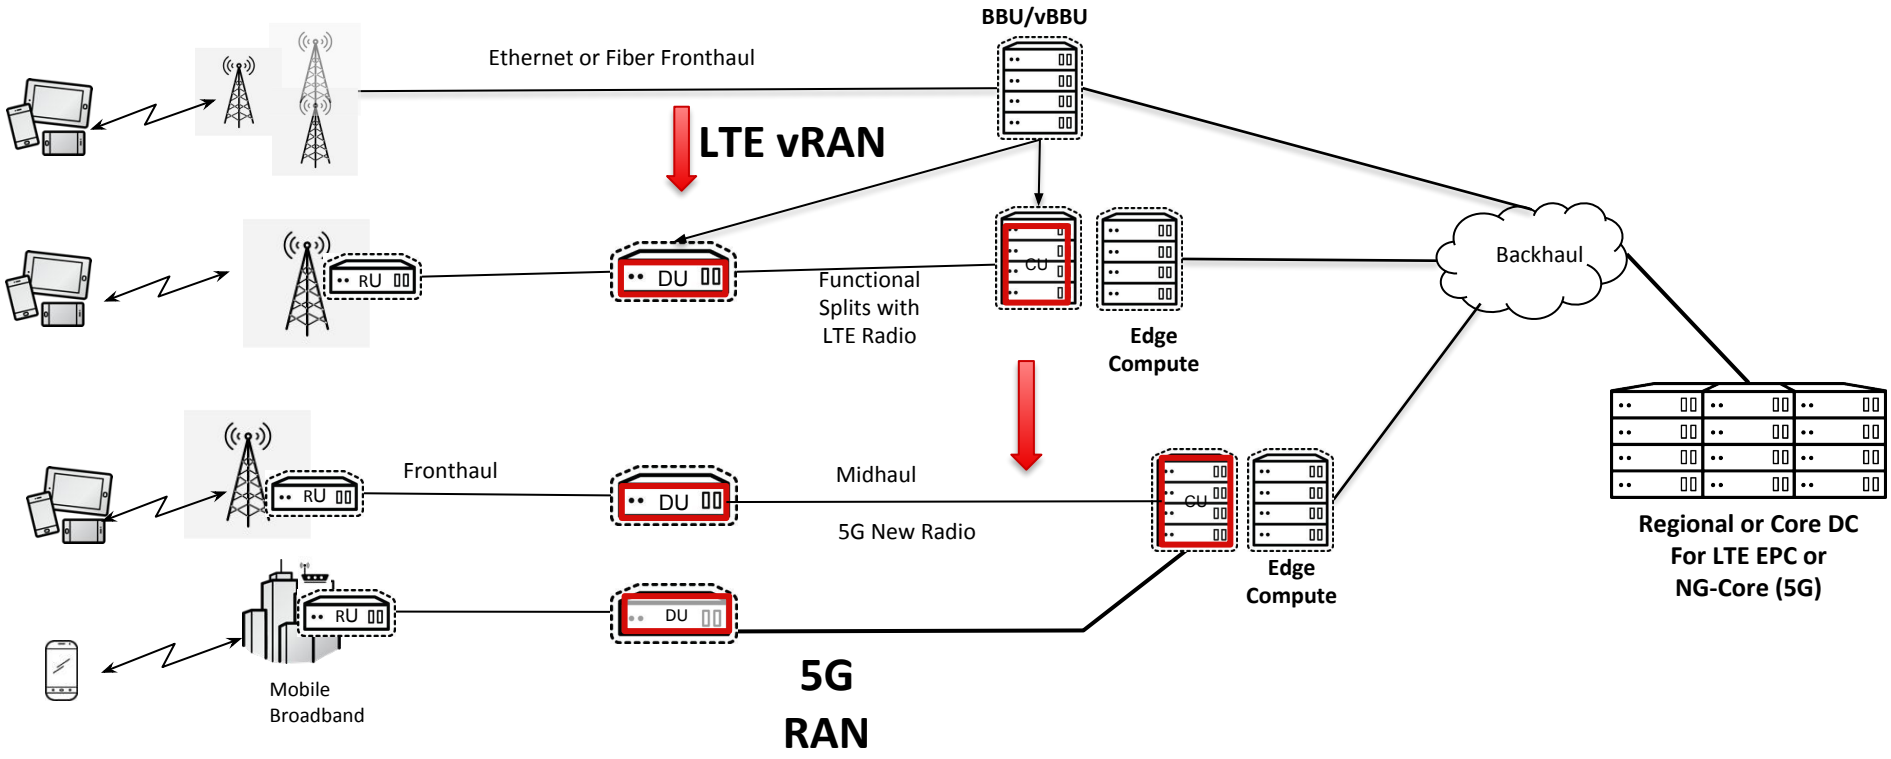

Compute Nodes at - RU (Radio Unit, DU (Distributed Unit), CU (Centralized Unit) and Edge Compute for Application in addition to packet core and DC

Red Hat intel

CLOUD NATIVE COMPONENTS & COMMON MICROSERVICES

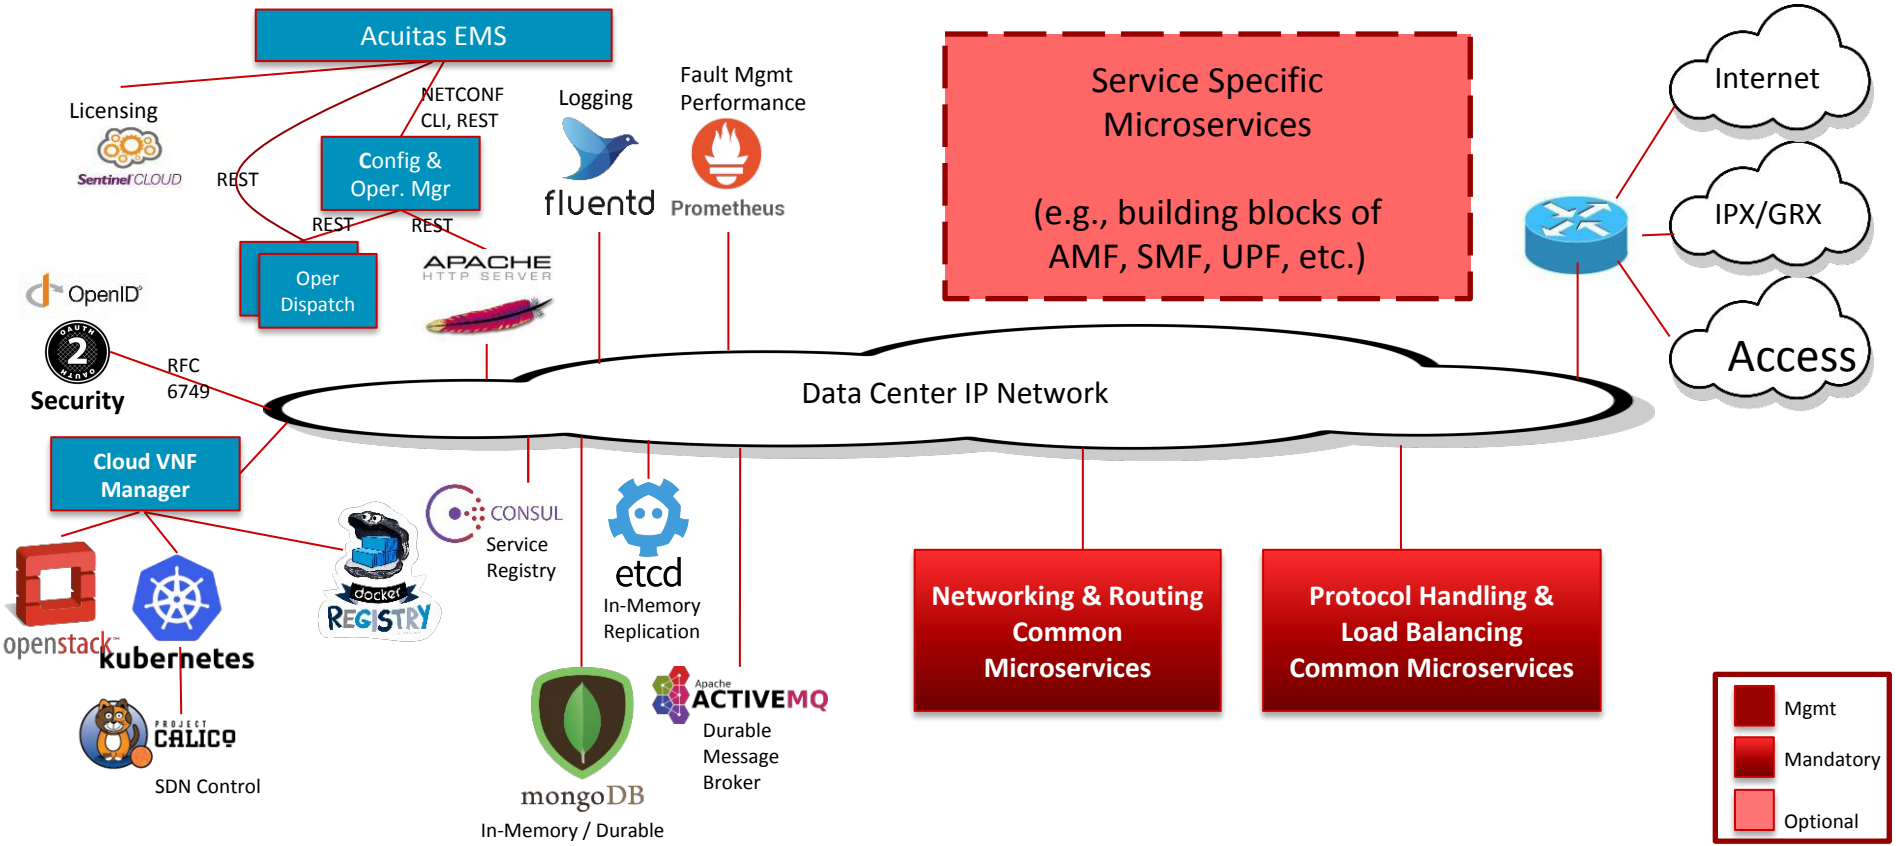

Courtesy : Affirmed Networks

Red Hat intel

Where is Edge Cloud and NEW servers growth?

China Mobile's Edge TICs

Located from city level to AP. Support services including mobile & residential/enterprise UP, MEC and CRAN. Based on open-source Virtualization and/or Container platform

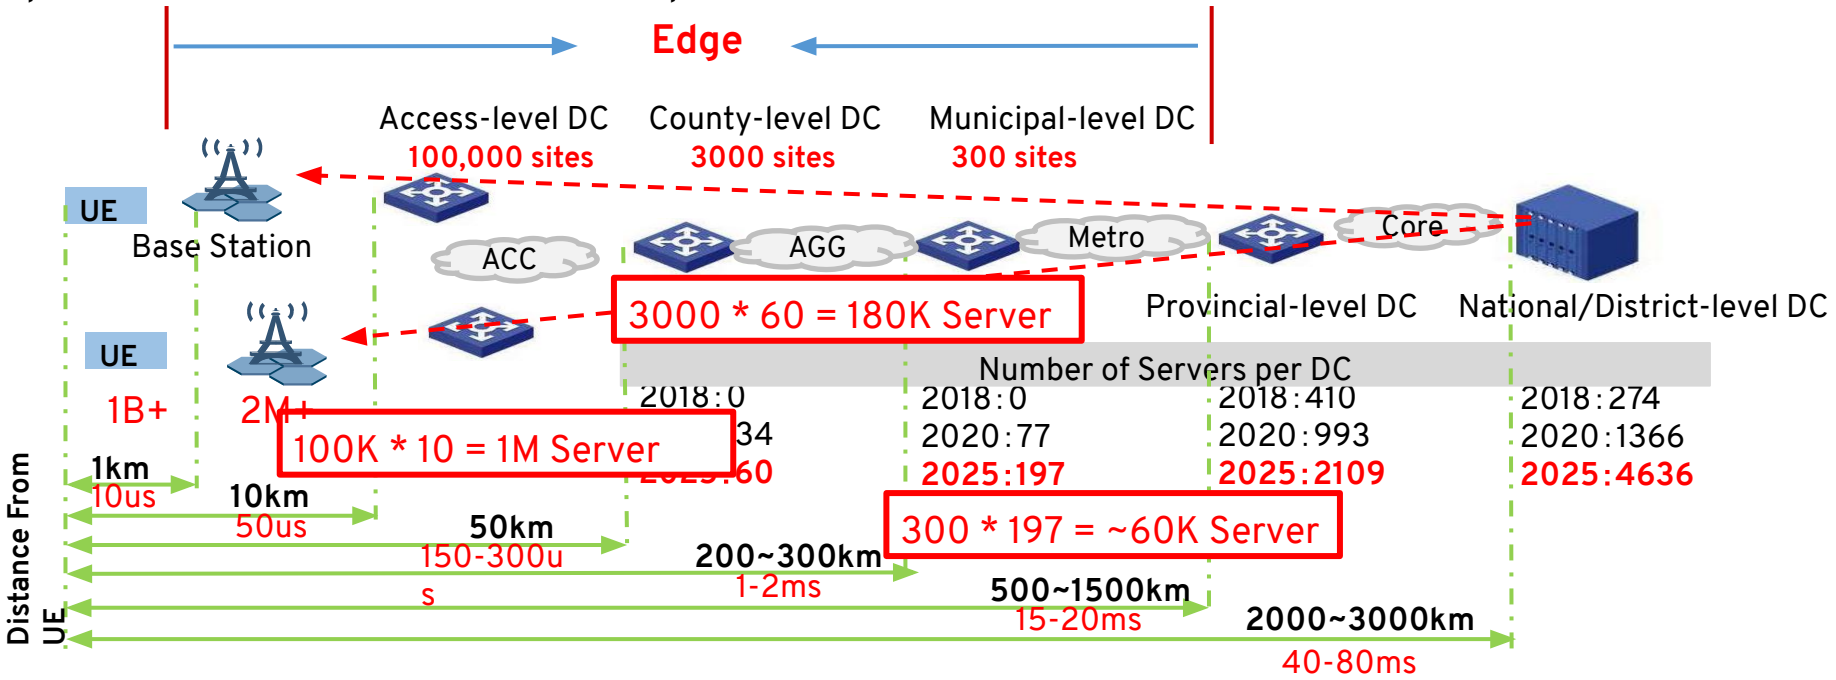

Source: VCO keynote Presentation – ONS

Red Hat | intel

EDGE

What does Edge mean?

