Going Towards The Edge / RAN / 5G

Red Hat @ Telco Networking

Open Tour 2022

Aki Pellinen, Senior Telco Solution Architect
Timo Jokiaho, Chief Technologist
Telecom, Media & Entertainment
Red Hat EMEA

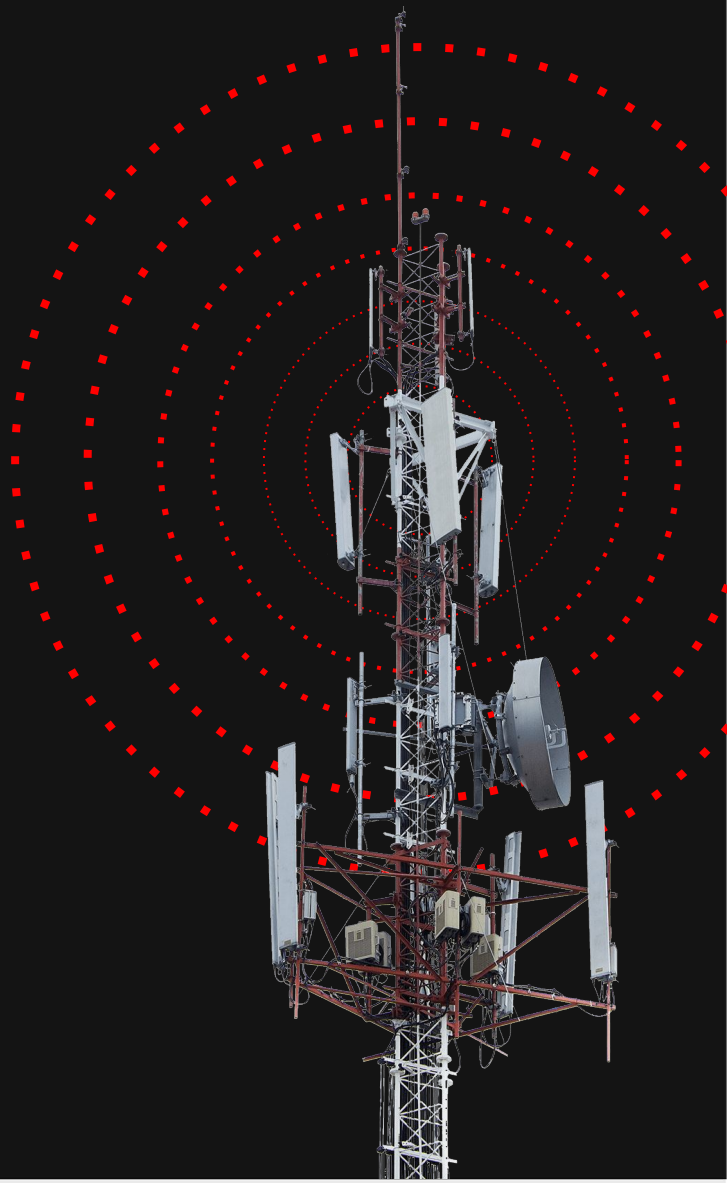

Agenda

- Introduction
- 5G Core
  - Overview
  - Service Based Architecture
  - Multivendor Models
- 5G (4G) RAN
  - Disaggregated RAN Overview
  - RAN Requirements to Cloud Platforms
- Technical Platform Update
- Red Hat Telco Ecosystem
- Summary

Red Hat

NFV Journey - Cloud Native

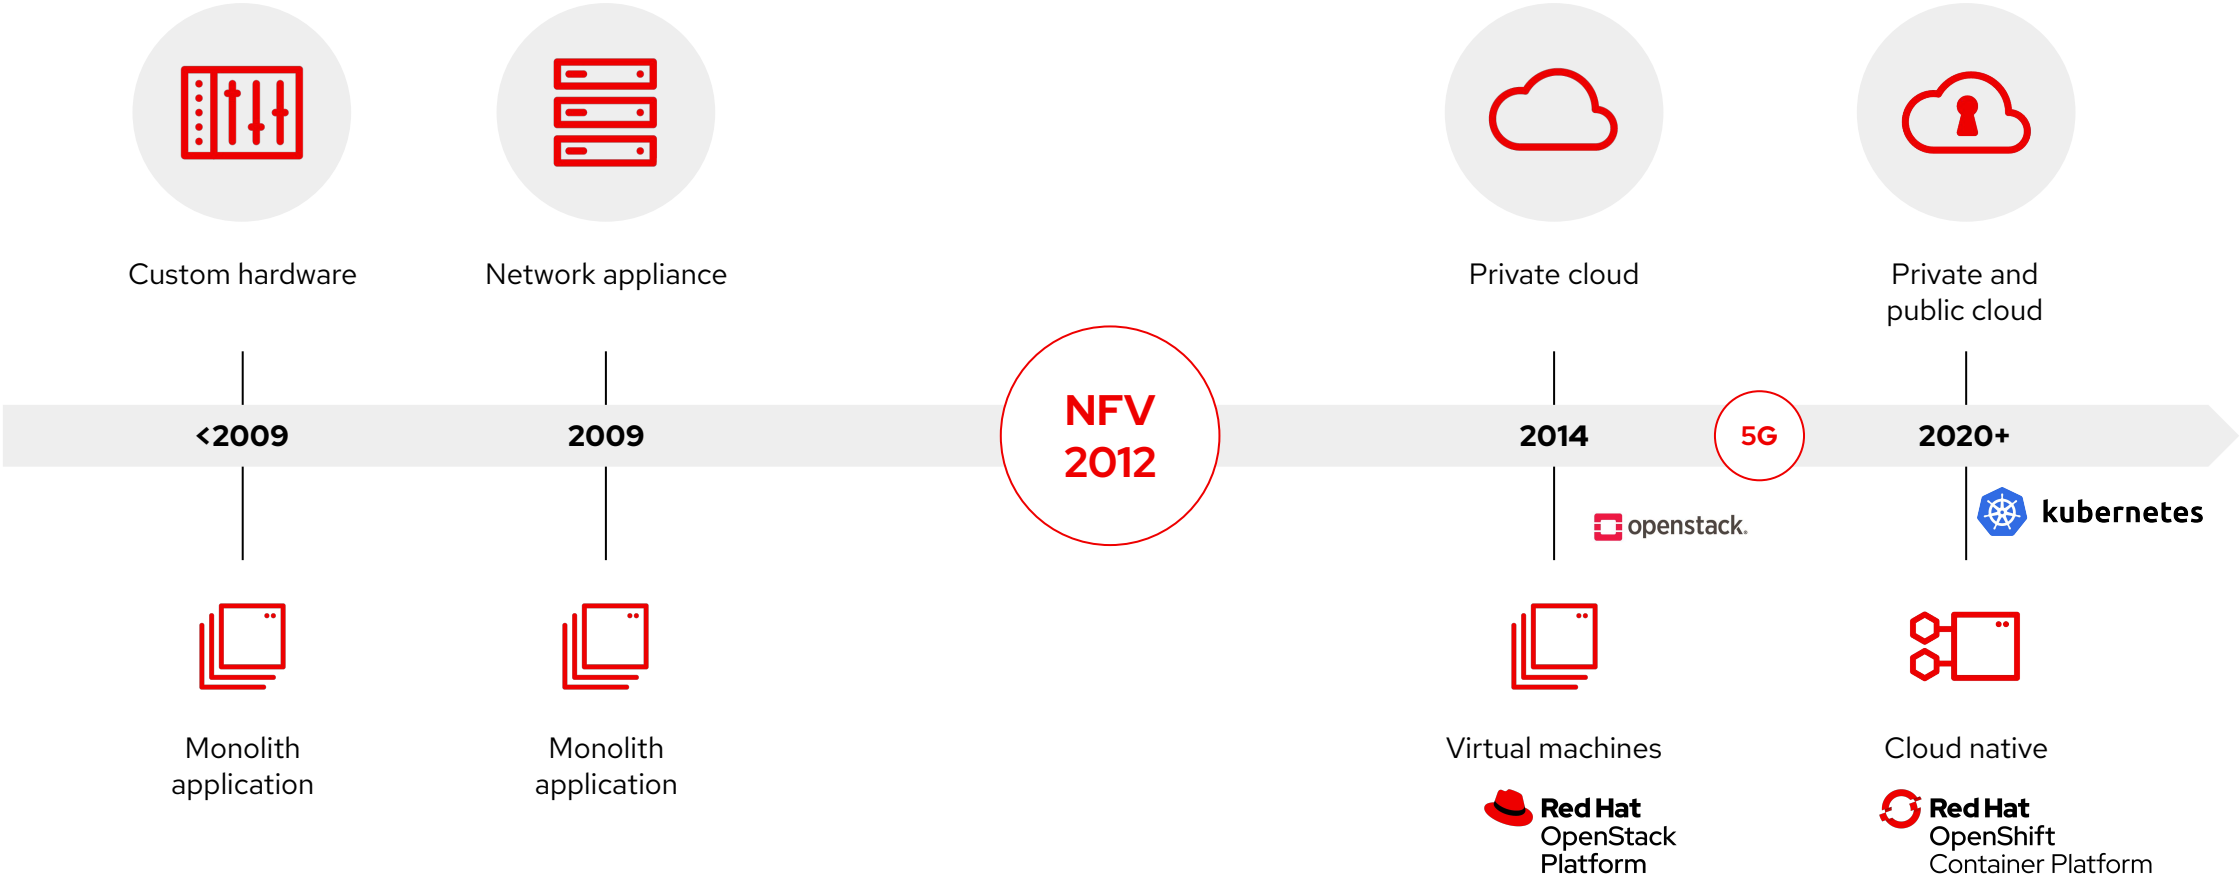

Red Hat

Open Telco Cloud

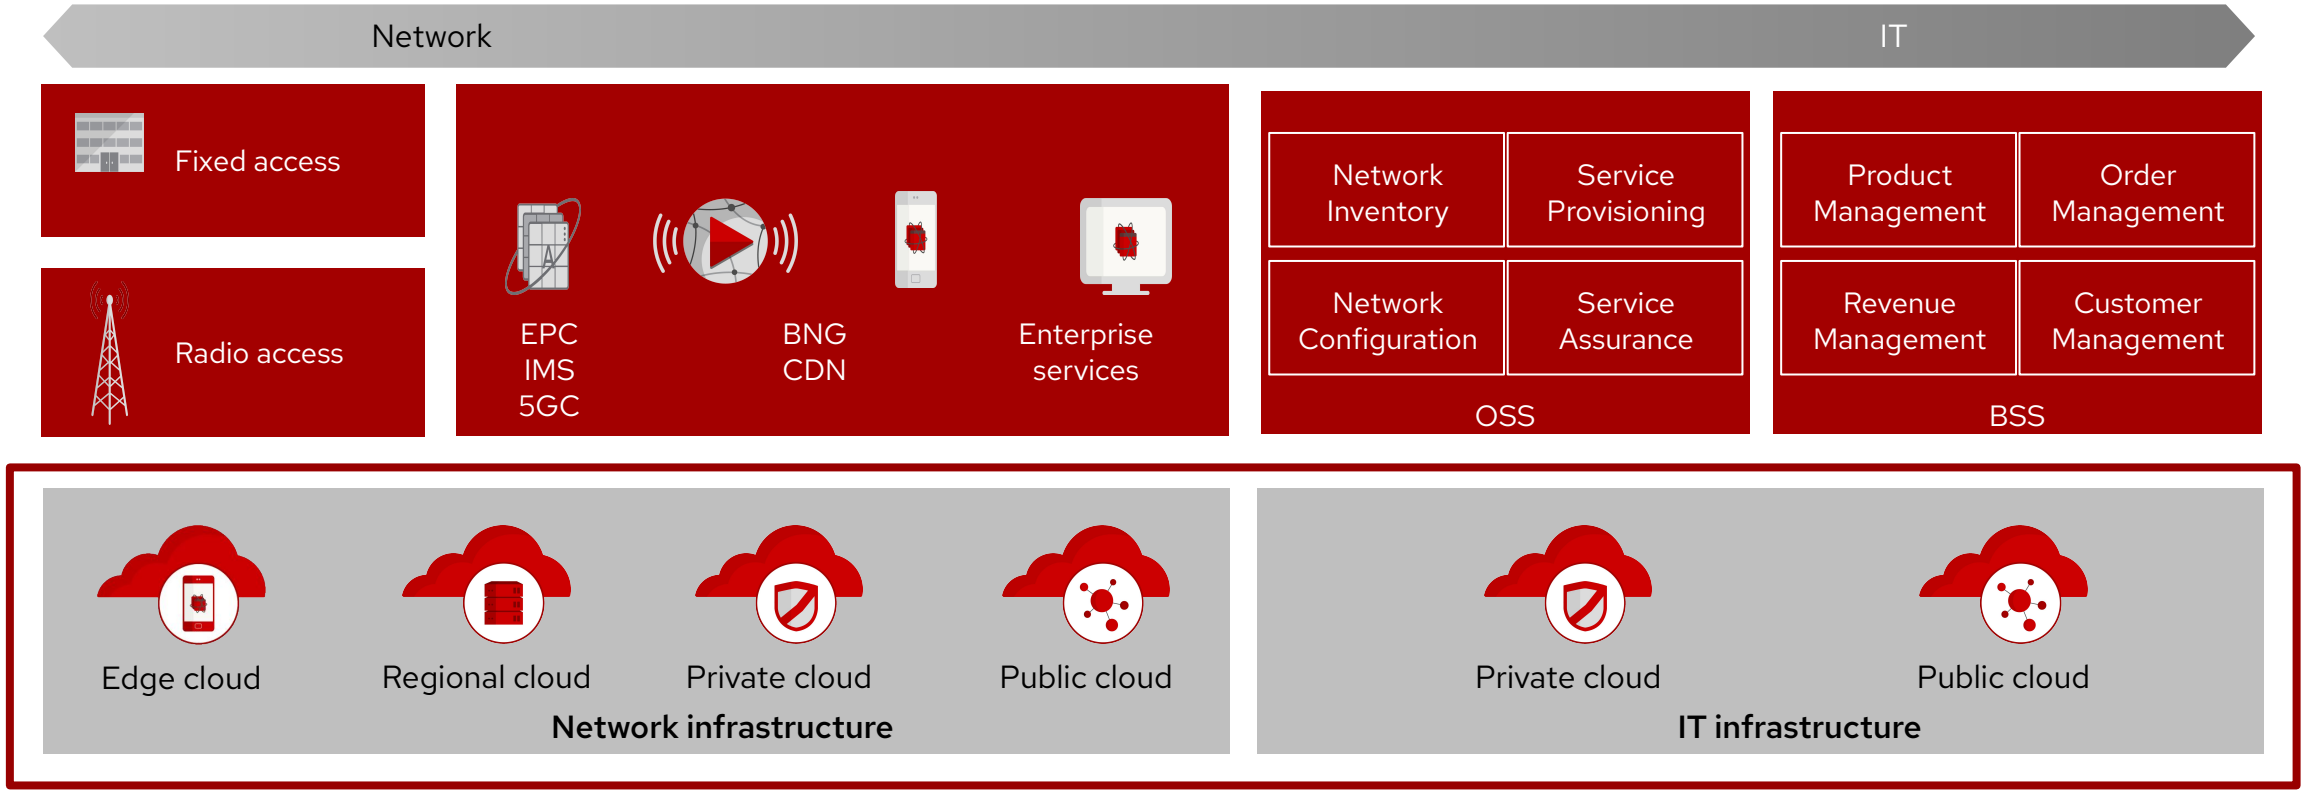

Red Hat
OpenShift
Container Platform

Red Hat

Enabling Uniform Open Telco Cloud

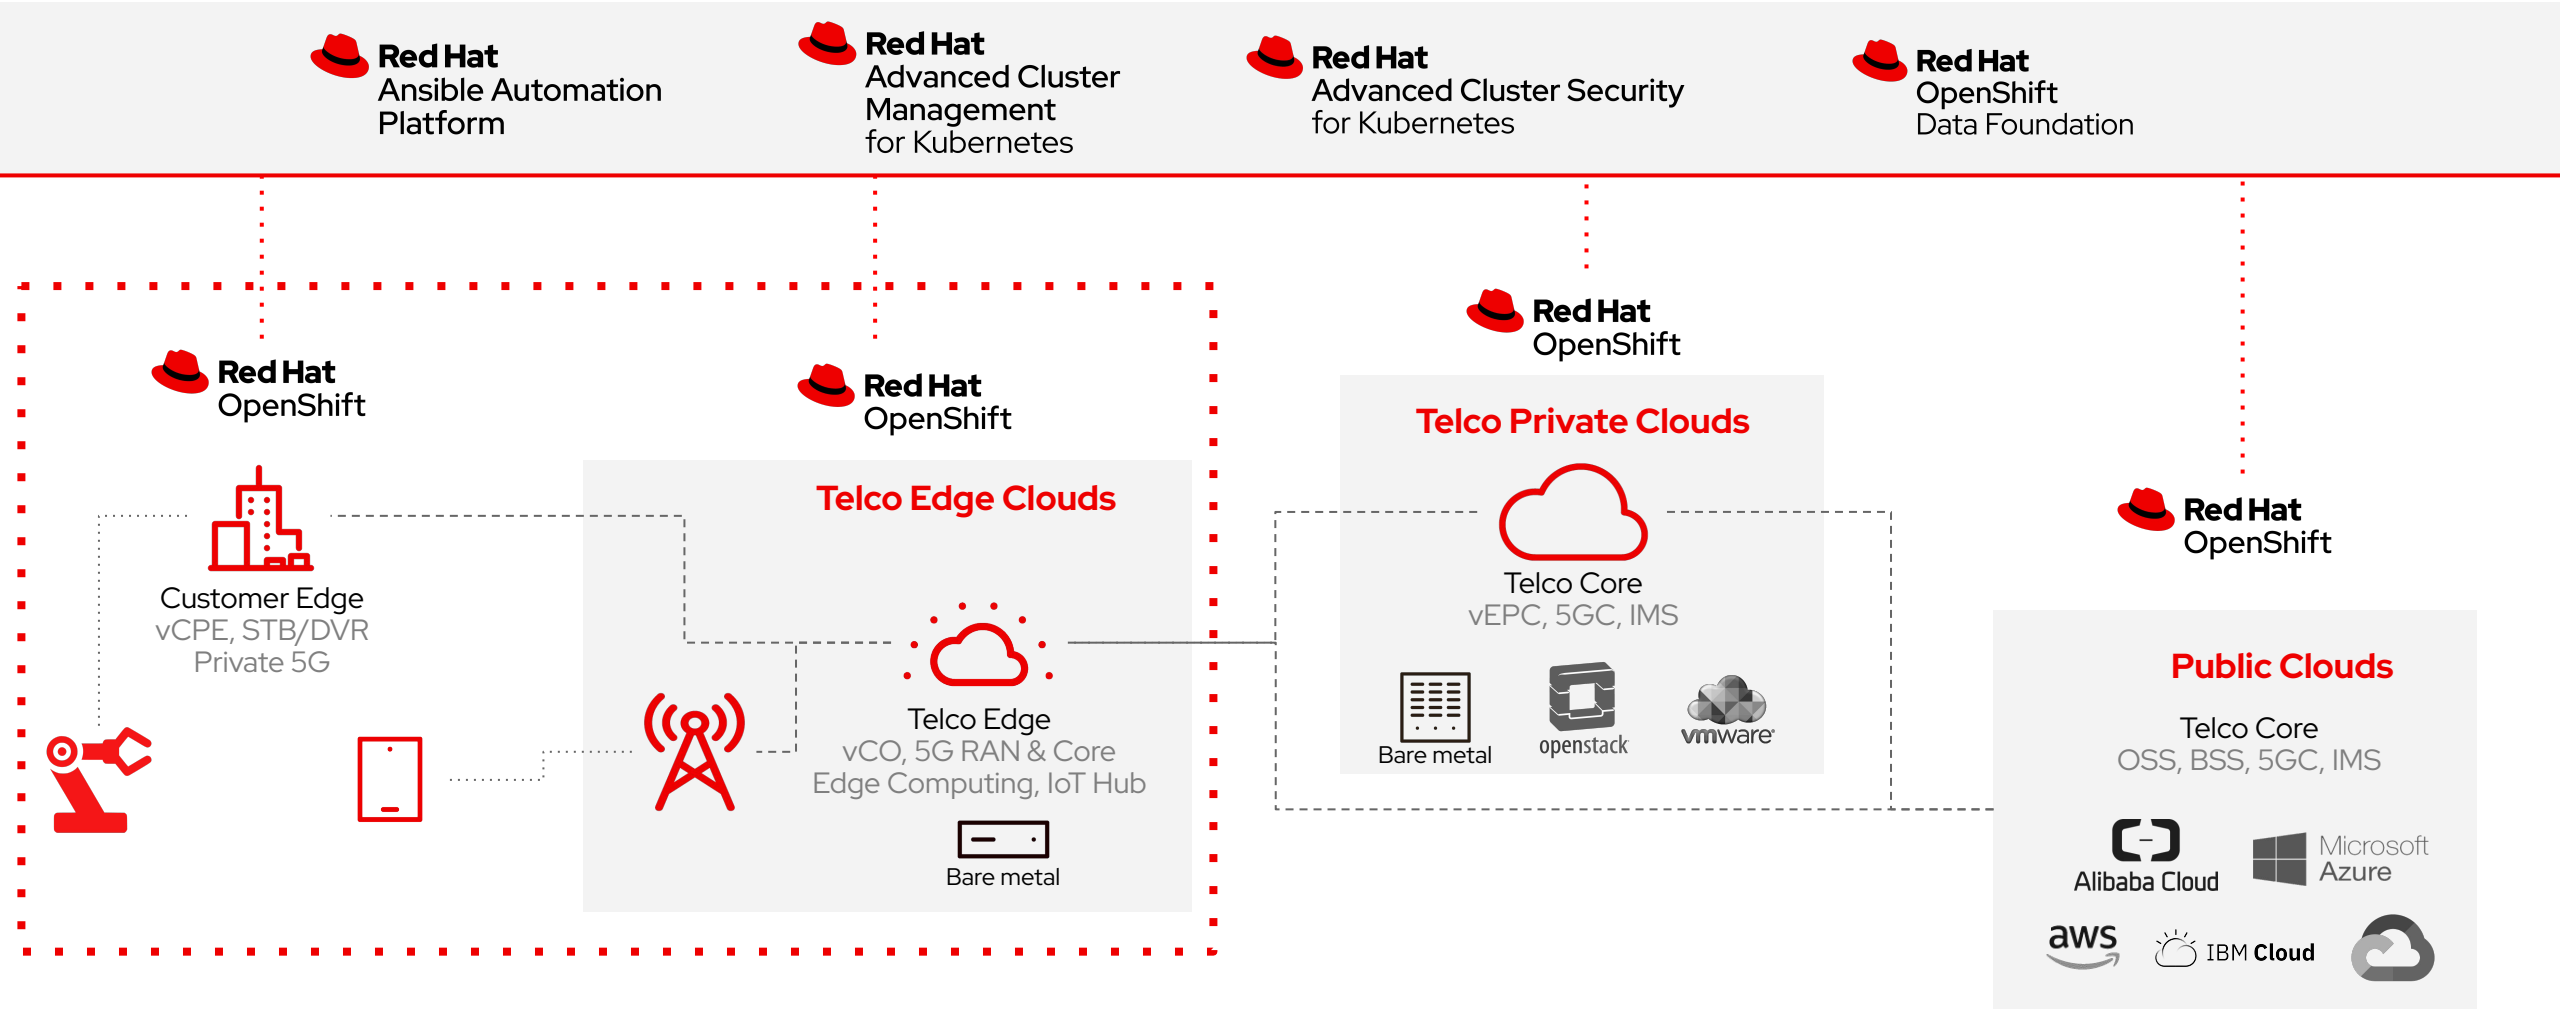

Red Hat

5G Core

Red Hat

5G is Enabler Technology
With Three Main Use Case Categories for 5G

Ultra-reliable, immersive experiences for people and objects when and where it matters most

Enhanced mobile broadband (eMBB)

5G enhances data speeds and experiences with new radio capabilities like mmWave frequency spectrum for higher bandwidth allocation and theoretical throughput up to 20Gbps.

Ultra-reliable, low-latency communication (uRLLC)

5G supports vertical industry requirements, including sub-millisecond latency with less than 1 lost packet in 10^5 packets.

Massive machine type communications (mMTC)

5G supports cost-efficient and robust connection for up to 1 million mMTC, NB-IOT*, and LTE-M** devices per square kilometer without network overloading.

* NB-IOT = Narrowband internet of things
** LTE-M = Long Term Evolution Machine Type Communication

Use Case Specialization with Slicing

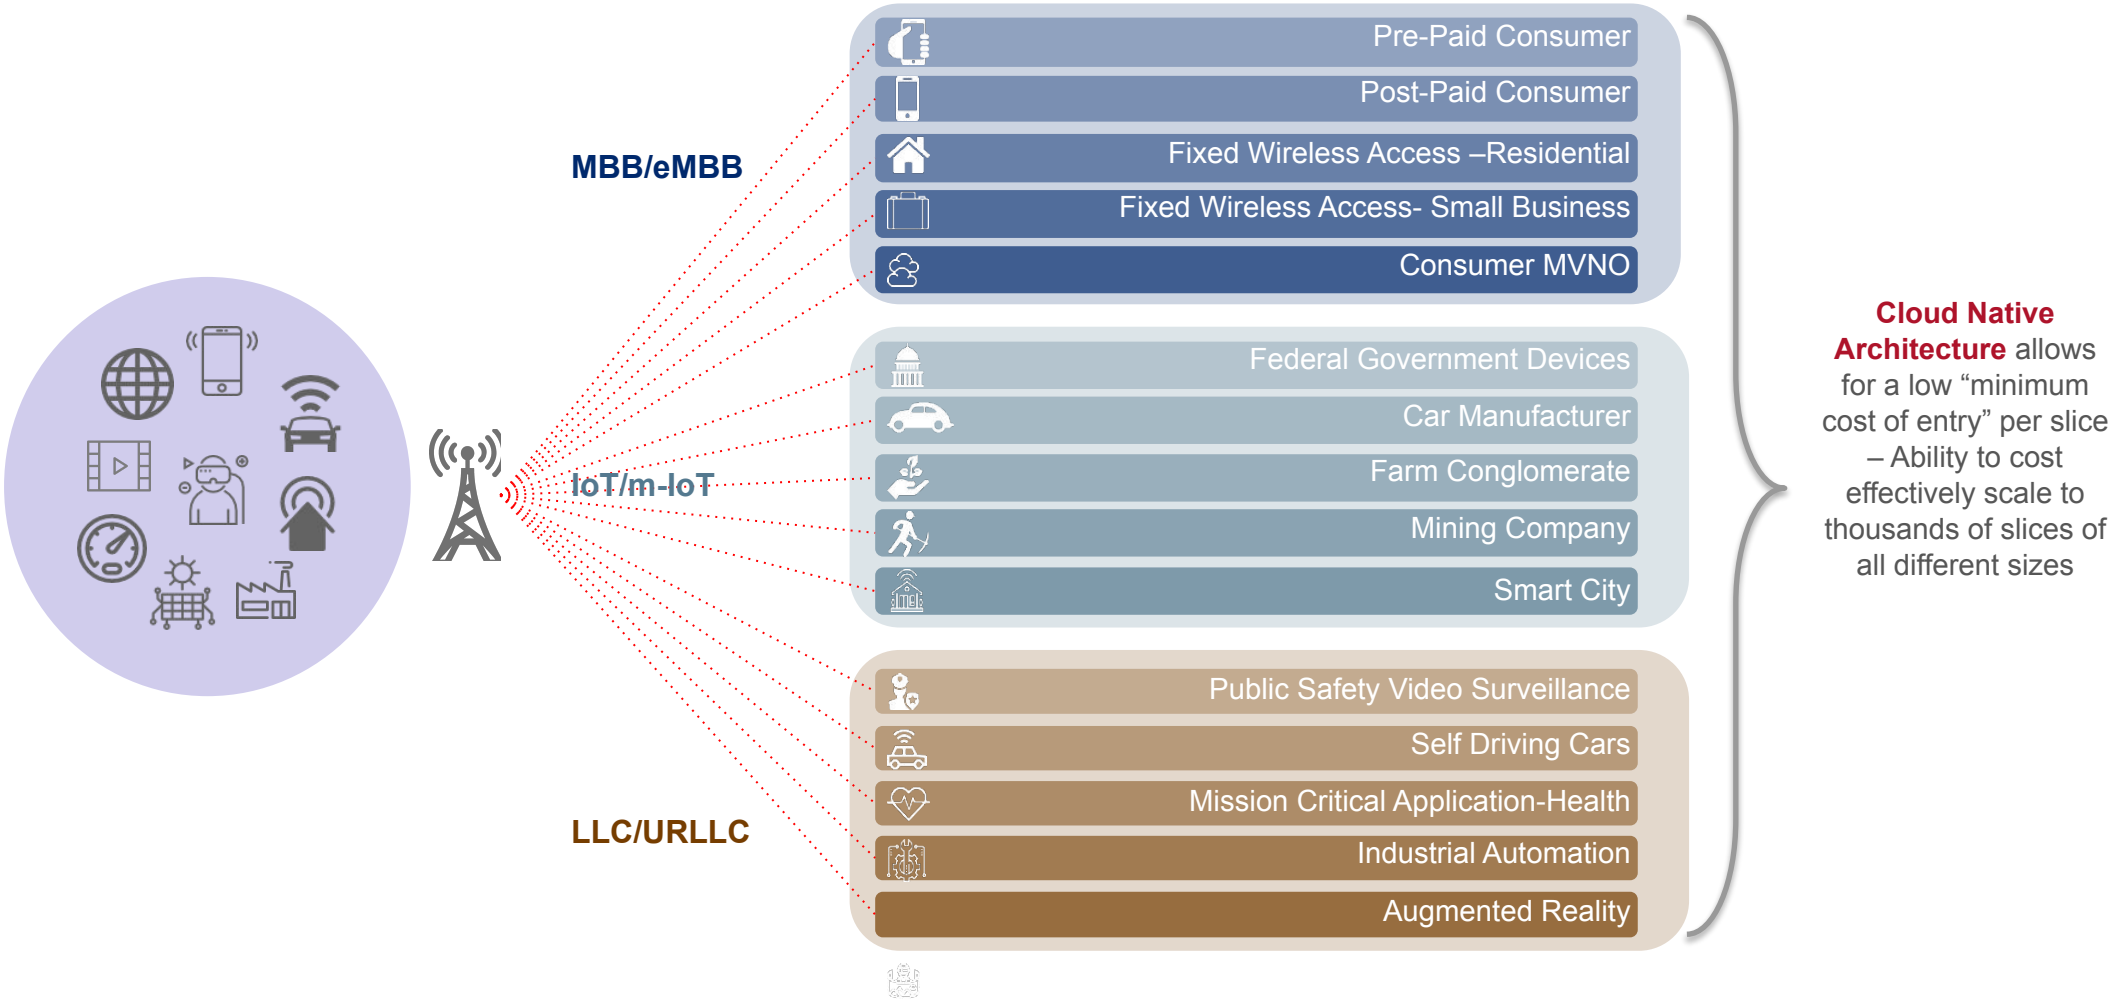

Red Hat

Network Telco Cloud Evolution to Cloud Native

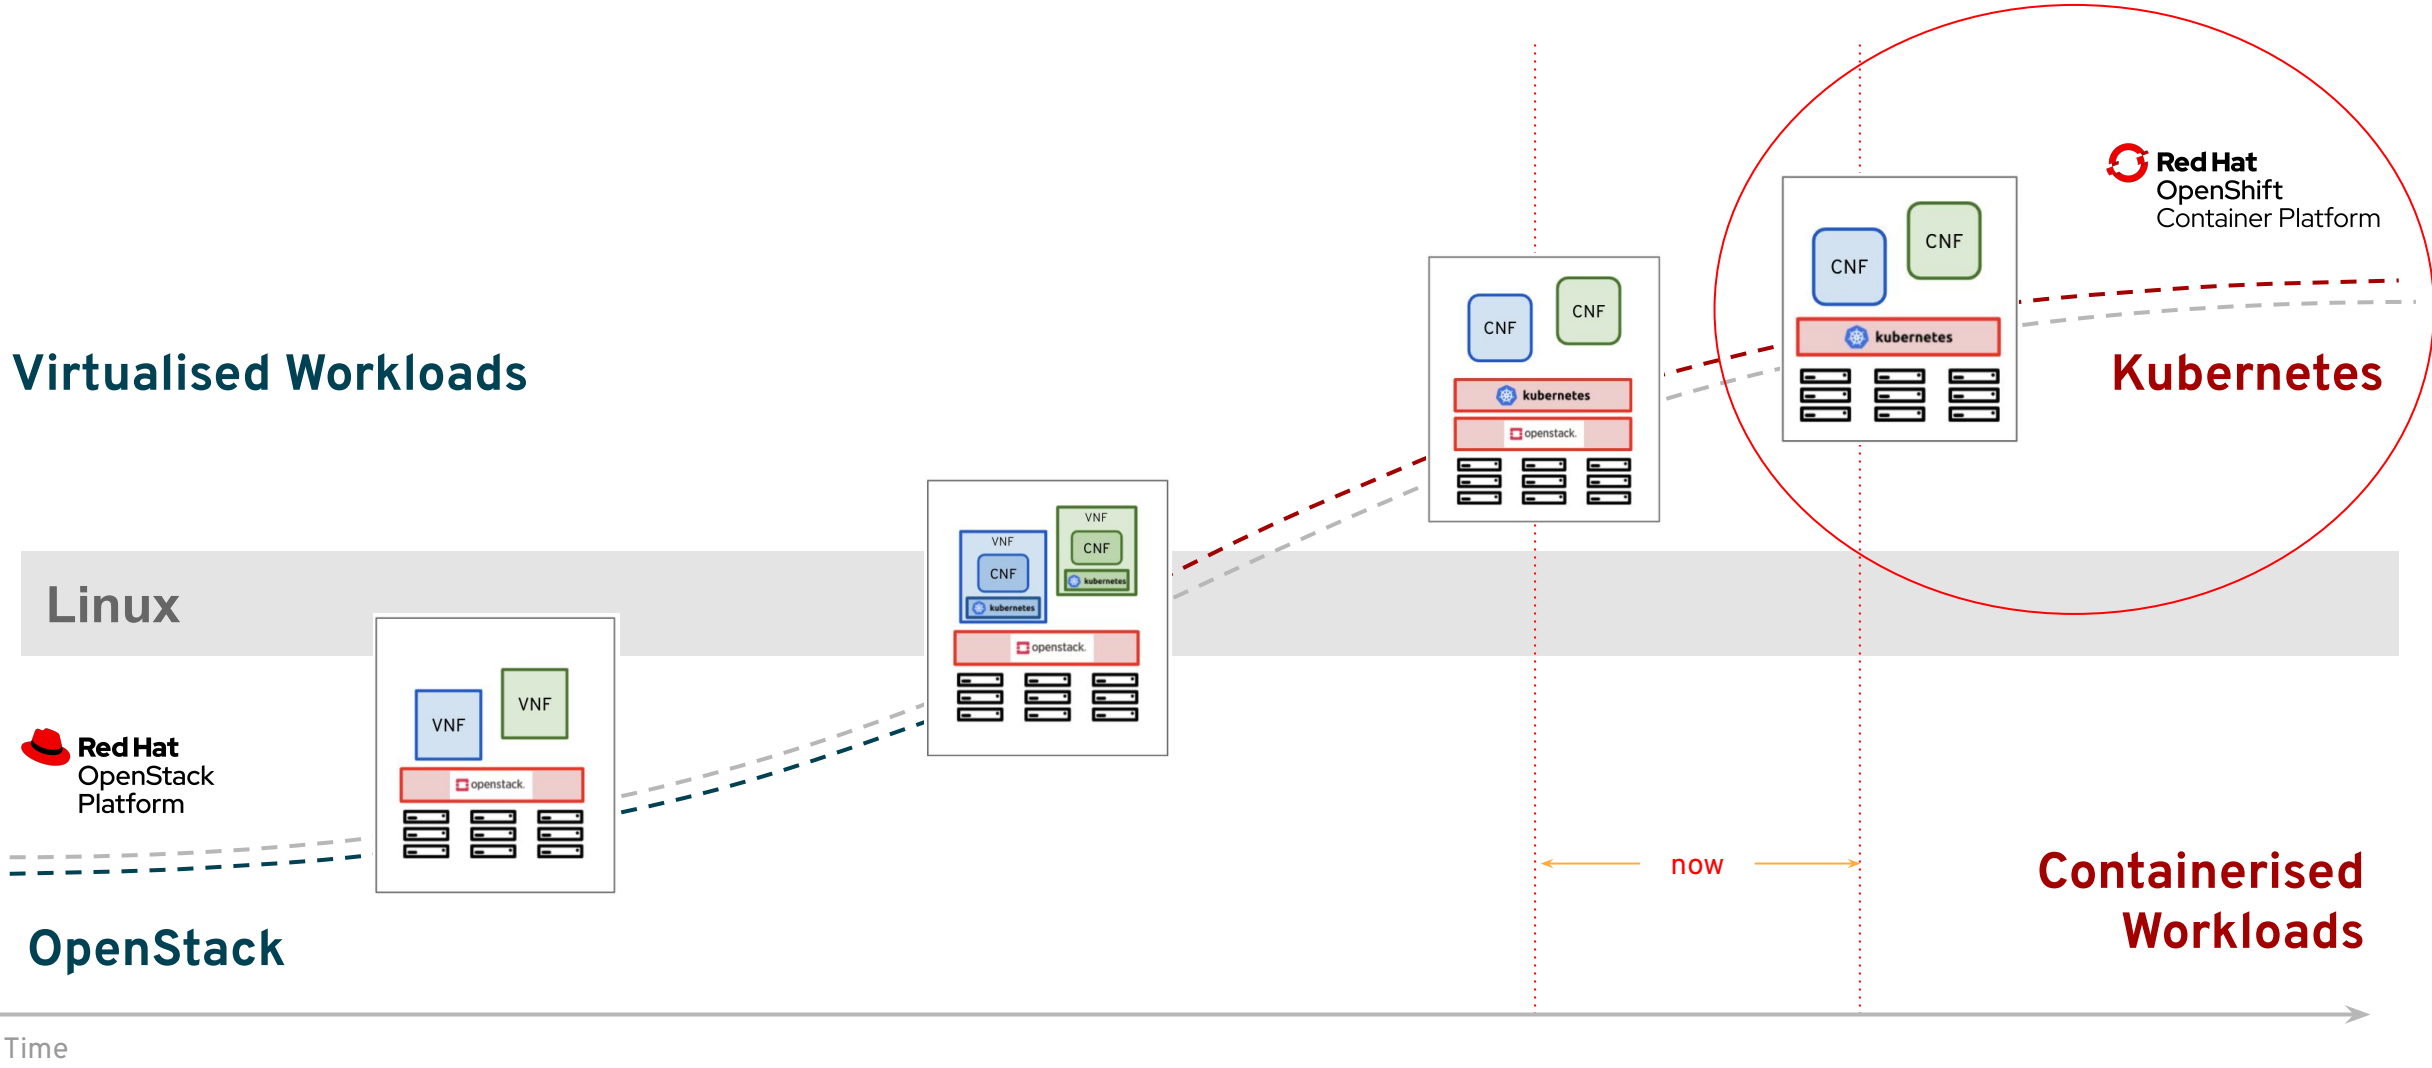

TelecomTV on Red hat continued support for Openstack.