"Edge is the next infrastructure paradigm for delivering applications and services closer to the user. Edge allows efficient processing and delivery of time sensitive data"

"Edge refers to the geographic distribution of compute nodes in the network"

"Edge is a distributing computing paradigm in which computation is largely performed on distributed devices sitting closer to end users"

Edge is

- Independent
- Elastic
- Massive scale

Edge must be

- Automated
- Resilient
- Be built using public, private or hybrid cloud

Edge Is about building ultra-reliable experiences for people and objects when and where it matters

Edge is anything that sits between the Subscriber /User/Endpoint and regional or core data center of the provider

Red Hat | intel

WHY EDGE?

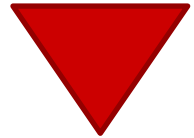

- Inverted Pyramid Model
  - Lots of functionality concentrated in few data centers => failures impact wider user base => High Risk
- Smaller instances => allows fencing of impact domain
  - Lots of smaller instances ; partition theory
  - May bring in other challenges – orchestration and placement => hopefully can be solved via automations and tools especially if the deployment models are repeatable
- Timely delivery of information and analytics (Ultra Reliable Low Latency Communications)
  - Smart vehicle
  - IoT – Environmental information processing
  - Local decision engines – “contextual computing”

Red Hat | intel

VARIETY OF EDGE DEPLOYMENTS

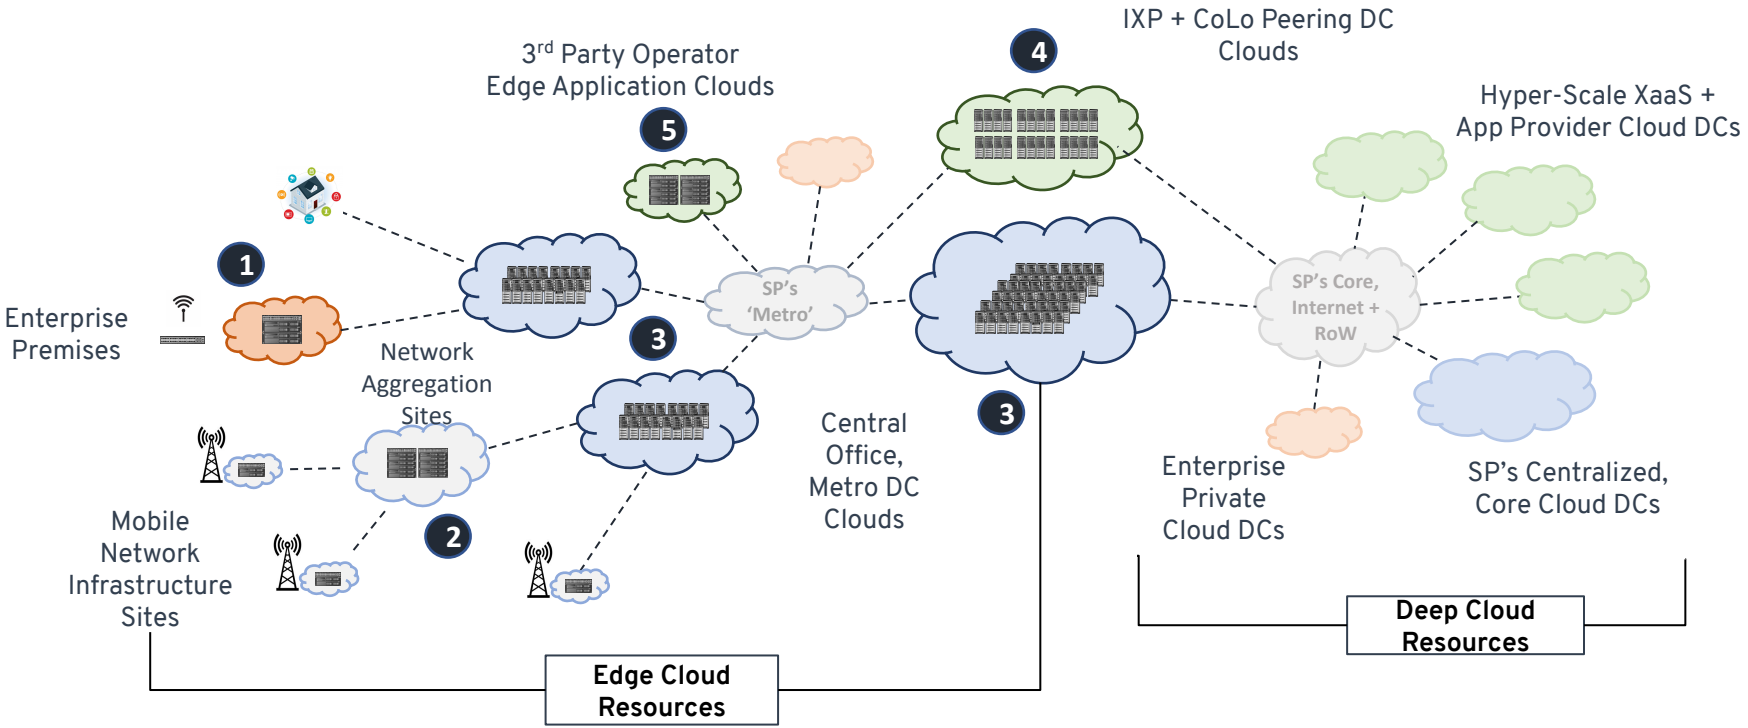

Courtesy: Paul Johnson - ACG

EDGE CLASSIFICATION

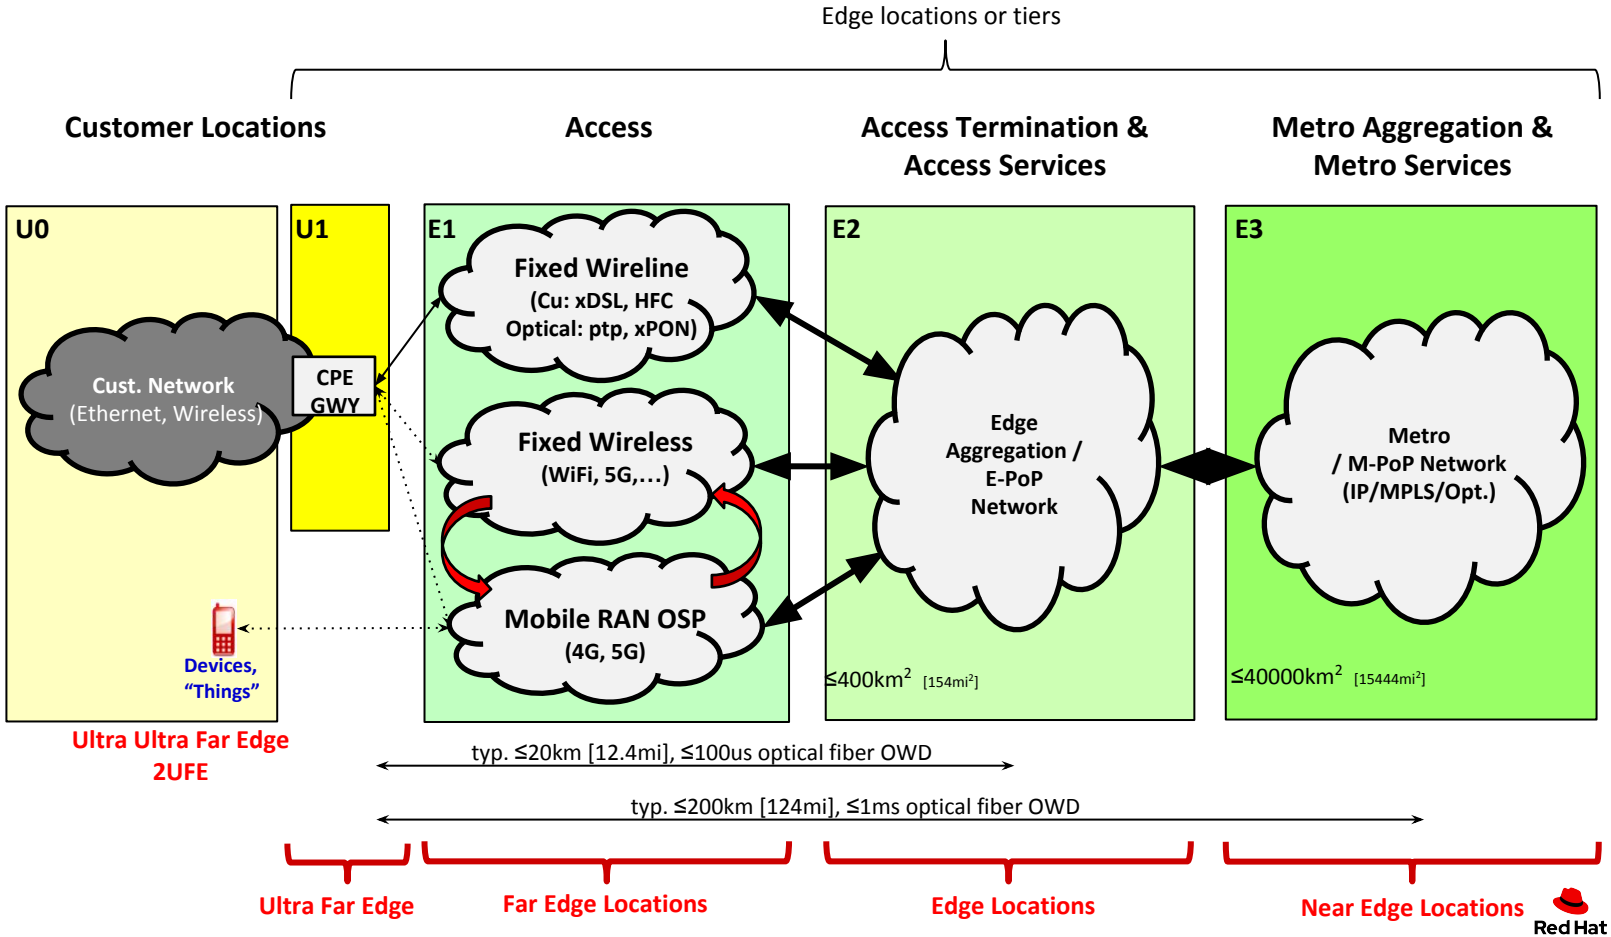

ULTRA/FAR EDGE USE CASES & SIZES - SUMMARY

Use Case	Layer	Descriptions	# of Nodes	Comments
Embedded Compute	U0	Managed Embedded compute	1	Embedded devices - 2UFE (Ultra Ultra Far Edge)
vCPE	U1	Branch Office Connectivity	1-2	Managed CPE service with VNFs for SD-WAN or Branch Office
IoT Gateway	U1/E1	IoT Gwy for SCADA, Protocol/Messaging conversion and data analytics, Atom class	1-2	Usually ruggedized devices collecting sensor data – low cost
Enterprise Edge	U1/E1	Xeon-D class servers	2+	Enterprise Edge – “Edge Cloud”
vRAN	E1	Intel Xeon Class with FPGA	1-2	vDU/DU on baremetal
SDR	U1	Software Radio or RIU	1	

Edge Use Cases

Red Hat | intel

CABLE DEPLOYMENTS

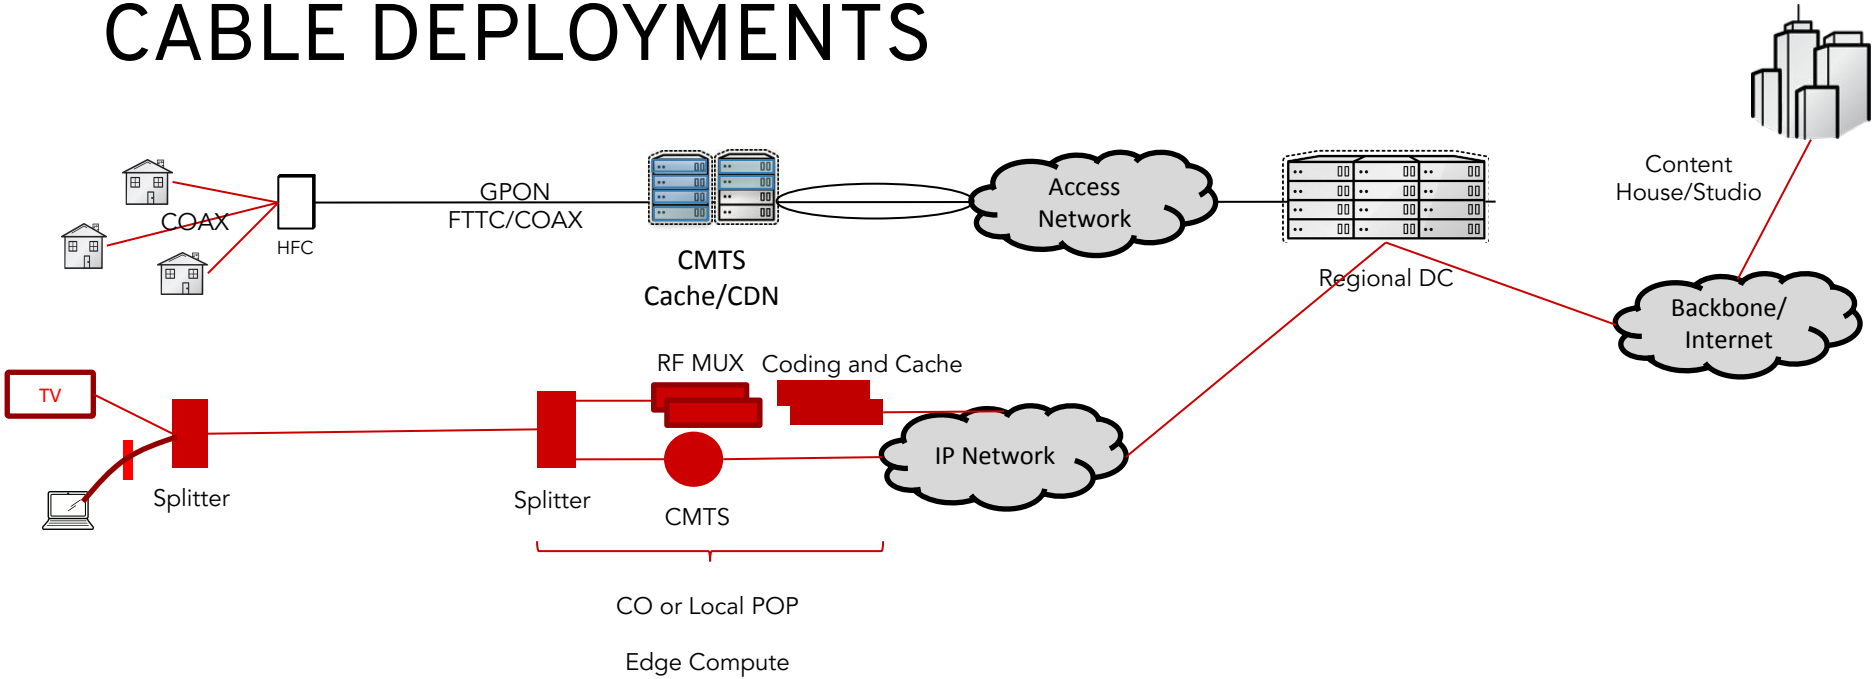

- Many deployment models
- Typical setup show above
- Local POP (aka Cable Head end) – Terminates cable subs and IP backend
- Authentication, Subscriber Management, Content Management, Billing
- Reach and access connectivity means SMB services in addition to residential

Red Hat intel

EDGE DEPLOYMENTS SUMMARY

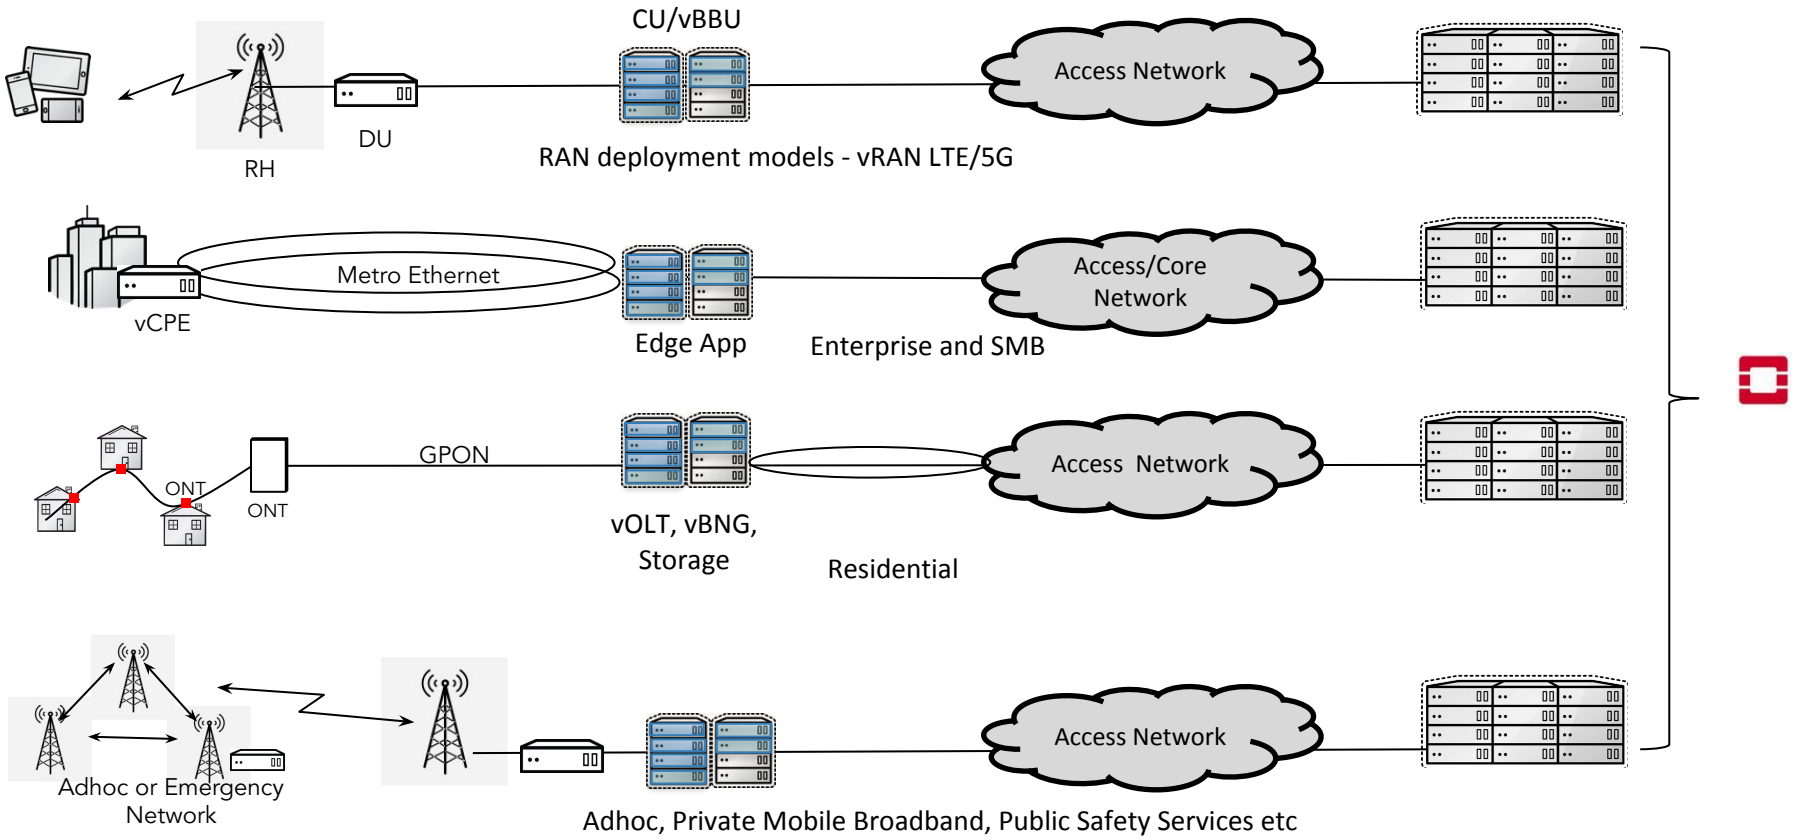

LOGICAL CONNECTIVITY OF THE VRAN (MACRO)

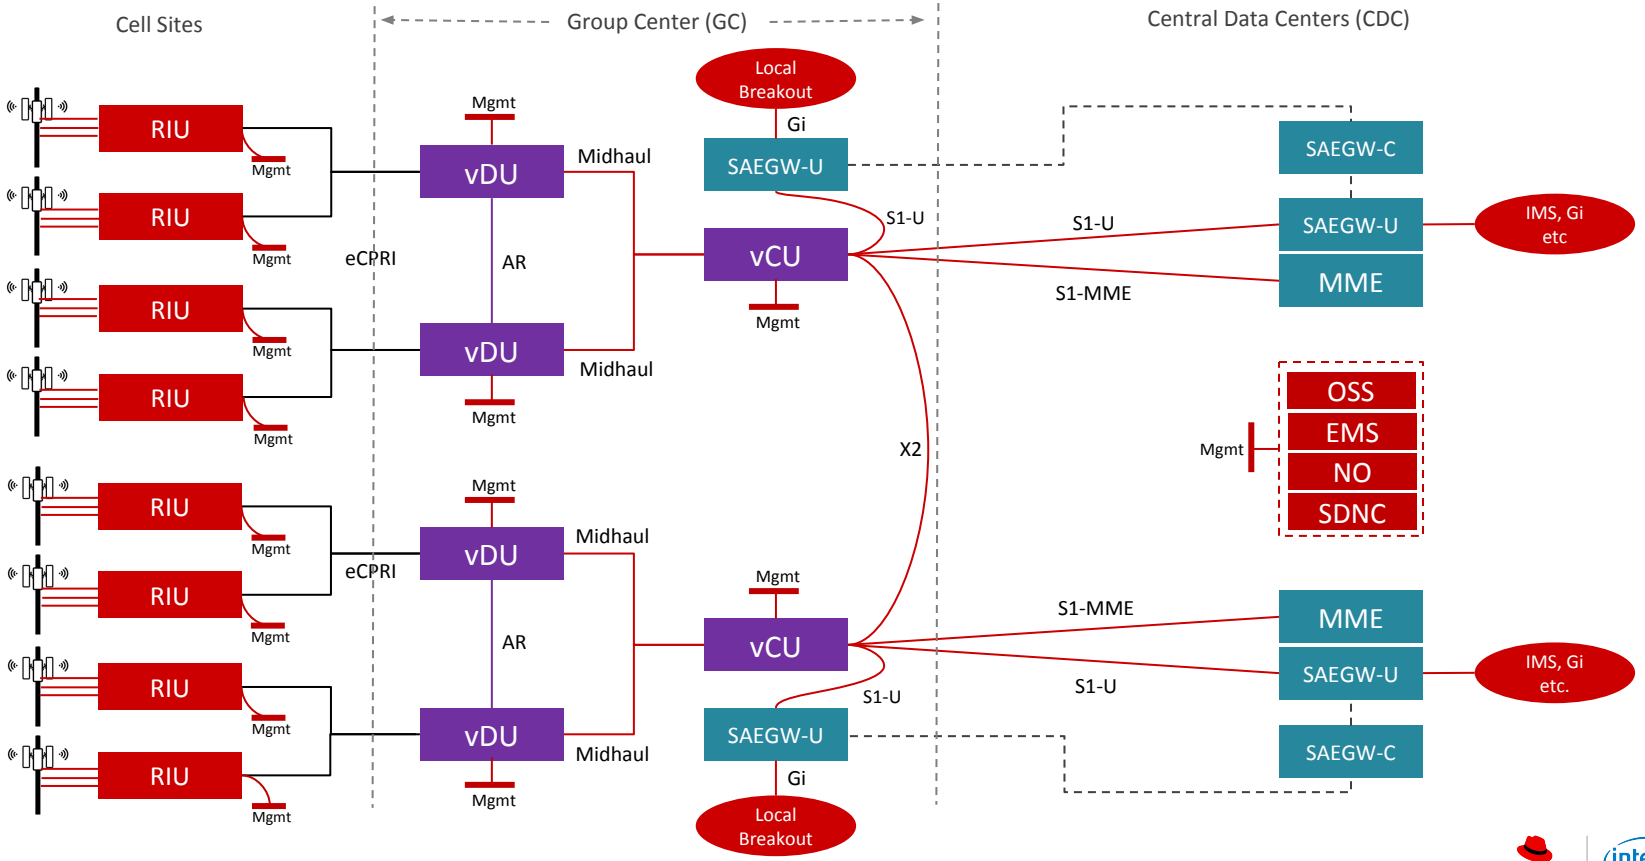

Hybrid workloads for Edge – VNFs and Applications

Containers, Virtual Machines, and Bare-metal “pick and mix”