Red Hat

Network Functions Deployment Scenarios

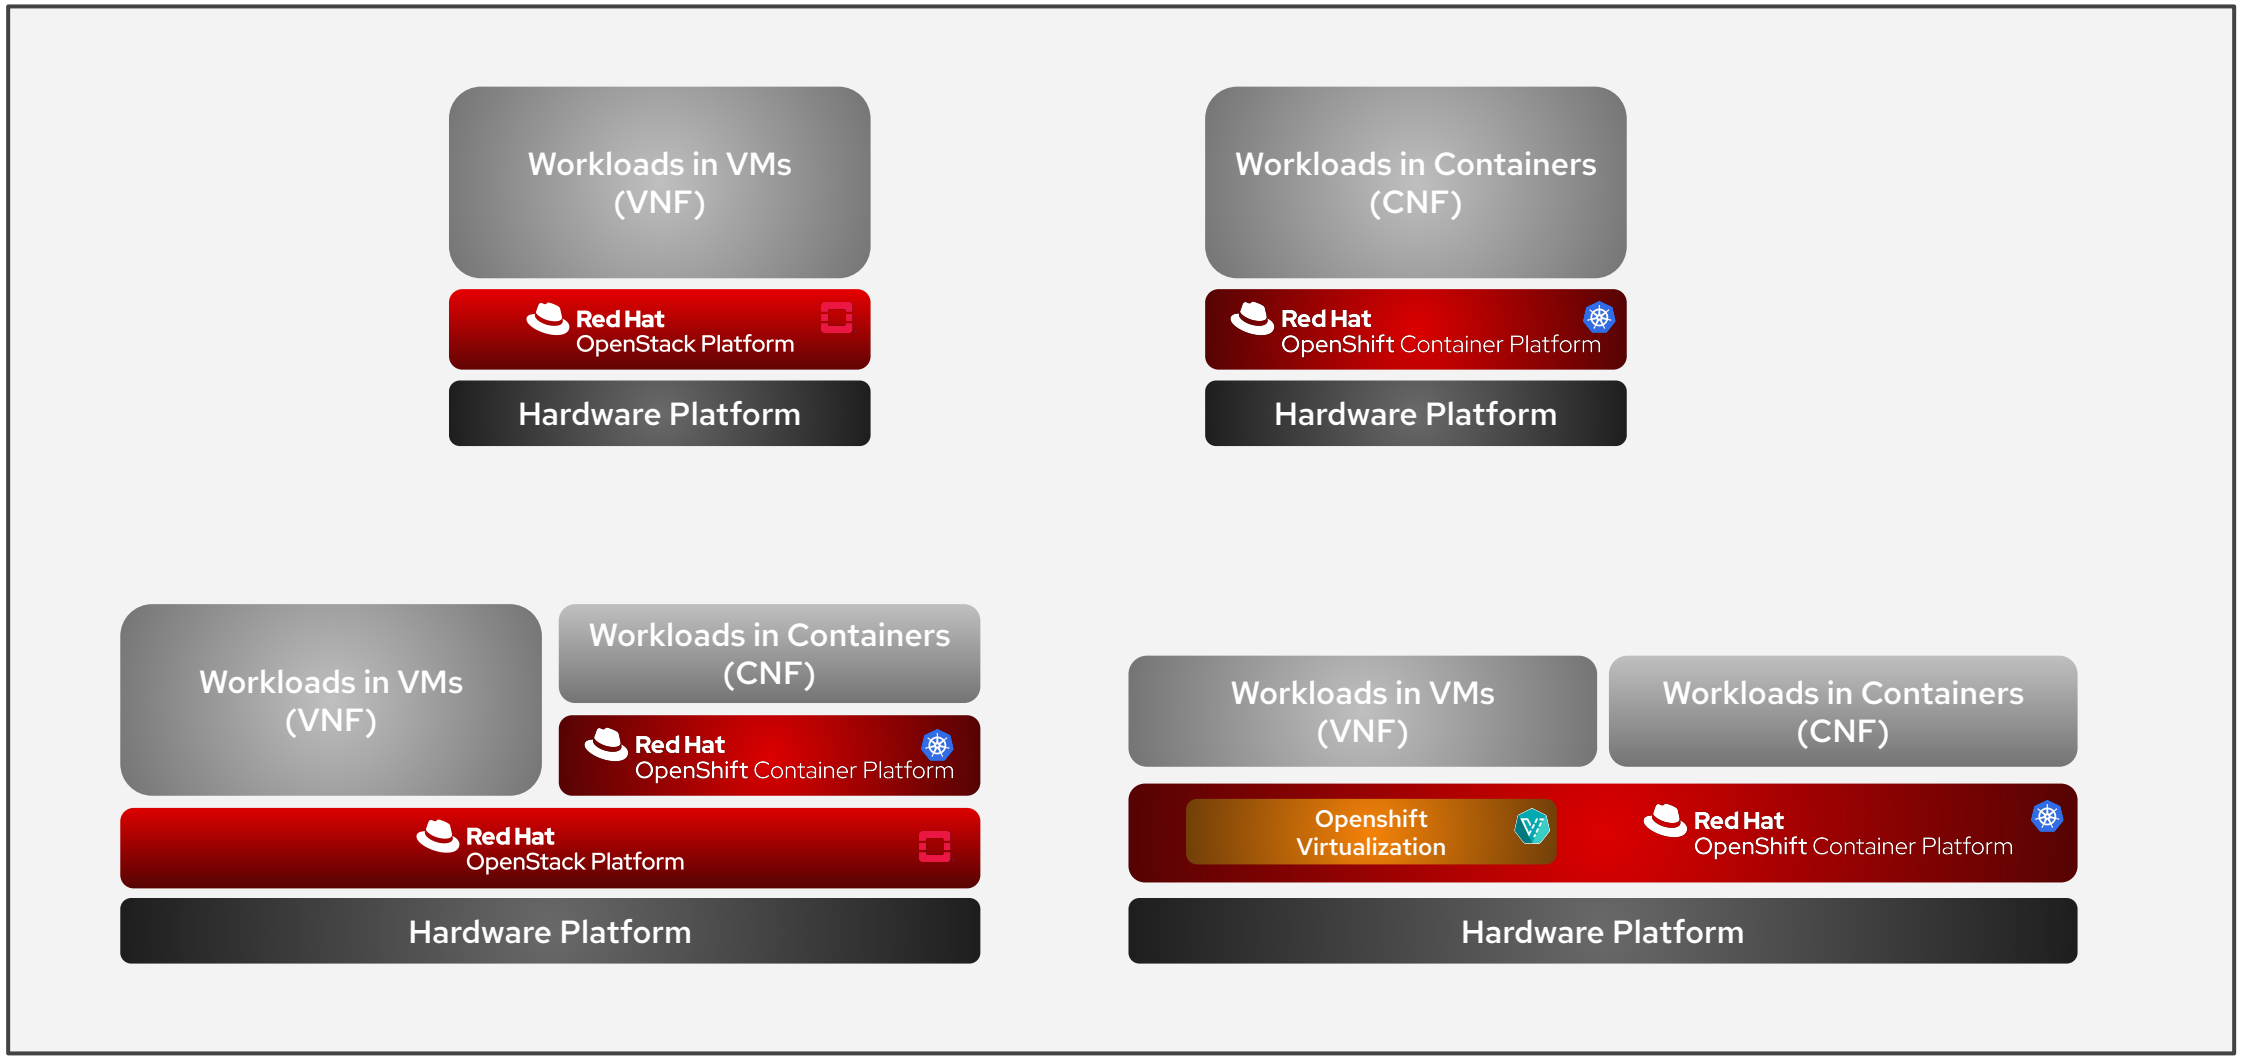

Red Hat

Service Based Architecture (SBA) - Software Based Approach

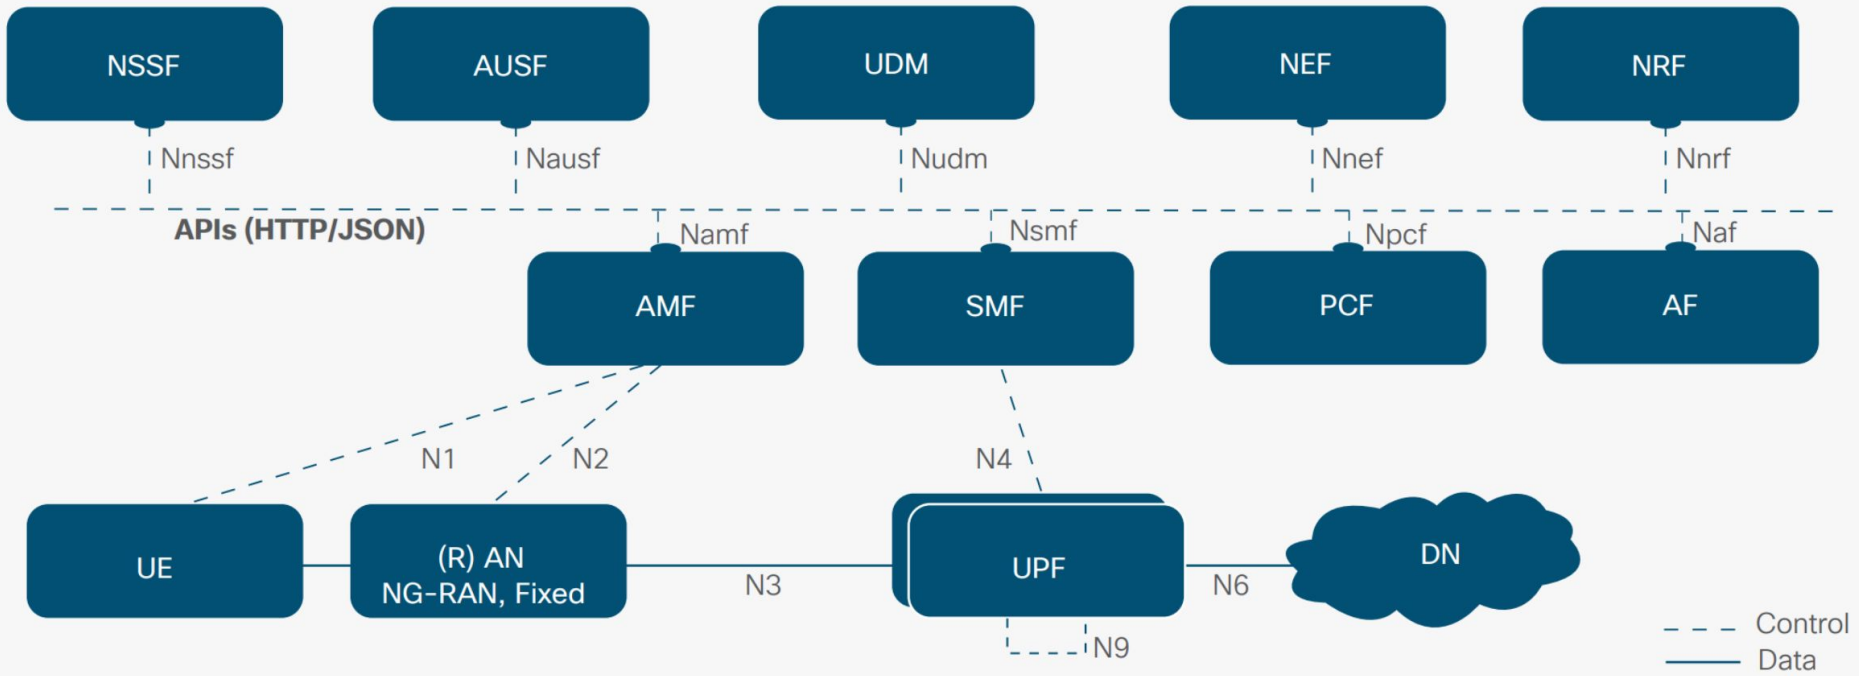

AUSF	Authentication Server Function
AMF	Access and Mobility Management Function
AF	Application Function
DN	Data Networks
SMF	Session Management Function
UDM	Unified Data Management

NEF	Network Exposure Function
NSSF	Network Slice Selection Function
NRF	Network Repository Function
PCF	Policy Control Function
UPF	User Plane Function
(R)AN	(Radio) Access Network

Diagram Source

Red Hat

Service Based Architecture (SBA) - Software Based Approach

(Open Source) Software focused

API Management

Security

Cloud Native / Containers

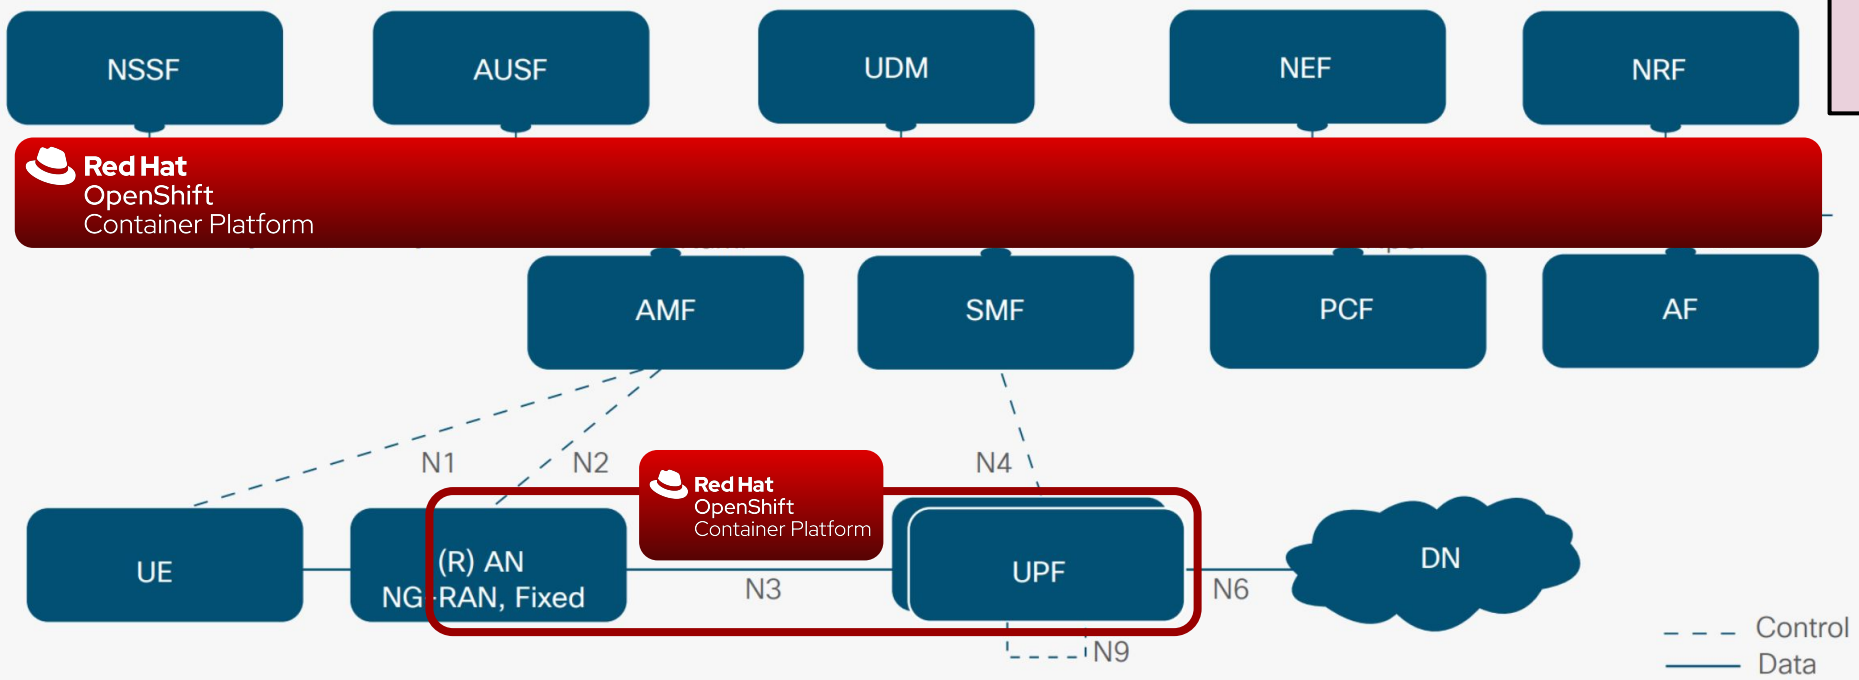

AUSF	Authentication Server Function
AMF	Access and Mobility Management Function
AF	Application Function
DN	Data Networks
SMF	Session Management Function
UDM	Unified Data Management

NEF	Network Exposure Function
NSSF	Network Slice Selection Function
NRF	Network Repository Function
PCF	Policy Control Function
UPF	User Plane Function
(R)AN	(Radio) Access Network

Diagram Source

Red Hat

Service Based Architecture (SBA) - Software Based Approach

(Open Source) Software focused

API Management

Security

Cloud Native / Containers

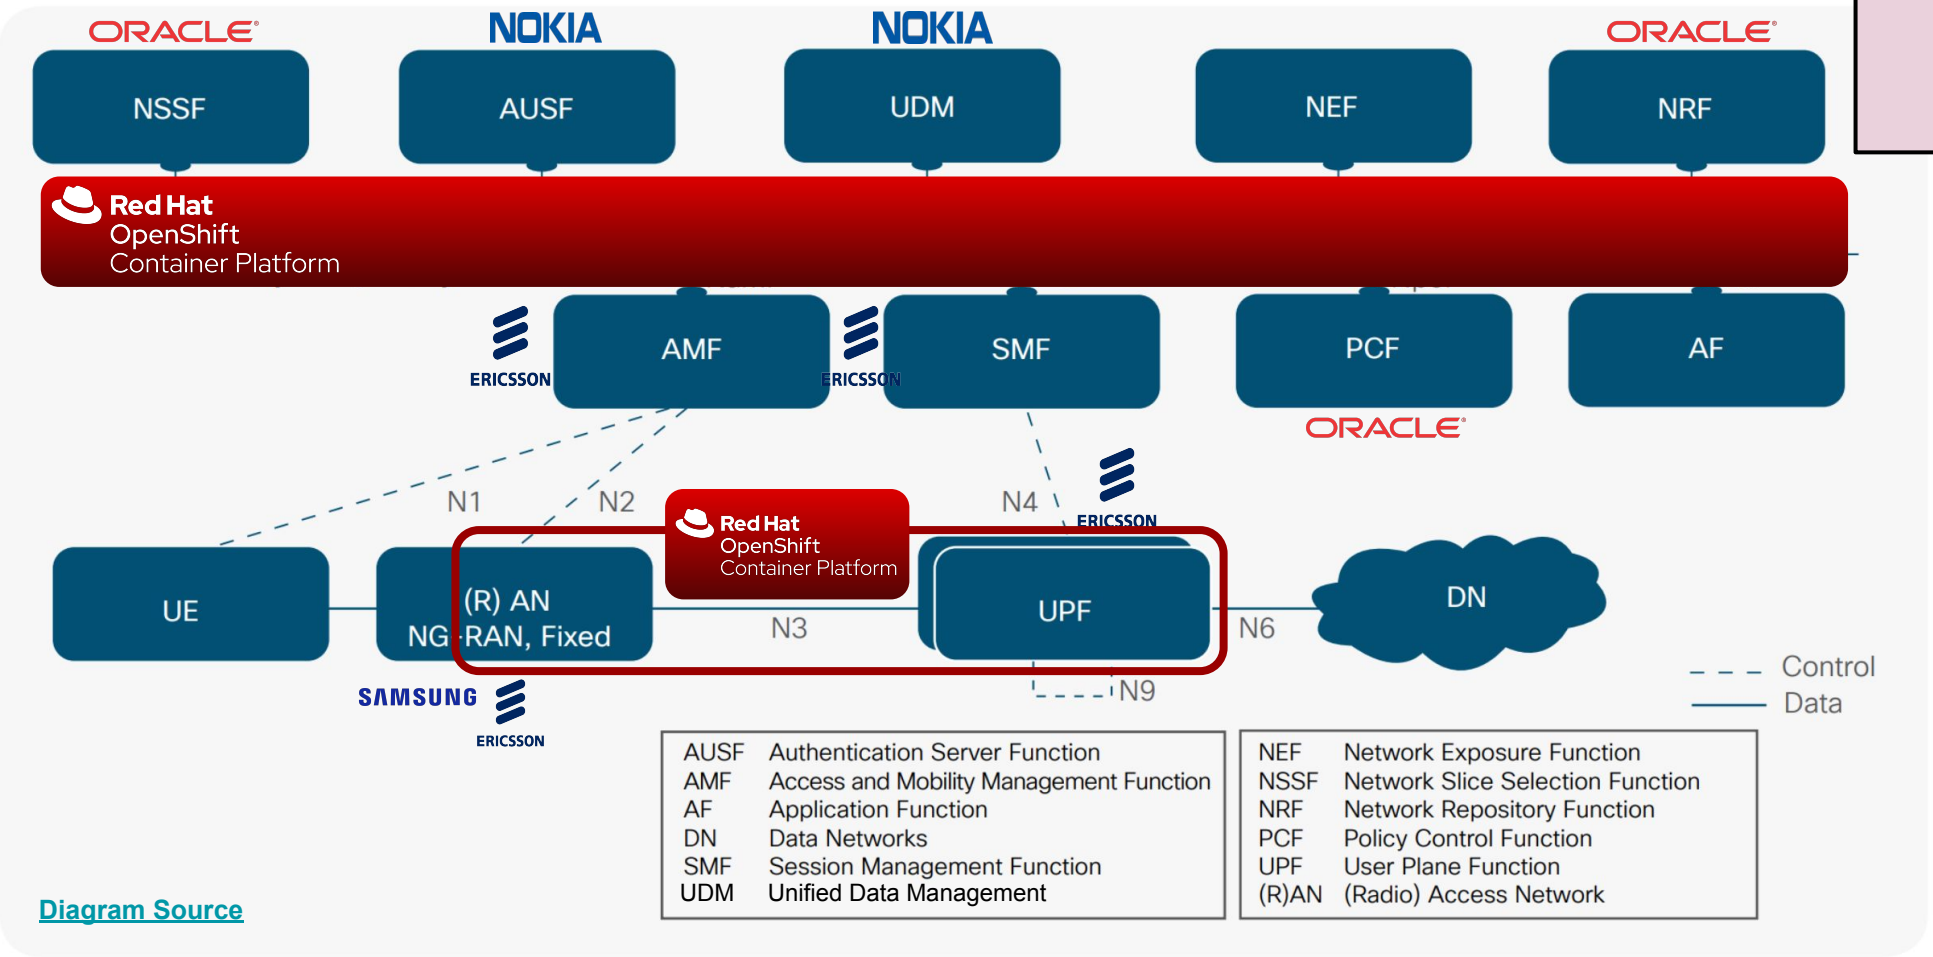

Diagram Source

Service Based Architecture (SBA) - Software Based Approach

(Open Source) Software focused

API Management

Security

Cloud Native / Containers

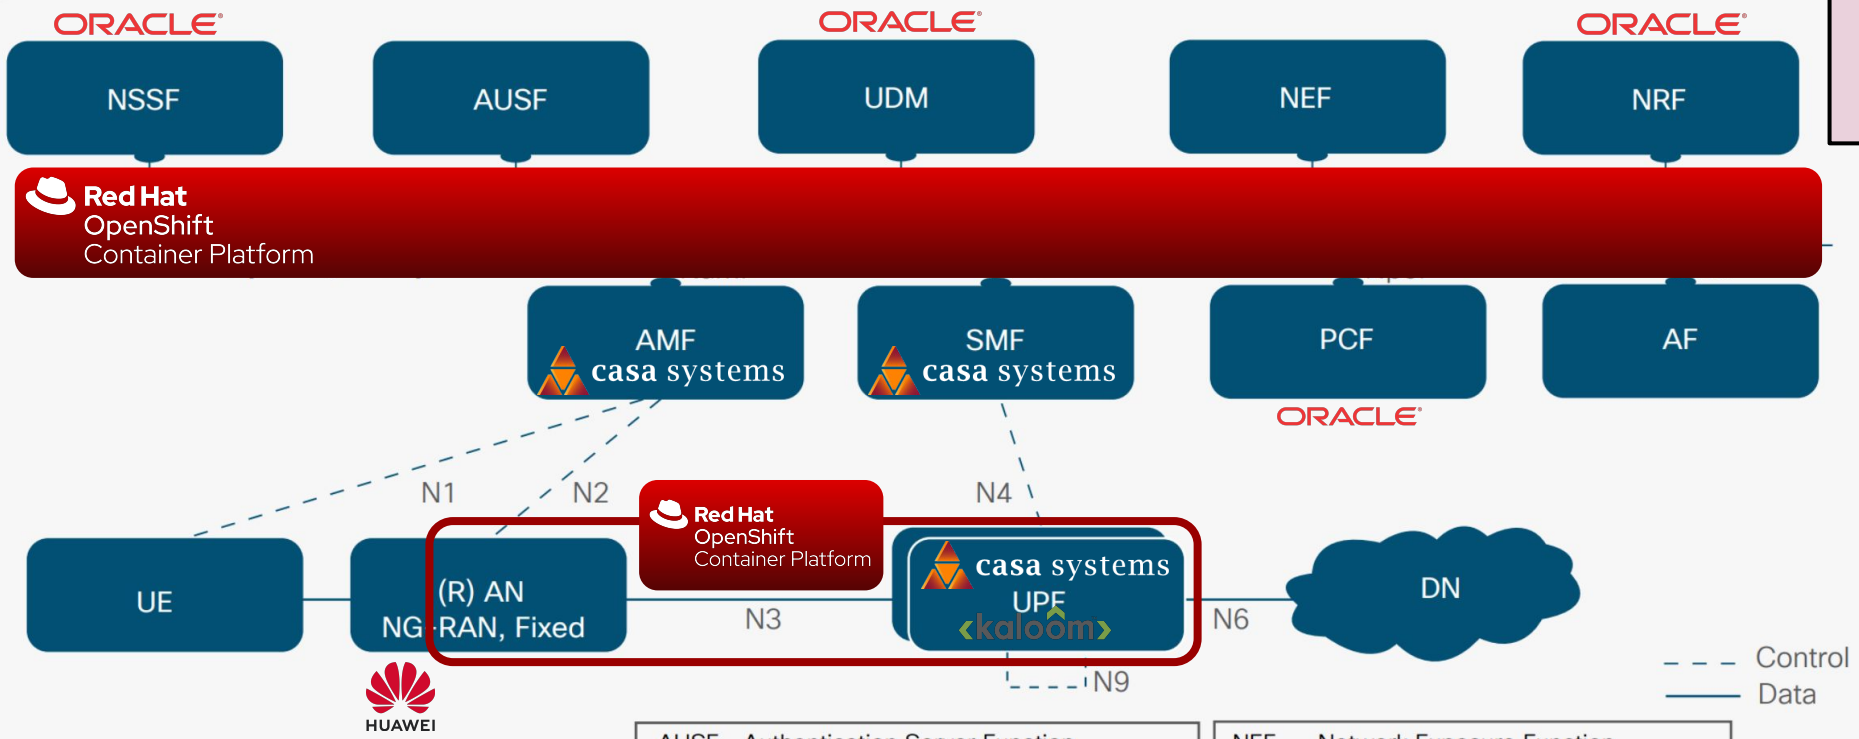

AUSF	Authentication Server Function
AMF	Access and Mobility Management Function
AF	Application Function
DN	Data Networks
SMF	Session Management Function
UDM	Unified Data Management

NEF	Network Exposure Function
NSSF	Network Slice Selection Function
NRF	Network Repository Function
PCF	Policy Control Function
UPF	User Plane Function
(R)AN	(Radio) Access Network

Diagram Source

Service Based Architecture (SBA) - Software Based Approach

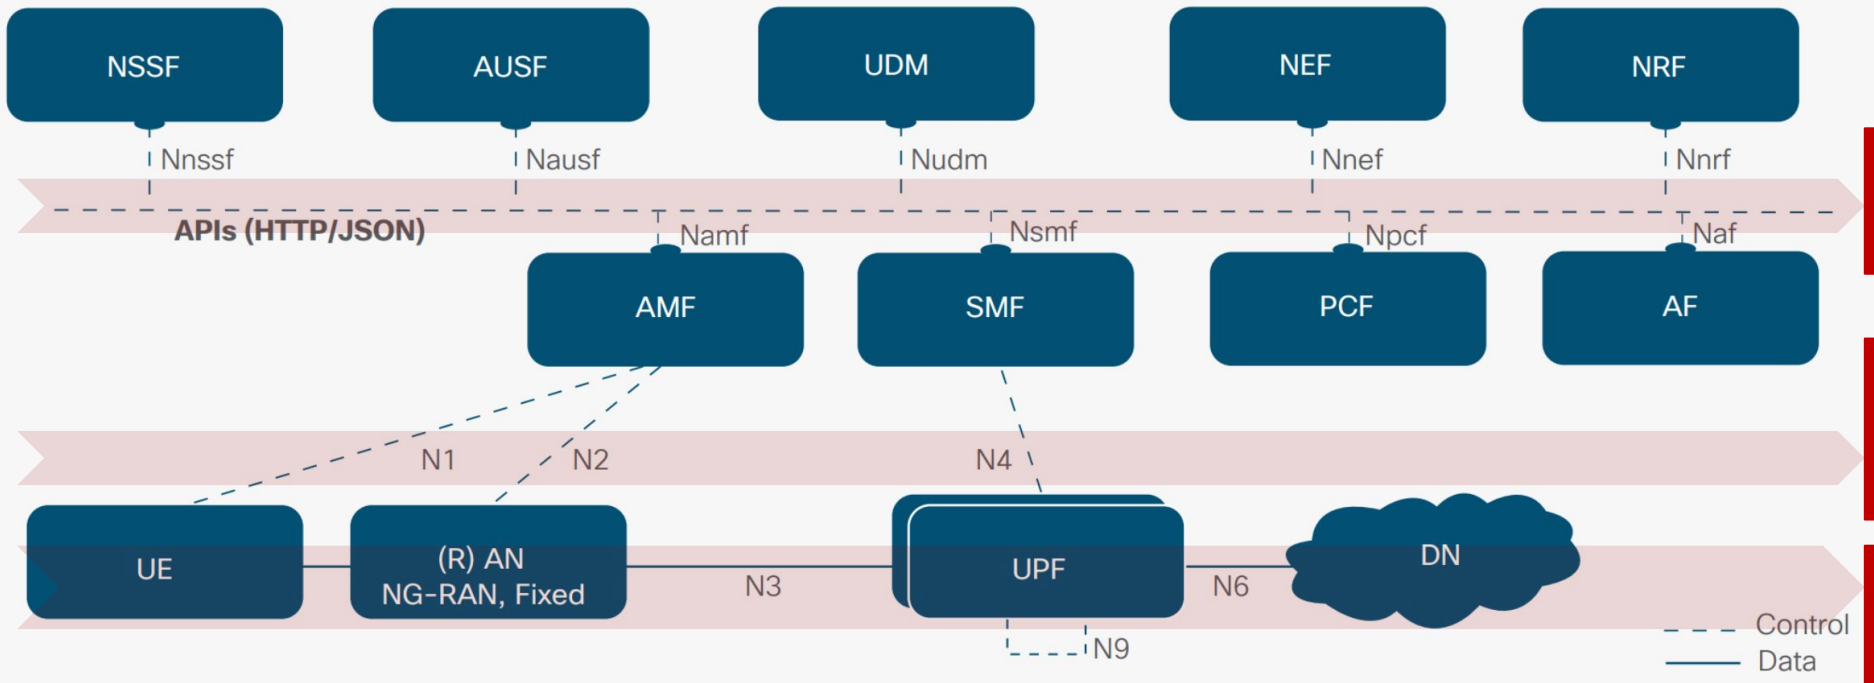

AUSF	Authentication Server Function
AMF	Access and Mobility Management Function
AF	Application Function
DN	Data Networks
SMF	Session Management Function
UDM	Unified Data Management

NEF	Network Exposure Function
NSSF	Network Slice Selection Function
NRF	Network Repository Function
PCF	Policy Control Function
UPF	User Plane Function
(R)AN	(Radio) Access Network

Kubernetes networking (REST): OVN-K, metalLB

Control-Plane networking OVN-K, MAC-VLAN, kernel SR-IOV

Dataplane networking: SR-IOV with DPDK (GTP, SCTP, ...)

Diagram Source

Red Hat

5G (4G) RAN

Red Hat

Mobile Network Radio Base Station

RAN Evolution: Traditional Model (2/3/4G)

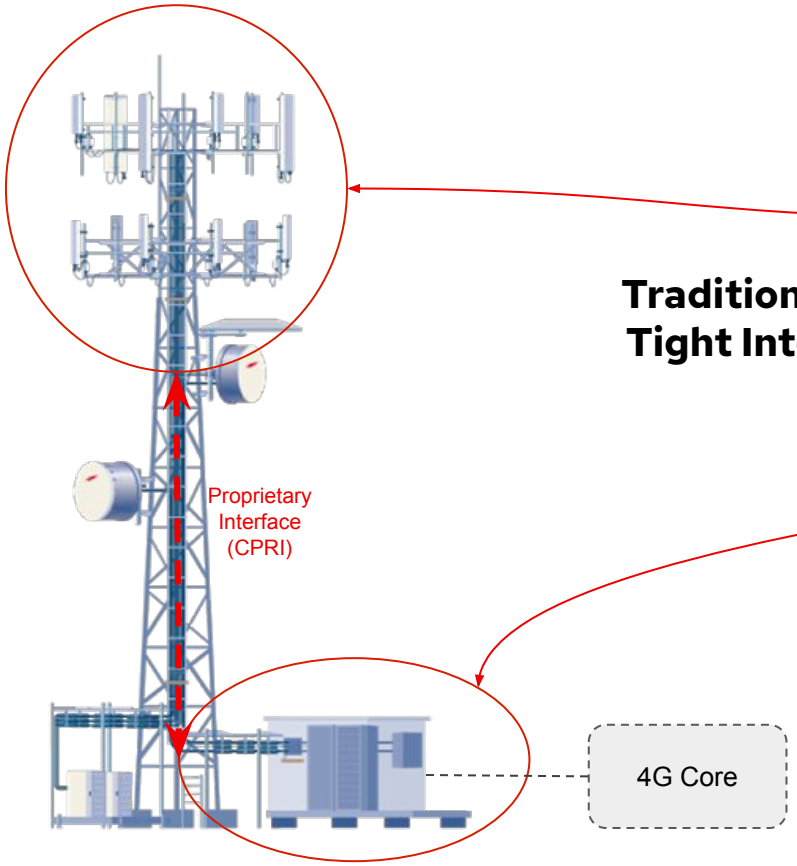

Traditional model
Tight Integration

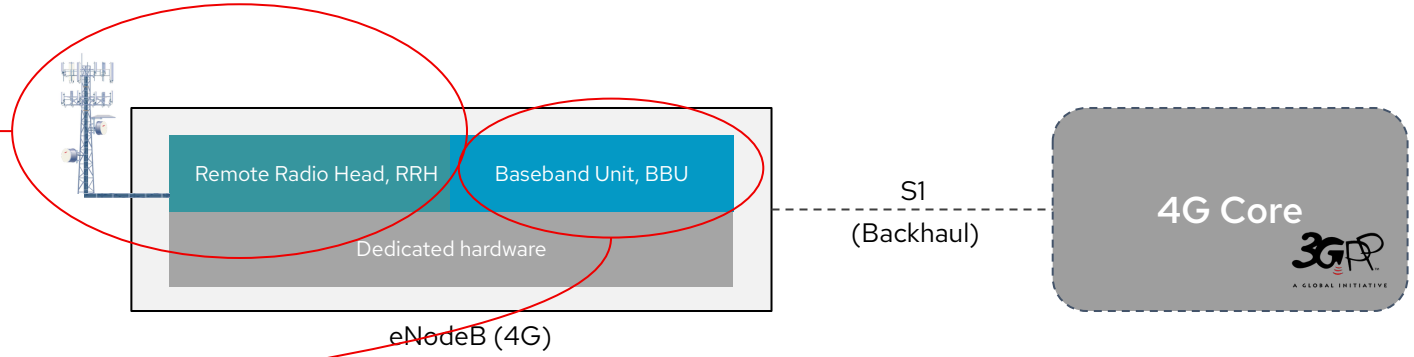

- Tight integration by single RAN vendor @ Cell Site
- Proprietary hardware and software (including platform software)
- Proprietary interface between Radios and Baseband
- Standard interface (Backhaul) towards Core Network

RRH = Remote radio head
BBU = Baseband unit
RU = Radio unit
DU = Distributed unit
CU = Centralized unit
UP = User plane
CP = Control plane
UPF = User Plane Function
S1 = 4G RAN and EPC interface
NG = Next-generation logical interface

Red Hat

RAN Evolution: Two Dimensions of Architecture Transformation

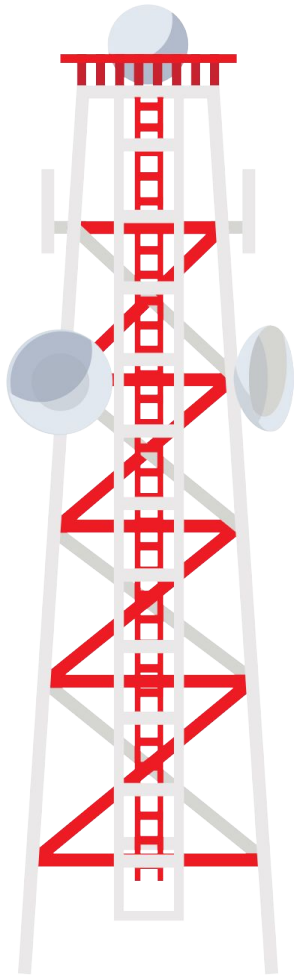

RAN Functional Split (disaggregation of functionality)

RAN Functional Split: disaggregating baseband functionality from radio functionality

RAN Cloudification (disaggregation of HW & SW)

RAN Cloudification: Deploy baseband functionality on consistent cloud platform

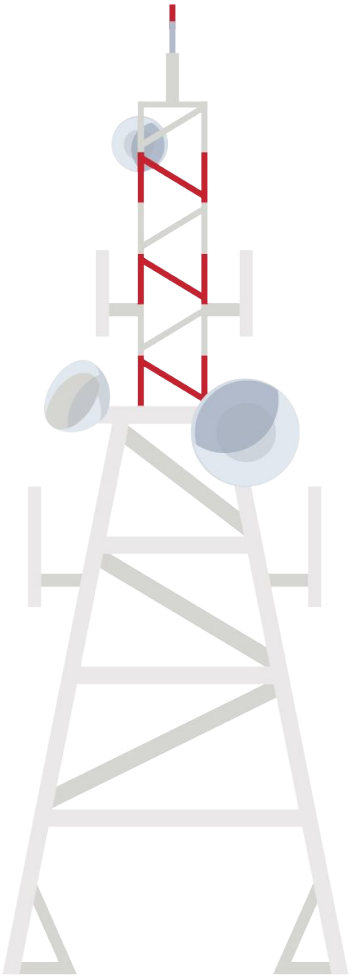

Red Hat

RAN Evolution: Industry Organizations

Most Visible Organisations Driving this Evolution & Transformation

3GPP™

A GLOBAL INITIATIVE

Standardization authority in Mobile Network space

O-RAN Alliance takes RAN related 3GPP specifications and complements those

O-RAN ALLIANCE

De Facto standards initiative within RAN space

TELECOM INFRA PROJECT

TIP runs practical workgroups and Community Labs to test Open RAN related work products and builds The Industry Ecosystem

Red Hat

RAN Evolution: Key Terminology Used

Open RAN, OpenRAN, O-RAN, vRAN, Cloud RAN , …

1. **Open RAN** is a generic term that refers to industry movement and open RAN architectures including open interfaces, virtualization / containerization and use of AI/ML, etc.
2. **OpenRAN** is a project initiated by the Telecom Infra Project (TIP). It's an attempt to realize the Open RAN concept on its own part. Its work covers 2G/3G/4G/5G. As inputs, OpenRAN uses 3GPP and O-RAN Alliance specifications.
3. **O-RAN** (ORAN) refers to the O-RAN Alliance or standards created by the O-RAN Alliance, which complements 3GPP specifications by defining interface profiles, new open interfaces and new nodes.
4. **vRAN (Virtualized RAN)**: Whereas Open RAN focuses on openness, vRAN is really about decoupling software from hardware.
5. **Cloud RAN (C-RAN)** is vRAN built on cloud native technologies, such as microservices, containers and CI/CD. Confusingly, C-RAN is also sometimes used to mean Centralized RAN where baseband processing is centralized and relocated out from antenna site to edge data center, but in most cases that is written as CRAN.