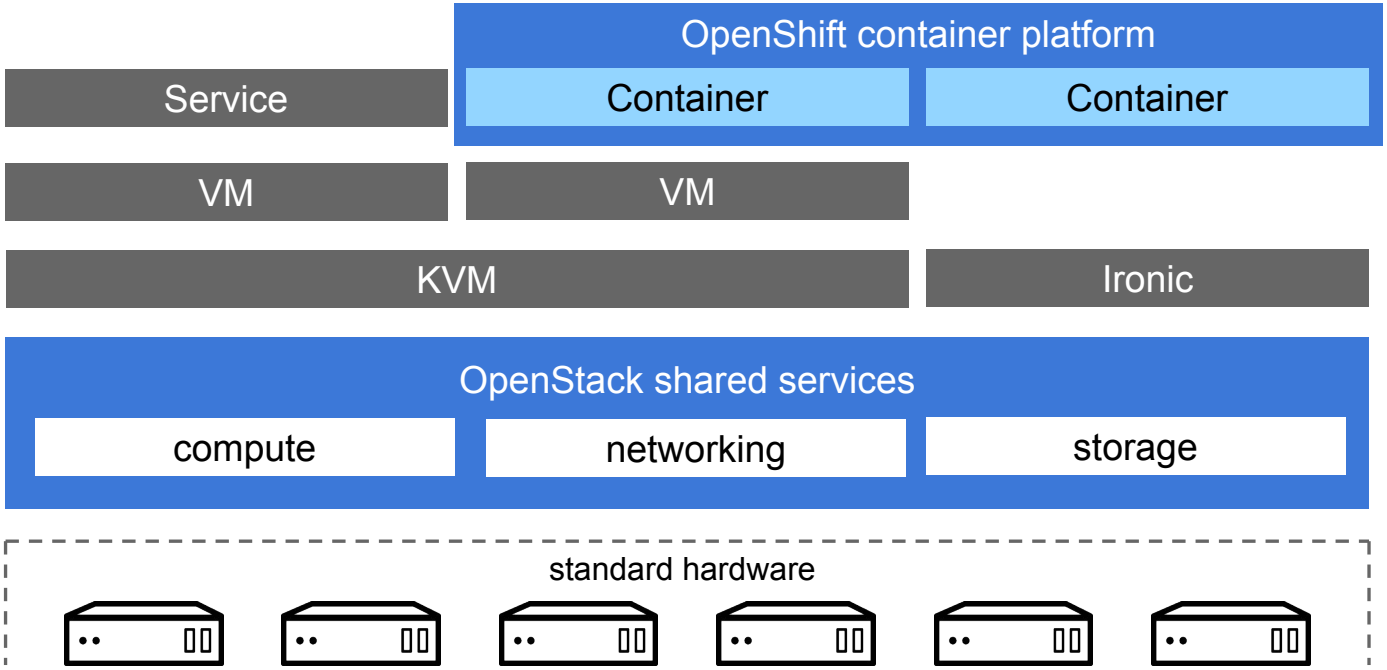

EDGE SCENARIO MAPPING

No Infrastructure control plane at Edge Site

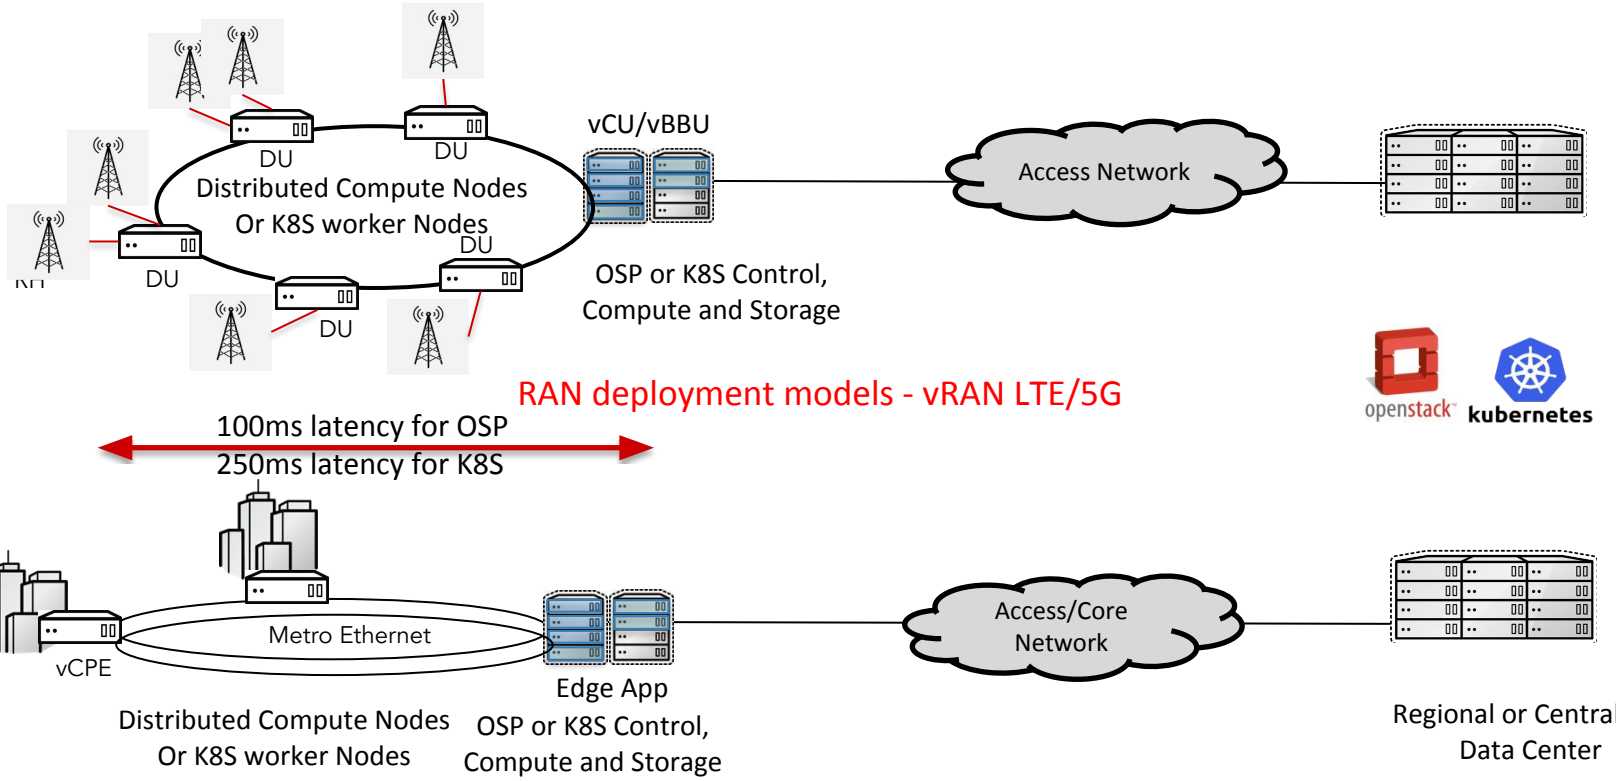

Enterprise and SMB

Edge Scenarios

Capabilities and constraints

OpenStack

- <100ms Latency for OSP Nova to OSP control Plane
- OSP Director deployable as of OSP13
- Ephemeral Storage
- Uses OpenStack Ironic to initialize and manage the node
- Ironic conductor
- Up to 300 remote nodes
- Support of real time linux and real time KVM

K8S/OpenShift

- Up to 250ms latency or more for remote worker nodes of K8S
- VMs via Kubevirt in future
- Node/remote cluster must be initialized and available with an IP address
- Up to 1000 nodes possible
- Support of real time coming in the future

SEPARATION OF SERVICES

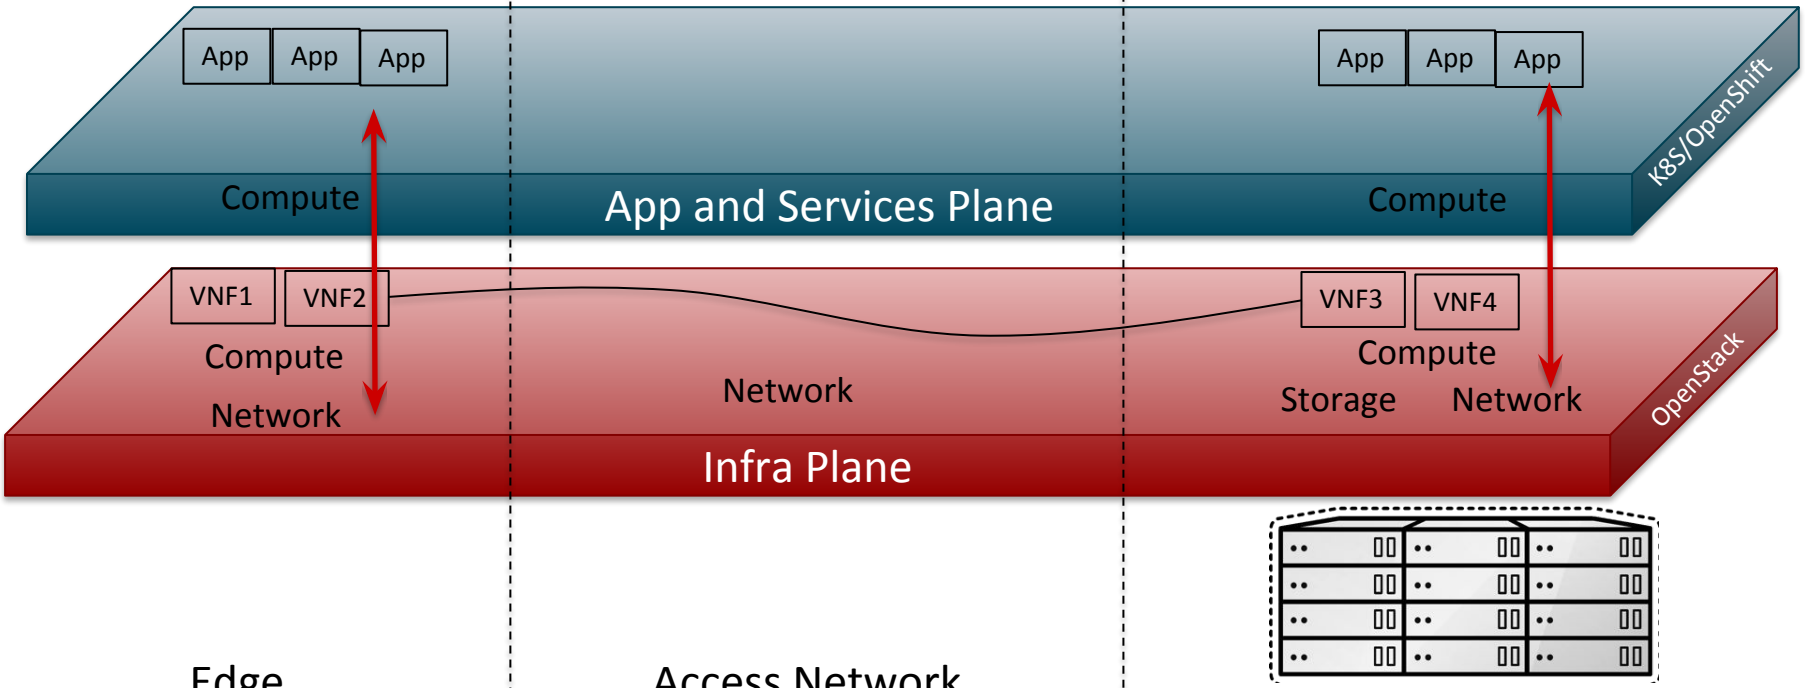

Edge

Access Network
(Physical)

Data Center or Agg location

Red Hat intel

MULTI-CLUSTER ENVIRONMENT

Independent K8S clusters (1 node to n nodes in each Far Edge site)