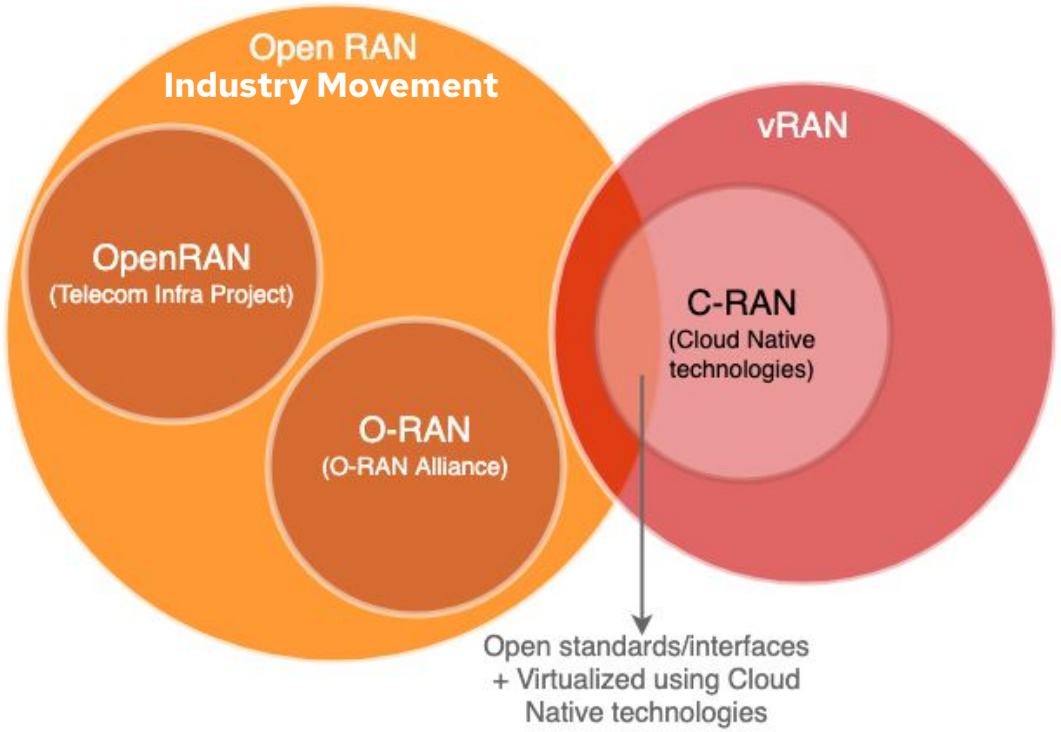

Source:Devopedia 2021.

Open RAN does not have to be virtual
Virtual RAN does not have to be open

Red Hat

Mobile Network Radio Base Station

RAN Evolution: Open RAN Model with Containerized RAN Workloads on Container Platform

Cloud RAN / vRAN

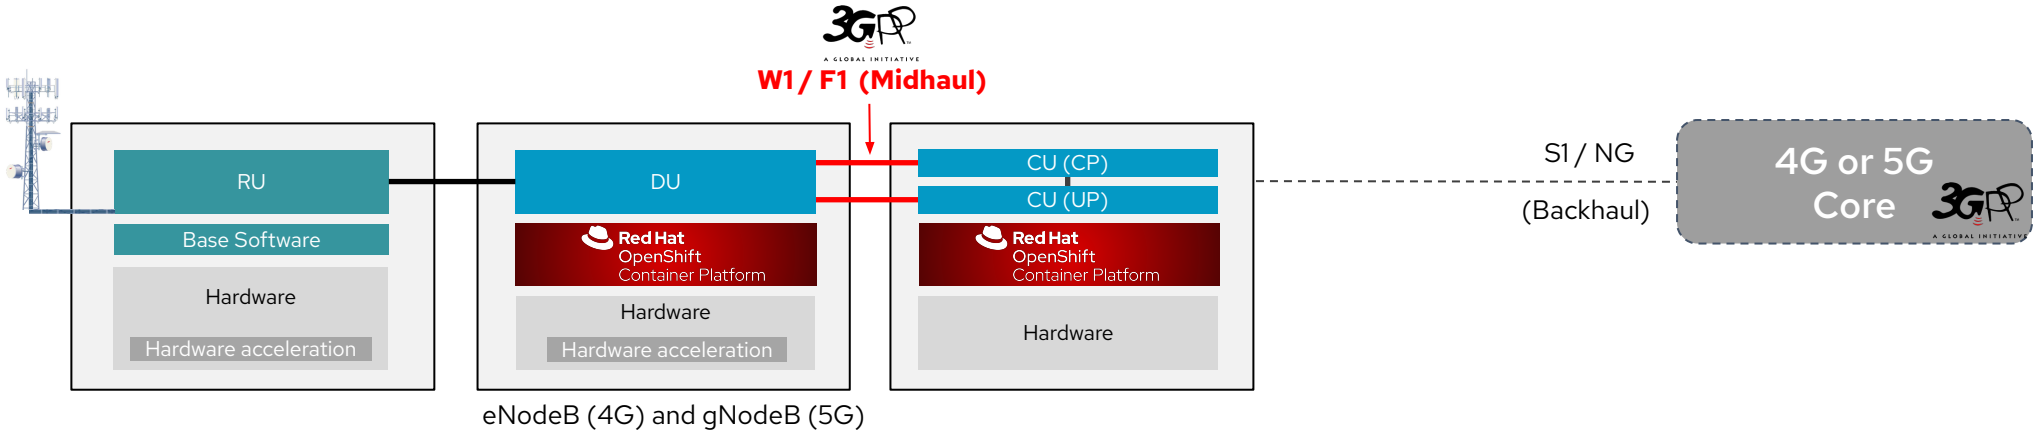

- Functional Split / Disaggregation per 3GPP
- Midhaul interfaces: W1 (Rel 16 for 4G) & F1 (Rel 15 for 5G)
- Standard interfaces (Backhaul) towards Core Network(s)
- Three entity model: Radio Unit (RU), Distributed Unit (DU), Centralized Unit (CU, Control and User Planes)
- Cloud Platform to host DU and CU workloads
- Single RAN vendor model

RRH = Remote radio head
BBU = Baseband unit
RU = Radio unit
DU = Distributed unit
CU = Centralized unit
UP = User plane
CP = Control plane
UPF = User Plane Function
S1 = 4G RAN and EPC interface
NG = Next-generation logical interface

Red Hat

Mobile Network Radio Base Station

RAN Evolution: Open RAN Model with Containerized RAN Workloads on Container Platform

Cloud RAN / vRAN

Open Workload Interfaces

(partially compliant to O-RAN Alliance)

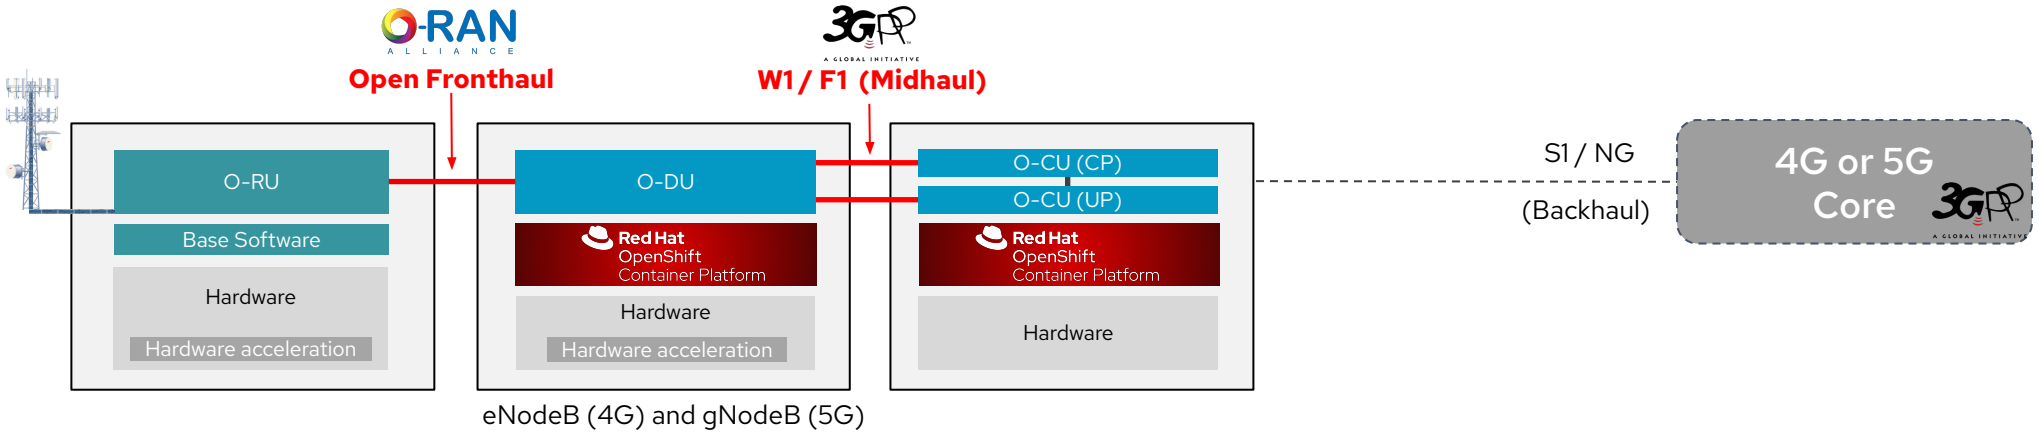

- Added: Open Fronthaul by O-RAN Alliance
- O-RAN Alliance nomenclature: O-RU, O-DU, O-CU
- Potential for multi RAN vendor model

RRH = Remote radio head
BBU = Baseband unit
RU = Radio unit
DU = Distributed unit
CU = Centralized unit
UP = User plane
CP = Control plane
UPF = User Plane Function
S1 = 4G RAN and EPC interface
NG = Next-generation logical interface

Red Hat

Mobile Network Radio Base Station

RAN Evolution: Open RAN Model aligned with O-RAN Alliance

O-RAN Alliance Compliant Model

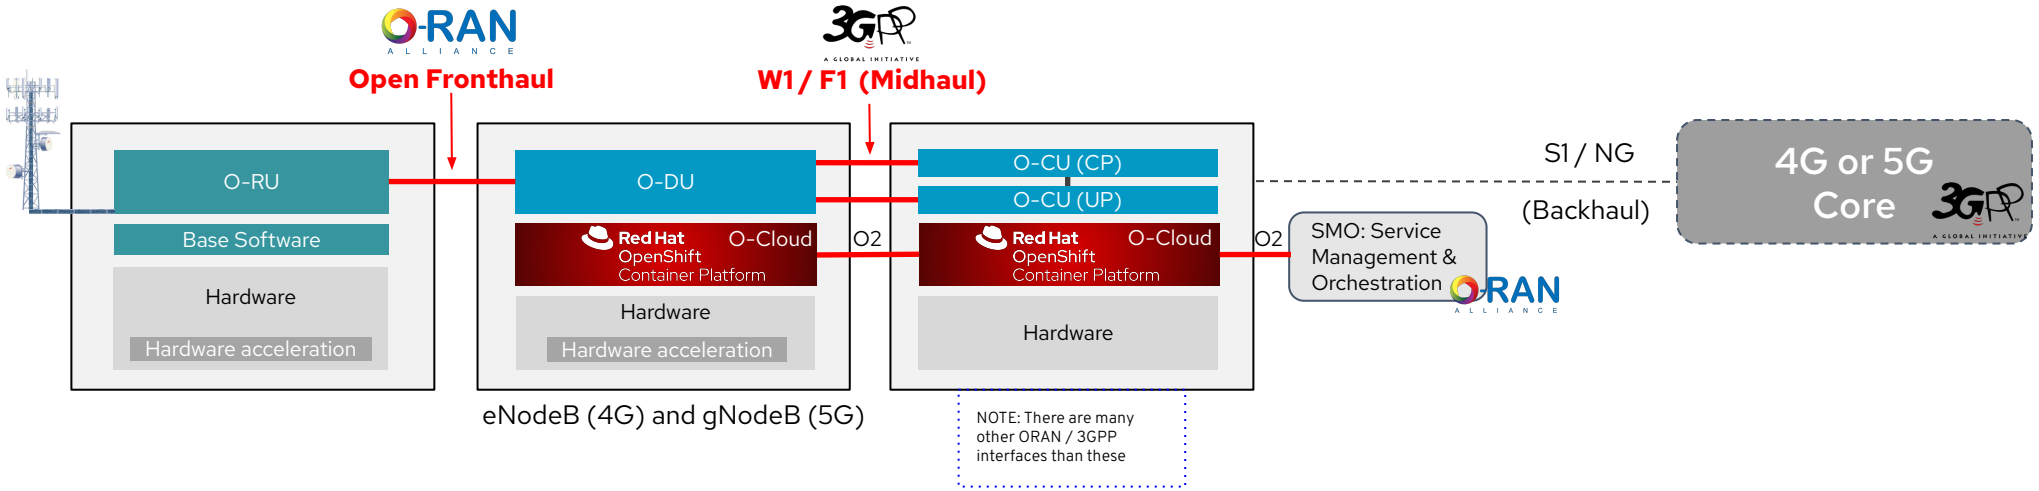

- O-RAN Alliance nomenclature for Cloud Platform: O-Cloud
- Goal: multi vendor

RRH = Remote radio head
BBU = Baseband unit
RU = Radio unit
DU = Distributed unit
CU = Centralized unit
UP = User plane
CP = Control plane
UPF = User Plane Function
S1 = 4G RAN and EPC interface
NG = Next-generation logical interface

Red Hat

High Level RAN Stack by O-RAN Alliance

O-RAN Interfaces:
- O1, O2
- A1, E2
- OFH

3GPP Interfaces:
- F1
- NG
- S1

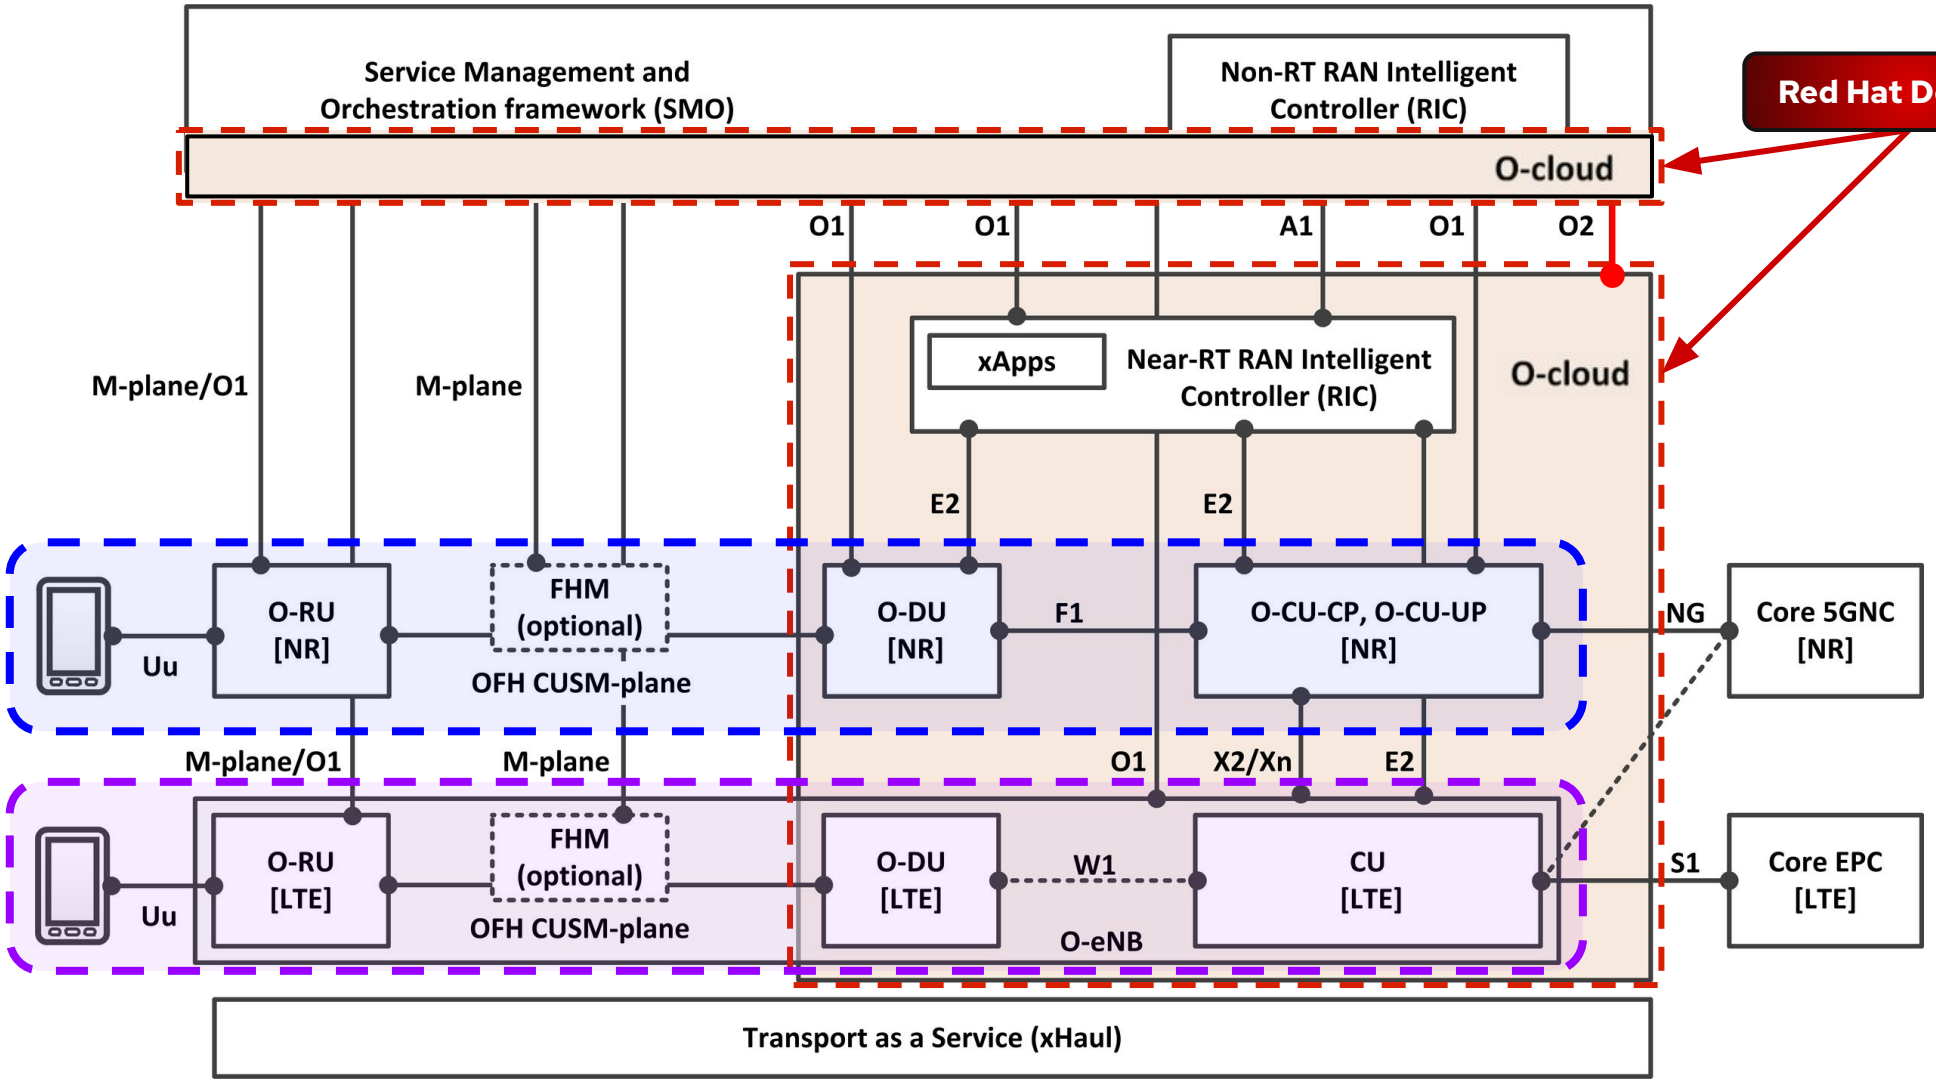

RRH = Remote radio head
BBU = Baseband unit
RU = Radio unit
DU = Distributed unit
CU = Centralized unit
UP = User plane
CP = Control plane
UPF = User Plane Function
S1 = 4G RAN and EPC interface
NG = Next-generation logical interface

Red Hat

RAN / Edge Evolution Adds New Requirements to Cloud Platforms

These are three of the most important new areas to cover …

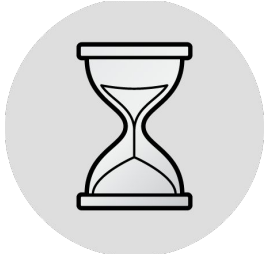

Realtime Kernel (RT Execution Environment)

Workloads stringent low-latency determinism requirements for core kernel features such as interrupt handling and process scheduling in the microsecond (μs) range.

Timing & Synchronization

Time synchronization via transport networks will be critical for 5G radios. Precision Time Protocol (PTP) remains the preferred method to deliver timing across packet-switched networks

Red Hat Timing & Sync work presented at:

- International Timing & Synchronization Forum (Nov 5, 2020), Virtual
- OpenAirInterface Fall 2020 Virtual Workshop (Nov 12, 2020), Virtual
  Video: OpenAirInterface Event
- Workshop on Synchronization and Timing Systems (Apr 1, 2021), Virtual
  Slides: WSTS 2021
- International Timing & Synchronization Forum (Nov 3, 2021), Brighton, UK
- Upcoming: IWPC 5G Network Architecture Evolution hosted by AT&T, (Oct 12, 2022), Austin, TX
- Upcoming: International Timing & Synchronization Forum (Nov 8, 2022), Düsseldorf, Germany

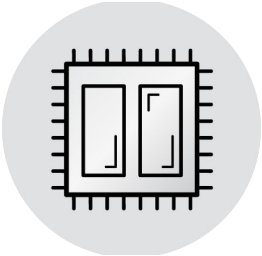

Hardware Acceleration

Field Programmable Gate Arrays (FPGA) , SmartNIC, and other hardware acceleration components will be vital for 5G virtualized infrastructure.

Red Hat

RAN / Edge Evolution Adds New Requirements to Cloud Platforms

CPU Management

CPU Manager manages groups of CPUs and constrains workloads to specific CPUs. CPU Manager is useful for workloads that have some of these attributes: require as much CPU time as possible or are low-latency network applications.

Zero touch provisioning

Provides all the tools required to install, upgrade and maintain the cloud infrastructure for the RAN workload with minimum user interaction in an “appliance” like deployment. Reduced complexity with increased flexibility of options and performance.

Topology Management

Topology Manager collects hints from the CPU Manager, Device Manager, and other Hint Providers to align pod resources, such as CPU, SR-IOV VFs, and other device resources, for all Quality of Service (QoS) classes on the same non-uniform memory access (NUMA) node.

Remote Management

Take full control of edge and RAN operation from a centralized single pane of glass. This includes, OpenShift installation and upgrade, application provisioning and monitoring.

Low Latency

A combination of multiple factors that allow the workload the maximum processing capacity and minimizes packet delivery latencies.

Reduce footprint

Remote Radio sites have limited space and power, therefore edge and RAN clouds would require a small cloud footprint.

Red Hat

Deployment Models for Disaggregated RAN

And Consistent Cloud Platform …

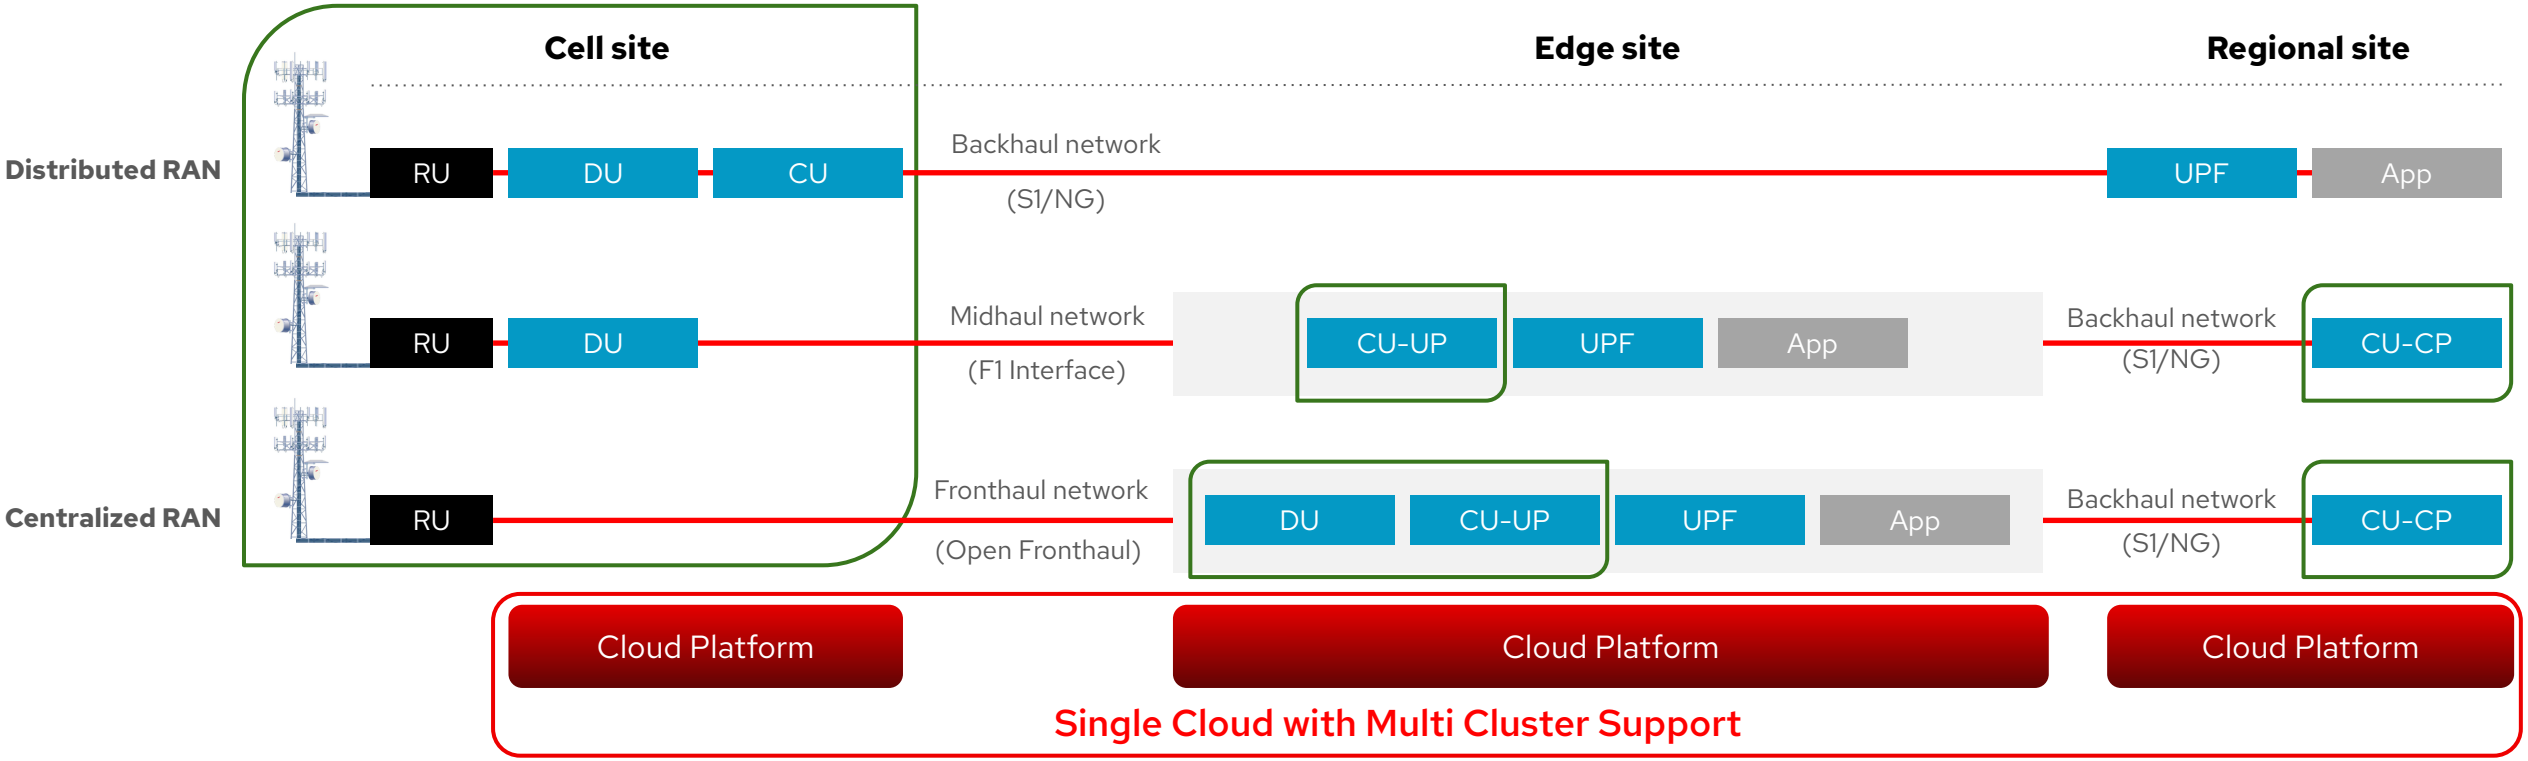

RRH = Remote radio head
BBU = Baseband unit
RU = Radio unit
DU = Distributed unit
CU = Centralized unit
UP = User plane
CP = Control plane
UPF = User Plane Function
S1 = 4G RAN and EPC interface
NG = Next-generation logical interface

Red Hat

Mobile Network Element UP NF Placement Locations;

5G Service Specific Slice Examples, Major Latency Components

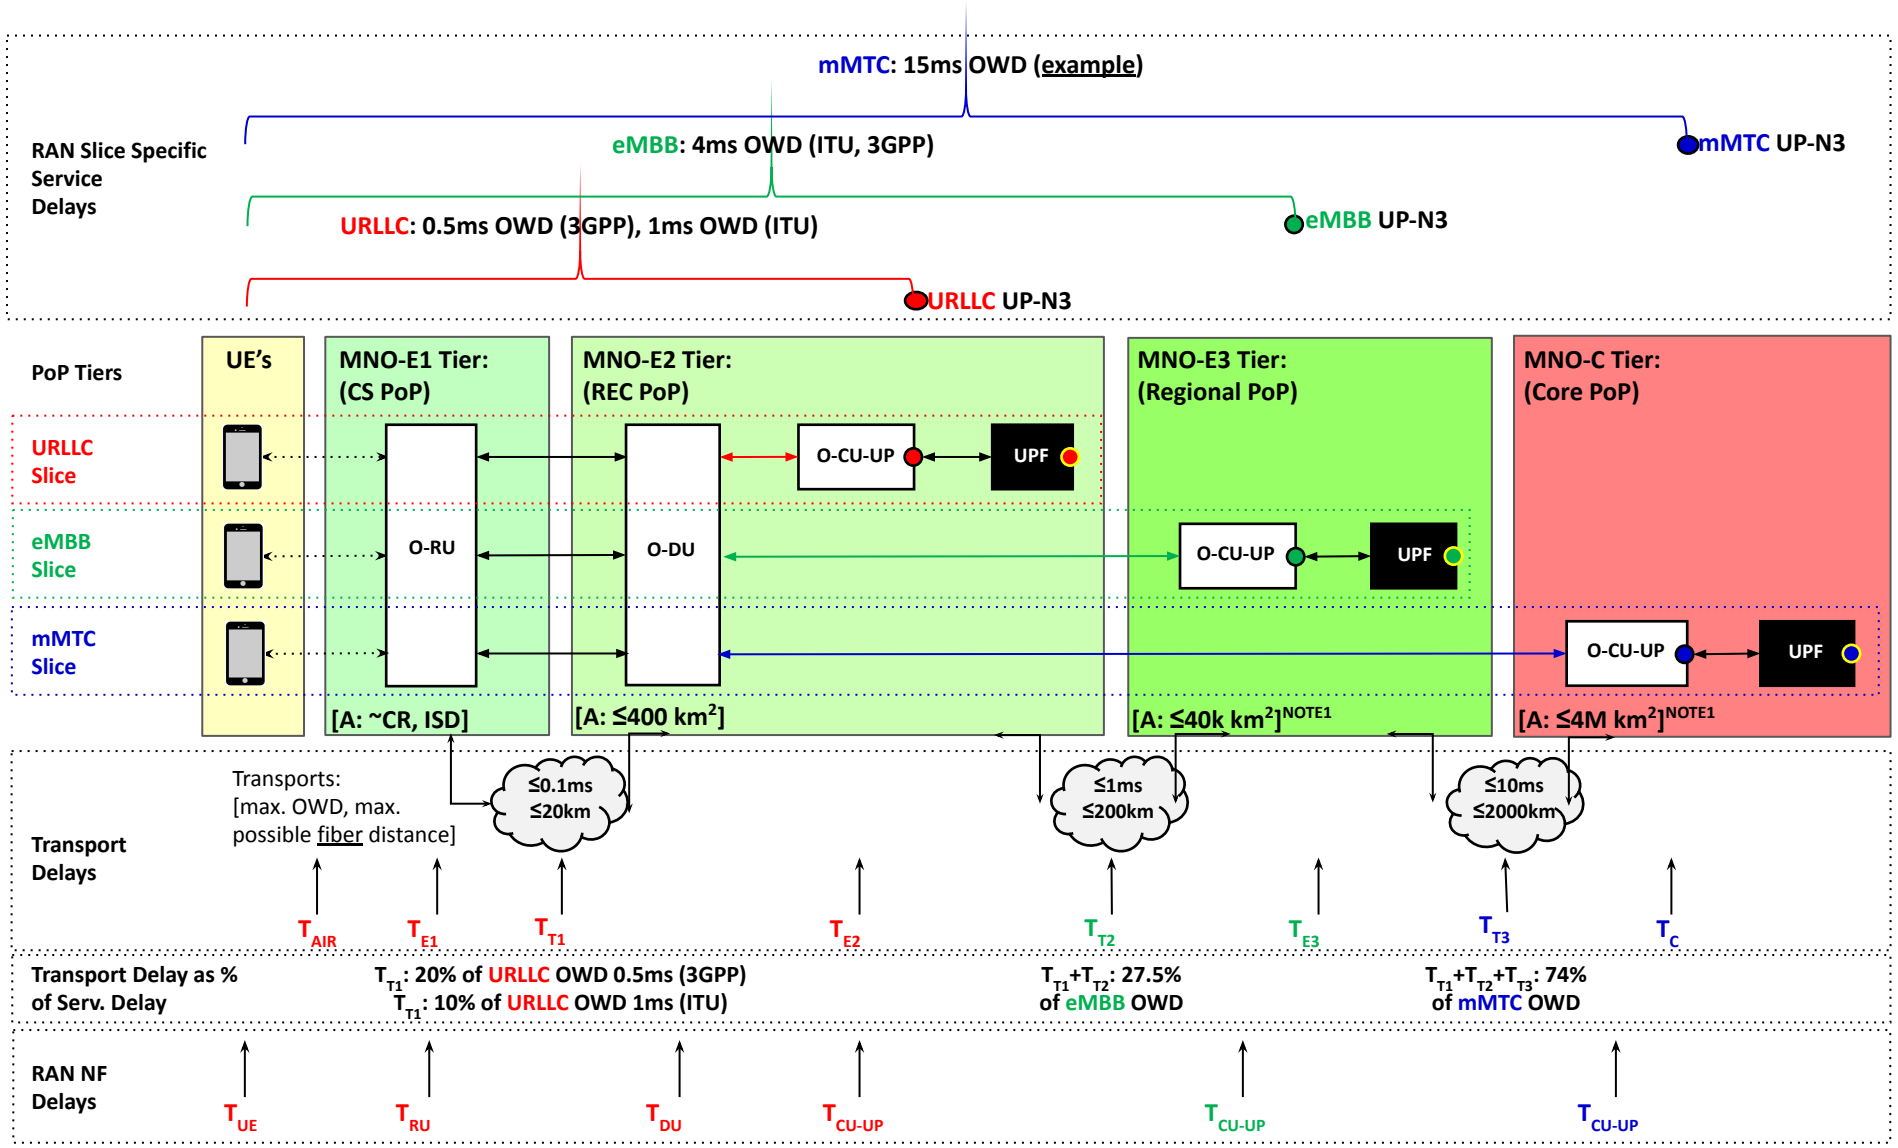

1. Speed of light is ~300 km in one millisecond (in vacuum)
2. IP packet can travel ~200 km in one millisecond (in fibre)

Red Hat

Red Hat Telco Platform Update

Red Hat

Telco Cloud Platform Requirements

Networking

- High Performance Data Plane
- Multiple Networking Interfaces
- High Throughput
- Low Latency
- Load Balancer
- Acceleration Hardware

Deployment & Scale

- Remote Distributed Deployments
- Single node and HA clusters
- Field upgradability
- Zero-touch-provisioning

Real-Time, Latency and Sync

- Time Synchronization
- Real-Time Determinism
- CPU Affinity and Isolation
- Topology Management
- Node Feature Discovery

Telco Cloud Platform Requirements

Networking

- High Performance Data Plane
- **Multiple Networking Interfaces**
- High Throughput
- Low Latency
- Load Balancer
- Acceleration Hardware

Deployment & Scale

- Remote Distributed Deployments
- Single node and HA clusters
- Field upgradability
- Zero-touch-provisioning

Real-Time, Latency and Sync

- Time Synchronization
- Real-Time Determinism
- CPU Affinity and Isolation
- Topology Management
- Node Feature Discovery

Traffic Isolation - Multiple Network Interfaces

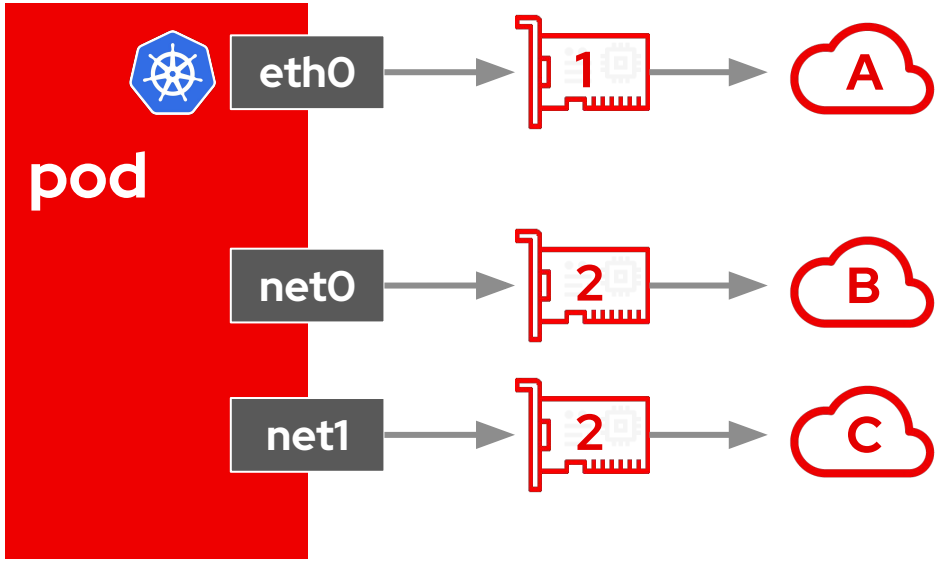

Multi-network Container

The Problem:
Kubernetes only supports one network interface, “eth0”, per pod.