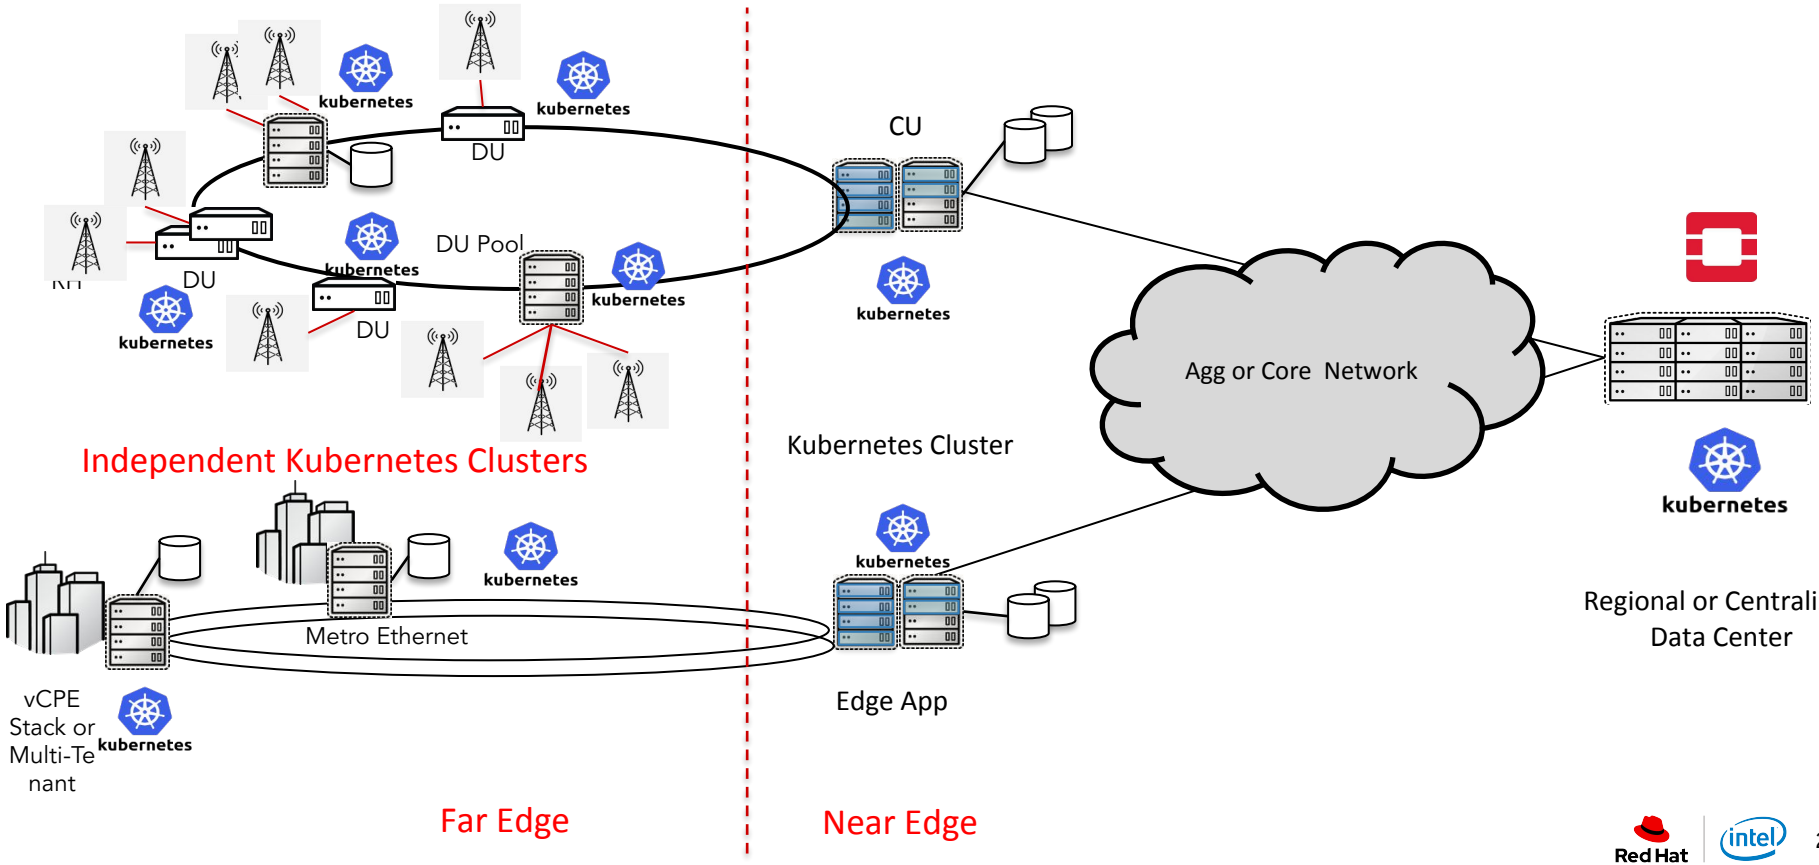

NEAR FUTURE

MULTI-CLUSTER – OPENSHIFT FOR HYBRID CLOUD
IN FUTURE FOR EDGE

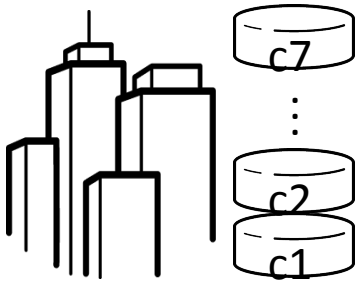

OpenShift Clusters c1 through c7

$ openshift-install launch

Cluster Registry CRD

Single Source of Truth

$ oc get clusters

Federated API

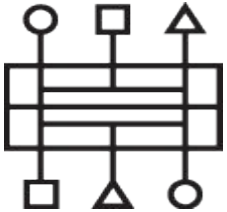

Base Federated Resources

FederatedDeployment
FederatedSecret
FederatedReplicaSet
FederatedConfigMap

Bonus: *Federate any CRD without writing code*

Schedule and Reconcile

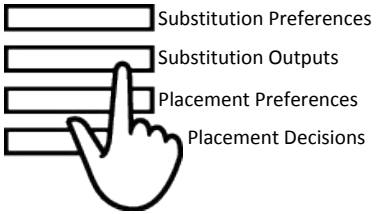

Auxiliary Resources

```
overrides:
  clusters:
  - clusterName: c1
    replicas: 5
  - clusterName: c3
    replicas: 10
  - clusterName: c7
    replicas: 15
```

https://github.com/kubernetes/community/tree/master/sig-multicluster

Red Hat | intel

NEAR FUTURE

OPENSHIFT HIVE

API Driven Multi-cluster Provisioning & Lifecycle Management

- Reliably provision/deprovision, upgrade, & configure OpenShift 4 clusters
  - 4.0: Internal only release
    - Initial support for OpenShift deployment on AWS only.
    - Primary focus supporting Dedicated for 4.0 clusters and the new UHC Portal/API.
    - May be used to drive cluster creation for CI.
- Leverages:
  - *openshift-install* - Uses CLI to launch clusters in the public cloud
  - Kubernetes Cluster API - Declarative, Kubernetes-style APIs for cluster creation, configuration, and management
  - Kubernetes Federation - Makes it easy to manage multiple clusters
- Working code & documentation now available:
  - *https://github.com/openshift/hive*

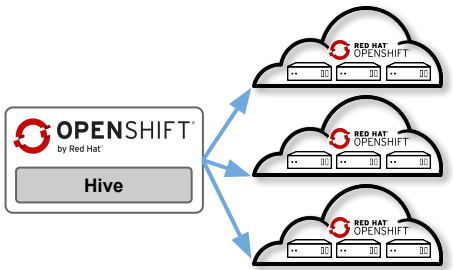


```yaml
- apiVersion: hive.openshift.io/v1alpha1
  kind: ClusterDeployment
  metadata:
    labels:
      controller-tools.k8s.io: "1.0"
    annotations:
      hive.openshift.io/delete-after: "8h"
      hive.openshift.io/try-install-once: "${TRY_INSTALL_ONCE}"
    name: ${CLUSTER_NAME}
  spec:
    platformSecrets:
      aws:
        credentials:
          name: "${CLUSTER_NAME}-aws-creds"
    images:
      hiveImage: "${HIVE_IMAGE}"
      hiveImagePullPolicy: "${HIVE_IMAGE_PULL_POLICY}"
      installerImage: "${INSTALLER_IMAGE}"
      installerImagePullPolicy: "${INSTALLER_IMAGE_PULL_POLICY}"
      releaseImage: "${OPENSHIFT_RELEASE_IMAGE}"
    sshKey:
      name: "${CLUSTER_NAME}-ssh-key"
    clusterName: ${CLUSTER_NAME}
    baseDomain: ${BASE_DOMAIN}
    networking:
      type: OpenshiftSDN
      serviceCIDR: "172.30.0.0/16"
      machineCIDR: "10.0.0.0/16"
      clusterNetworks:
        - cidr: "10.128.0.0/14"
          hostSubnetLength: 9
    platform:
      aws:
        region: us-east-1
    pullSecret:
      name: "${CLUSTER_NAME}-pull-secret"
```

Red Hat | intel

Multi-Cluster Federation

Multiple Independent Clusters

- Each Cluster is independent of each other
- Each Cluster can be installed from a central site
  - With the site and topo descriptions (yaml declarative topology) - Push model
  - From the site - Pull model
- Each Cluster manages local resources including storage
- Cookie cutter approach to site install and management
- Federation across clusters
  - From any cluster in the federation, workloads can be scheduled to any other member cluster including replica scheduling
  - Workloads can be moved from one member in the federation to another member in the federation
- Federation of API any type including CRDs

EDGE - RED HAT PRODUCT OFFERS

Edge Attribute	Result	Technology Mapping => Red Hat Impact
vRAN	Distributed Computing	DCN for OpenStack – 100ms latency constraint with OSP 13 vRAN in Containers – Remote worker Nodes for OpenShift Hybrid Cloud OpenShift Models
Distributed Install	Installation of many many sites	OC multi-cluster Install – OpenShift Hive Ansible playbook bundles Director Install of DCN over Layer 3
CUPS	Control and User Plane Separation - allows placement of applications and Functionality closer to user Local offload of traffic Various workload types	Scaling control plane via OpenShift/OpenStack and Data Plane via Smart NIC, FPGA and acceleration support on RHEL, OpenStack and OpenShift Efficient and Flexible architectures – Containers or VMs
Distributed Computing and massive scale	Orchestration Challenges Hybrid Cloud Models	Automation necessary to efficient orchestration – Ansible OpenShift Hybrid Cloud
IoT and Industry 4.0	IoT Services => Hybrid Cloud and IoT Gateways	Middleware, Messaging, OCP, JBOSS and 3Scale
Assurance for Edge	Assurance Framework – Distributed Monitoring and Management model	OSP, OCP Assurance Framework – Prometheus, Operator Framework, EFK, ELK, Collectd – Centralized Monitoring with distributed agents
Low Latency Services for Edge compute including URLLC	Drones, Autonomous Vehicles and Holographic calling	Low latency processing of information => RT Kernel, RT-KVM, PTP (Precision Time Protocol) Accuracy

EXTENDING ASSURANCE MODEL TO ULTRA/FAR EDGE

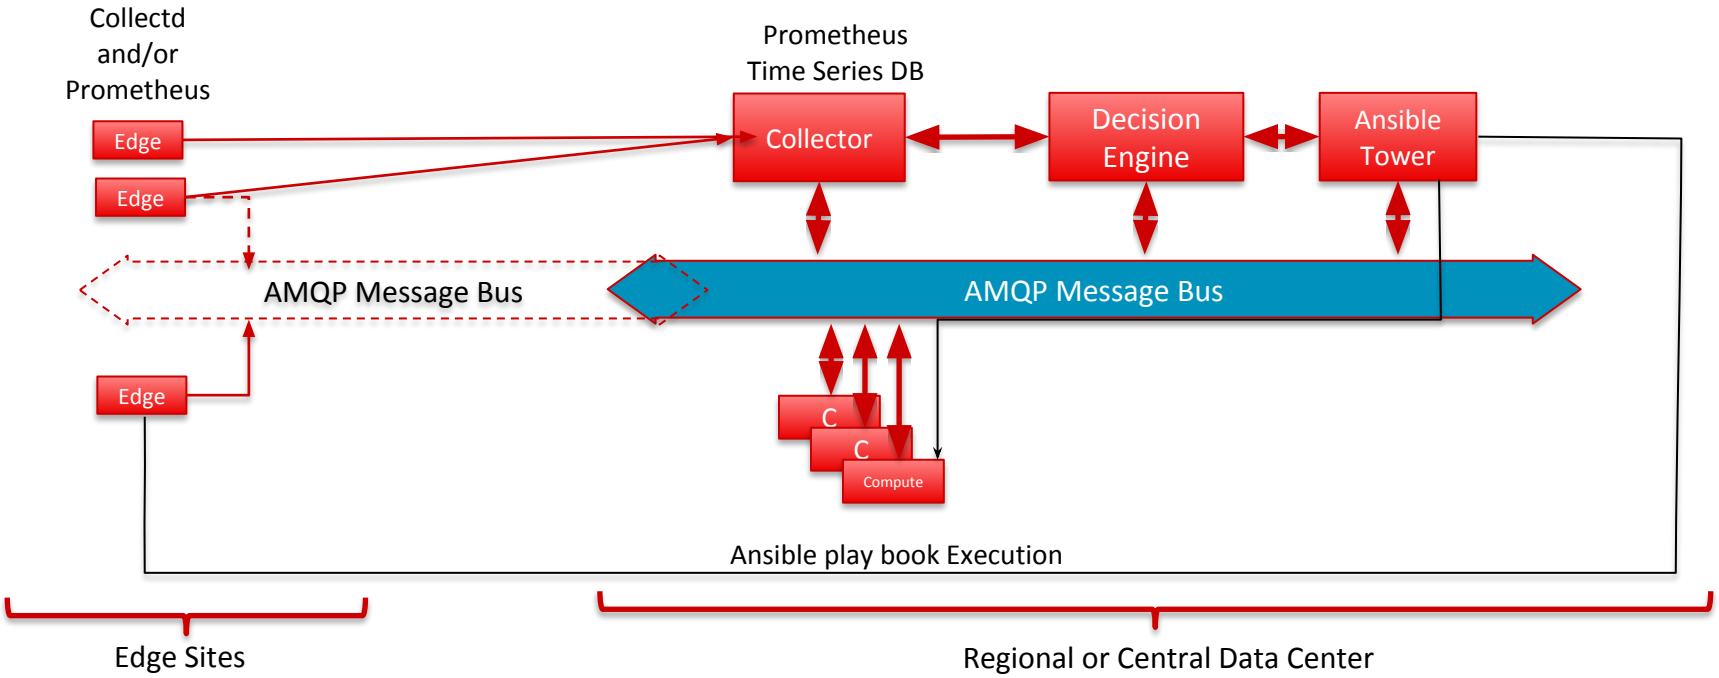

- Automating placement of workloads that meet the criteria
- Monitoring Change
- Proactive and reactive action

- Rules based closed feedback loop
- Requires all layers of stack to communicate
- Correlation

Red Hat intel

OPENSTACK, KUBERNETES, FUTURE?

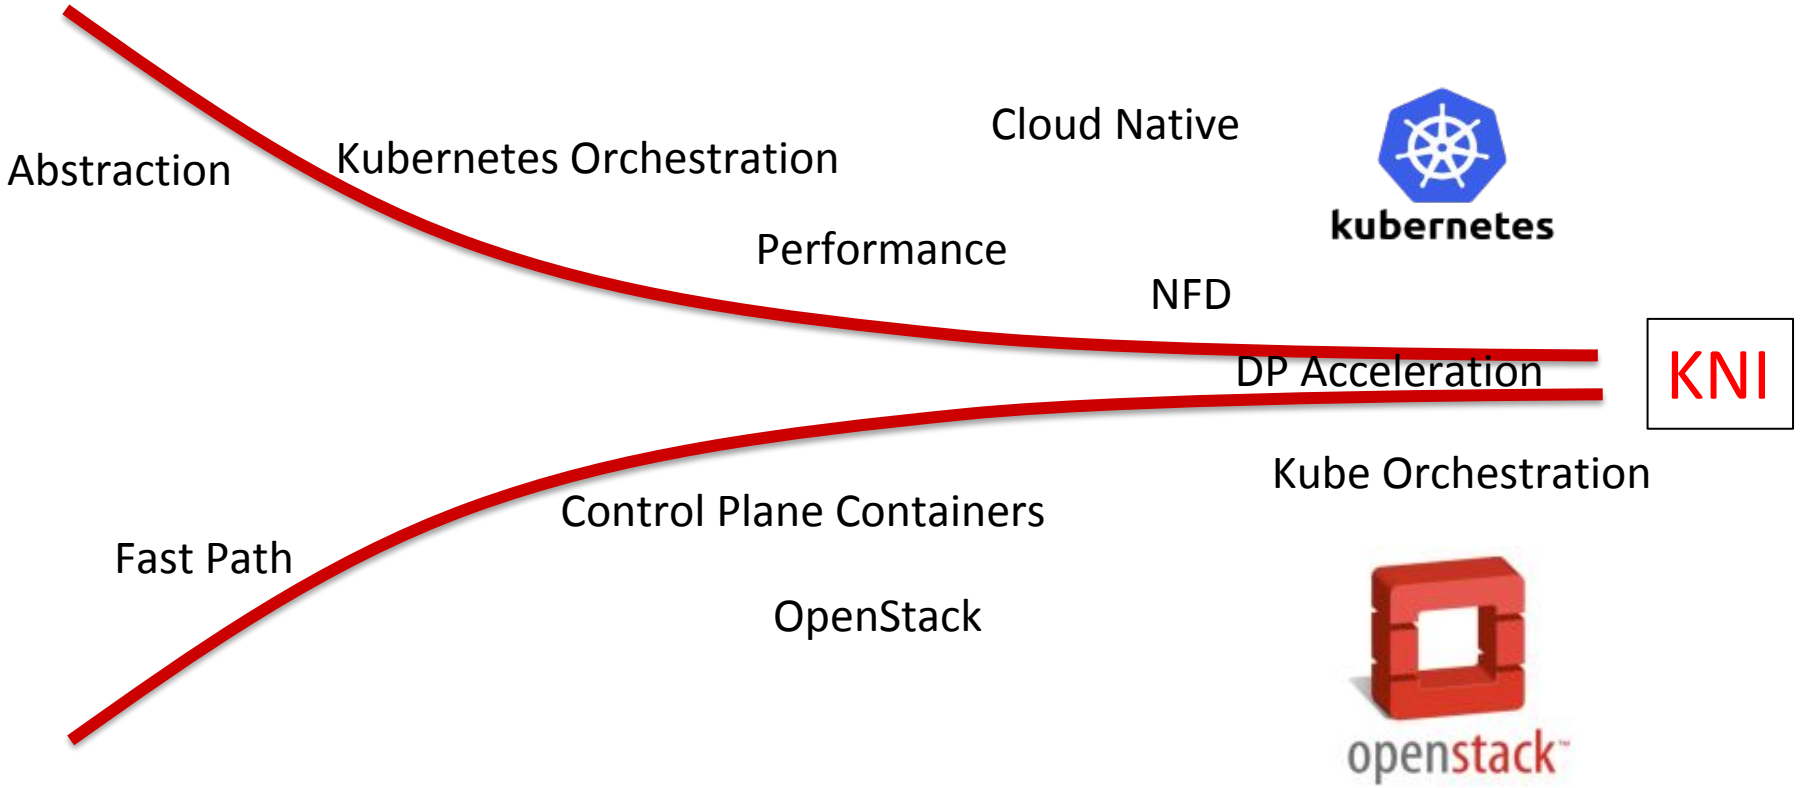

Red Hat intel

Kubernetes-Native Infra for Edge (KNI-Edge) Family

The KNI-Edge Family unites edge computing blueprints sharing the following characteristics:

- Implement the Kubernetes community's Cluster API
  - declaratively configure and consistently deploy and lifecycle manage Kubernetes clusters on-prem or public cloud, on VMs or bare metal, at the edge or at the core.
- Leverage the community's Operator Framework for app LCM
  - applications lifecycle managed as Kubernetes resources, in event-driven manner, and fully RBAC-controlled
  - more than deployment + upgrades, e.g. metering, analytics
  - created from Helm Charts, using Ansible or Go
- Optimize for Kubernetes-native container workloads
  - but allow mixing in VM-based workloads via KubeVirt as needed.

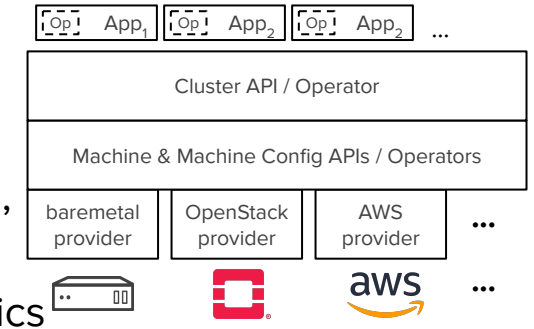

Red Hat intel

OTHER SOLUTIONS

- Third Party solutions for Kube cluster management for multisite and distributed deployment
- KubeEdge Project accepted into CNCF
  - Targeted towards edge devices -
  - MQTT for IoT
  - New release v0.2 out today

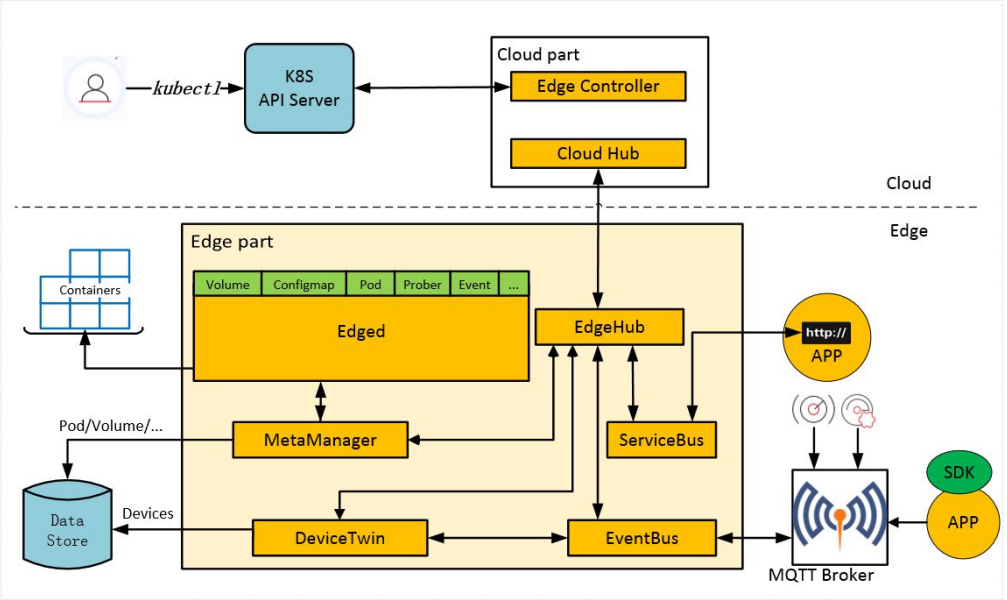

Red Hat intel

FUTURE TOPICS

In search of AI/ML

What can AI/ML offer

- Self Optimizing Infrastructure - Autonomy in operations (within constraints) for large scale – such as the “Edge Challenge”
  - Closed loop feedback control system that takes input analyzes and acts upon the result
- Workload placement and graph partition problems – as we grow to larger clouds and multiple clouds
- Network, DC troubleshooting
- Traffic hotspot detection and re-routing especially in massively distributed systems
- Security : Anomaly detection, Bot and denial of service attack detection

Intent and AI/ML for self Optimizing Infra

- Specifying Intent vs. Achieving a specific outcome
- Translation of Intent to actionable items
  - Automating placement of workloads that meet the criteria
  - Monitoring Change
  - Pro active and re-active action
- Rules based closed feedback loop
- Requires all layers of stack to communicate
  - Correlation

Hey Shadowman!!! – Mutate my Infrastructure from A to B with the set of C constraints

{A and B current and future states}

C = { rate of R, partition tolerance = null, floor = k nodes, ceiling = n nodes}

SUMMARY – RED HAT VALUE PROPOSITION

- Red Hat has the Infrastructure and tools to build the Cloud Ready and/or Cloud Native Platform
- Red Hat is enhancing its infrastructure to run VMs, containers, storage and networking seamlessly
- Red Hat has special functionality required to deploy an infrastructure for 5G
  - Real Time Linux
  - Timing Synchronization support
  - RT-KVM
  - Hardware Acceleration (Smart NICs and FPGA support)
  - Massive scale assurance infrastructure
  - Microservices catalog
  - Distributed deployment of Edge Compute
    - Distributed Compute Node - OpenStack
    - OCP-clustering and Kubernetes Federation

redhat.

THANK YOU

plus.google.com/+RedHat

facebook.com/redhatinc

linkedin.com/company/red-hat

twitter.com/RedHatNews

youtube.com/user/RedHatVideos

redhat®

RED HAT VCO MOBILE SERVICES

EDGE EXAMPLE APPLICATION

Streaming Video Optimization

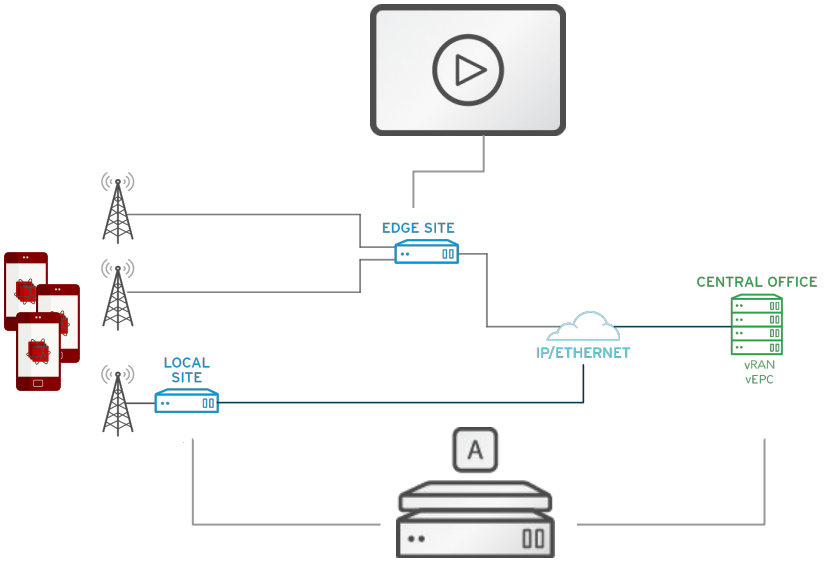

Multi-View Video Feeds

Red Hat | intel

LATENCY TESTING – WITH SPECTRE AND MELTDOWN

Attribute		RHEL -RT	Centos + Starling X Patches = Wind River	Comments
CVE ON	Min (us)	9	9	Worst case scenario with CVE mitigation on and background stress traffic
	Average (us)	10	10	
	Max (us)	26	25	
CVE OFF	Min (us)	4	4	CVE mitigation Off and background stress traffic
	Average (us)	4	4	
	Max (us)	17	19	

Guest with 2 vCPU – running stress and cyclic test – 24 Hour Duration

Host Level RT latency

- All tests are 8 hour duration
- CVE mitigation:
- ON - spectre/meltdown CVE enabled
- OFF- spectre/meltdown CVE disabled

CVE mitigation	cyclictest mode	kernel	1-cpu			8-cpu		
			min	avg	max	min	avg	max
ON	with background stress	RHEL-RT 7.5	4	5	7	4-5	5-6	21-26
		Centos-RT + WR	4	4	7	4-5	5	21-25
OFF	no background stress	RHEL-RT 7.5	1	1	4	1-2	2	4-6
		Centos-RT + WR	1	1	4	1-2	2	3-6

Agenda:

- Industry Surveys/Trends
- Hot use cases/Solutions
  - -VCO
  - -vCDN/Media & Entertainment Examples
- 5G & Edge Solutions
  - -Current Offers
  - -vRAN & Edge - Rakuten Use case
  - -KNI and Akraino
- Untapped spaces/conversations
  - -Automation, AI/ML
  - -Pricing
- Summary