The Solution:
Multus “meta plug-in” for Kubernetes CNI* which enables multiple physical network interfaces (nic’s) per pod

- virtual network functions are assigned to a different CNI plugins
- each function is defined in different CRD** object
- the first nic remains for the control plane traffic
- the secondary nic is for the data plane traffic

* CNI - Container Network Interface
** CRD - Custom Resource Definition

Red Hat

Traffic Isolation - Multiple Network Interfaces

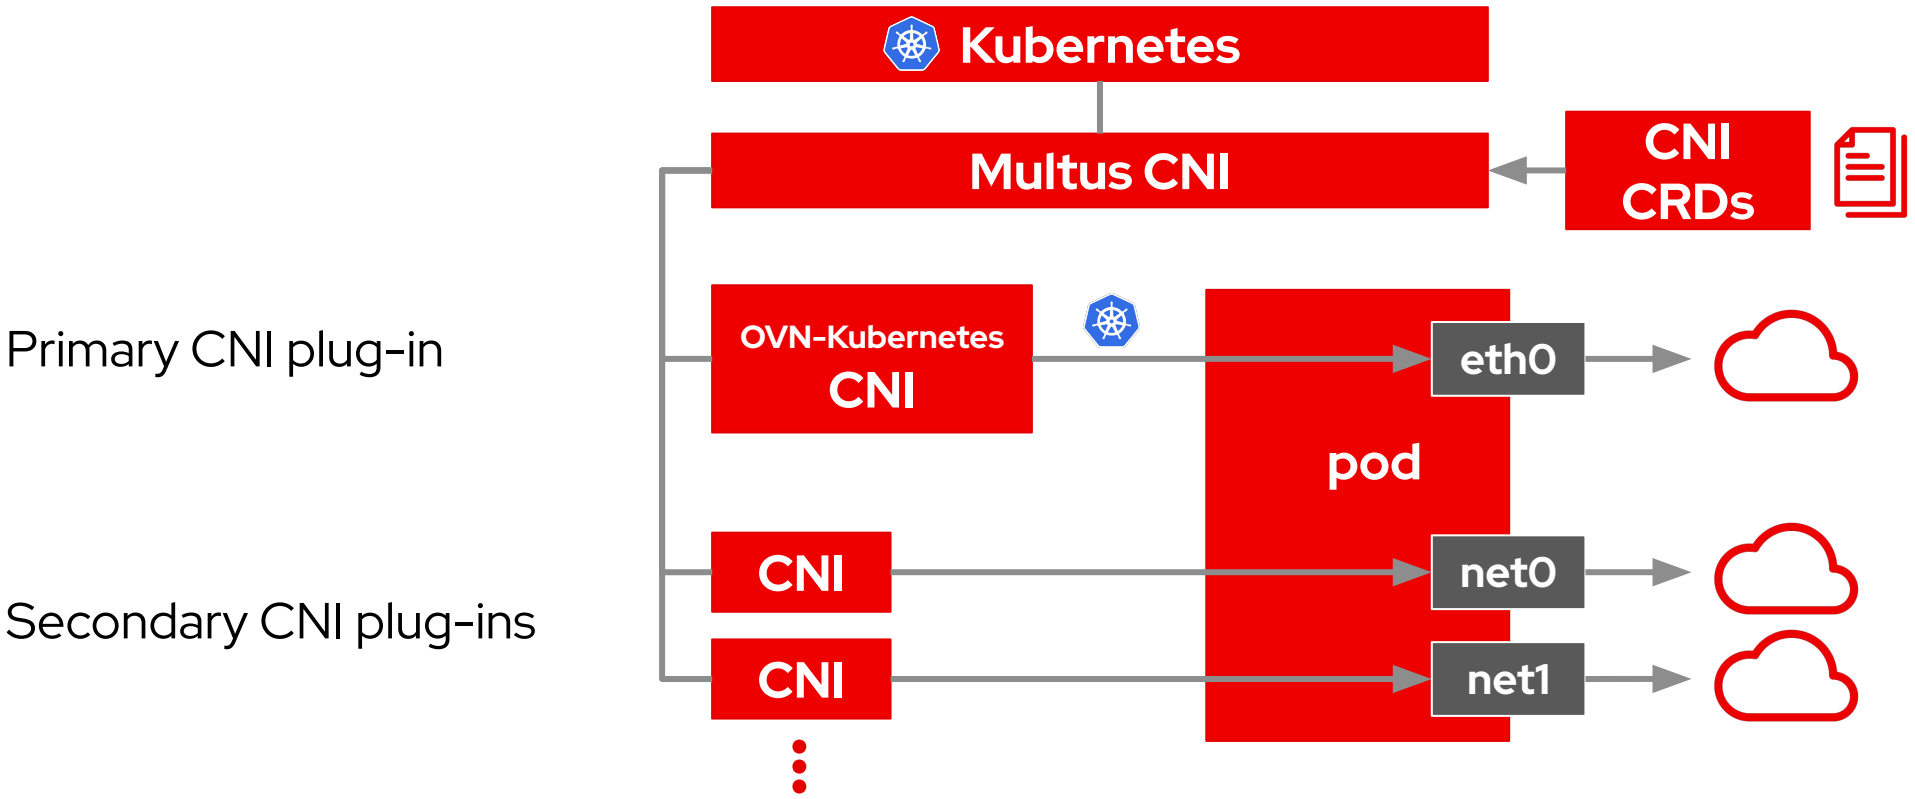

SR-IOV support enables physical nic to be used across multiple virtual functions

- MACVLAN (container MAC addressing)
- host device
- IPVLAN (IP control for both ipv4 and ipv6)
- Bridge with VLAN

- Static IPAM, (Static IP address management)
- DHCP IPAM (DHCP IP address management)
- Whereabouts (CNI IPAM plugin that assigns IP addresses cluster-wide)
- Route Override (override IP route given by previous CNI plugins)

IPv4/IPv6 Dual Stack Support Configuration

- IPv6 single/dual stack is supported in OpenShift 4.8 (k8s 1.21) with OVN.
- **Single Stack**
  - Either an IPv4 or IPv6 address is assigned to the pod interface
- **Dual Stack**
  - Both IPv4 and IPv6 addresses assigned to the interface

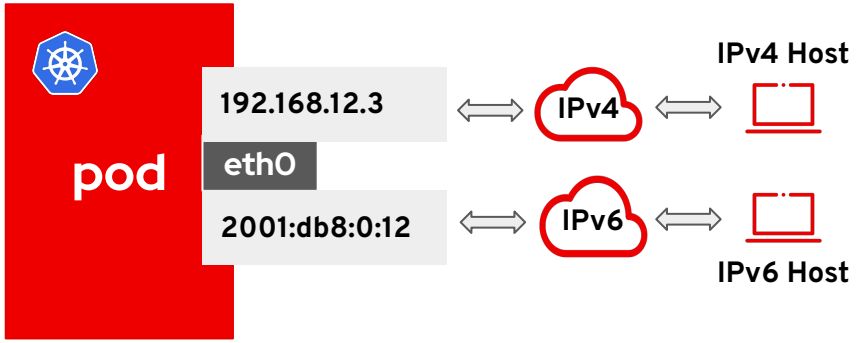

- Simple install-time configuration
  - Modify "`install-config.yaml`" to specify IPv6 subnets in addition to IPv4.
- Post-install configuration:
  - Edit "`network.config.openshift.io`" config to add secondary "`(machine|cluster|service)Network`" values, and they will get rolled out correctly.
- Restrictions / Caveats / Notes
  - OVN *only*, no plans to support in openshift-sdn (3rd-party plugins may have some support)
  - Supported platform at GA: Bare Metal IPI (UPI 4.11*, vSphere 4.12*, other platforms TBD)
  - Multus secondary interfaces have always been IPv6 ready

* = targeting

Telco Cloud Platform Requirements

Networking

- High Performance Data Plane
- Multiple Networking Interfaces
- High Throughput
- Low Latency
- **Load Balancer**
- Acceleration Hardware

Deployment & Scale

- Remote Distributed Deployments
- Single node and HA clusters
- Field upgradability
- Zero-touch-provisioning

Real-Time, Latency and Sync

- Time Synchronization
- Real-Time Determinism
- CPU Affinity and Isolation
- Topology Management
- Node Feature Discovery

Types of Kubernetes Services

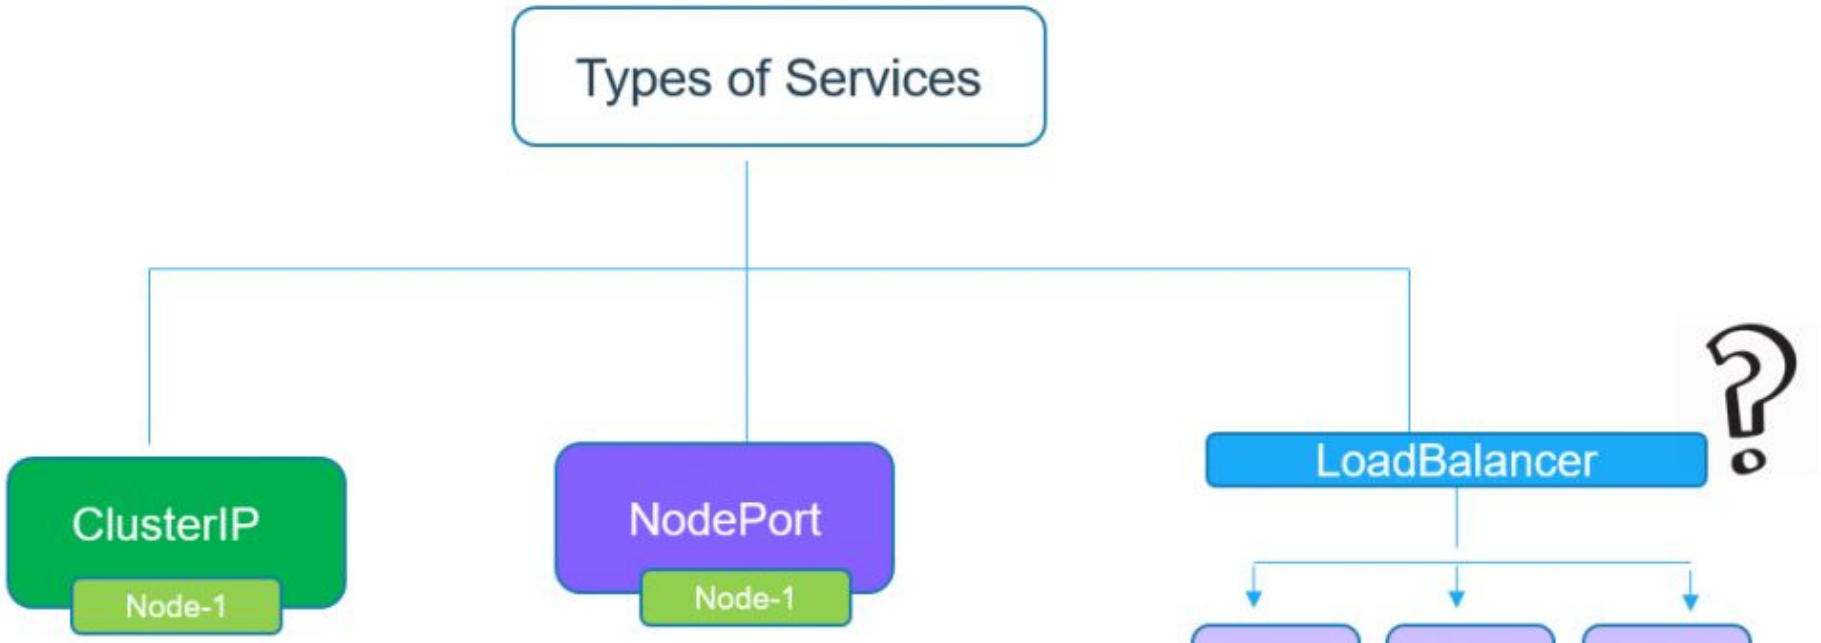

- Reachable within the cluster.
- Connects Frontend Pods to Backend Pods

- Exposing Frontend app to external world

- Equally distribute the loads

Load Balancer for Bare Metal?

Load Balancer Service

Exposes the Service externally using a cloud provider's load balancer. NodePort and ClusterIP Services, to which the external load balancer routes, are automatically created.

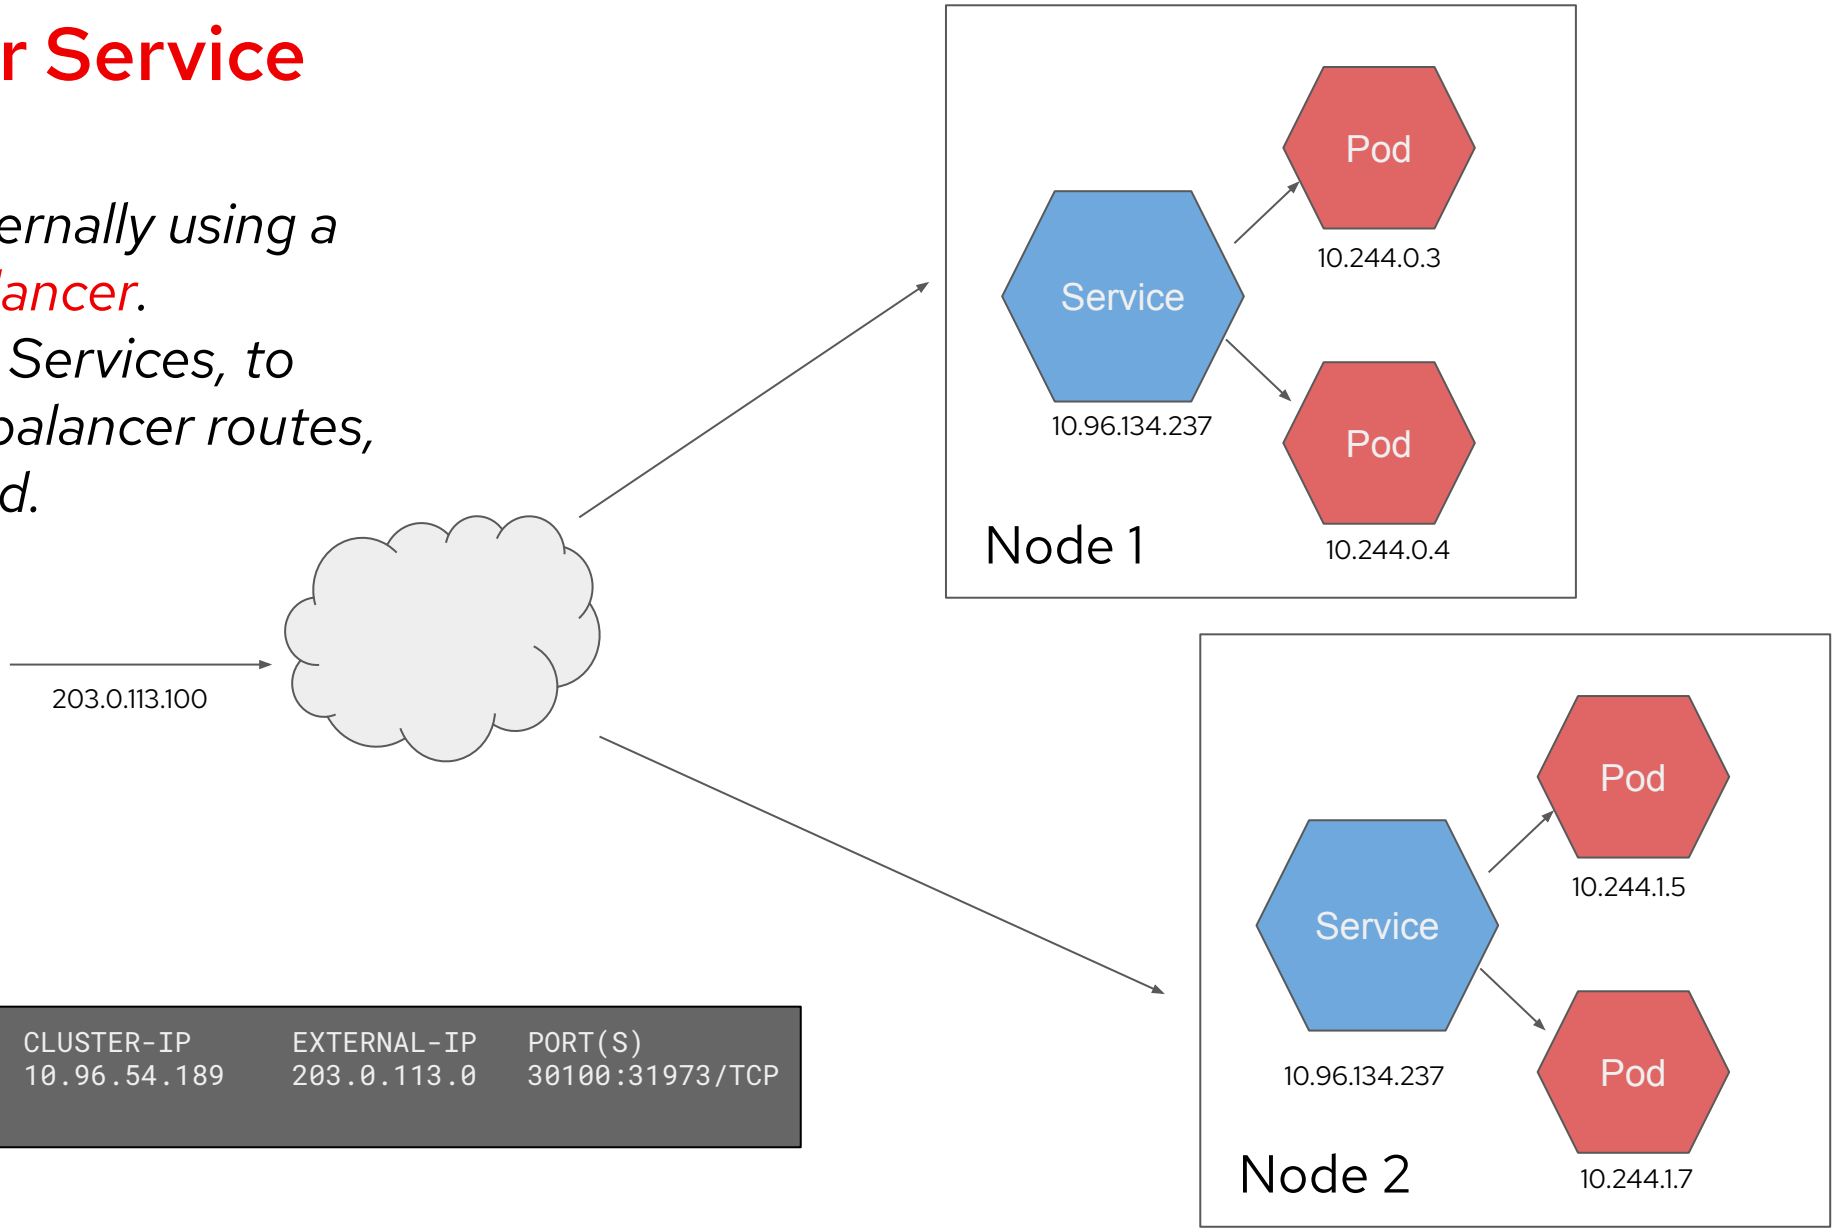


```
NAME          TYPE           CLUSTER-IP      EXTERNAL-IP    PORT(S)
lbservice     LoadBalancer   10.96.54.189    203.0.113.0    30100:31973/TCP
```

Load balancer service

Stable IP to reach our application

Load Balancing across the nodes

Load Balancer Service (On bare metal)

```yaml
type: LoadBalancer
status:
  loadBalancer:
    ingress:
    - ip: 203.0.113.100
```

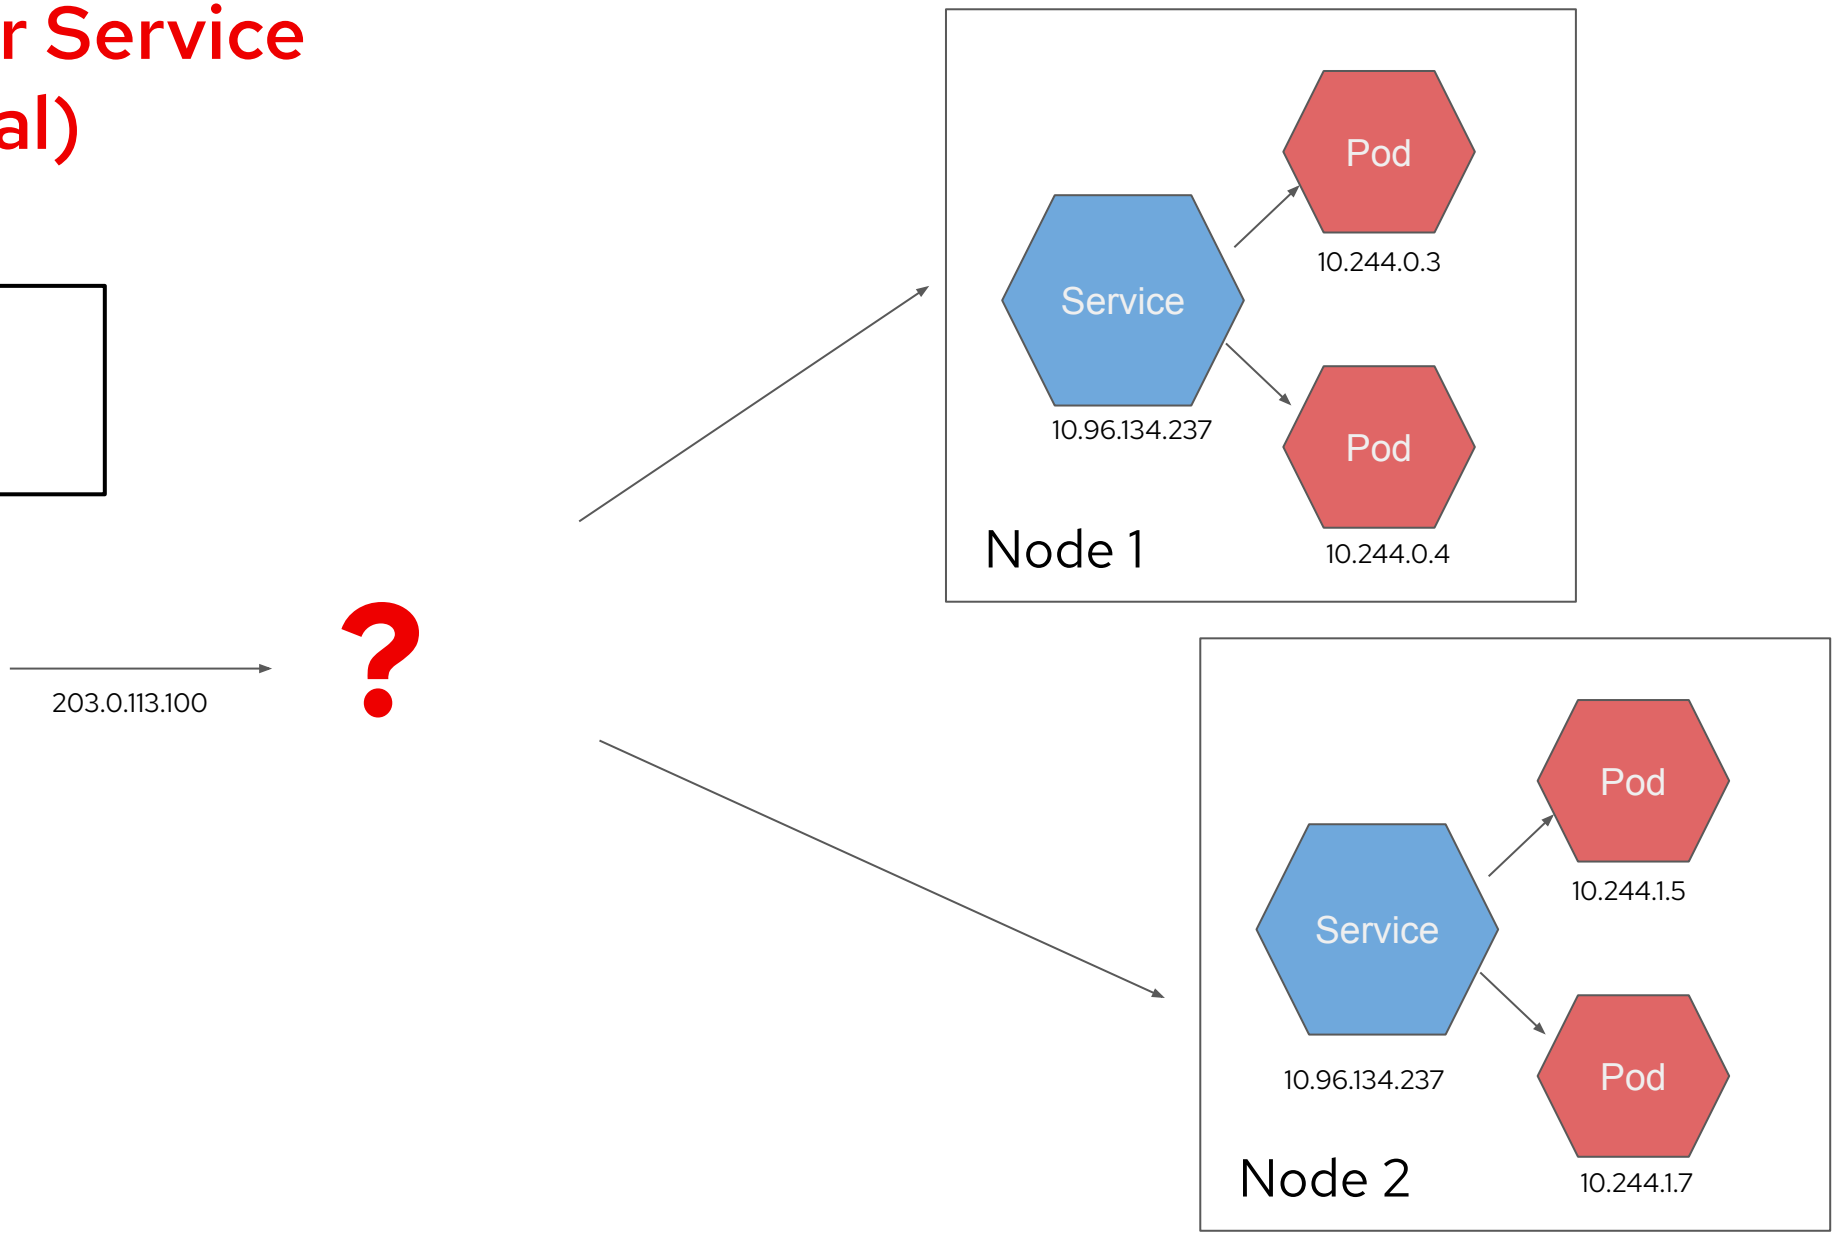

MetalLB - BGP (L3) Mode

Leverage the ECMP load balancing capabilities of the Site Gateway (router) coupled with dynamic BGP announcements.

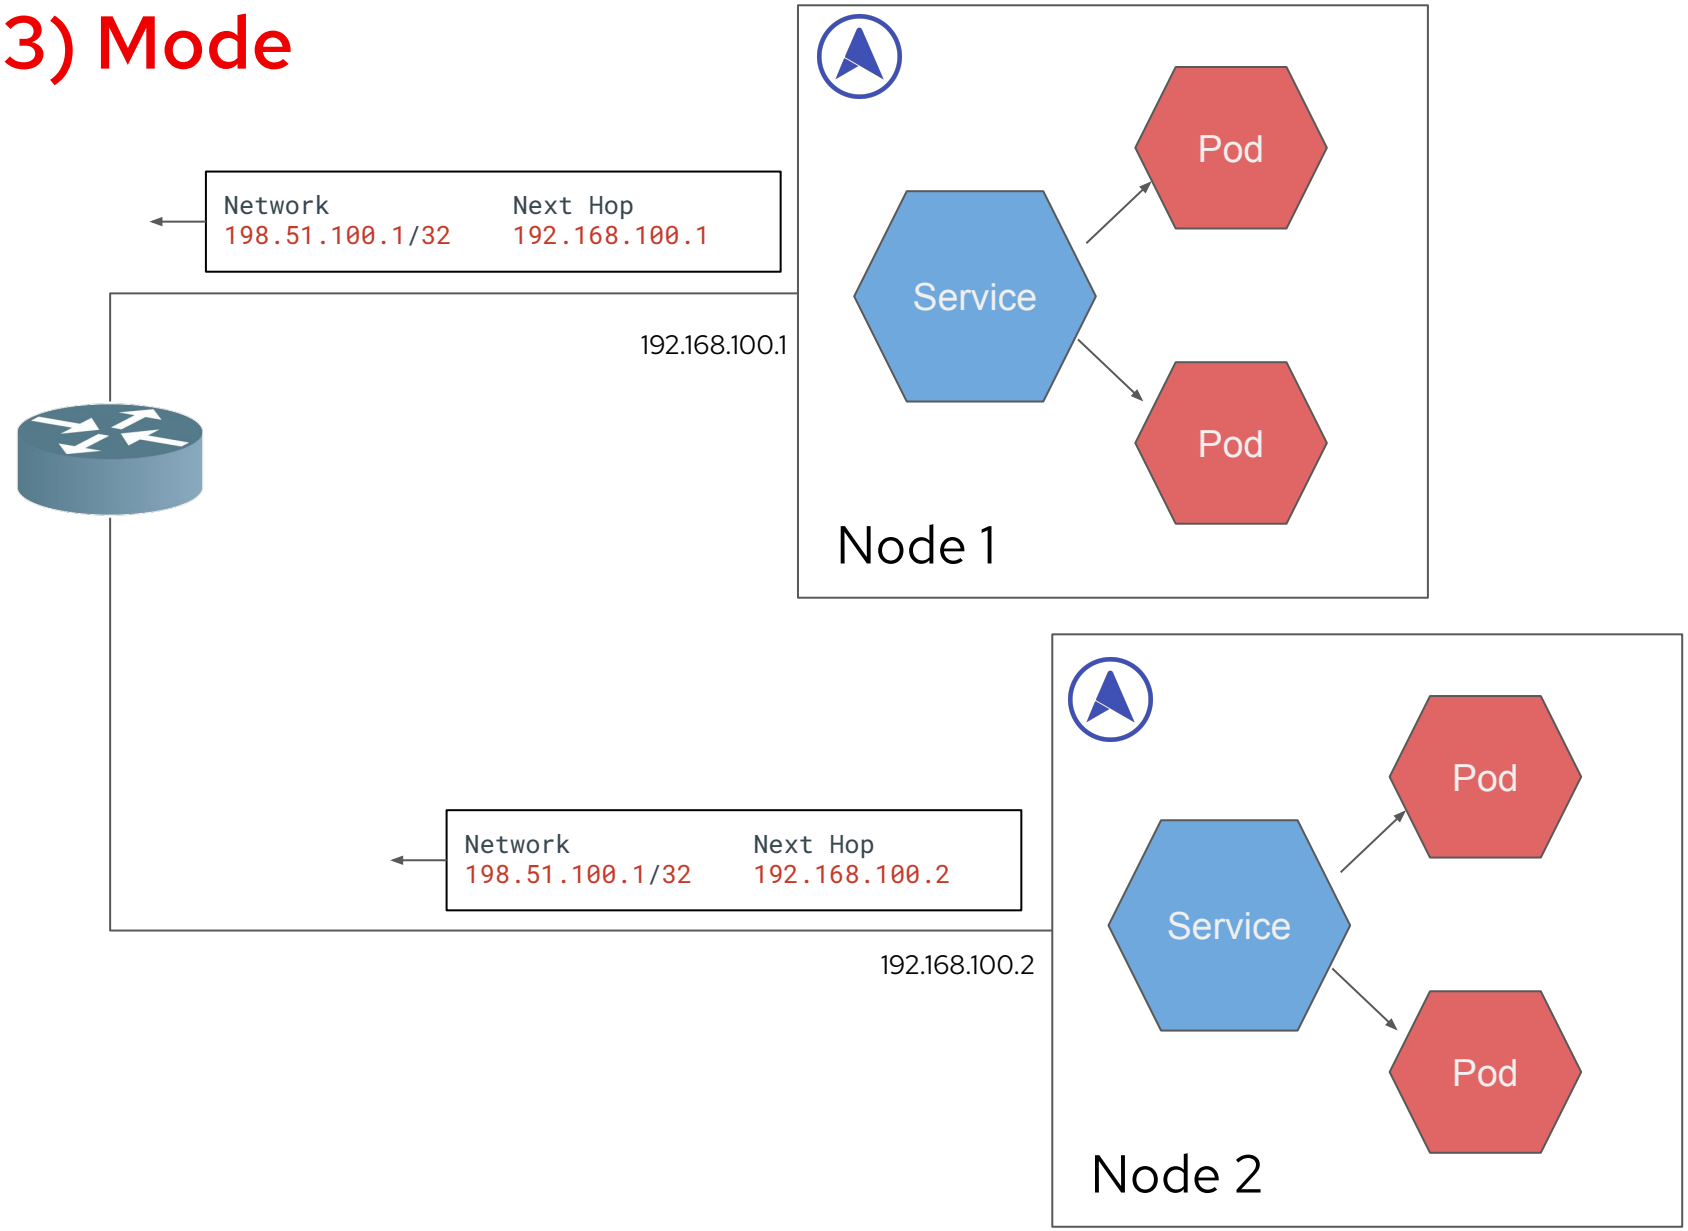


```yaml
Service:

  type: LoadBalancer
  status:
    loadBalancer:
      ingress:
      - ip: 198.51.100.1
```

MetalLB Architecture

- Controller
  - Single Instance
  - Handles the IP pooling and allocation
- Speaker
  - One per node
  - Hostnetworked pod
  - Handles the IP announcement (both L2 and BGP)

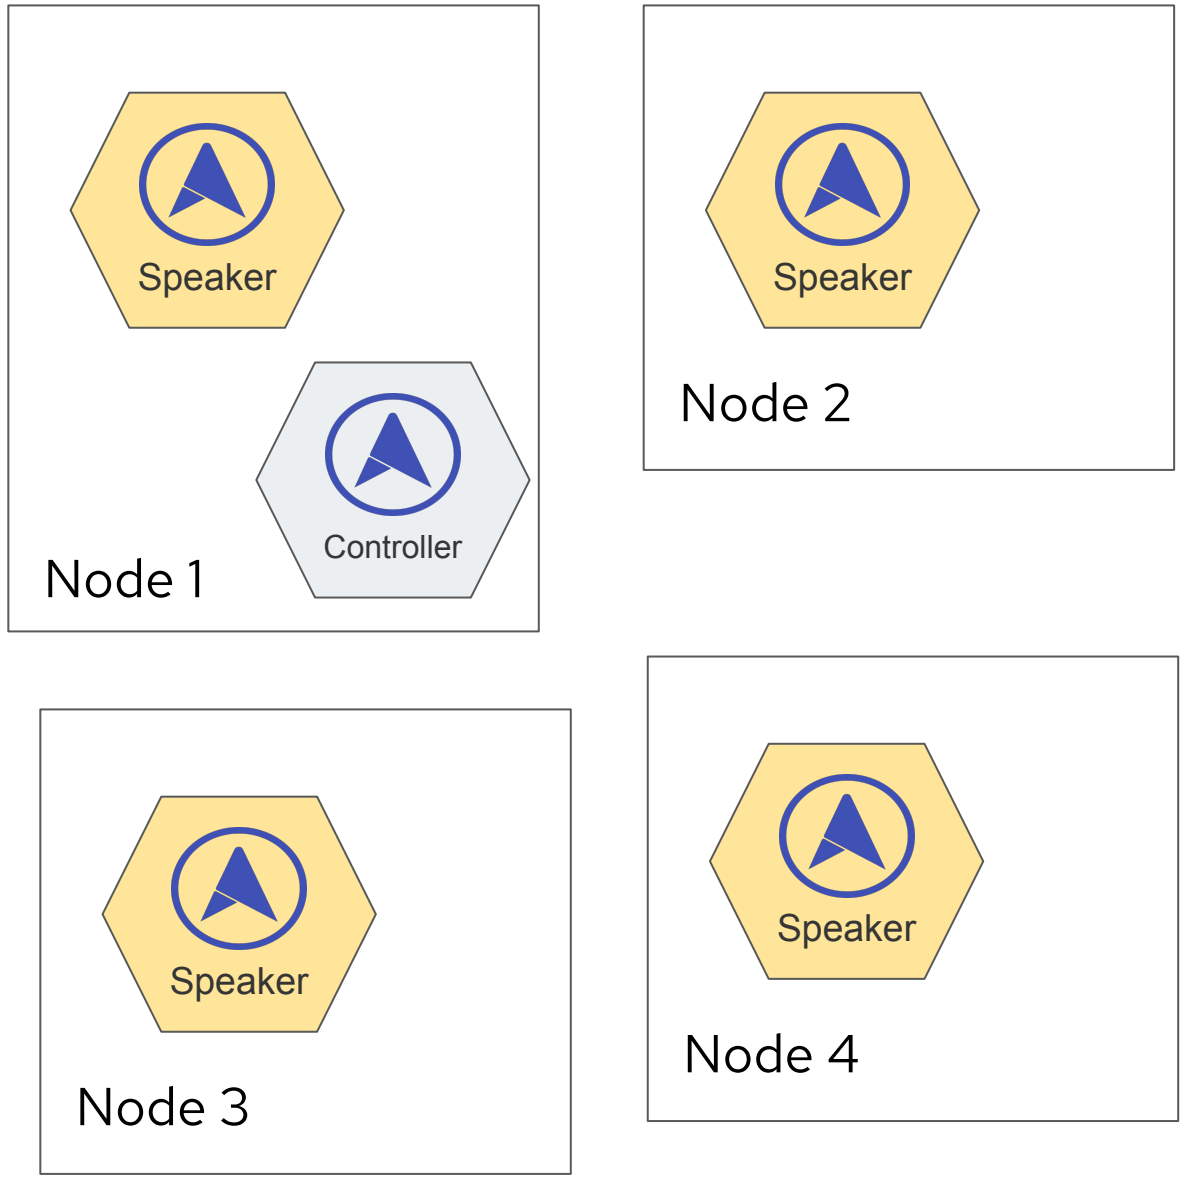

Speaker BGP mode (FRR)

- FRR Container
  - https://frrouting.org/
- Speaker Container
  - Listens for services + MetalLB configuration
  - Based on the K8s API, renders the FRR configuration
  - Reloads the FRR configuration on the fly
- Reloader Container

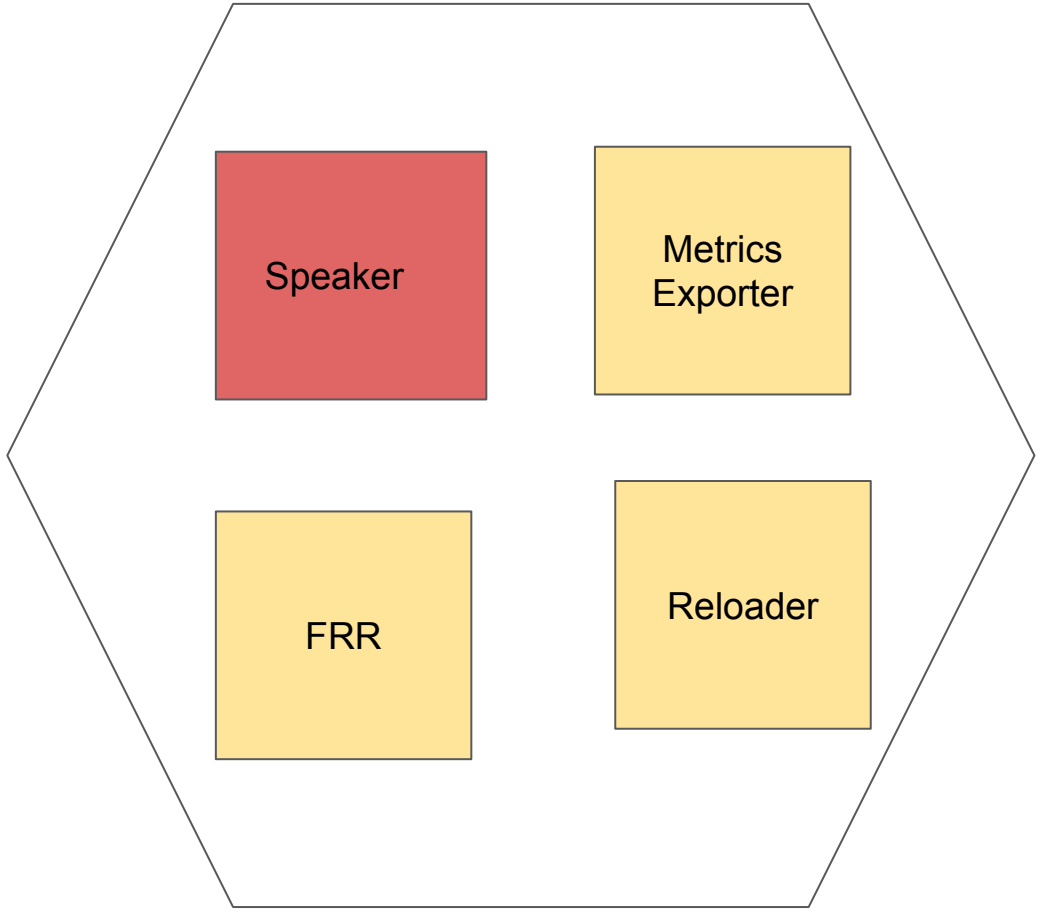

OpenShift MetalLB (BGP) documentation

[About MetalLB and the MetalLB Operator - Load balancing with MetalLB | Networking | OpenShift Container Platform 4.10]

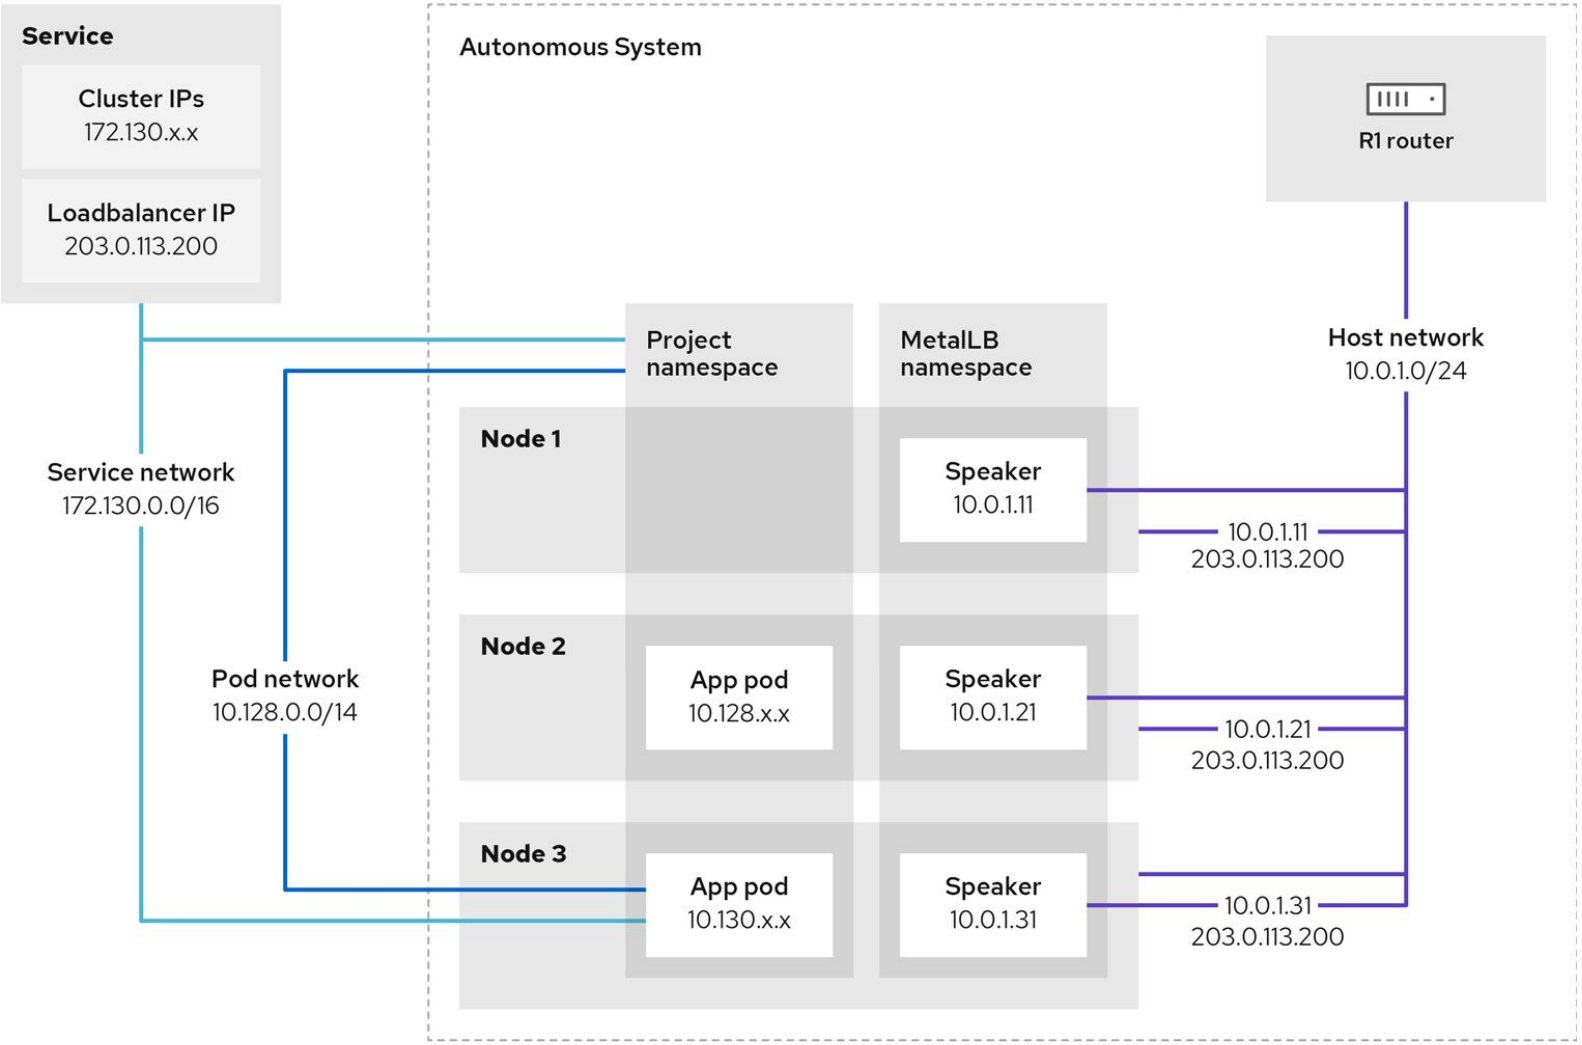

Telco Cloud Platform Requirements

Networking

- High Performance Data Plane
- Multiple Networking Interfaces
- High Throughput
- Low Latency
- Load Balancer
- Acceleration Hardware

Deployment & Scale

- Remote Distributed Deployments
- Single node and HA clusters
- Field upgradability
- Zero-touch-provisioning

Real-Time, Latency and Sync

- **Time Synchronization**
- Real-Time Determinism
- CPU Affinity and Isolation
- Topology Management
- Node Feature Discovery

Frequency, Phase and Time Synchronization

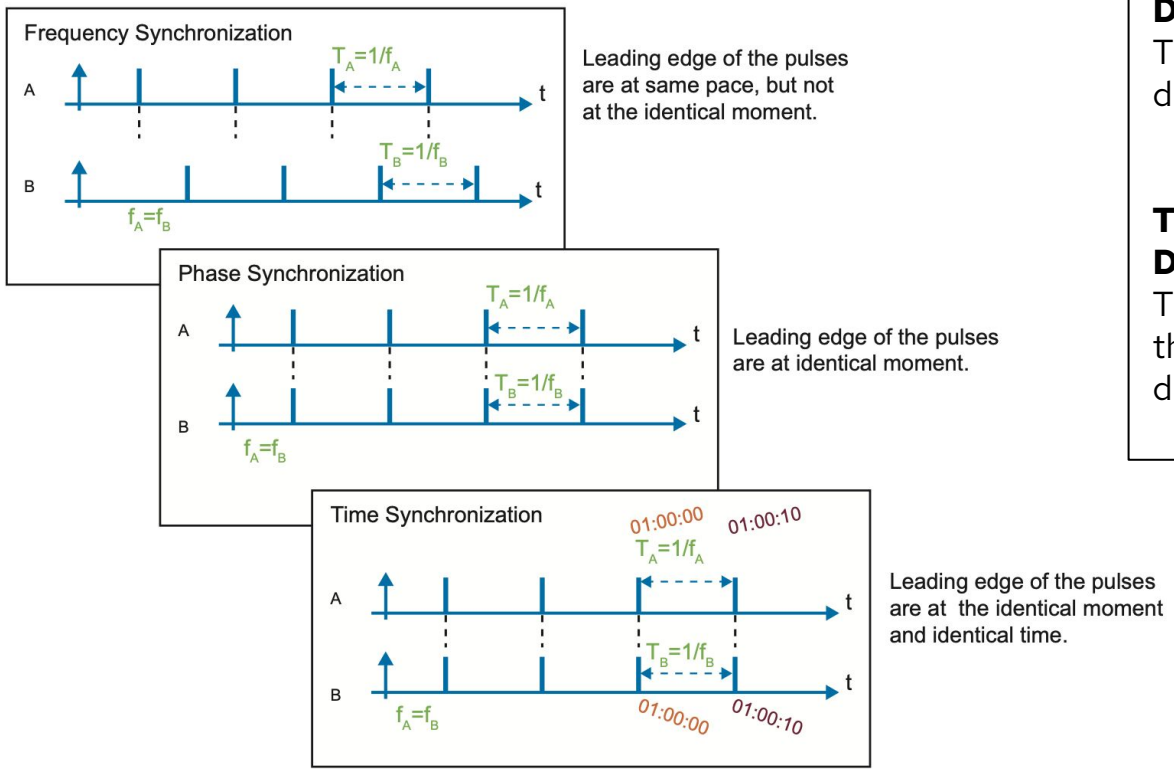

Figure 1: Frequency, Phase and Time Synchronization

Microsemi

FDD - Frequency Division Duplex:
Transmit/Receive use different frequencies

TDD - Time Division Duplex:
Transmit/Receive use the same frequency at different times

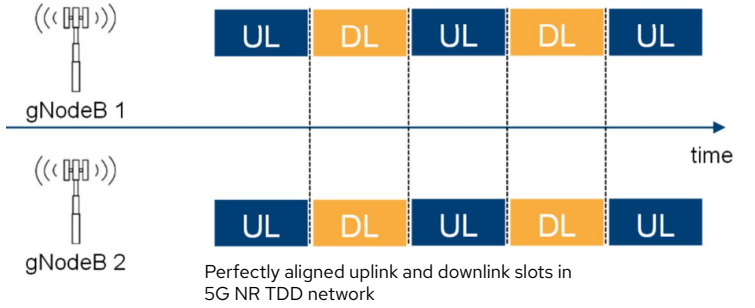

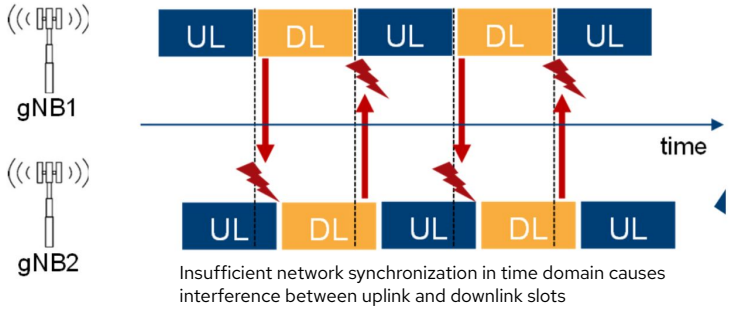

Rohde & Schwartz

Red Hat

RAN Synchronization Challenge on One Slide

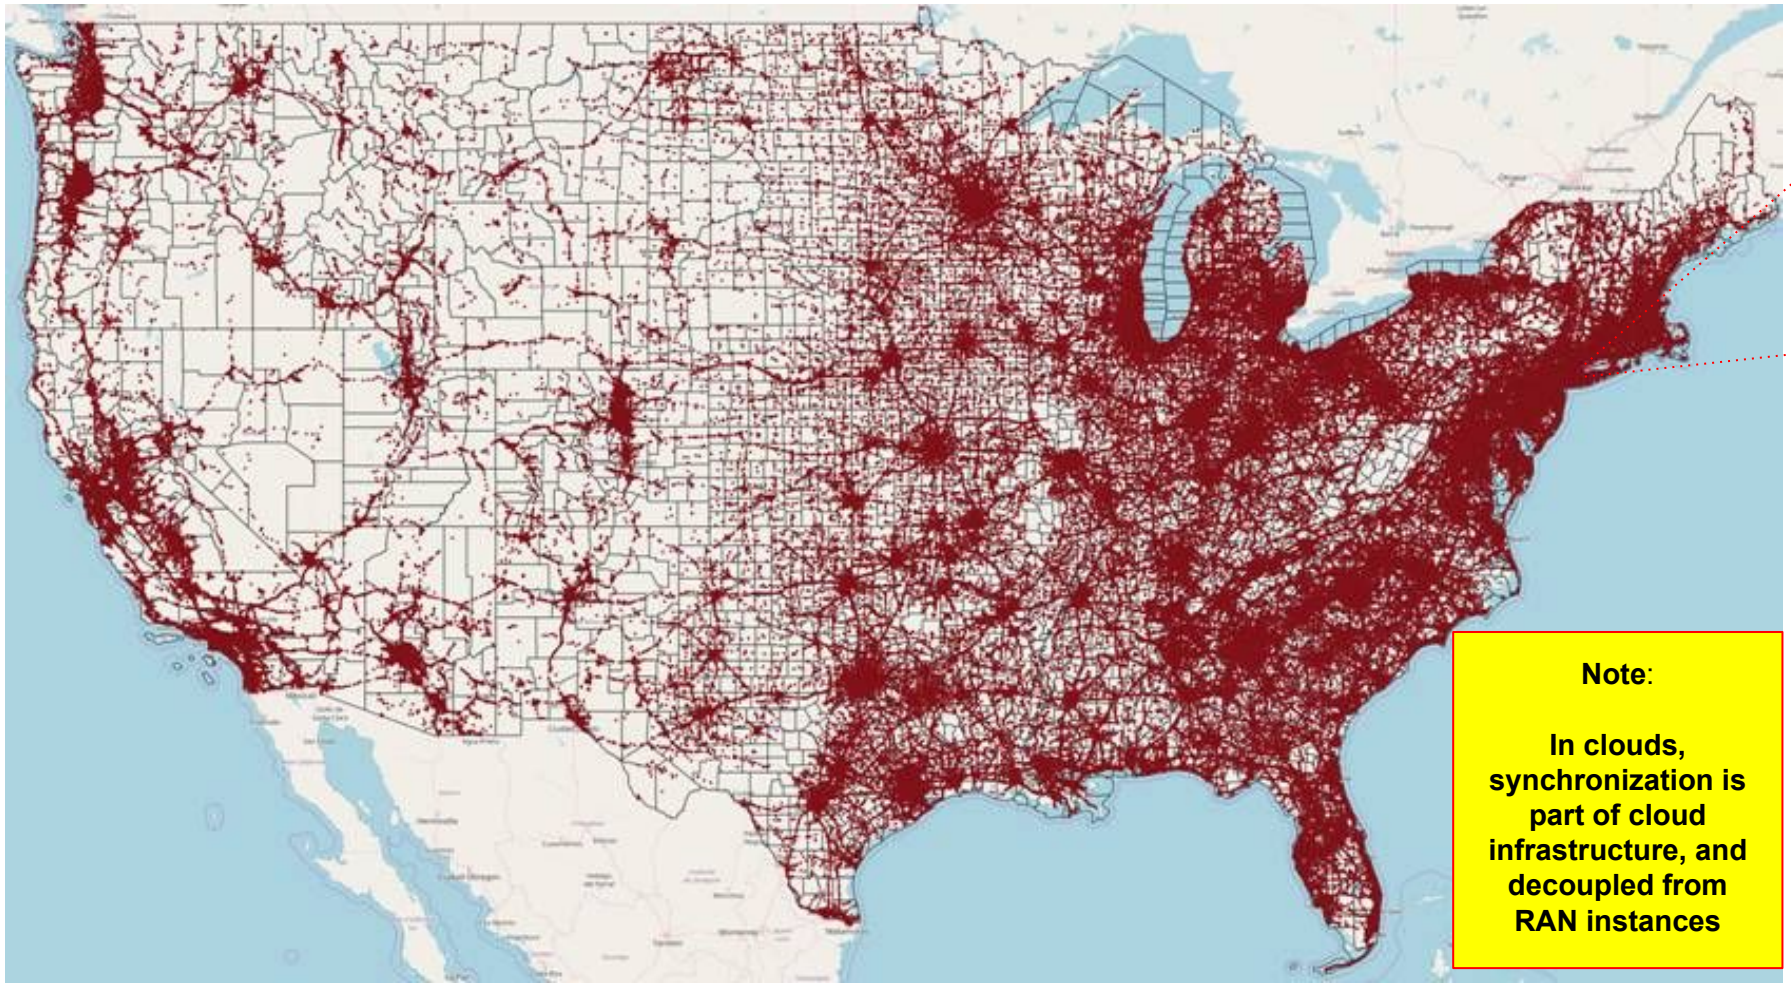

Note:

In clouds, synchronization is part of cloud infrastructure, and decoupled from RAN instances

One provider: >500k cell sectors, >60k cell sites; radio & site numbers will increase for 5G, esp. with utilization of FR2 "mmwave" spectrum

NYC: >10k+ cells, >2k+ sites

Synchronize

- 100's of k - M's of distributed Network Elements to 1.5μs absolute |TE|
- 100's to thousands of physical and cloud NEs within a synchronization cluster to O(≦100's) ns relative |TE|,

while

- Minimizing CapEx, Installation and OpEx costs, and
- Maximizing the synchronization performance and availability

and also

- Across widely diverse transports
- Interoperability in multi-vendor & multi-technology environment
- With tight freq. stability reqs.

Absolute |TE| (TE=Time Error): maximum absolute time difference against a common reference standard
Relative |TE| (rel. |TAE|=Time Alignment Error): max. relative time difference between two entities (within a sync. cluster)

Red Hat

O-RAN "LLS-Cx" Synchronization Reference Configs

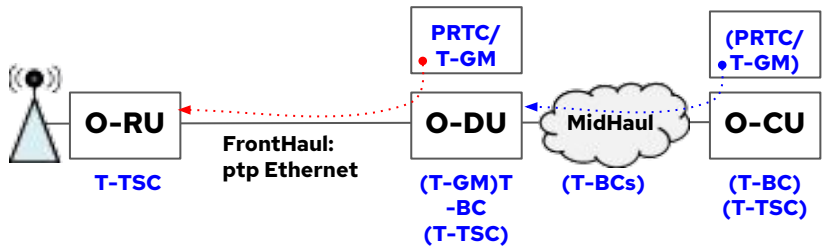

WG4 Configuration LLS-C1

FH Direct connection to O-RU(s) from O-DU; sync source in O-DU

WG9 C1, Option A: T-GM Embedded in O-DU
WG9 C1, Option B: T-GM Directly Connected to O-DU
WG9 C1, Option C: T-GM connected to O-DU via chain of network nodes

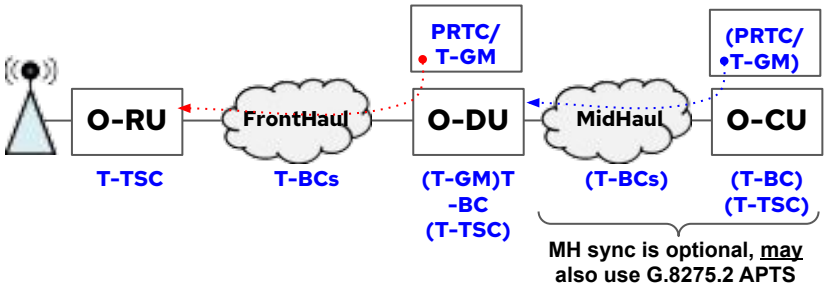

WG4 Configuration LLS-C2

FH Network connection to O-RU(s) from O-DU; sync source in O-DU

WG9 C2, Opt. A: O-DU is the nearest common T-BC
WG9 C2, Opt. B: nearest common T-BC not O-DU

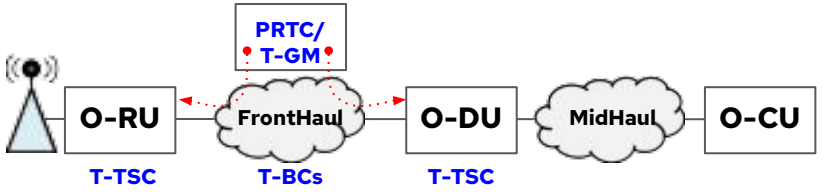

WG4 Configuration LLS-C3

Network connection to O-RU from O-DU & sync source in FH network

WG9 C3, Option A: T-GM is the nearest common master
WG9 C3, Option B: nearest common master is not T-GM
WG9 C3, Option C: T-GM in Mid/Back-haul
WG9 C3, Option D: T-GM in Mid/Back-haul with T-BC chain

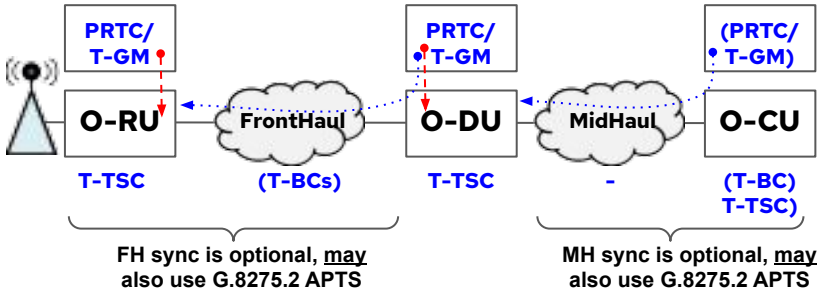

WG4 Configuration LLS-C4

Network connection to O-RU from O-DU & local sync sources

WG9 C4, Option 1: GNSS at Cell Site (e.g. in O-RU / xNB)
WG9 C4, Option 2: GNSS at Cell & Edge + APTS network

Note:

In clouds, synchronization is part of cloud infrastructure, and decoupled from RAN instances

Red Hat

PTP in OpenShift (Linux) Nodes

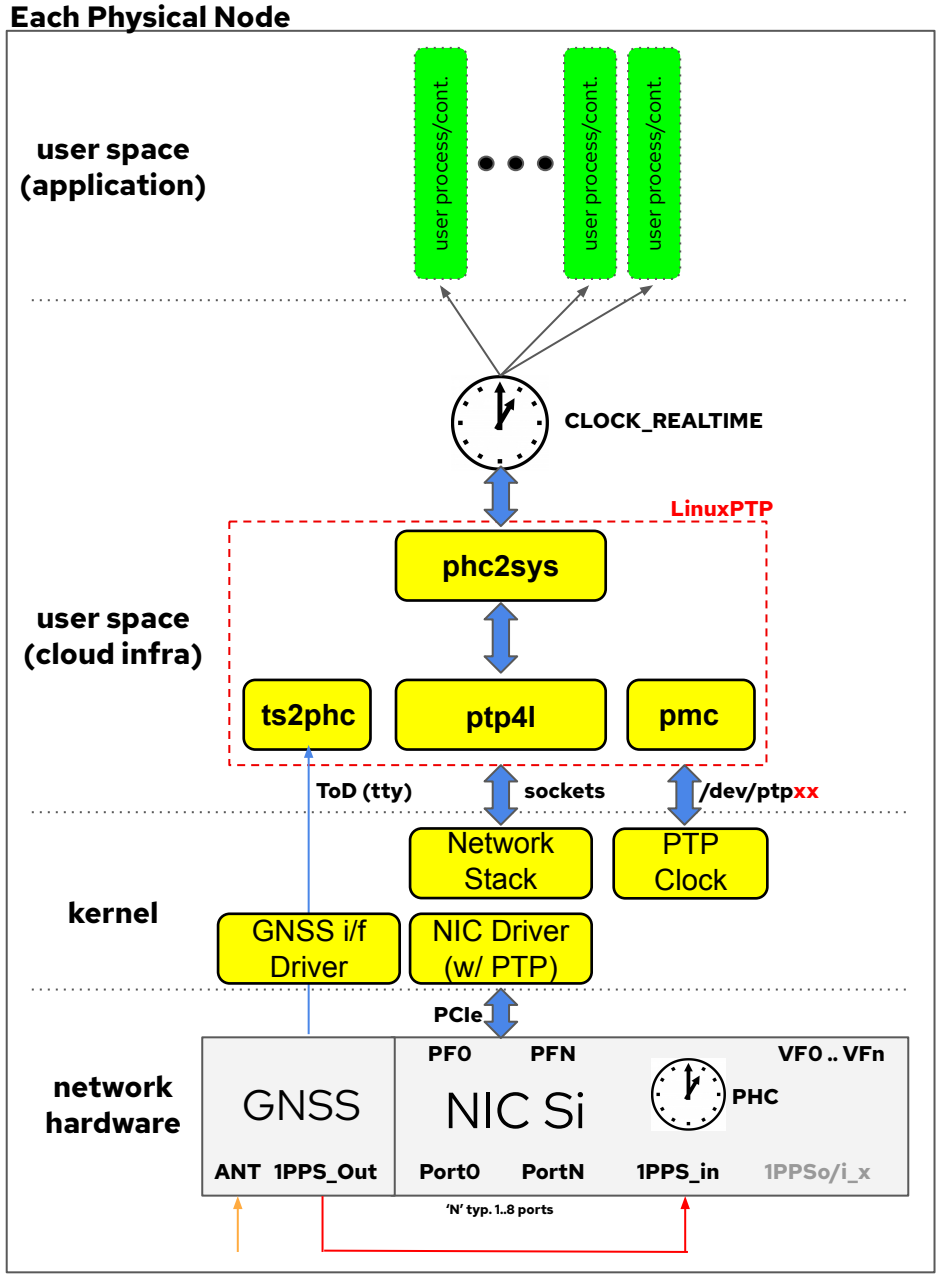

Key Components of the node PTP implementation

- HW specific synchronization SW support features are implemented in HW device drivers
- HW Clock (PHC) support in NIC Si is required for high accuracy
- **Linuxptp** is an Open Source project implementation of the PTP SW stack for Linux
- **ptp4l** implements Boundary Clock (BC) and Ordinary Clock (OC), it synchronizes PTP hardware clock (PHC) to remote master clock
- ptp4l is very flexible, and can be configured to support specific profiles, assuming that HW & driver supports associated features (e.g. PHC, L3 vs. L2 transport, accuracy targets)
- **phc2sys** synchronizes two or more clocks in the system, typically used to synchronize the system clock from PTP / PHC
- **pmc** - PTP management client; 1588 basic management access for ptp4l
- **ts2phc** synchronizes PHC(s) from external reference signals, such as 1PPS_in and ToD messages - used in certain HW assisted T-BC, and GNSS driven T-GM configurations
- In k8s clusters, synchronization functions are configured and monitored with k8s, and associated general CM, PM and FM event and metrics tools.