SOLUTIONS

VIRTUAL CENTRAL OFFICE

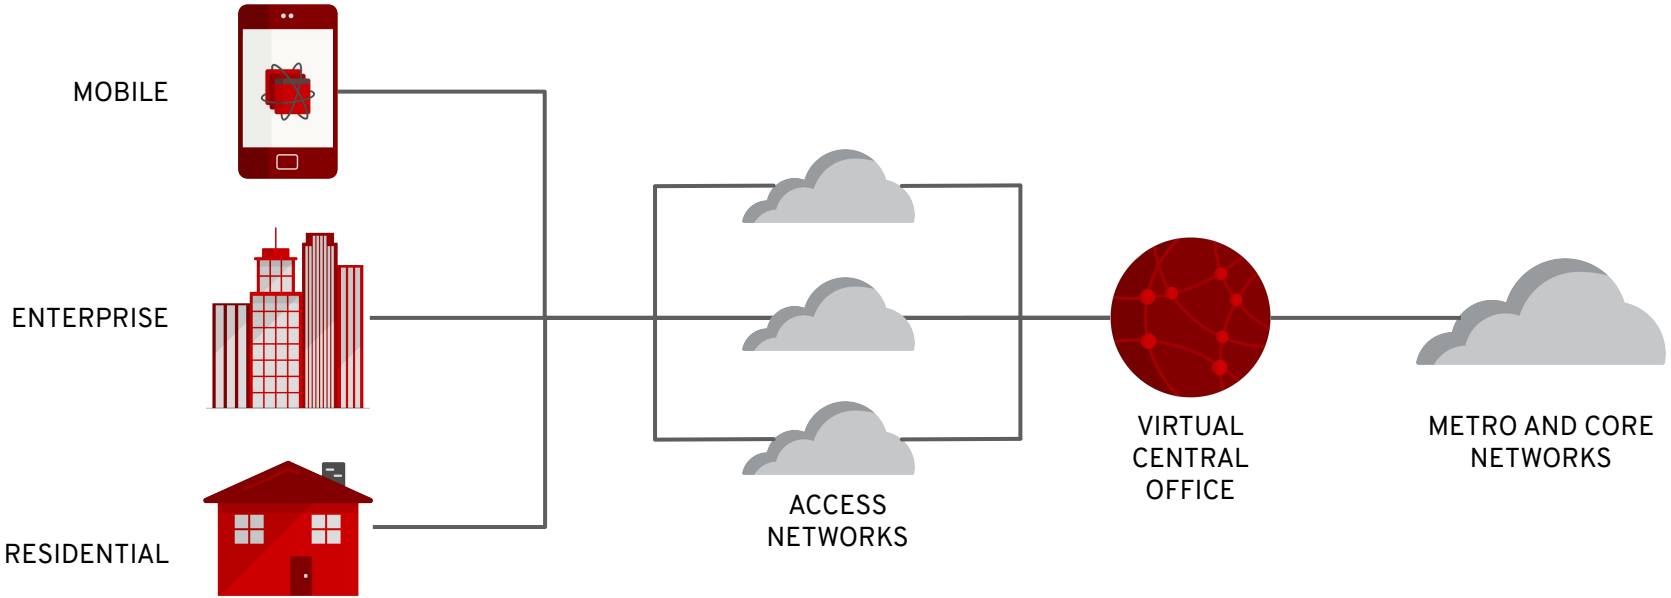

Red Hat | intel

VCO 2.0 – ONS POC

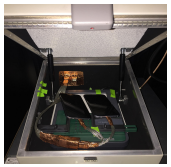

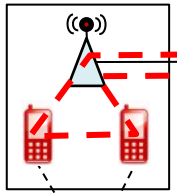

SDR

RU

DU

CU

Baremetal Nodes
Managed by OpenStack

OpenStack

SGW

MME

HSS

SD

PGW

PGW

GiLAN

Jumphost

Internet

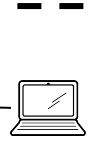

Amsterdam Stage

VPN

OPENVPN®

Internet

MOBILE SERVICES

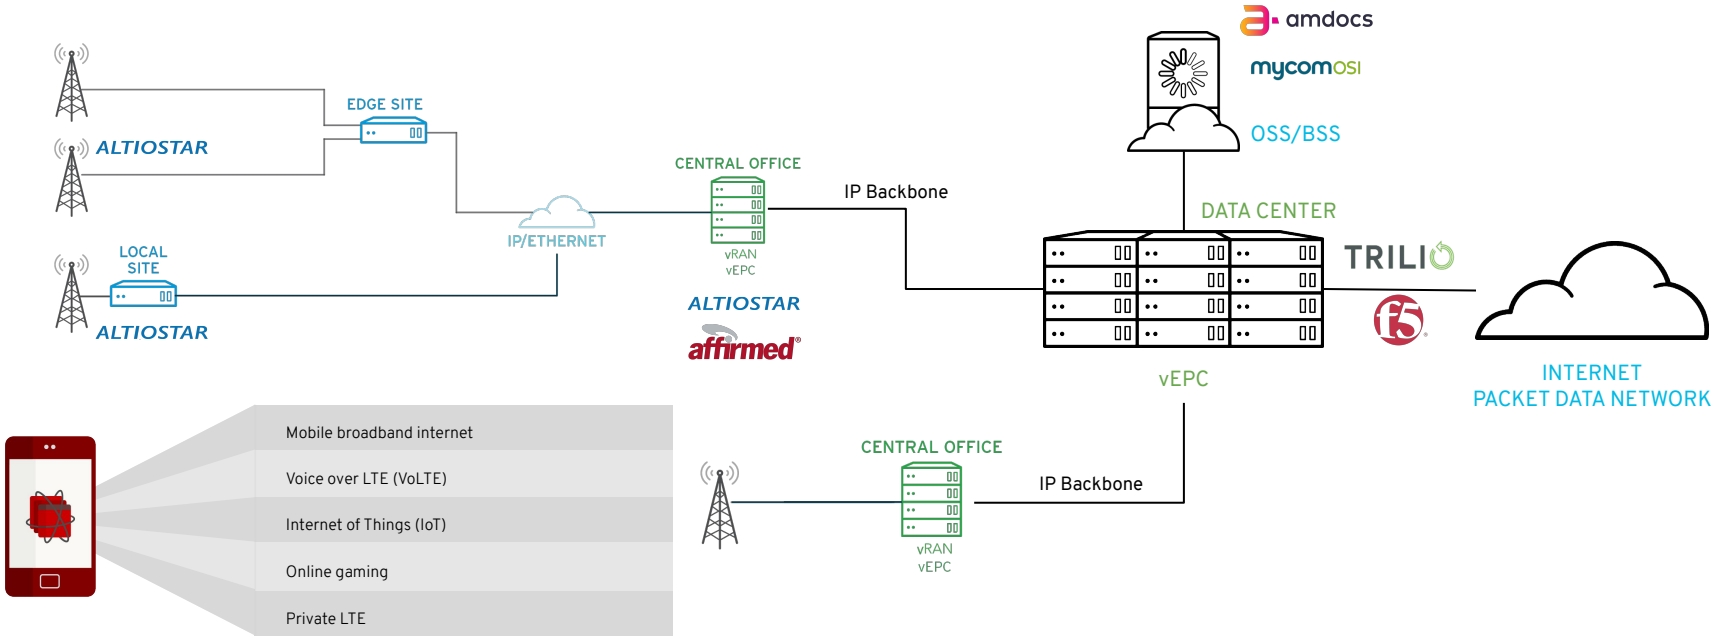

Red Hat | intel

PACKET CORE EVOLUTION 4G/LTE TO 5G

MOVE TO CONTAINERS

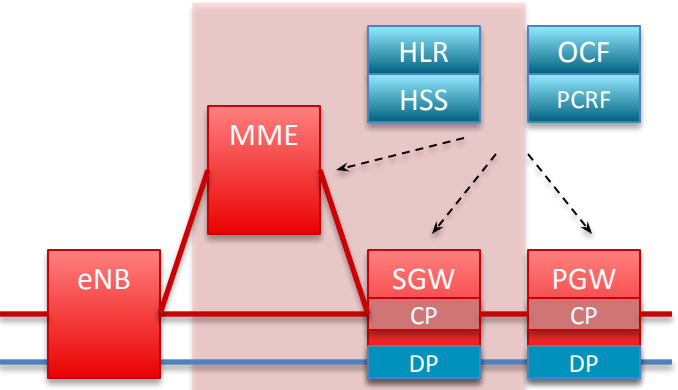

Box / Device centric
LTE/4G

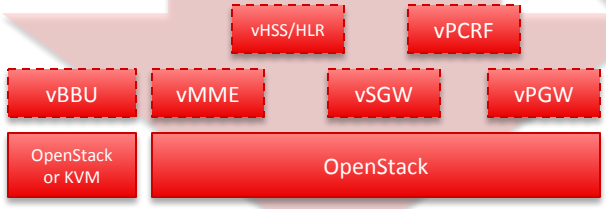

- CP-DP Separation
- UPF is controlled by AMF and SMF
- Data plane extensibility

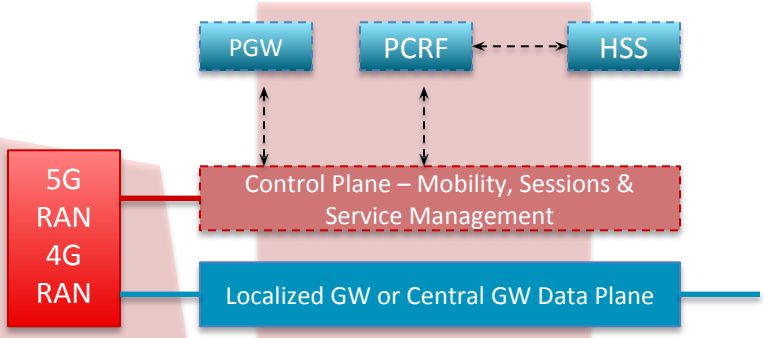

5G - Cloud Based

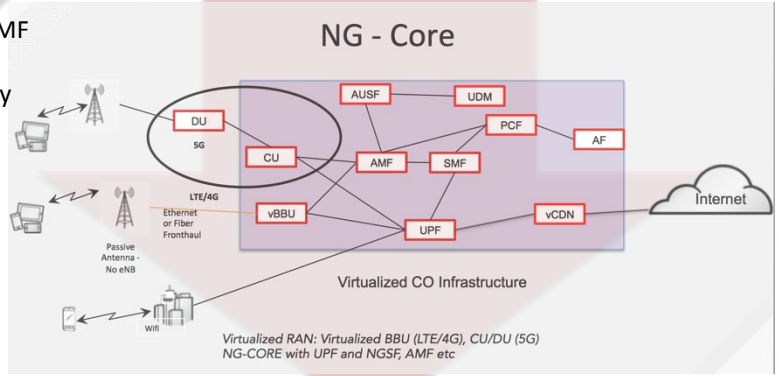

Red Hat intel

Micro-segmentation with Slicing

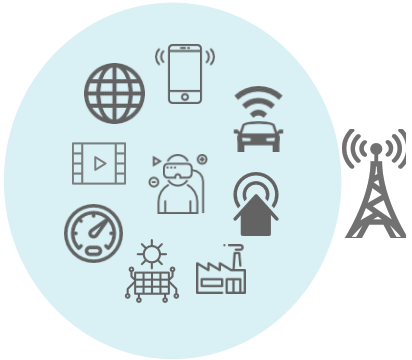

MBB/eMBB

- Pre-Paid Consumer
- Post-Paid Consumer
- Fixed Wireless Access –Residential
- Fixed Wireless Access- Small Business
- Consumer MVNO

IoT/m-IoT

- Federal Government Devices
- Car Manufacturer
- Farm Conglomerate
- Mining Company
- Smart City

LLC/URLLC

- Public Safety Video Surveillance
- Self Driving Cars
- Mission Critical Application-Health
- Industrial Automation
- Augmented Reality

Cloud Native Architecture allows for a low “minimum cost of entry” per slice – Ability to cost effectively scale to thousands of slices of all different sizes

Red Hat | intel

SCALE AND SIZE OF MARKET

- RAN
  - Customer Examples
    - US has 350K cell tower locations
    - Verizon 150K Cell Sites
    - AT&T ~150 K Cell Sites
    - Reliance Jio – 200K Cell Sites
    - China Mobile 2M+
  - Edge Compute
    - Additional compute requirements over an above RAN sites depending on applications
  - Add – small cell deployments for higher frequency spectrum
  - Ratio of small cell to macro cell – min 10:1 – dense urban
  - Coverage and penetration rate 30-40% of macro base station market

Multiply 40% of current base station market 10 fold to get a number for Small Cell market size

5G – WHY DOES RED HAT CARE?

5G Attribute	Result	Technology Mapping => Red Hat Impact
FWA (Fixed Wireless Access)	High Speed Broadband Service based on mmWave band spectrum	Performance, throughput – OSP and OCP, Smart NIC, FPGA support
NG-Core	Next Generation Packet Core for 5G – based on Microservices and Containers	OpenShift and KNI/Foundation
High Frequency Spectrum with more spectrum bandwidth	High Speed Residential Broadband service => Instant Subscriber acquisition	Hardware deployment at CO and local pops => OpenStack and KNI
Massive Scale	IoT Services => Hybrid Cloud	Hybrid Cloud => Middleware, Messaging, OCP
Control-User Plane Separation	Flexible placement of traffic exit points allows Edge compute for application optimization	Edge Compute foot prints => CNV, OSP, OCP etc
Low Latency Services (URLLC)	Drones, Autonomous Vehicles and Holographic calling	Low latency processing of information => RT Kernel, RT-KVM, PTP (Precision Time Protocol) Accuracy
Slicing	Virtual Operators, Private Mobile Broadband, Distinct grades of service	Distributed Cloud deployment models, Infrastructure partitioning, multi-tenancy, Networking and VPN support
4x4 MIMO, Beam Forming	Higher throughput per user on Macro cell and small cell	Higher Aggregate throughput => Smart NIC, FPGA etc

EDGE DEPLOYMENT– REAL CUSTOMER EXAMPLE

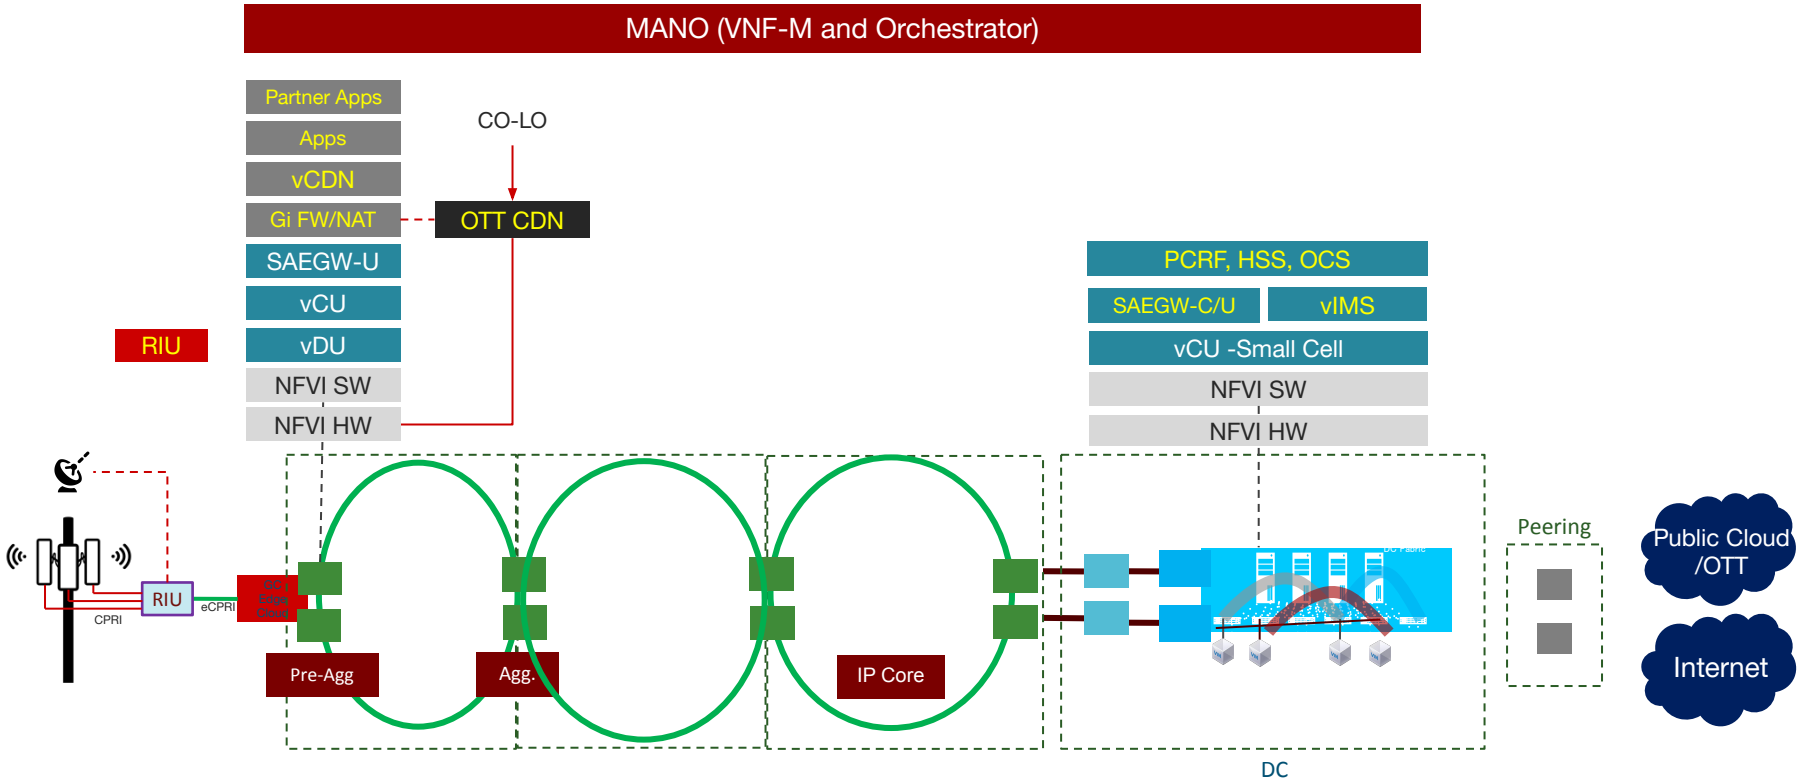

SUMMARY

Key Messages

- 5G Enables Edge Compute through Control and User Plane Separation (CUPS)
- 5G provides 100x more bandwidth through EMB Services (mm Wave broadband service)
- 5G is all about being Cloud Native
- 5G an enabler to new services with Ultra Low Latency (URLLC) – Autonomous Vehicles, AR/VR and Holographic calling or gaming
- 5G can scale IoT to billions of devices
- 5G puts stringent constraints on infrastructure – low latency, high performance, cloud native and massive scale
- Edge is necessary to address URLLC use cases
- Edge is applicable outside of 5G
- Edge has specialized requirements
- Edge Orchestration and provisioning – big challenges – Automation is key
- Requirements being addressed upstream

WHAT IS DIGITAL TRANSFORMATION?