And many other Si & HW Partners

OpenShift 4.9 PTP for RAN Workloads

- OCP Node as an Ordinary Clock or as a Boundary Clock
- [O-RAN Approved] Low-latency, Node-local Event Bus w/ PTP Events and sidecar image for easy CNF (vDU) consumption

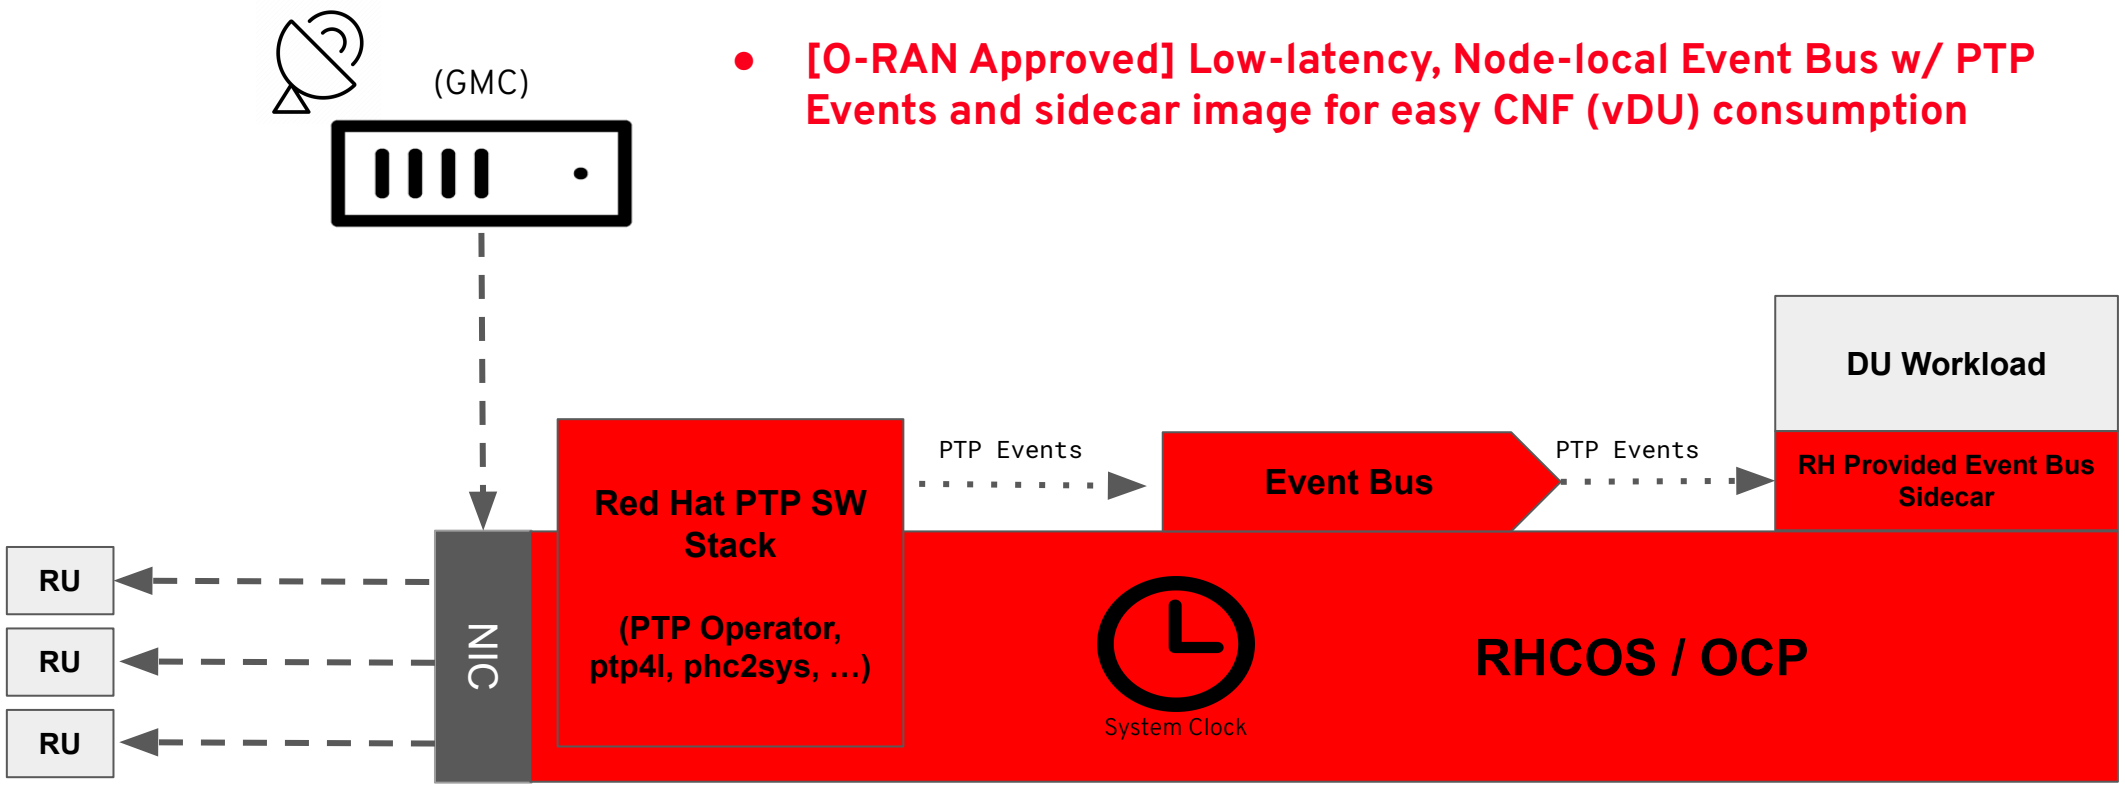

- Cell Site Router (CSR) GMC - Grandmaster Clock BC - Boundary Clock OC - Ordinary Clock

Telco Cloud Platform Requirements

Networking

- High Performance Data Plane
- Multiple Networking Interfaces
- **High Throughput**
- **Low Latency**
- Load Balancer
- Acceleration Hardware

Deployment & Scale

- Remote Distributed Deployments
- Single node and HA clusters
- Field upgradability
- Zero-touch-provisioning

Real-Time, Latency and Sync

- Time Synchronization
- **Real-Time Determinism**
- **CPU Affinity and Isolation**
- Topology Management
- Node Feature Discovery

OpenShift Performance Add-On Operator

- Provides deterministic application performance for latency sensitive workloads
- Intel EPA features
- HugePages Support
- CPU Pinning and Isolation
  - Provides CPU Manager with *'isolcpus'*
  - Intel CMK plugin functionality
- SR-IOV, DPDK and RDMA
- Real-Time Kernel installation and configuration
- Additional Operators Supported:
  - Node Topology Manager (NUMA)
  - Node Feature Discovery
    - GPU, FPGA, SR-IOV

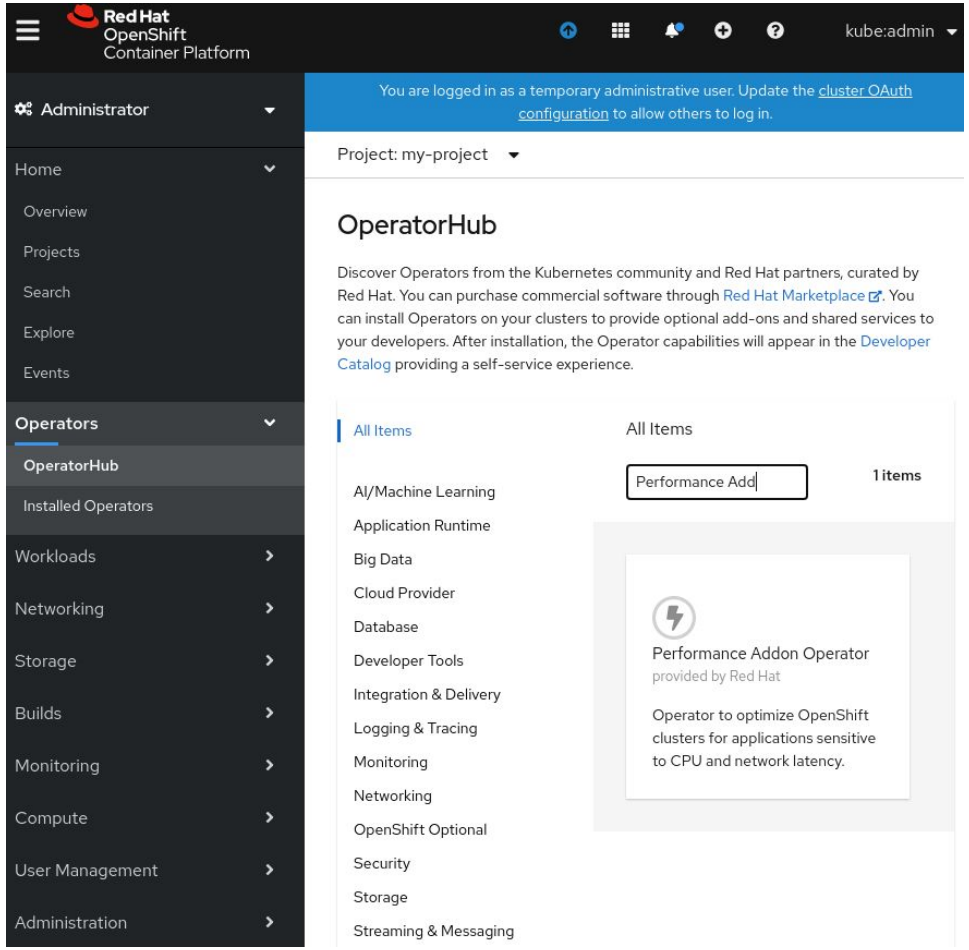

OpenShift Performance Add-On Operator

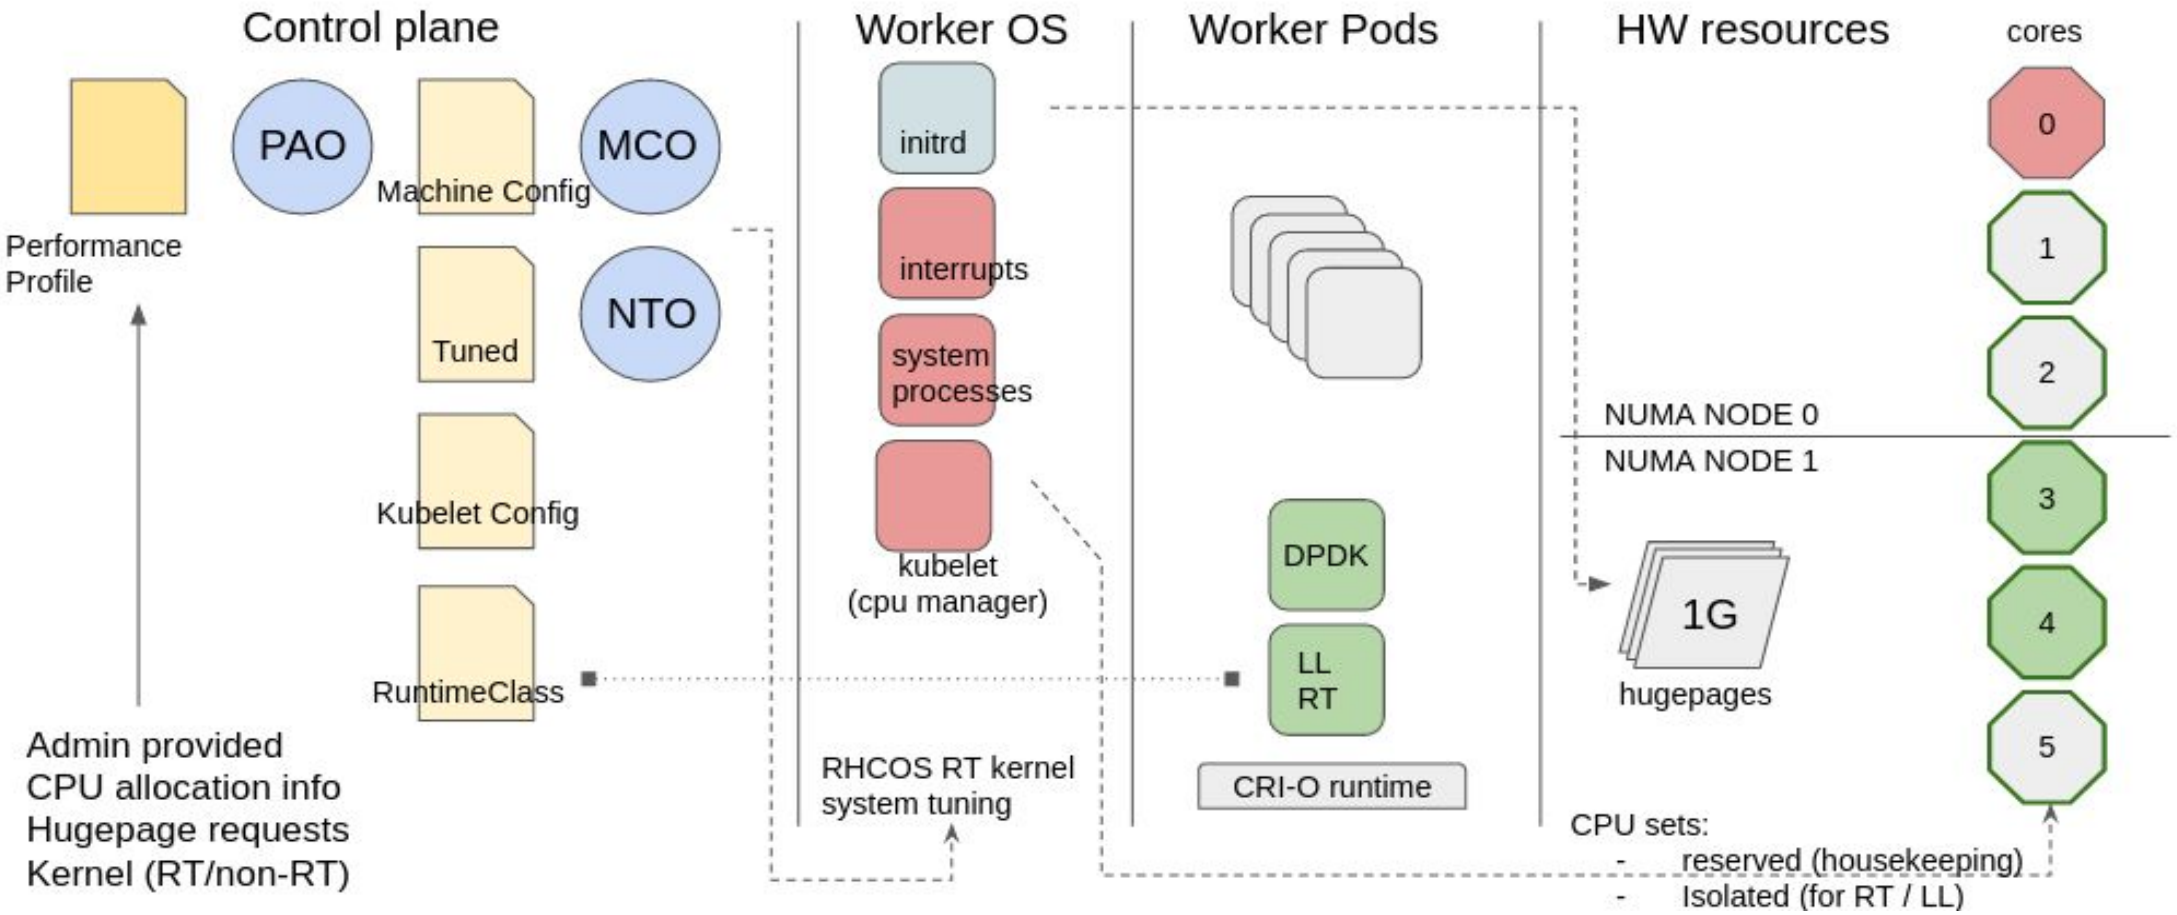

Telco Cloud Platform Requirements

Networking

- High Performance Data Plane
- Multiple Networking Interfaces
- High Throughput
- Low Latency
- Load Balancer
- Acceleration Hardware

Deployment & Scale

- Remote Distributed Deployments
- Single node and HA clusters
- Field upgradability
- Zero-touch-provisioning

Real-Time, Latency and Sync

- Time Synchronization
- **Real-Time Determinism**
- CPU Affinity and Isolation
- Topology Management
- Node Feature Discovery

Real Time RHEL CoreOS

A Real Time Kernel is a Red Hat Enterprise Linux kernel that is modified to maintain **low latency**, **consistent response time** and **workload determinism**.

This feature allows workloads to run uninterrupted by the Operating System.

- Allow the installation of the Real Time Kernel on RHEL CoreOS nodes.
- Allow the cluster administrator to provide a PerformanceProfile that defines:
  - A number of CPU cores dedicated to “housekeeping” tasks.
  - A number of CPU cores dedicated for workloads (CPU Pinning).
- NUMA alignment for devices, memory and cores used by Low Latency Workloads.

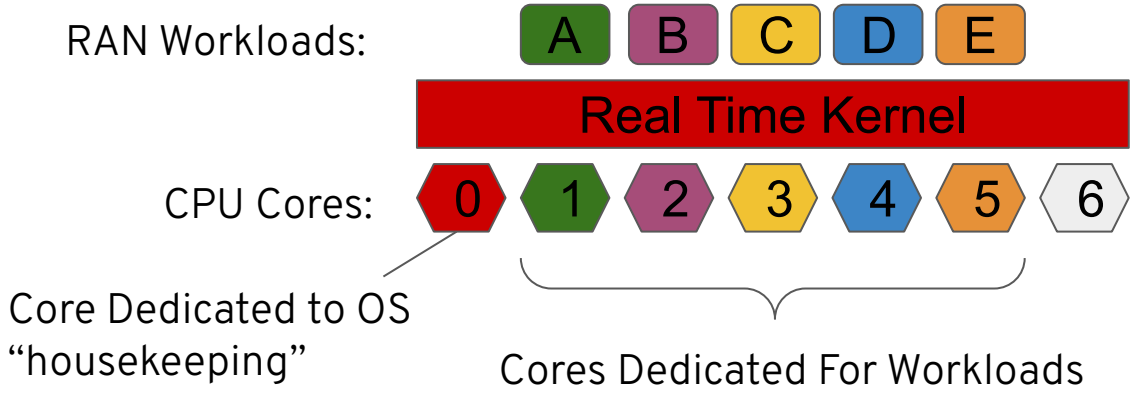

OpenShift Real Time Latency

Core OS Real-Time Host Kernel Support

- OpenShift 4.6
- Max Latencies
  - HT enabled: 10µsec (average 3µsec)
  - HT disabled: 8µsec (average 4µsec)
- Results based on 24 hour cyclict test across 8 cores w/ stress-ng running on cores
- Dell PowerEdge R740 - Intel(R) Xeon(R) Gold 6126 CPU @ 2.60GHz
  - 2 sockets*12 cores/sockets = 24 cores
  - 192G, 2666 MHz, Quad channel, 2 installed DIMMs/channel
- CoreOS 44.81.202003100230

```
# Total: 086400000 086399970 086399940 086399909 086399879 086399847
086399818 086399787
# Min Latencies: 00002 00002 00002 00002 00002 00002 00002 00002
# Avg Latencies: 00003 00003 00003 00003 00003 00003 00003 00003
# Max Latencies: 00010 00009 00009 00010 00009 00009 00010 00010
# Histogram Overflows: 00000 00000 00000 00000 00000 00000 00000 00000
```

```
Disabled HT
Disabled VT and VT-d in BIOS (virtualization)
Disabled “memory patrol scrubbing” (memory)
Disabled “embedded hw monitoring” (system)

# Min Latencies: 00004 00004 00004 00004 00004 00004 00004 00004 00004 00004 00004 00003
# Avg Latencies: 00004 00004 00004 00004 00004 00004 00004 00004 00004 00004 00004 00003
# Max Latencies: 00008 00007 00007 00007 00007 00007 00007 00007 00008 00008 00008 00008
```


Red Hat

Telco Cloud Platform Requirements

Networking

- High Performance Data Plane
- Multiple Networking Interfaces
- High Throughput
- Low Latency
- Load Balancer
- **Acceleration Hardware**

Deployment & Scale

- Remote Distributed Deployments
- Single node and HA clusters
- Field upgradability
- Zero-touch-provisioning

Real-Time, Latency and Sync

- Time Synchronization
- Real-Time Determinism
- CPU Affinity and Isolation
- Topology Management
- Node Feature Discovery

OpenShift Hardware Acceleration Support for DU

- Intel N3000 (FPGA PAC)
  - Look-aside architecture
  - L1 baseband physical layer processing
  - FlexRAN
  - FEC offload
  - OpenNESS
- NVIDIA GPU Operator
  - Inline architecture
  - Combined with linux kernel RDMA support and Mellanox NIC card
  - L1 baseband physical layer processing
- Additional Intel SoC and FPGA support Future

Inline Acceleration Model

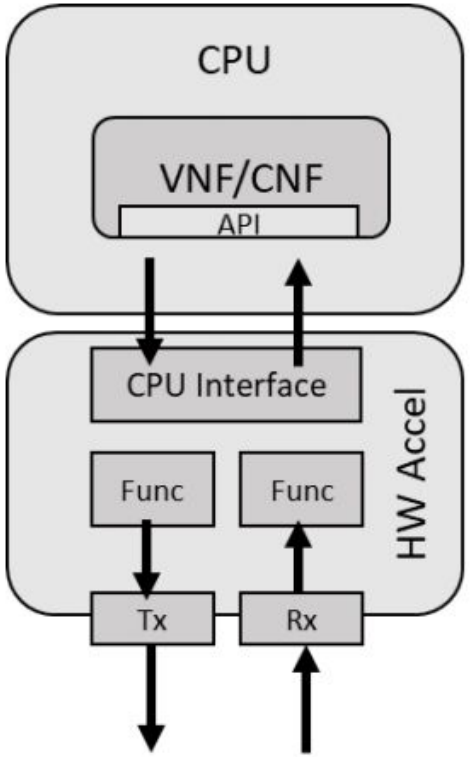

Look-aside Acceleration Model

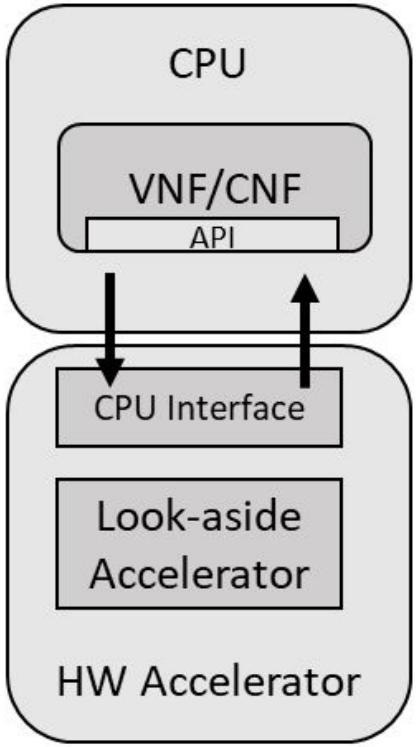

Intel Accelerator Boards and OpenShift Operators

Intel FPGA PAC N3000

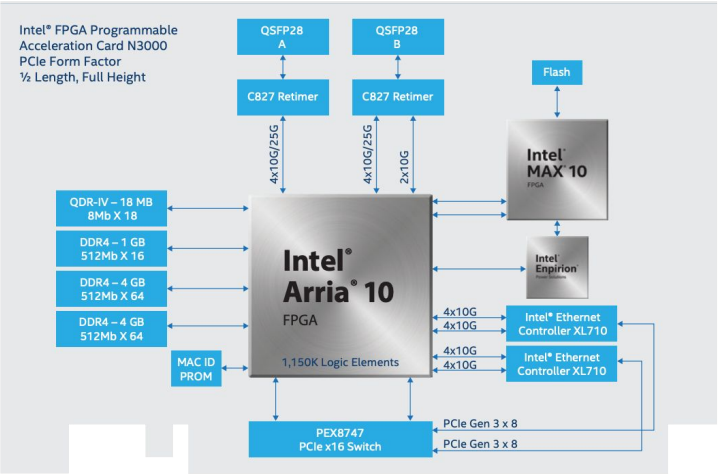

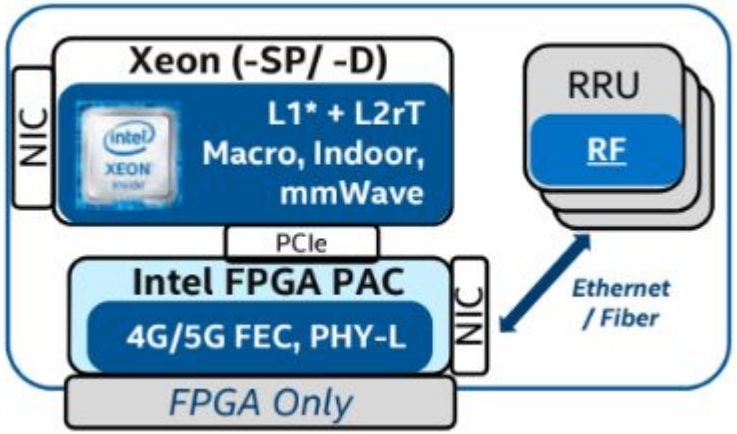

Intel eASIC Mount Bryce (ACC100)

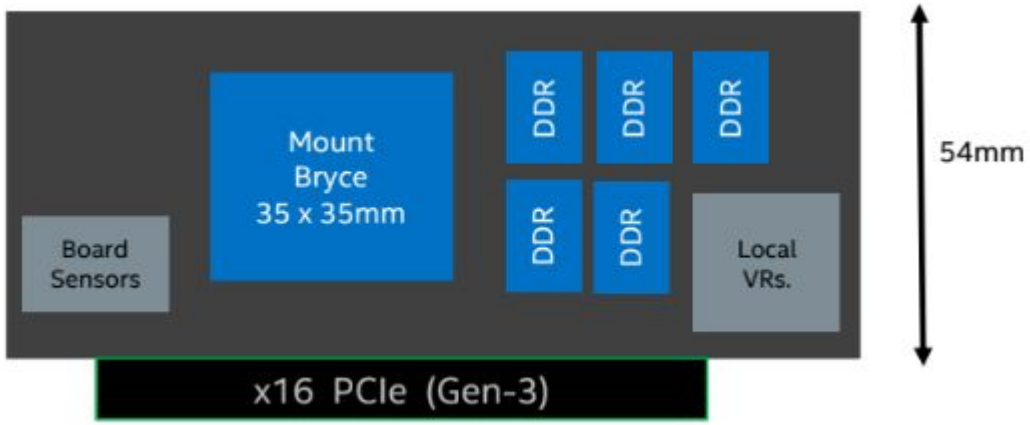

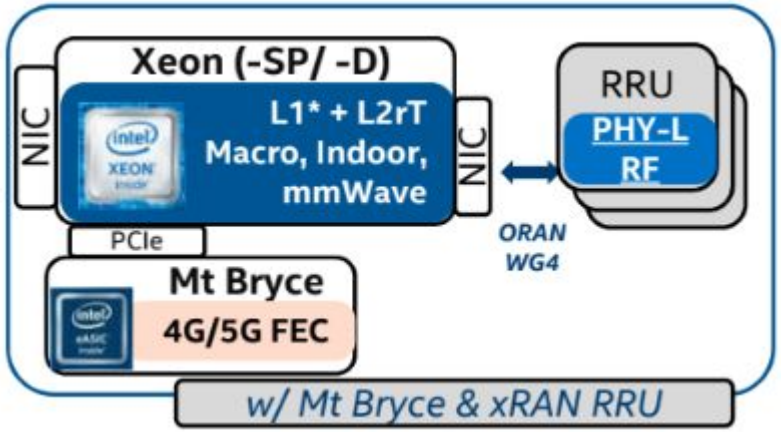

Red Hat

Intel Operator for Wireless FEC Accelerators

https://catalog.redhat.com/software/operators/search?q=Intel%20Operator&p=1

intel.

Intel Operator for Wireless FEC Accelerators

by Intel Corp

A collection of container images for the Intel Wireless FEC Accelerators based on the SDK Operator Framework. This is a complete package for automating installation, the configuration of SR-IOV FEC...

Red Hat Ecosystem Catalog | Platforms | Hardware | Software | Cloud & service providers

Home > Software > OpenShift operators > Browse products > Intel Operator for Wireless FEC Accelerators

Intel Operator for Wireless FEC Accelerators

Provided by Intel Corp

A collection of container images for management of SR-IOV Forward Error Correction (FEC) resources of the Intel's Wireless FEC Accelerators on OpenShift.

Deploy and use

Overview | Features & benefits | Documentation | Deploy & use | FAQs | Components | Get support

Red Hat

Telco Cloud Platform Requirements

Networking

- High Performance Data Plane
- Multiple Networking Interfaces
- High Throughput
- Low Latency
- Load Balancer
- Acceleration Hardware

Deployment & Scale

- Remote Distributed Deployments
- Single node and HA clusters
- Field upgradability
- Zero-touch-provisioning

Real-Time, Latency and Sync

- Time Synchronization
- Real-Time Determinism
- CPU Affinity and Isolation
- Topology Management
- Node Feature Discovery

OpenShift Edge Deployments

Low bandwidth or disconnected sites

Low bandwidth or disconnected sites.

~1,500 Single Node Servers provisioned and managed by one RHACM entity

Space constrained environments

Space constrained environments.

Up to 2,000 remote worker nodes per Cluster

Small footprint with high availability

Small footprint with high availability.

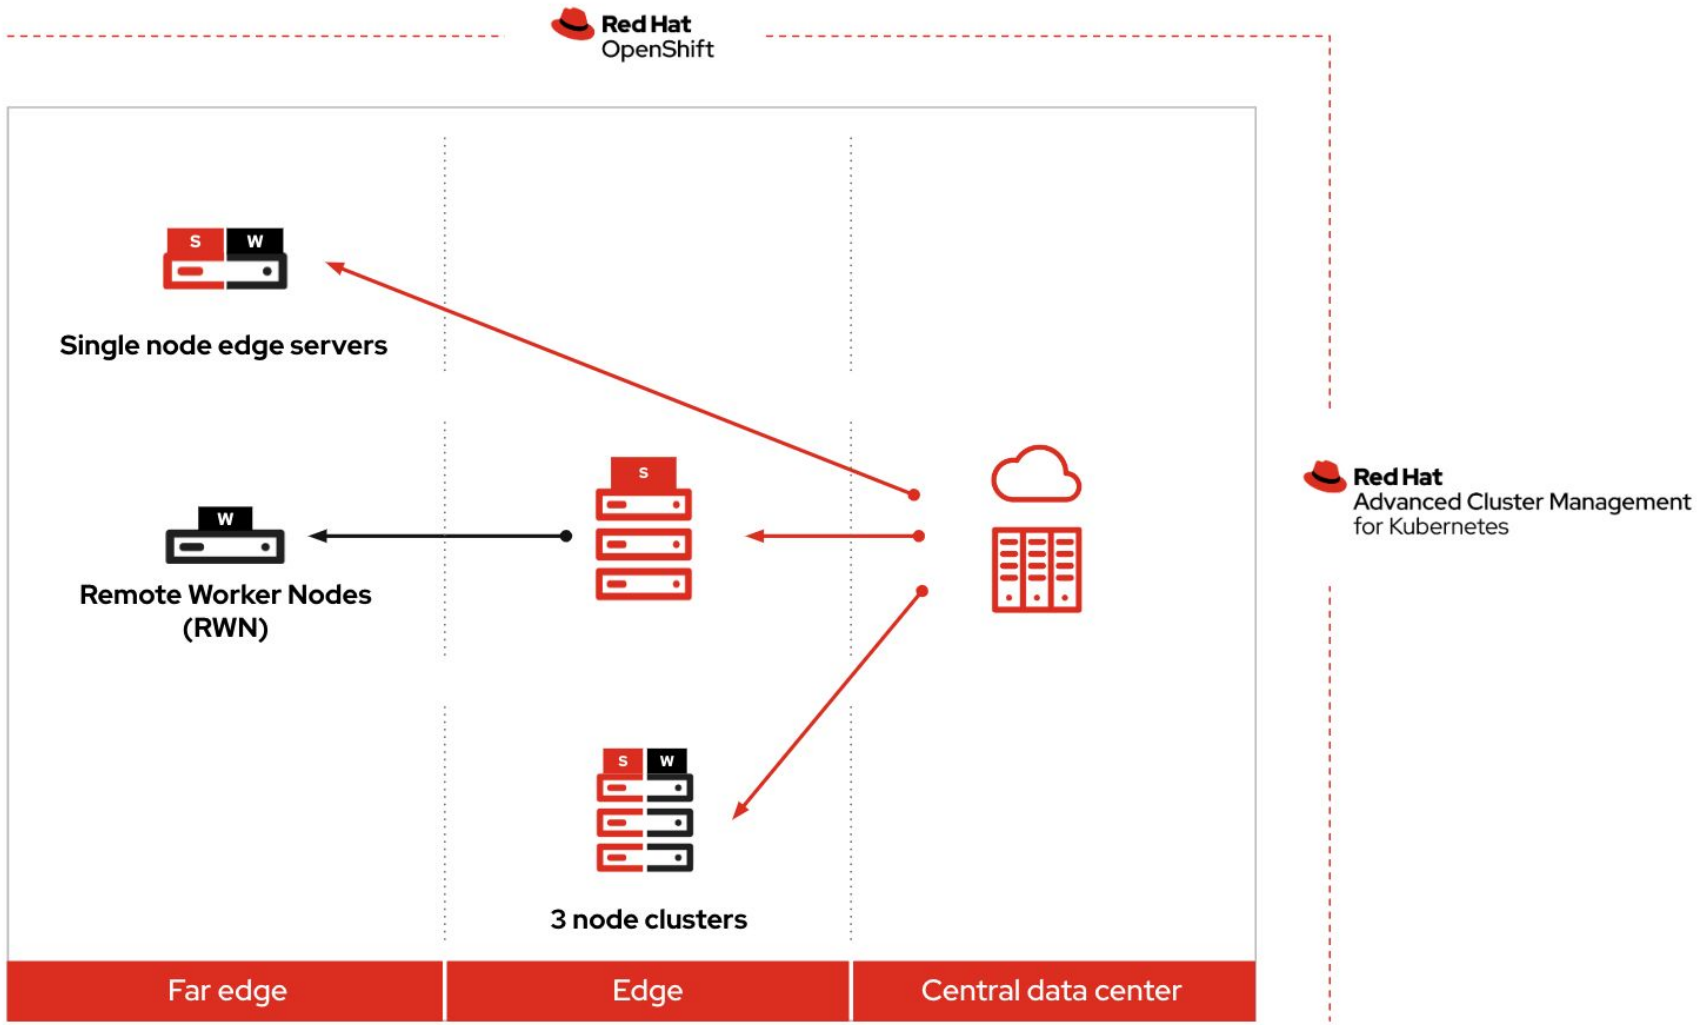

Cluster management and application deployment

Kubernetes node control

OpenShift from Core to the Edge

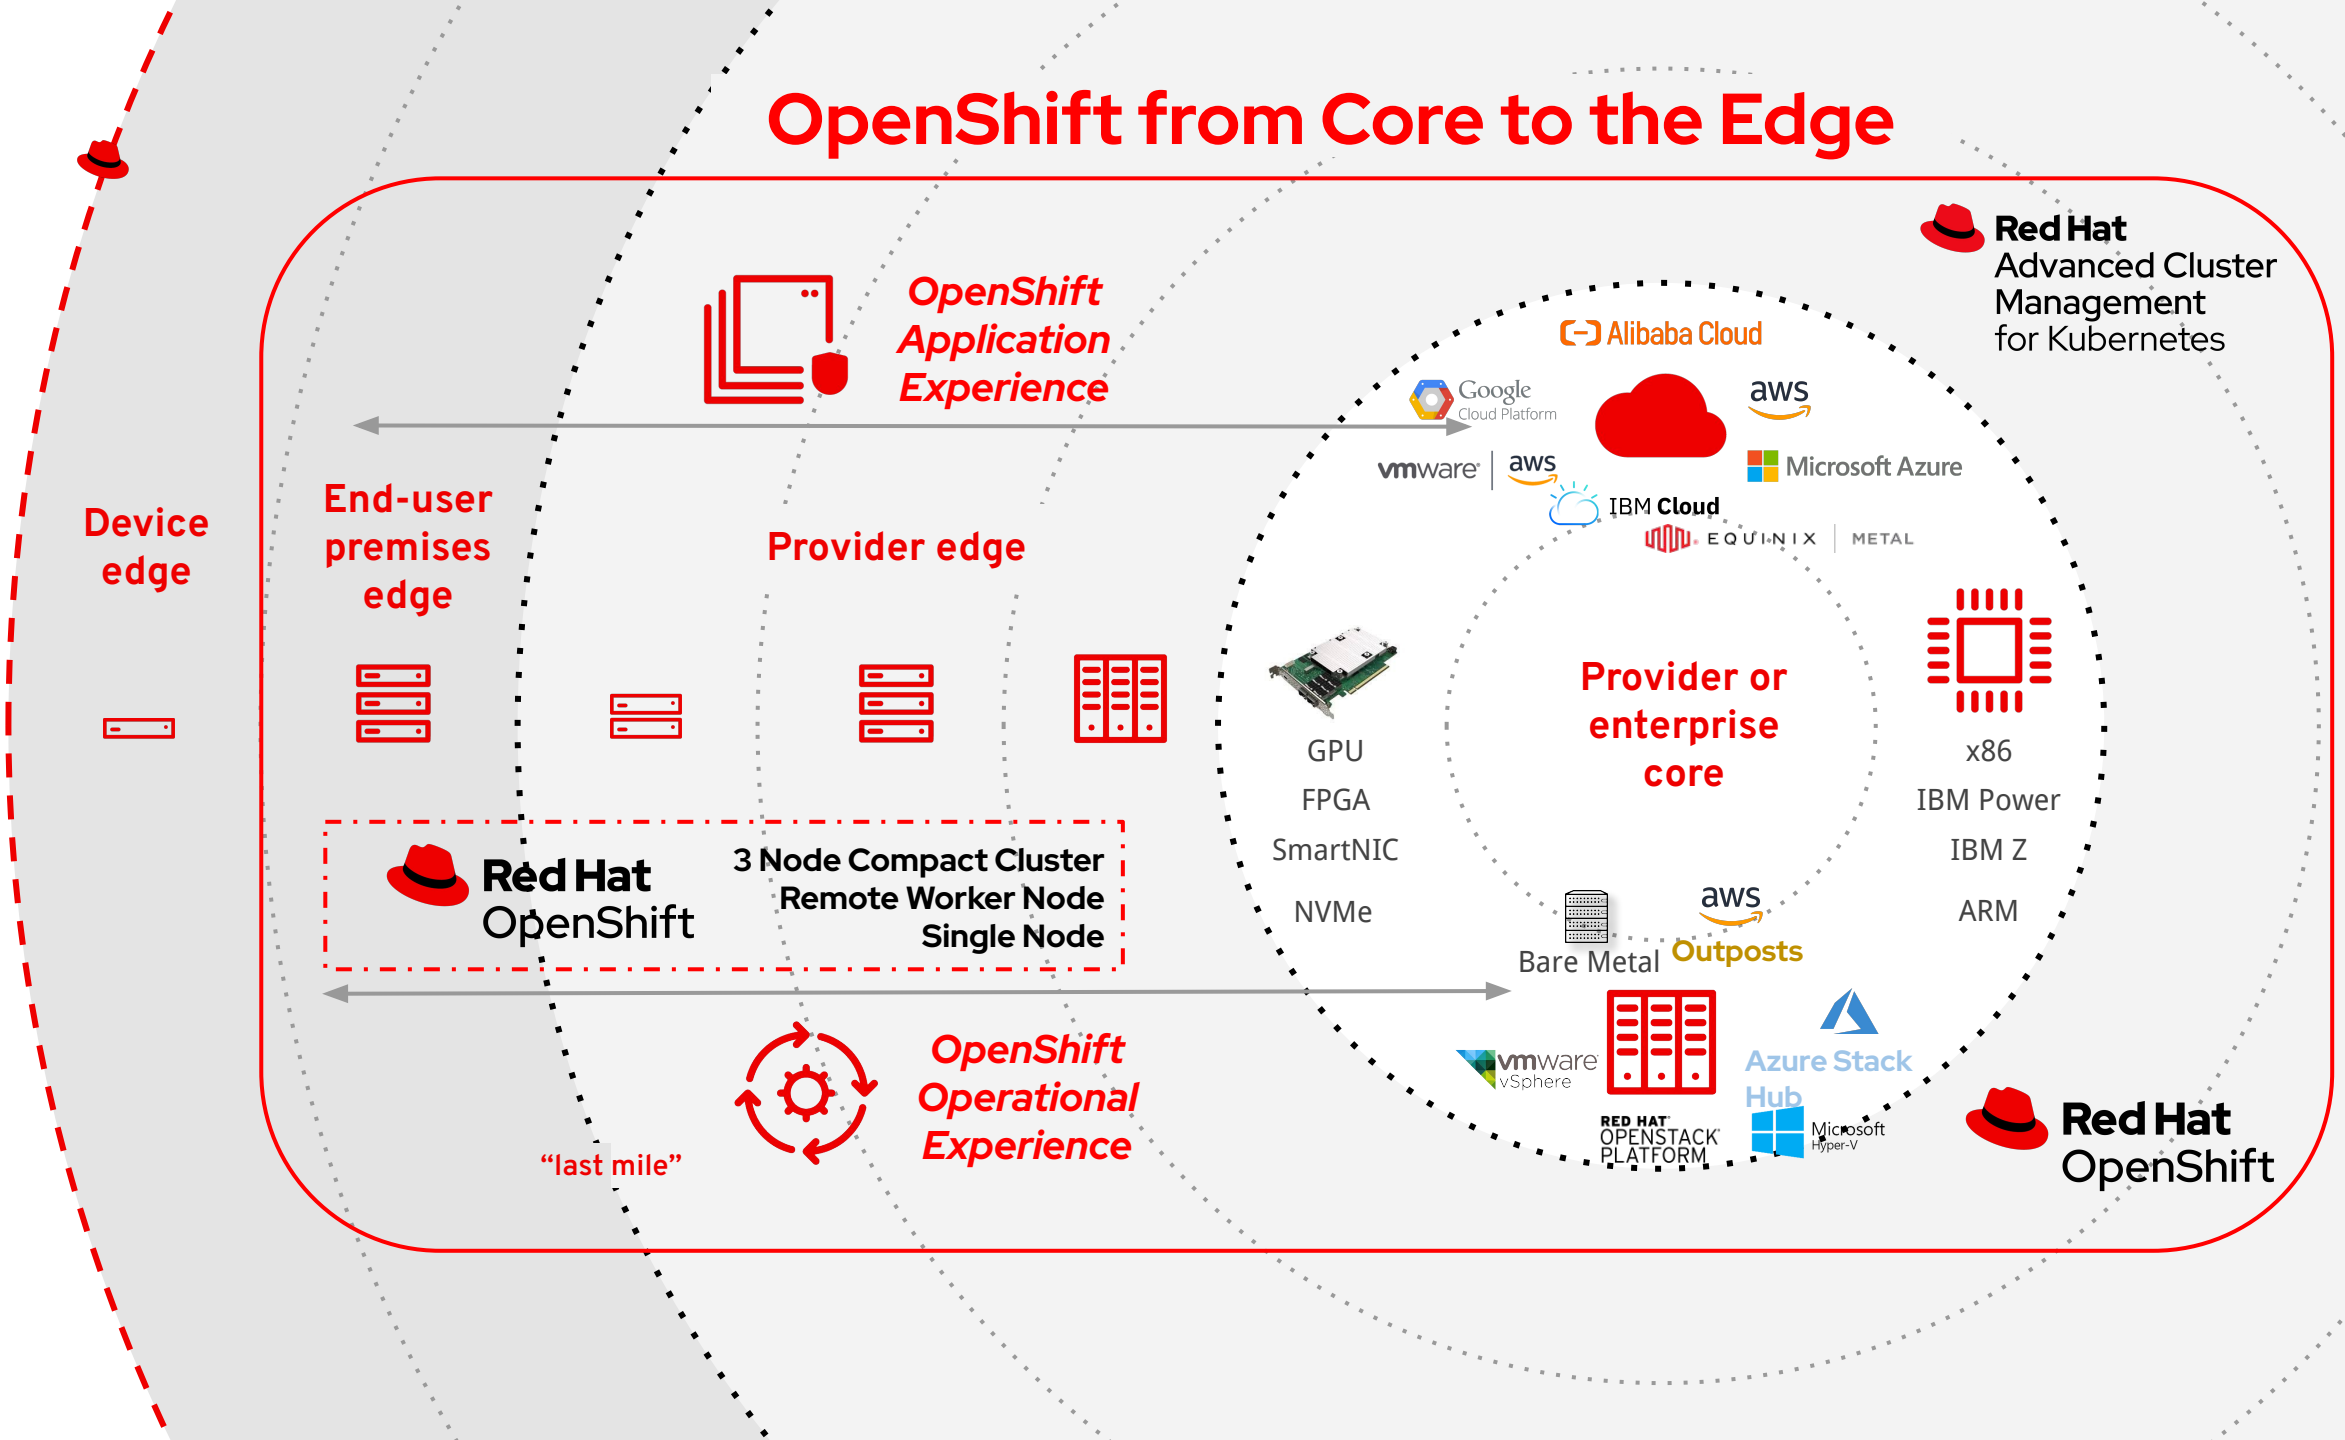

Manage clusters and deployments, all the way to the edge with ACM

CAPABILITIES

- Manage up to 2000 clusters with the RHACM hub
- Zero Touch Provisioning (Tech Preview): Leverage RHACM with Assisted Installer
- Hub-side Policy Templating: Reduce the number of policies for high scale management scenarios

BENEFITS

- Provides management capabilities for the entire OpenShift deployment to help ease the operational strain of managing a scaled out architecture
- Enables the scale needed for edge use cases, including SNO
- Ensures ease of deploying and provisioning of clusters at scale

What is a validated pattern?

Red Hat Edge's validated patterns are pre-defined edge computing configurations that bring together the Red Hat portfolio and technology ecosystem to help you stand up your edge architecture faster.

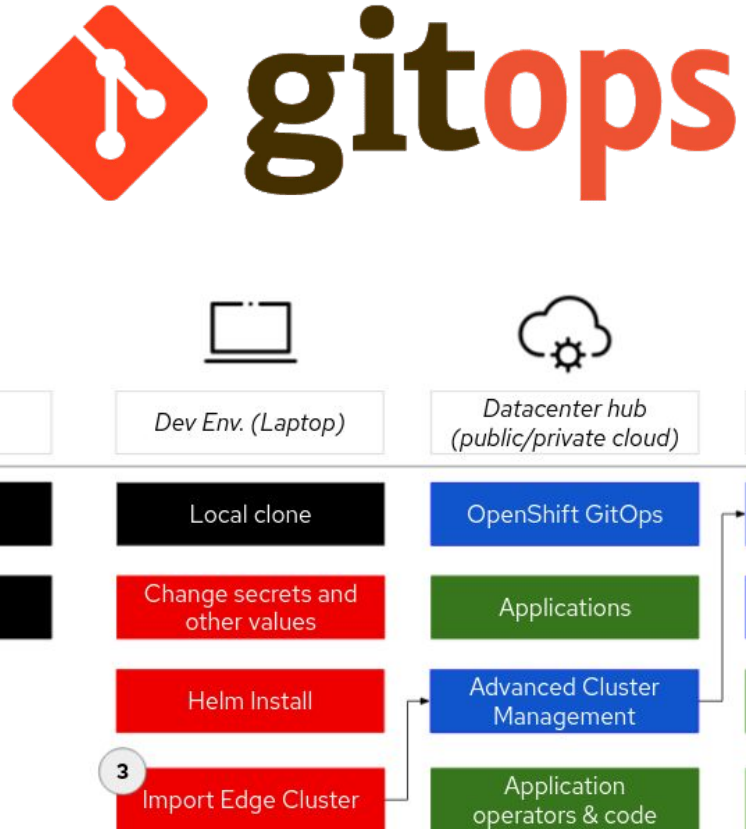

Red Hat

Git repository example: https://github.com/hybrid-cloud-patterns/industrial-edge

Validated Patterns : Simplifying the creation of edge stacks

Bringing the Red Hat portfolio and ecosystem together - from services to the infrastructure

Blueprint as code

Go beyond documentation using GitOps process to simplify deployment

From POC to production

Ensure your teams are ready to operate at scale

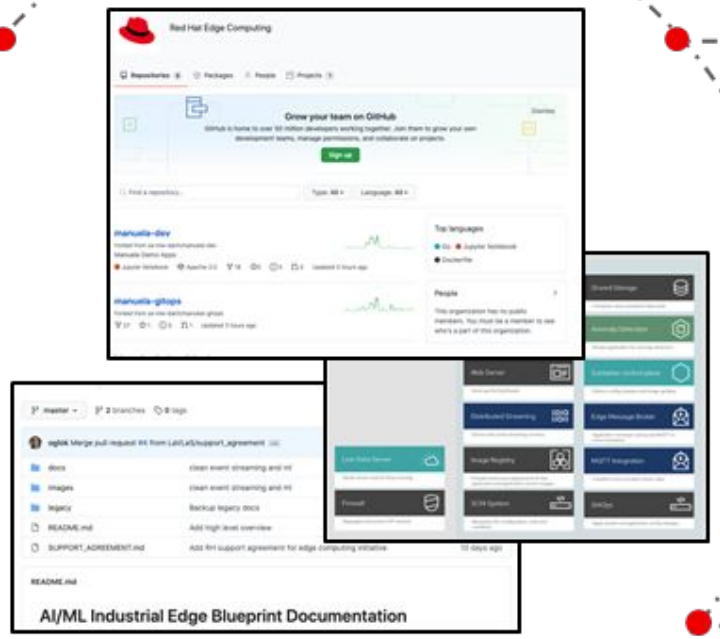

Highly reproducible

So that you can scale out your deployments with consistency

Open for collaboration

Anyone can suggest improvements, contribute to it

Telco Cloud Platform Requirements

Networking

- High Performance Data Plane
- Multiple Networking Interfaces
- High Throughput
- Low Latency
- Load Balancer
- Acceleration Hardware

Deployment & Scale

- Remote Distributed Deployments
- Single node and HA clusters
- **Field upgradability**
- **Zero-touch-provisioning**

Real-Time, Latency and Sync

- Time Synchronization
- Real-Time Determinism
- CPU Affinity and Isolation
- Topology Management
- Node Feature Discovery

Supporting application modernization

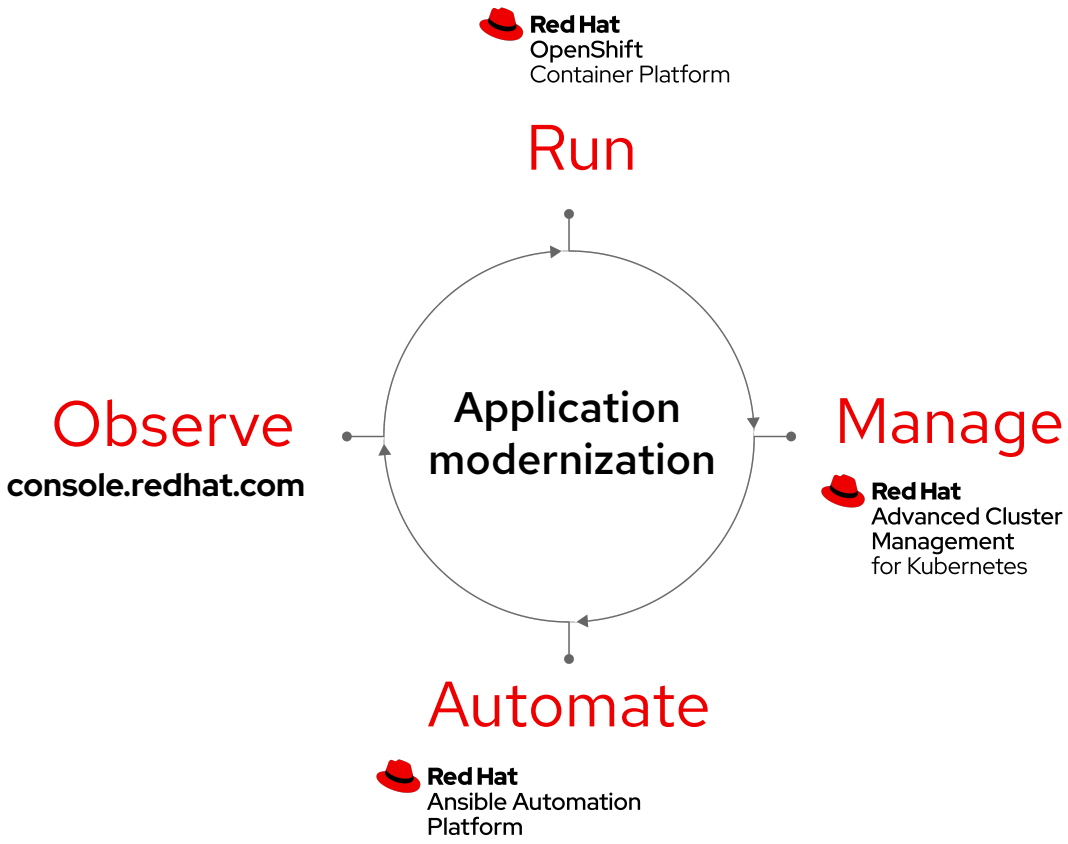

Run
- Cluster services: metrics, chargeback, registry, logging
- Advanced infrastructure functionality
- Dev services: dev tools, automated builds, CI/CD, IDE

Manage
- Multicluster and Kubernetes lifecycle management
- Policy-based governance, risk, and compliance
- Application lifecycle management
- ODF integration for Disaster Recovery of stateful apps

Automate
- Configuration management
- Workflow orchestration
- Network and security automation
- Automation analytics
- Certified content
- Automation services catalog

Observe
- Red Hat Insights for OpenShift
- Connected customer experience
- Subscription watch

Red Hat

Red Hat

Advanced Cluster Management for Kubernetes

Advanced Cluster Management runs on Red Hat OpenShift.

Hub Cluster

The OpenShift cluster where Advanced Cluster Management is installed is called the hub cluster. It is the central hub for creating, managing, and monitoring Kubernetes clusters.

Managed Cluster

A Kubernetes cluster managed by Advanced Cluster Management is called a managed cluster.

With Advanced Cluster Management, you can view managed cluster health, deploy applications to the managed clusters, and enforce policies on the managed clusters. You can easily upgrade as well as destroy any managed clusters created from Advanced Cluster Management.

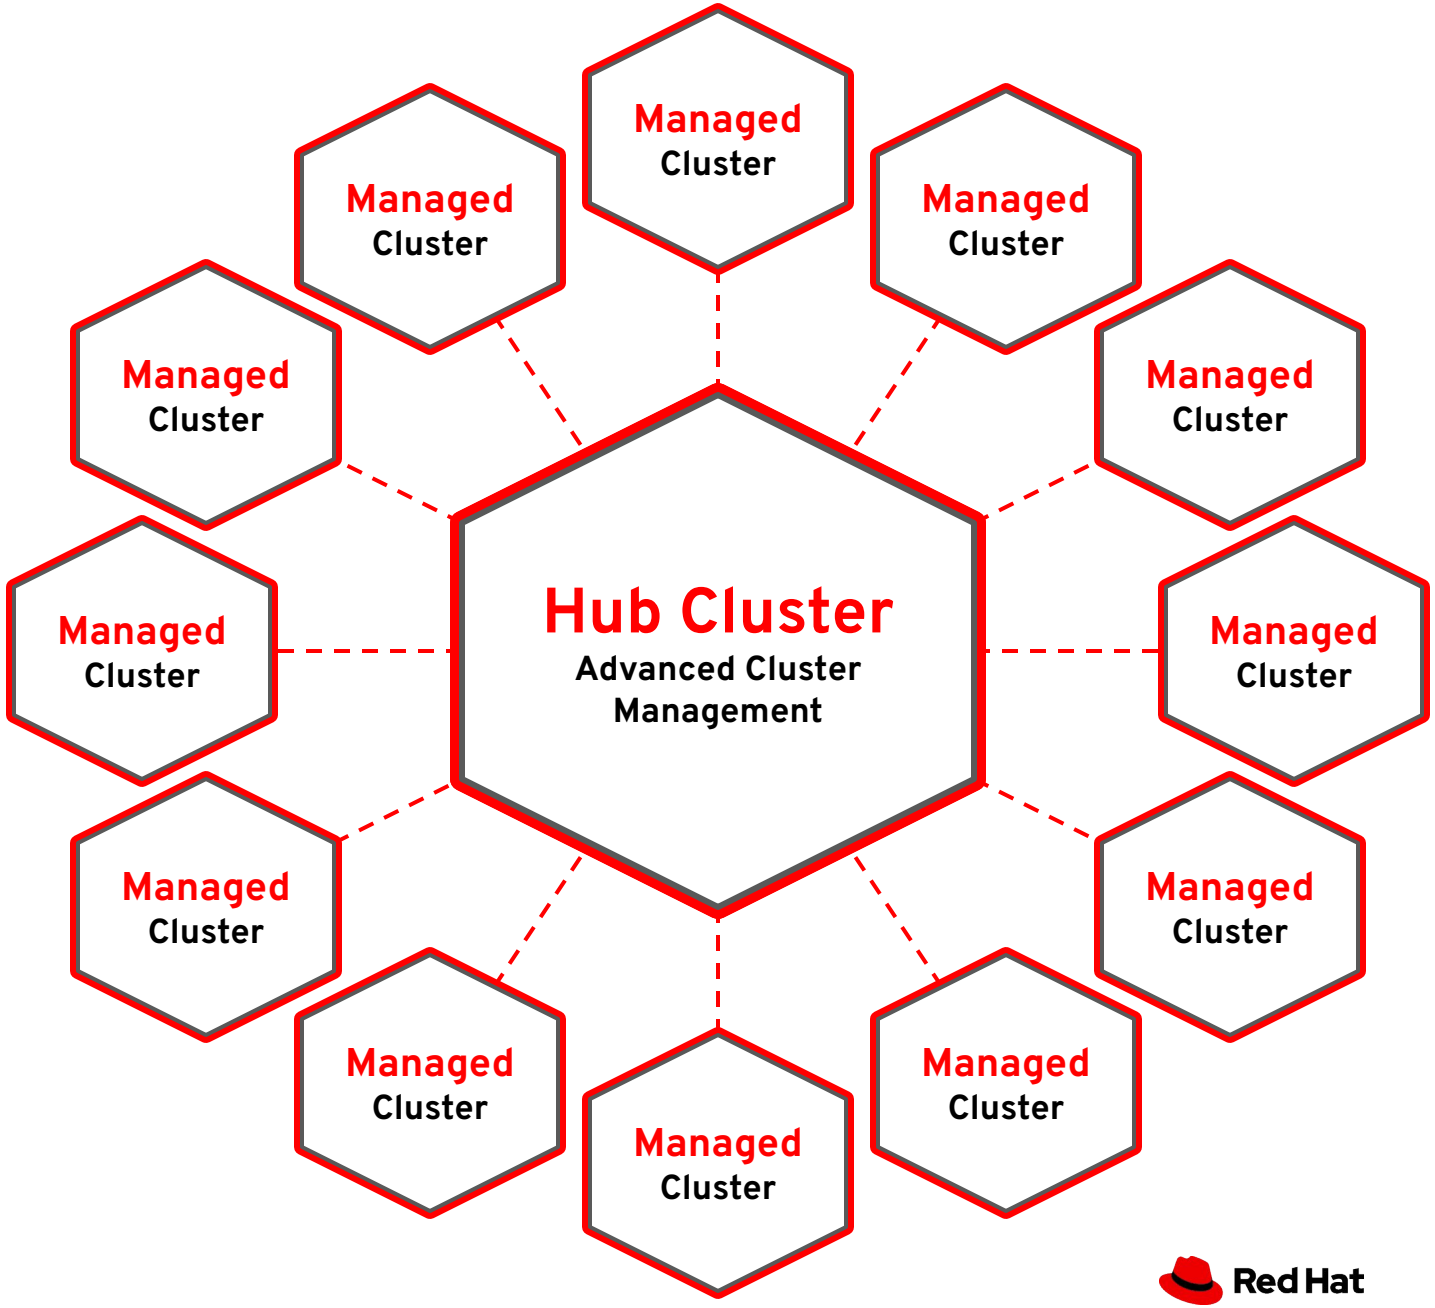

Red Hat

OpenShift 4: The NFV Edge Platform

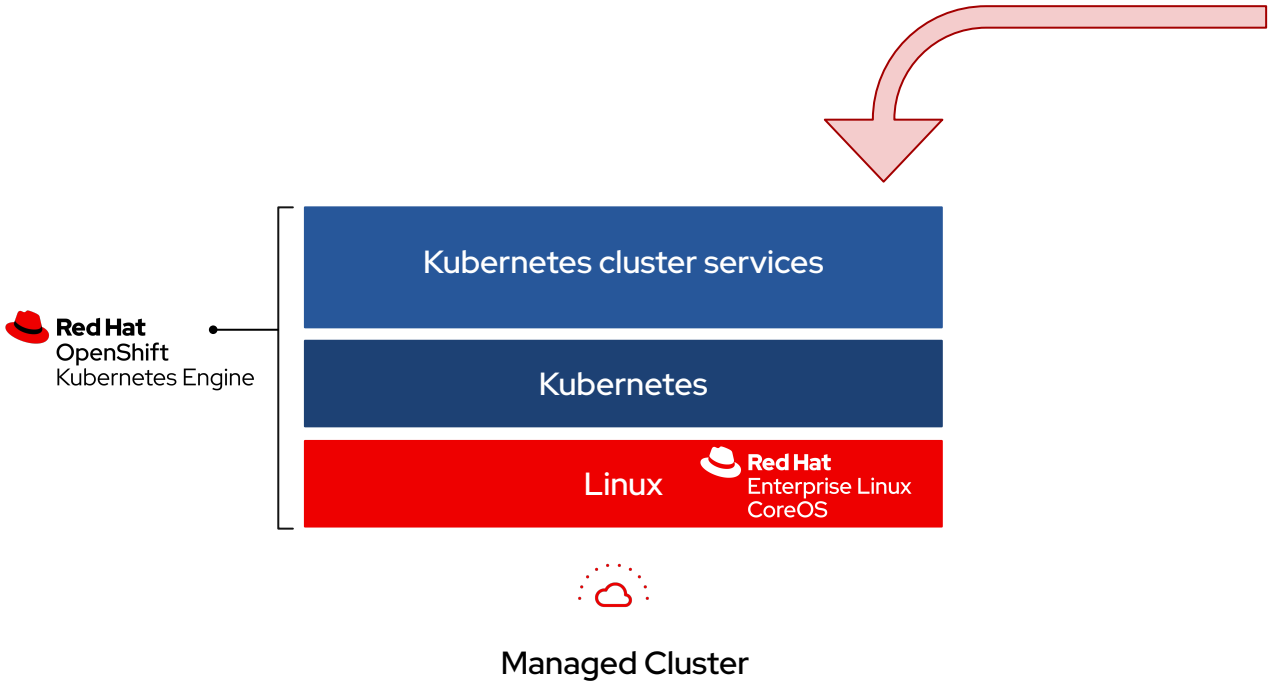

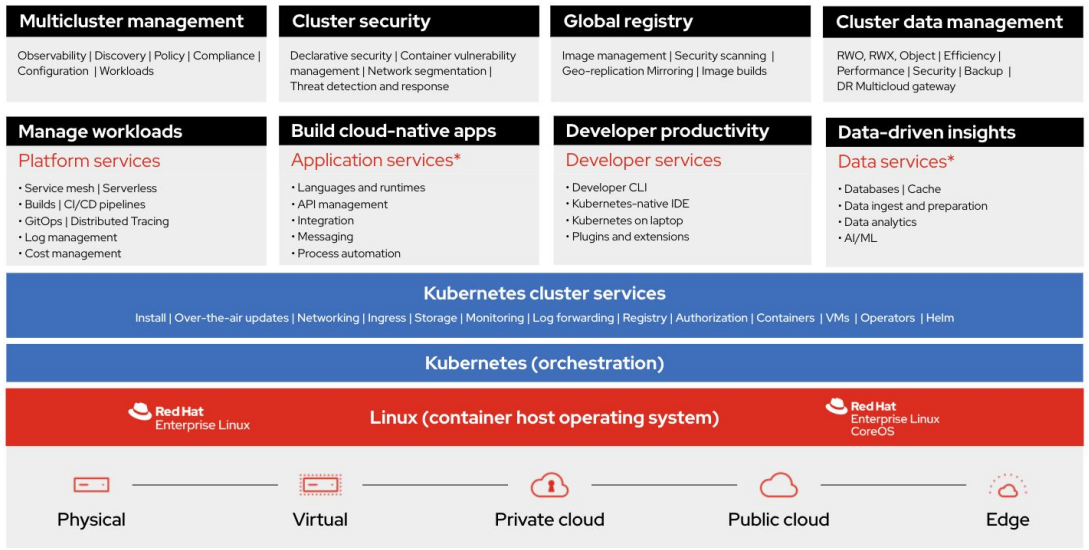

Hub Cluster

Red Hat

ZTP - High Level Provisioning Flow (GitOps)

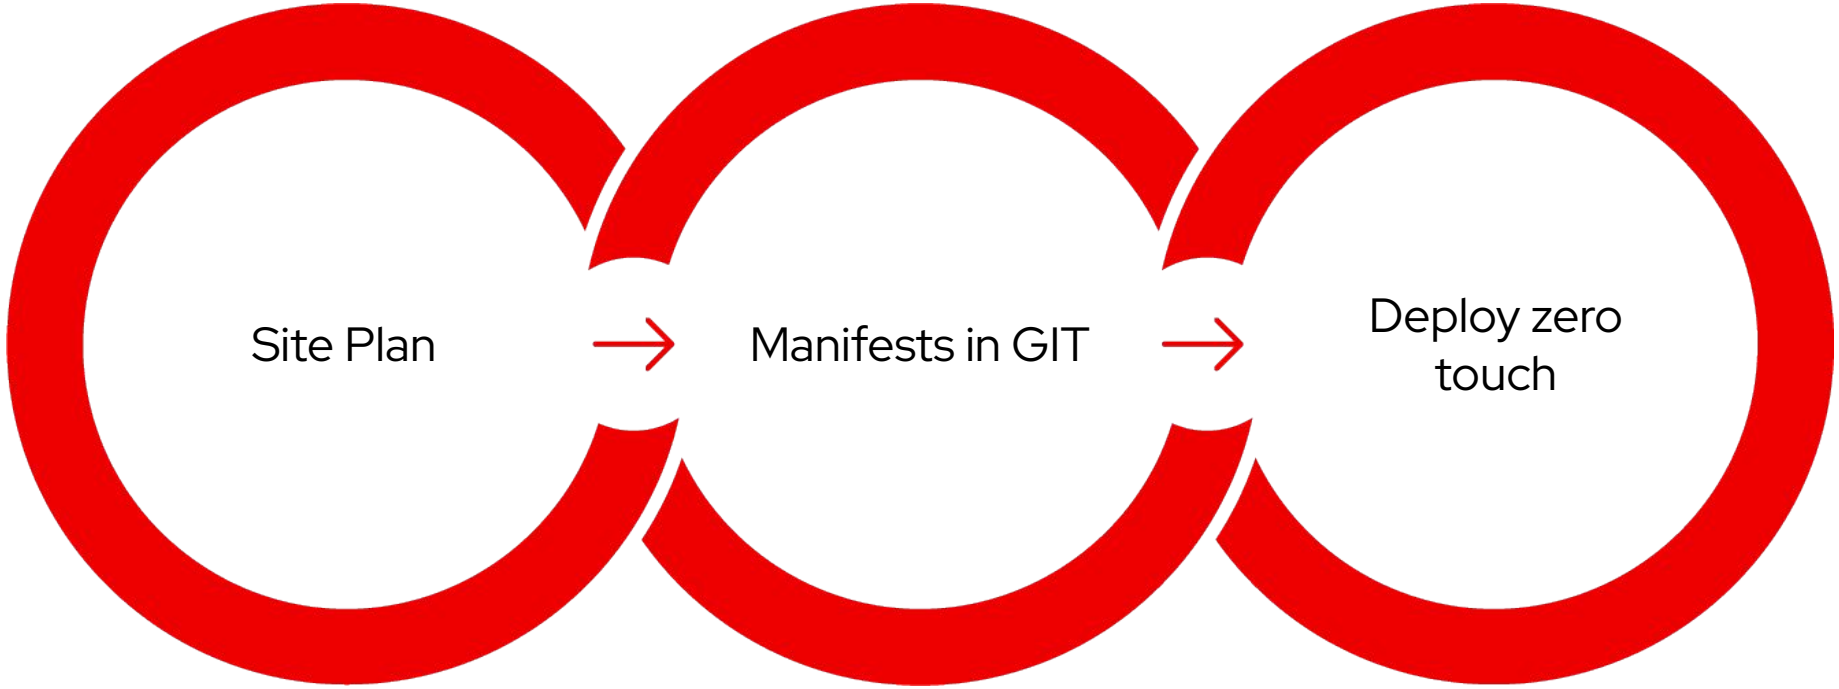

Red Hat

Site Plan

All managed Clusters can be added to site plan, despite connection / location limitations

GITOPS tool or Red Hat Quay

Git: will host the ingredients of the managed cluster code information "Single Point of Truth"
Quay: A private container registry for distributed, secure development

Automated cluster creation with ZTP (Zero Touch provisioning)

OS configuration is stored and applied across the cluster via the Machine Config Operator (Inside Machine Configuration Pool).

- Standard k8s YAML/JSON manifests
- Desired state of nodes is checked/fixed regularly
- Can be paused to suspend operations

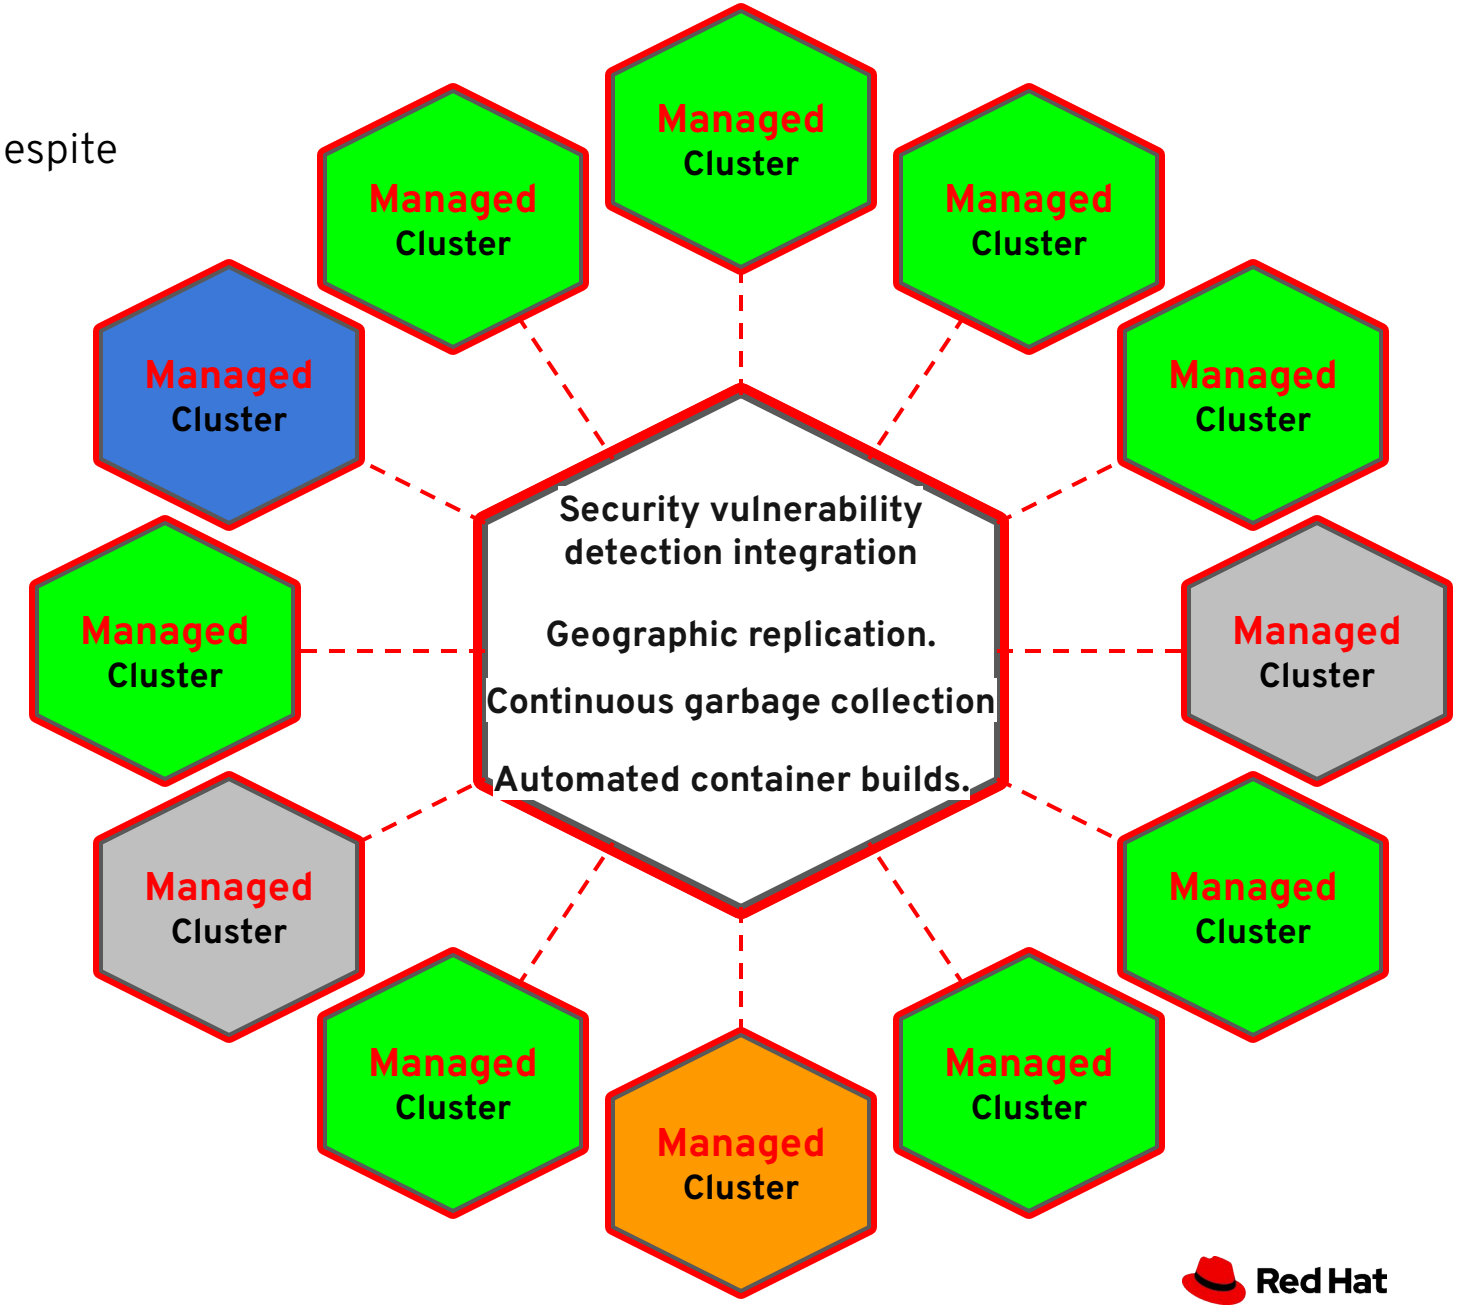

Red Hat

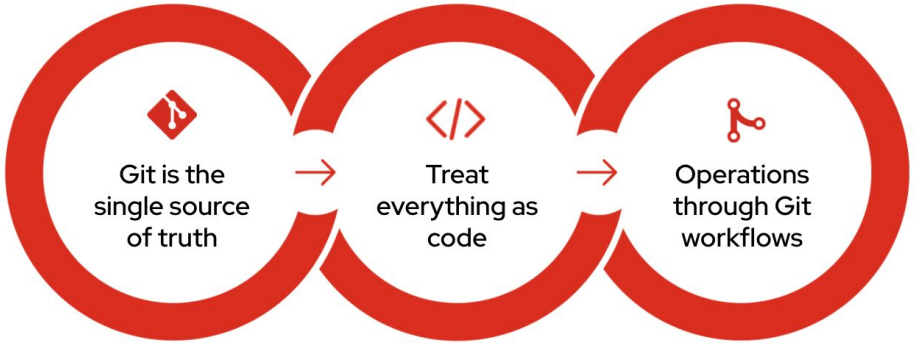
Git is the single source of truth
Treat everything as code
Operations through Git workflows

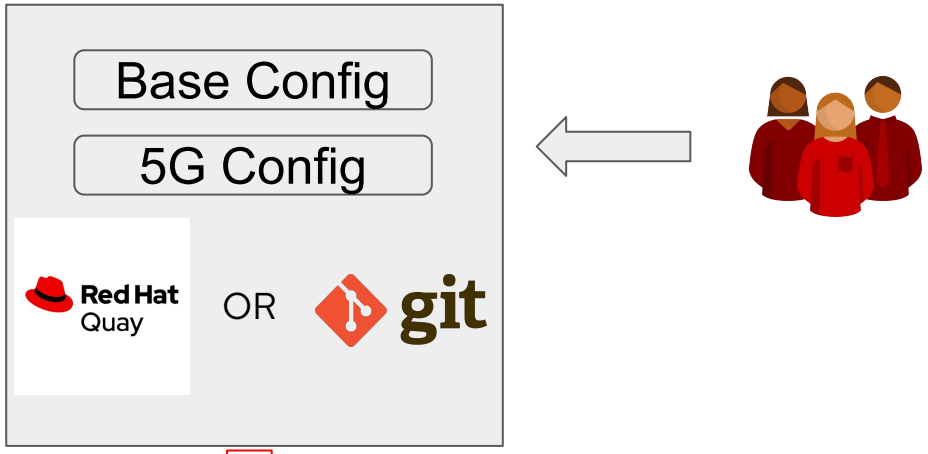
Base Config
5G Config
Red Hat Quay
OR
git

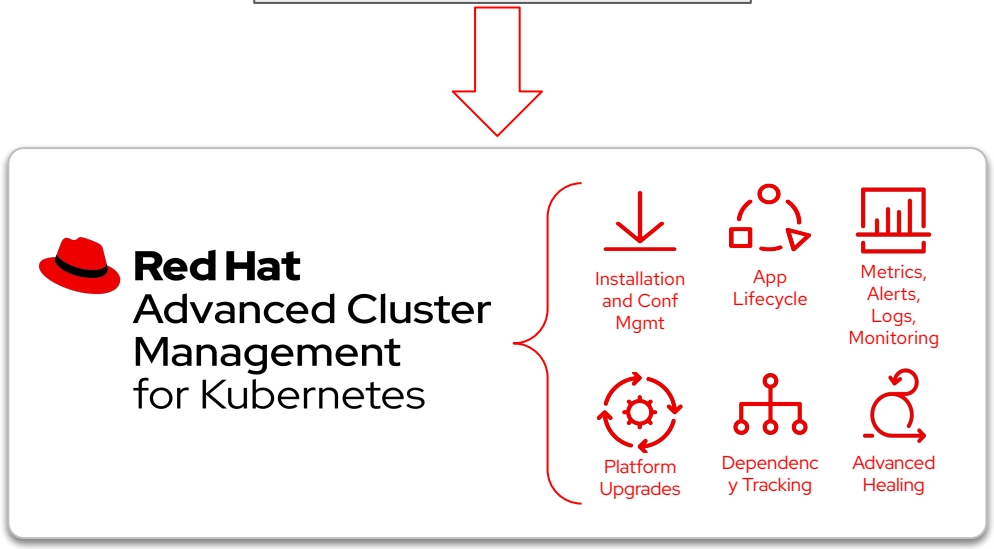
Red Hat Advanced Cluster Management for Kubernetes
Installation and Conf Mgmt
App Lifecycle
Metrics, Alerts, Logs, Monitoring
Platform Upgrades
Dependency Tracking
Advanced Healing

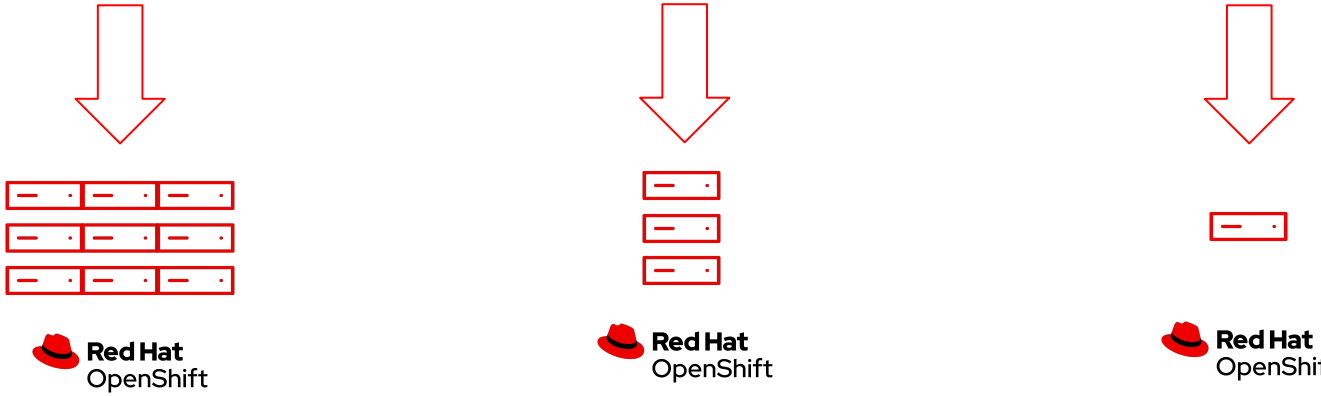
Red Hat OpenShift
Red Hat OpenShift
Red Hat OpenShift

Red Hat

Our Security Approach Does Not Change

Continuous Integration(CI) & Continuous Delivery (CD)

Security must be continuous

And integrated throughout the lifecycle

Identify security requirements & governance models

FedRAMP

NIST National Institute of Standards and Technology U.S. Department of Commerce

...