Digital transformation is the integration of digital technology into all areas of a business, fundamentally changing how you operate and **deliver value** to customers. It's also a **cultural change** that requires organizations to continually challenge the status quo, experiment, and get comfortable with failure.

Digital Transformation

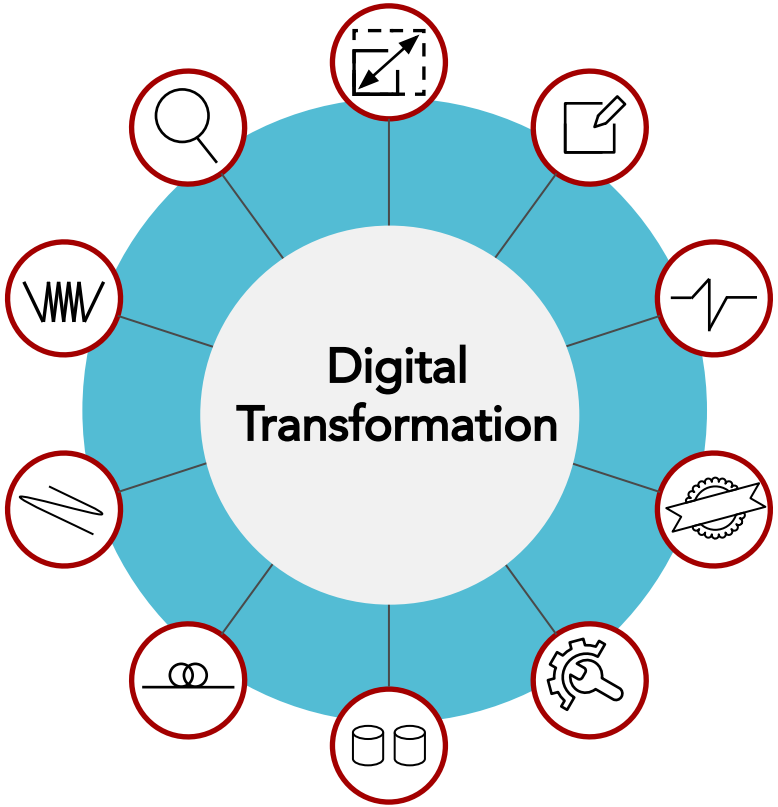

Infrastructure Transformation

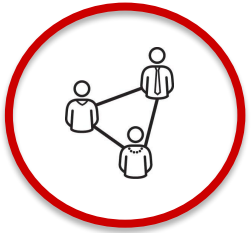

People Transformation

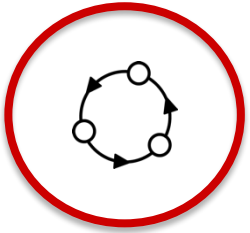

Process Transformation

Red Hat | intel

Infrastructure Transformation

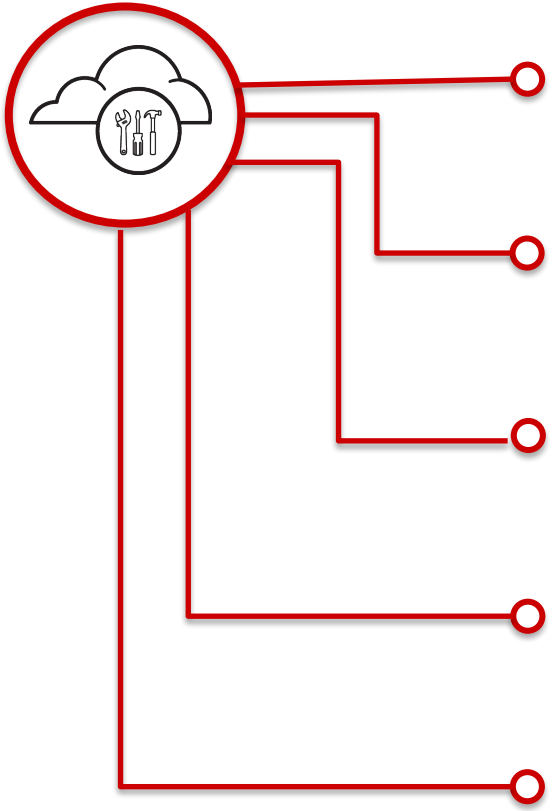

- Programmable
- Scale to massive usage and deployment
- Flexible
- Re-usable
- OPEN

Red Hat intel

People Transformation

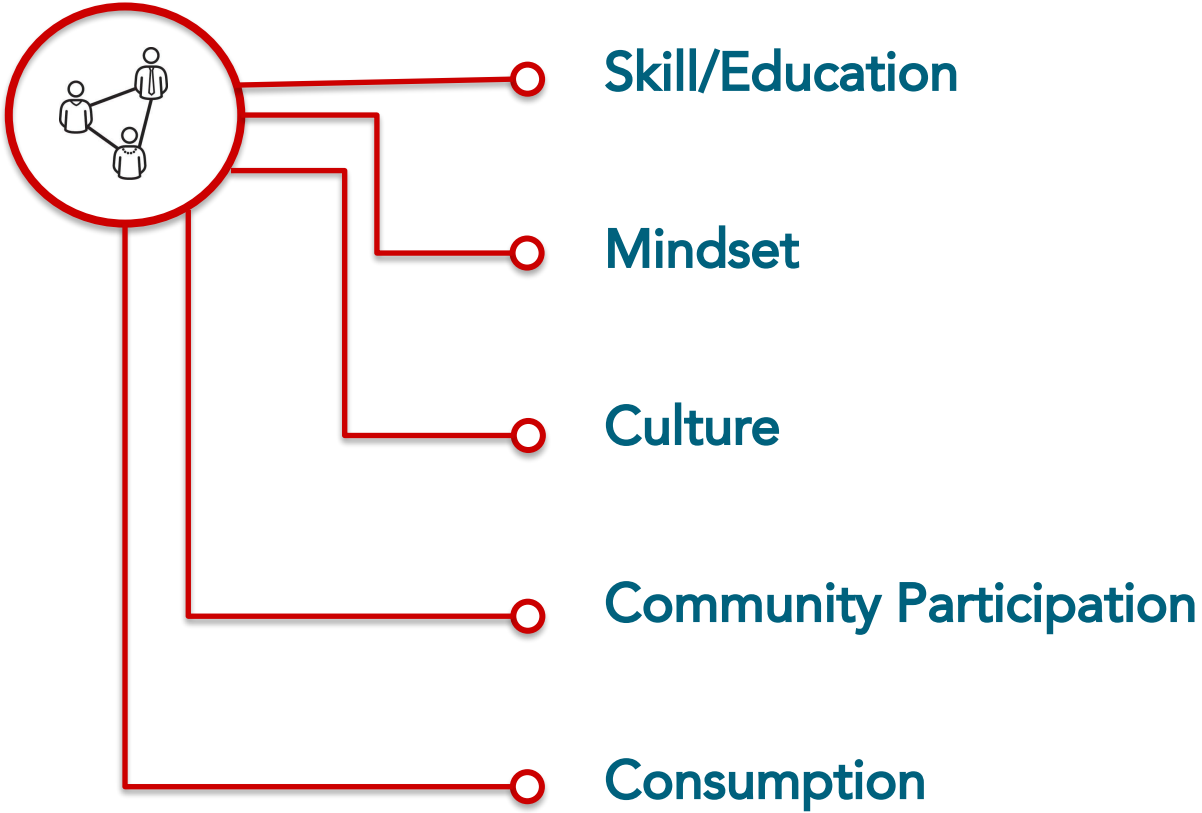

Process Transformation

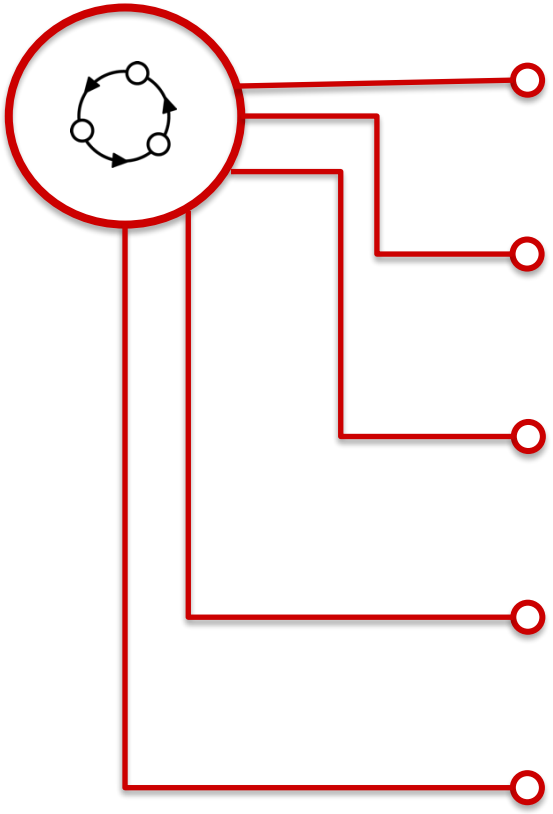

Ease of doing business / Barrier to Entry

Regulation

Data Privacy & Security

Transparency

Market place / Clearing House

CHALLENGES

Culture
Open Approach

Ecosystem
Critical mass for a thriving ecosystem

Technology
Readiness and adoption with Risk Mitigation

KNI

Kubernetes-Native Infra for Edge (KNI-Edge) Family

The KNI-Edge Family unites edge computing blueprints sharing the following characteristics:

- Implement the Kubernetes community's Cluster API
  - declaratively configure and consistently deploy and lifecycle manage Kubernetes clusters on-prem or public cloud, on VMs or bare metal, at the edge or at the core.
- Leverage the community's Operator Framework for app LCM
  - applications lifecycle managed as Kubernetes resources, in event-driven manner, and fully RBAC-controlled
  - more than deployment+upgrades, e.g. metering, analytics
  - created from Helm Charts, using Ansible or Go
- Optimize for Kubernetes-native container workloads
  - but allow mixing in VM-based workloads via KubeVirt as needed.

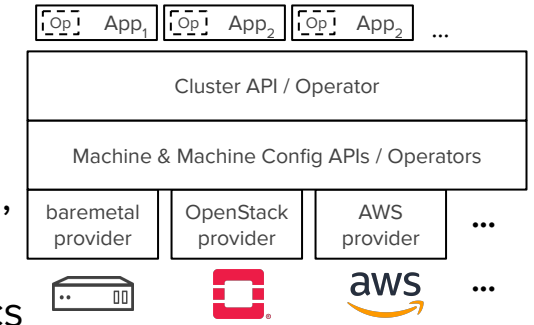

KNI-Edge Family - Proposal Template

Case Attributes	Description
Type	New
Blueprint Family - Proposed Name	Kubernetes-Native Infrastructure for Edge (KNI-Edge)
Use Case	various, e.g.: ● Provider Access Edge (Far/Near), MEC ● Industrial Automation ● Enterprise Edge ● ...
Blueprint proposed	various; initially: ● Provider Access Edge (PAE) ● Industrial Edge (IE)
Initial POD Cost (capex)	(depends on blueprint)
Scale	1 to hundreds of nodes, 1 to thousands of sites.
Applications	any type of workloads: ● containerized or VM-based ● real-time, ultra-low latency or high-throughput ● NFV, IoT, AI/ML, Serverless, ...
Power Restrictions	(depends on blueprint)
Preferred Infrastructure orchestration	End-to-end Service Orchestration: depends on use case; e.g. ONAP App Lifecycle Management: Kubernetes Operators Cluster Lifecycle Management: Kubernetes Cluster API/Controller Container Platform: Kubernetes (OKD) Container Runtime: CRI-O w/compatible backends VM Runtime: KubeVirt OS: CentOS, CentOS-rt, or CoreOS
Additional Details	

Thank you

Red Hat is the world's leading provider of enterprise open source software solutions. Award-winning support, training, and consulting services make Red Hat a trusted adviser to the Fortune 500.


linkedin.com/company/red-hat


youtube.com/user/RedHatVideos


facebook.com/redhatinc


twitter.com/RedHat

Red Hat | intel