CIS Center for Internet Security

National Cyber Security Centre

PCI Security Standards Council

HIPAA Health Insurance Portability & Accountability Act

Revise, update, remediate as the landscape changes

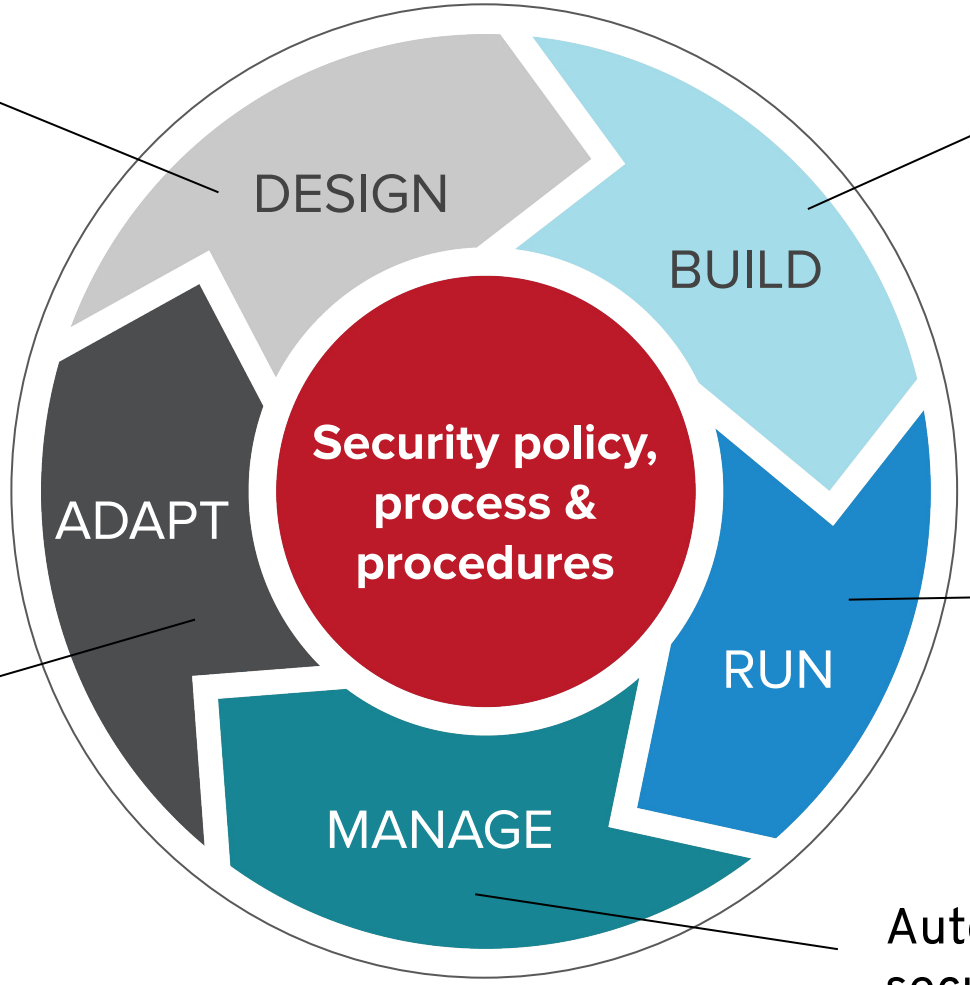

Built-in from the start; not bolted-on

Deploy to trusted platforms with enhanced security capabilities

Automate systems for security & compliance

CONTROL
Application Security

DEFEND
Infrastructure

EXTEND

Red Hat

Edge Security: Influencing Factors

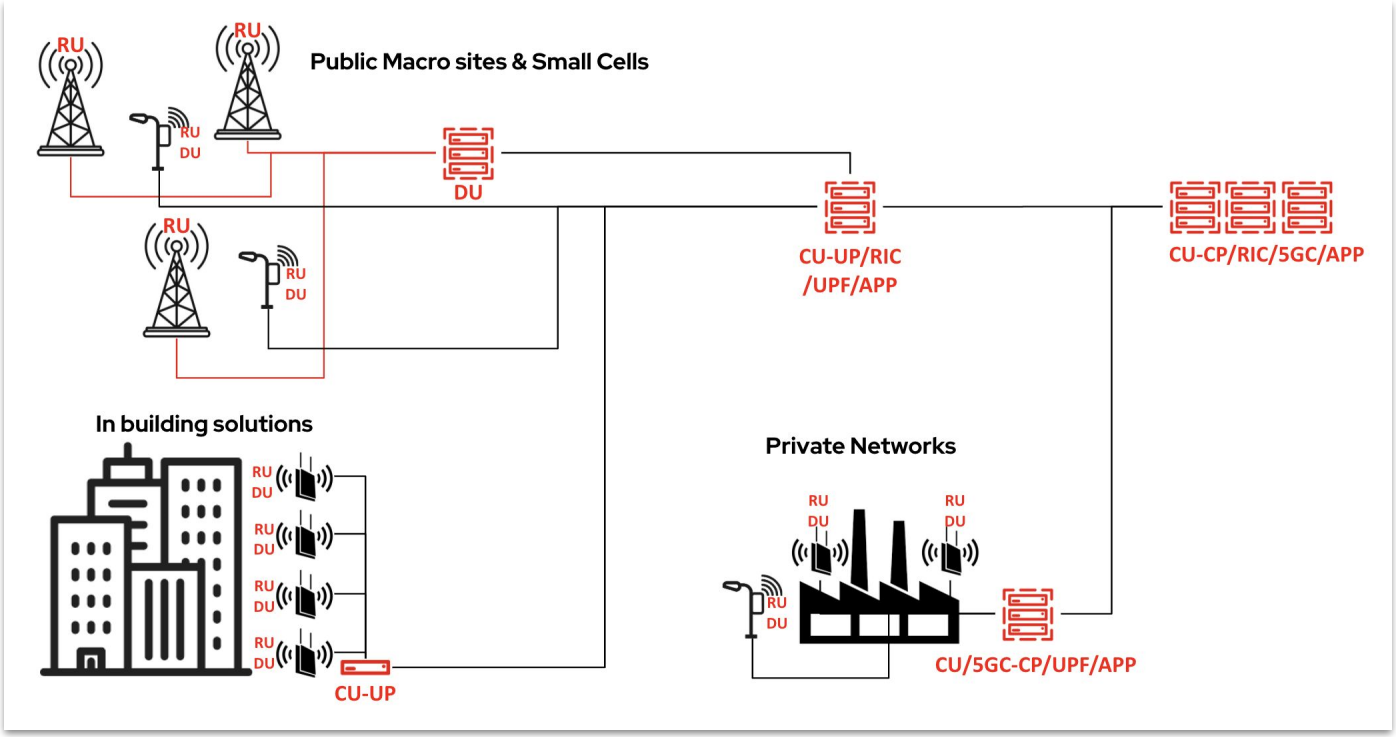

Reduced Physical Security

- No trust by default
- Need for more confidentiality & integrity

Multitenancy

- Workload isolation
- Software Supply Chain risk

Reduced Footprint

- Physical & Resources wise
- Less security implementation options

  e.g. security zones,...

Remote Management

- Risk of traffic interception
- Hardening of remote management tools

Red Hat

Security in Open RAN

Red Hat Products Solutions Learning & support Resources Red Hat & open source

ANALYST MATERIAL

Security in Open RAN

Open RAN is an open radio access network (RAN) architecture standardized by the O-RAN Alliance based on 3GPP and other standards. 5G and Open RAN enable new capabilities and control points that allow suppliers, test equipment manufacturers, wireless carriers, and network operators to assess, mitigate and manage security risks efficiently. This paper details how O-RAN enables operators with full visibility and control of their network's end-to-end security.

Last Updated: February 4, 2021

Download

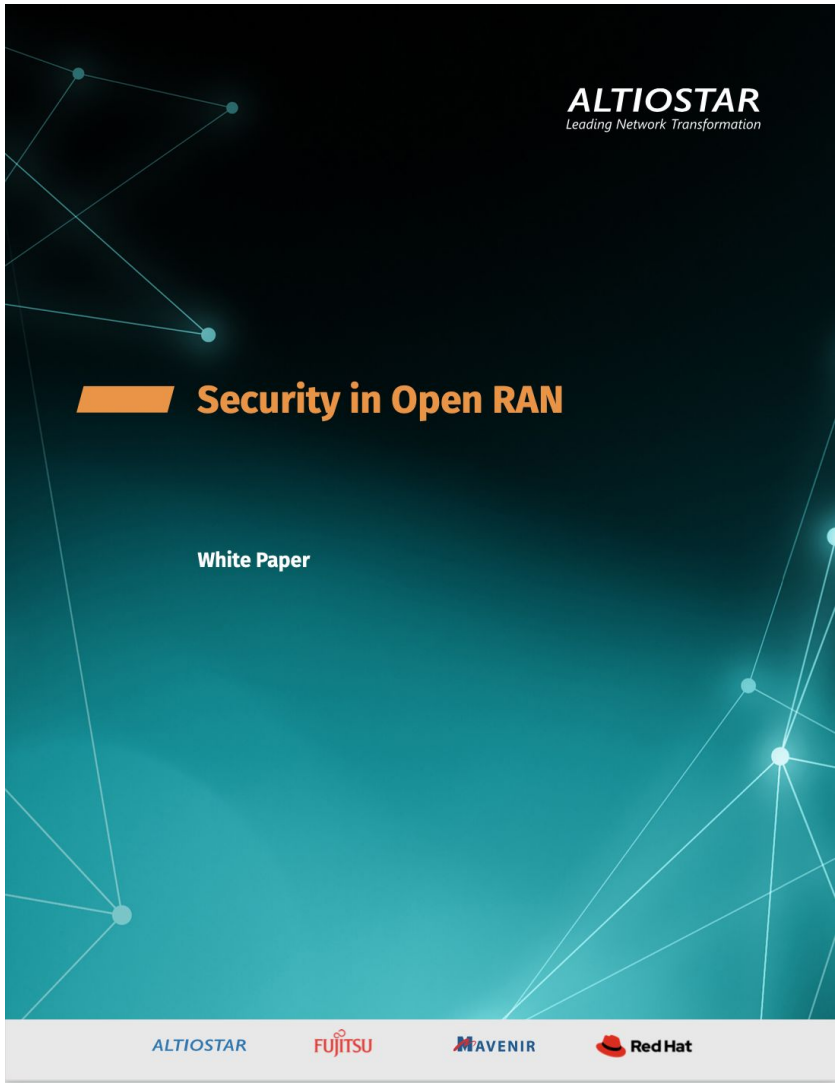

https://www.redhat.com/en/resources/open-ran-security-analyst-material

Red Hat

OPENSHIFT REQUIRES TO SECURE A FULL SW STACK

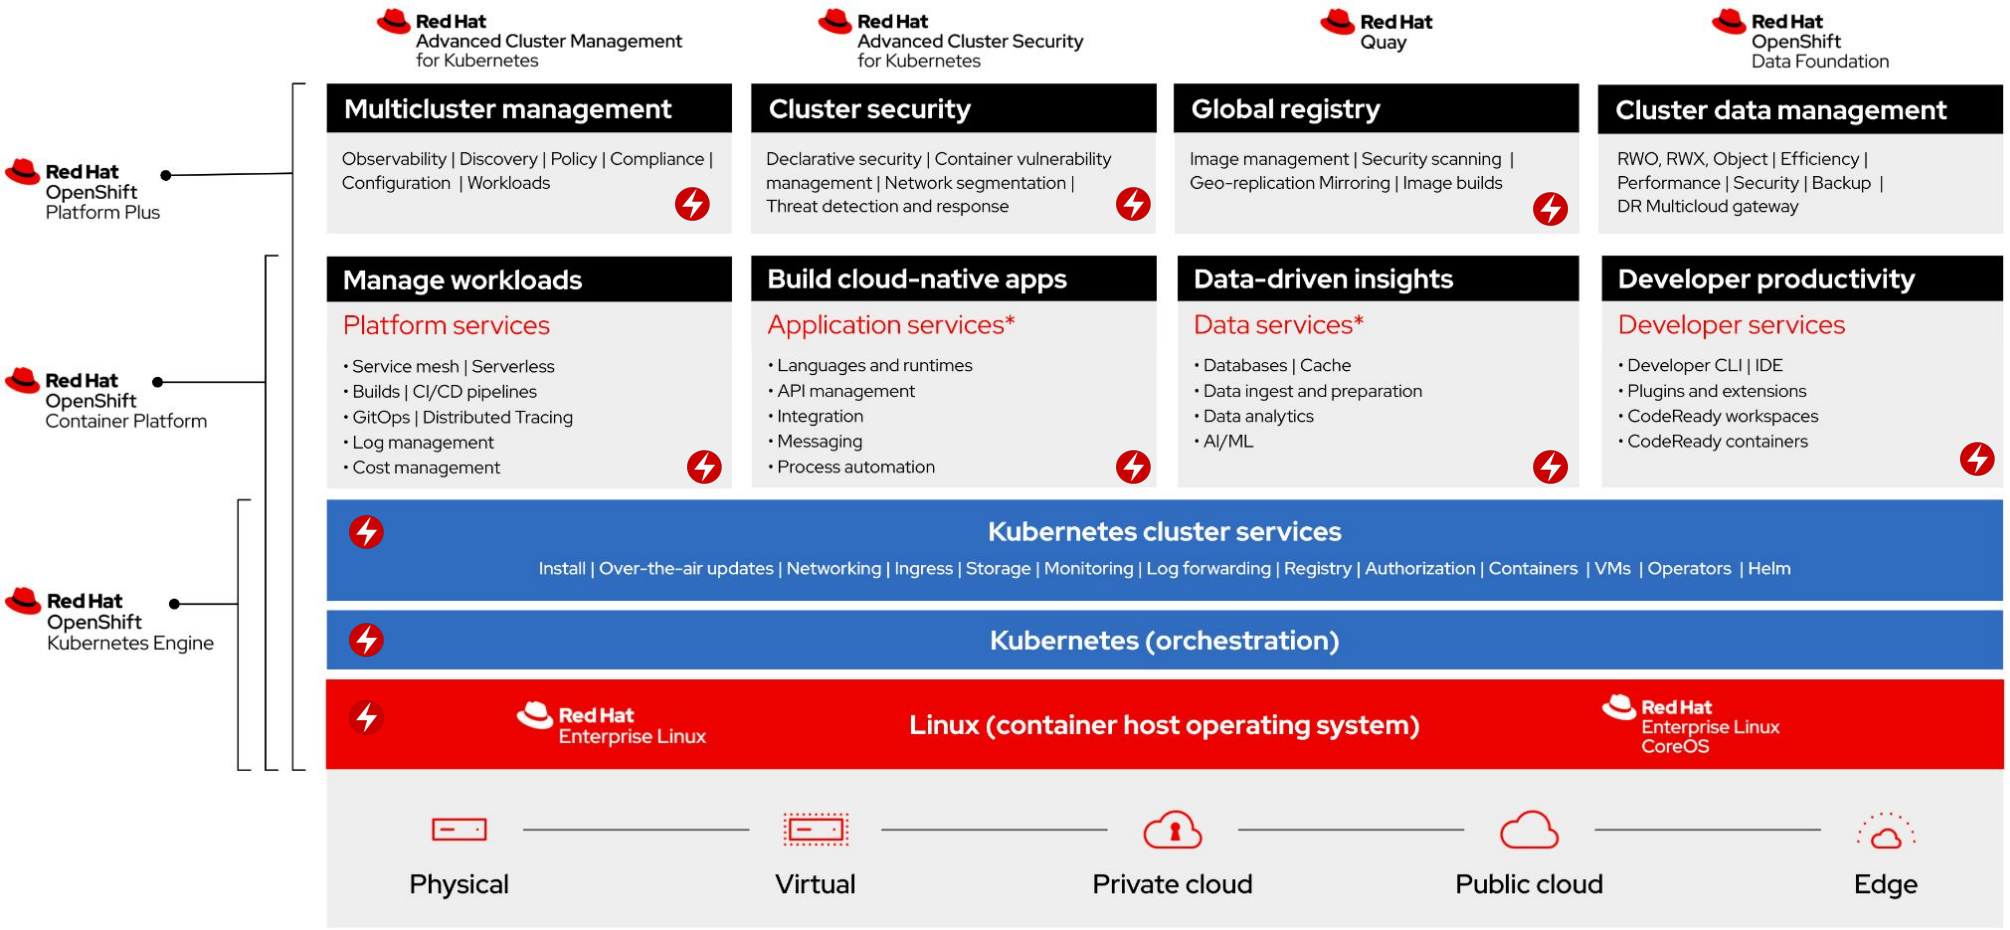

Multi-clusters

Multi-tenant

Over The Air updates, Automated Compliance

Container specific Host OS & CRI-O

Red Hat Security Supply chain

Support SW/HW integration (TPM, Vaults,...)

One platform with a unified operational model and consistent experience wherever it is deployed

* Red Hat OpenShift® includes supported runtimes for popular languages/frameworks/databases. Additional capabilities listed are from the Red Hat Application and Data Services portfolio.

Red Hat
CONFIDENTIAL CLARO - RED HAT

Why Red Hat to secure 5G ?

1	**Containers are Linux**	Securing containers require deep Linux expertise
2	**Container Security**	Rootless containers, SELinux, Security Context Constraint (SCC),...
3	**Container OS**	Red Hat CoreOS for minimal surface of attack & easy LCM
4	**Container Confidentiality**	Different encryption mechanisms supported
5	**Product Security**	Track-record product support for security bugfix & backports
6	**Compliance at scale**	Compliance Operator & RH ACM for multi-cluster compliance
7	**Security Certifications**	FIPS 140-2, FedRAMP,...
8	**Security Automation**	Ansible Security, RH ACM, DevSecOps pipelines,...
9	**Platform Integrity**	Secure Boot, File Integrity Operator, Container scanning,...
10	**Security Partner Ecosystem**	Rich ecosystem of security partners

Red Hat

Red Hat Telco Ecosystem

Red Hat

Red Hat's Ecosystem Strategy

Bring together vendors, technologies, and approaches across industries

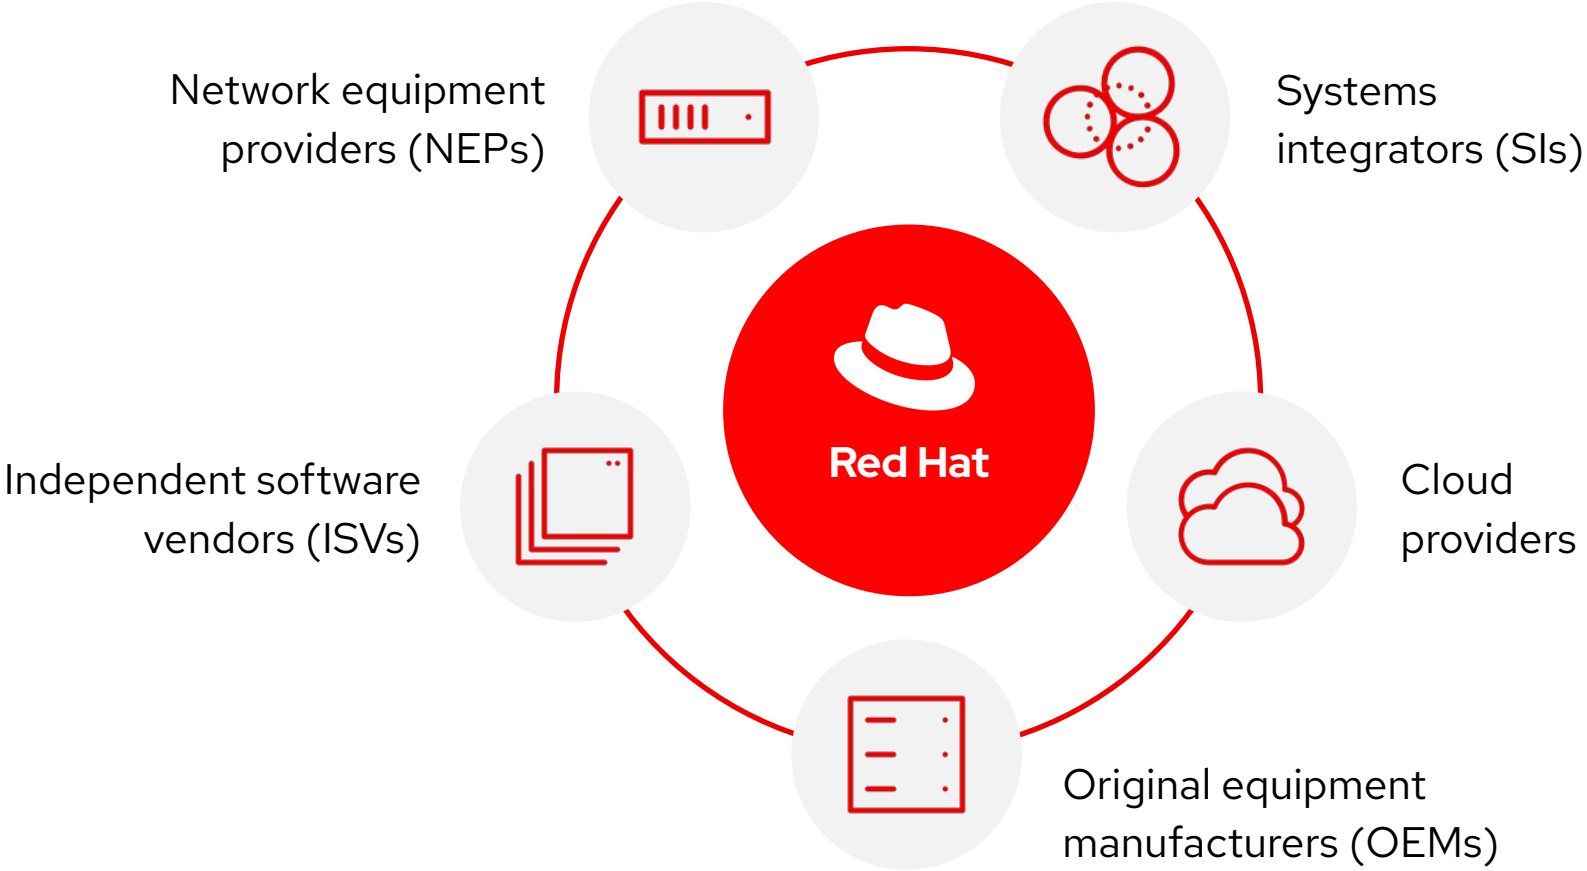

Strategy & Investment

Engineering

250 New dedicated resources for telco engineering

Partner Enablement

$$$ Investment in Joint Partner Labs

Labs & CapEx

$$$ Telco Solution Lab

Red Hat is committed to Telco market and as part of that, is **building dedicated lab** to deal with Telco deployment-specific configurations for Edge/RAN and Core.

Partner Categorization:

- **E2E** management and orchestration, service assurance, OSS/BSS, ...
- **Core SW** partners; closed source / open source; single/multi-vendor
- **RAN SW** partners (DUs, CUs, management & control, near-RT RIC, xAPPs, ...)
- **Radio Unit** (RU) partners to bring in their units to the lab
  (Walk-in Faraday shielded room - non-anechoic, focus on containment to stay within FCC limits & basic OTA testing only)
- **SW Technology Partners** which have parts of solution for integration, silicon vendors ref implementations & SW vendors with specific parts of the stack
- **HW Partners (NEs/servers/storage)**: switches/routers, synchronization gear, servers/storage, RUs/DUs
- **HW Partners (silicon and boards)**: Intel, ARM, NVIDIA, Xilinx, Silicom, Broadcom...

Red Hat

Red Hat Service Provider and Open Source

Mature the ecosystem

Engage in service provider specific projects and industry initiatives

Package the pieces

Secure technical capabilities for service providers in open source projects

Build the platform

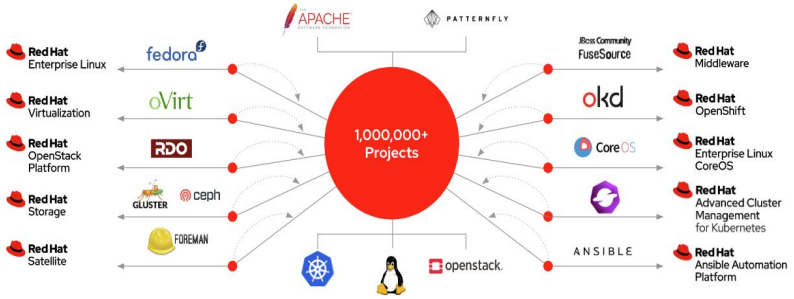

Upstream first development model

Red Hat

We are collaborating on CSP workloads with the Ecosystem

vEPC/5GC	vIMS	vRAN	SDN	AI/ML	OSS	BSS	SD-WAN	vDPI	SERVICES
Cisco	Nokia	Altiostar	Juniper Networks	hazelcast	amdocs	amdocs	Cisco	f5	accenture
affirmed	Mavenir	Nokia	nuagenetworks	ProphetStor	IBM	IBM	nuagenetworks	sandvine	Capgemini engineering
Nokia	Ericsson	Mavenir	Cisco	H_2O	Atos	Atos	Juniper Networks	affirmed	IBM
Ericsson	ZTE	Parallel Wireless	NEC	GIGASPACES	Ericsson	Ericsson	paloalto Networks	Cisco	Fujitsu
casa systems		Ericsson		CognitiveScale	SevOne	SevOne	A10 Networks	Citrix	Tech Mahindra
Mavenir		ZTE		NVIDIA	mycomosi	mycomosi			NEC
ZTE		Samsung							wipro
Hewlett Packard Enterprise		NEC							kyndryl
NEC									
Samsung									

COMMON MANAGEMENT	Red Hat ACM	Red Hat ACS	Red Hat Satellite	Red Hat Insights	Red Hat Ansible
COMMON INFRASTRUCTURE	Red Hat Enterprise Linux	Red Hat OpenShift	Red Hat OpenStack Platform	Red Hat OpenShift Data Foundation	

OEM	Dell	Hewlett Packard Enterprise	NVIDIA	Lenovo	Supermicro	IBM	intel

Note: *Partners listed represent the primary market drivers for each use case*

Red Hat

Building Red Hat's Open RAN Ecosystem

vRAN CNF vendor

ERICSSON, NOKIA, SAMSUNG, Parallel Wireless, MAVENIR, NEC, Baicells, ALTIOSTAR, Juniper Networks, Capgemini engineering

OEM servers, accelerators, nics

DELL, Hewlett Packard Enterprise, Lenovo, intel, NVIDIA, SUPERMICRO, Silicom Ltd. Connectivity Solutions, Juniper Networks

Red Hat Openshift platform

System integrators

IBM, accenture, NEC, Tech Mahindra, Atos, kyndryl, Capgemini engineering

Orchestrator service assurance

IBM, amdocs, Netcracker

RU vendor – Open FH

Comba, NEC, NOKIA

Red Hat

Summary

Red Hat

Looking Back … Developments to Make Openstack NFV-Ready

Absolutely relevant and necessary in **vRAN / Edge / 5G** space as well

PERFORMANCE	AVAILABILITY	SECURITY	MANAGEABILITY	LIFECYCLE
Enhanced Platform Awareness (EPA) SR-IOV, OVS/DPDK, vCPU pinning, Huge pages, RT-KVM …	High Availability Fault Tolerant Design Enterprise Hardened Code Instance Availability	End-to-End SELinux sVirt Neutron Security Groups Block Encryption SSL/TLS on APIs	Logging Performance Monitoring Operational Visibility Policy and Compliance	OpenStack Lifecycle Updates/Patches

Doing the same to make

Red Hat OpenShift Kubernetes Container Platform

NFV-ready

Red Hat

OpenShift from Core to the Edge

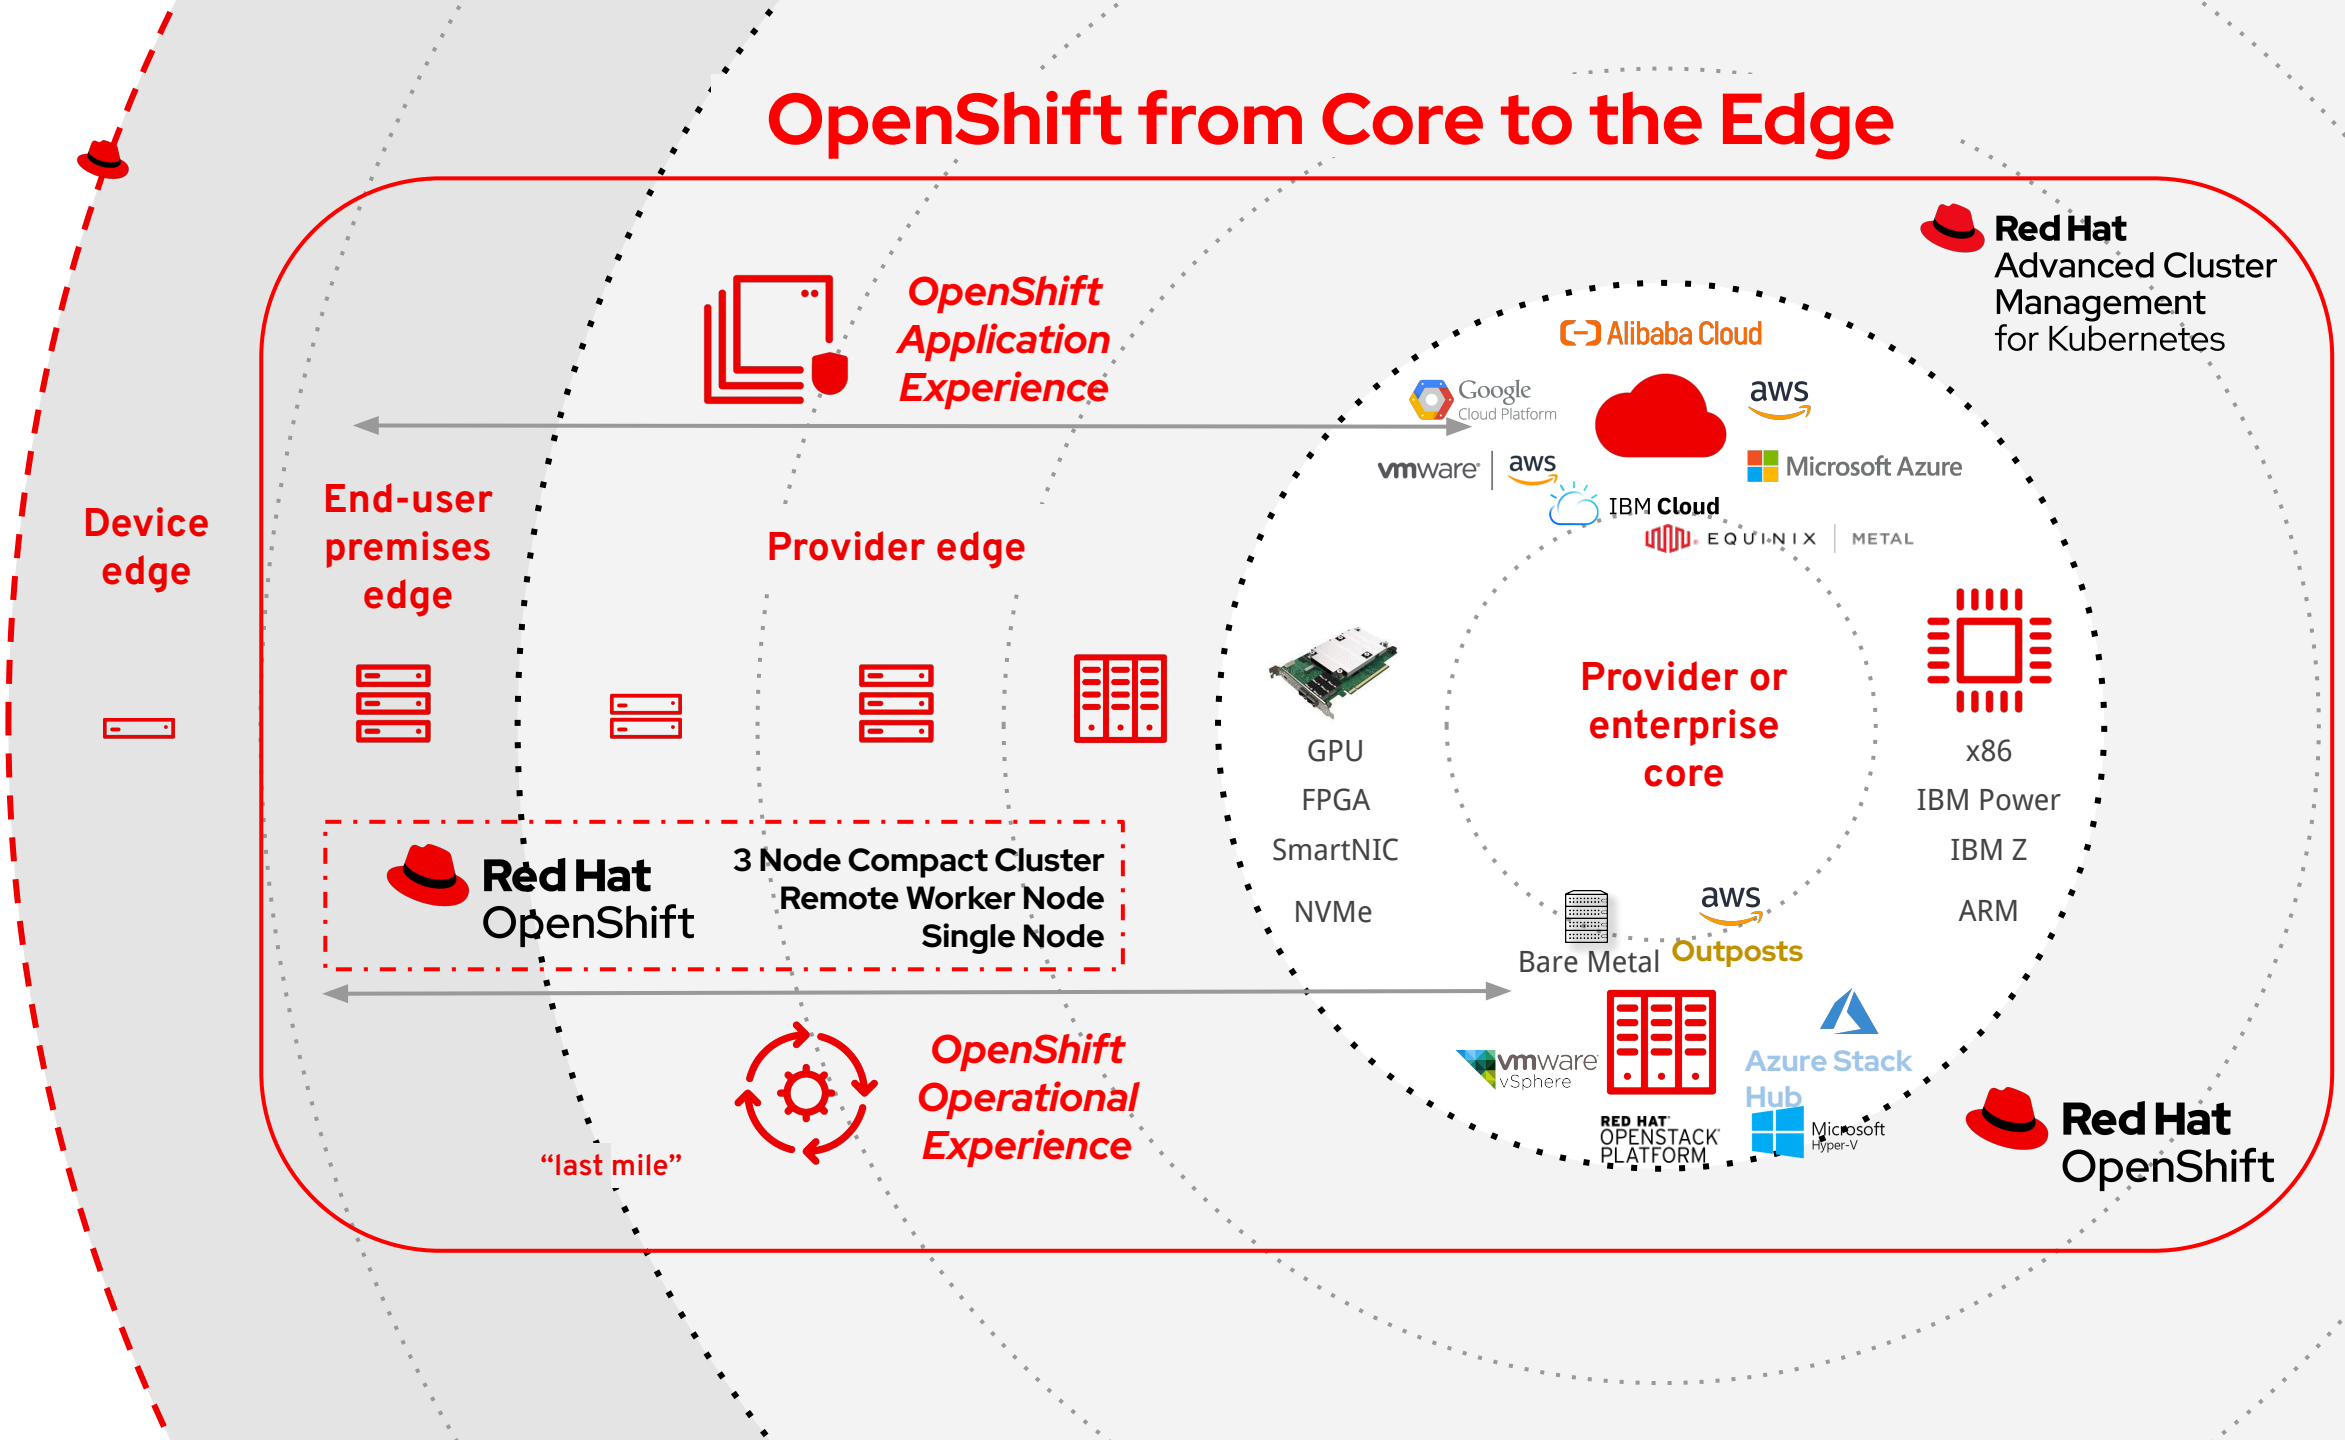

Thank you

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.

Red Hat