Disaggregated RAN Models & Energy Savings Options by Cloud Platform

Red Hat Open Tour 2023

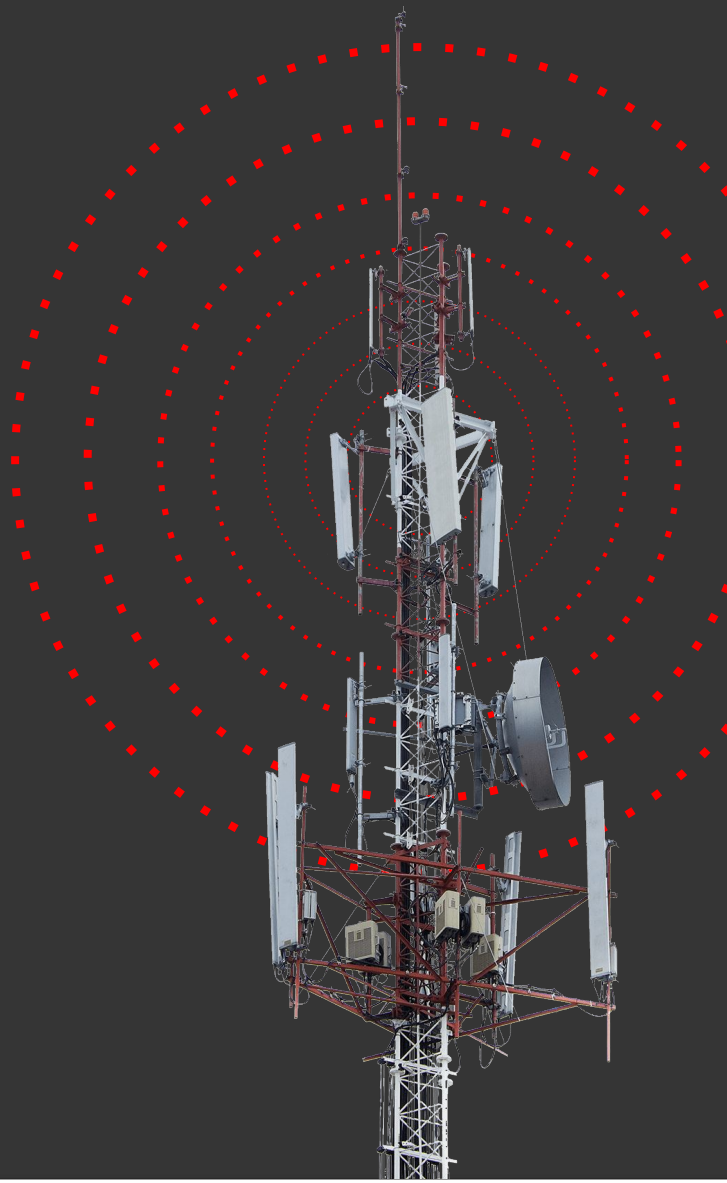

Aki & Timo
Telecom, Media & Entertainment
Red Hat EMEA

What we will discuss today

1- Open and Disaggregated RAN Models

2- Dimensions of Telco Power Saving & Energy Efficiency

3- Our Partners

4- Open Source Projects in this Domain

RAN Evolution: Two Dimensions of Architecture Transformation

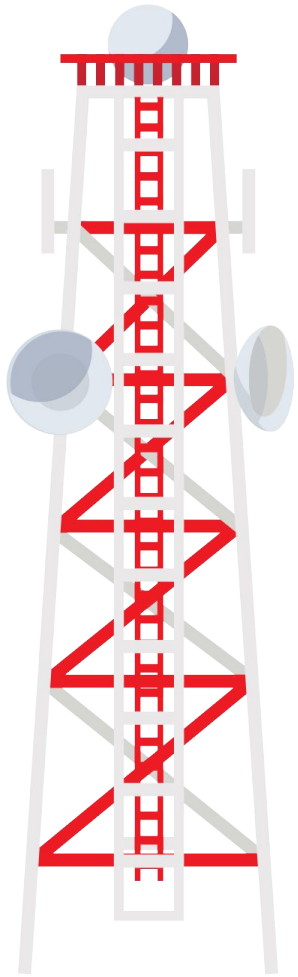

RAN Functional Split (disaggregation of functionality)

RAN Functional Split: disaggregating baseband functionality from radio functionality

RAN Cloudification (disaggregation of HW & SW)

RAN Cloudification: Deploy baseband functionality on consistent cloud platform

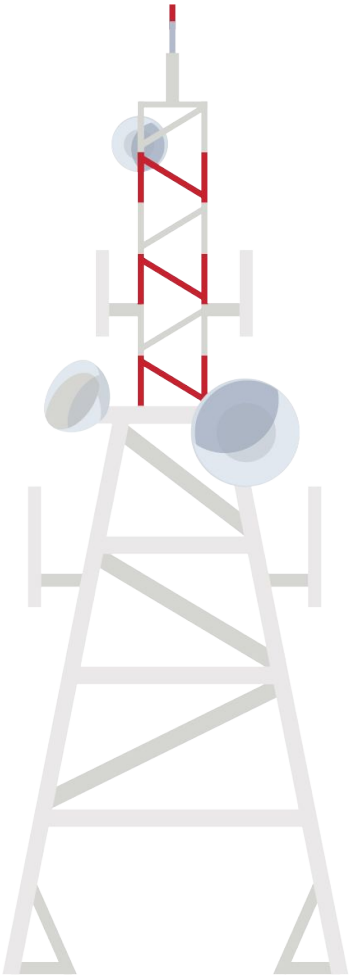

Red Hat

Mobile Network Radio Base Station

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

vRAN / Cloud RAN

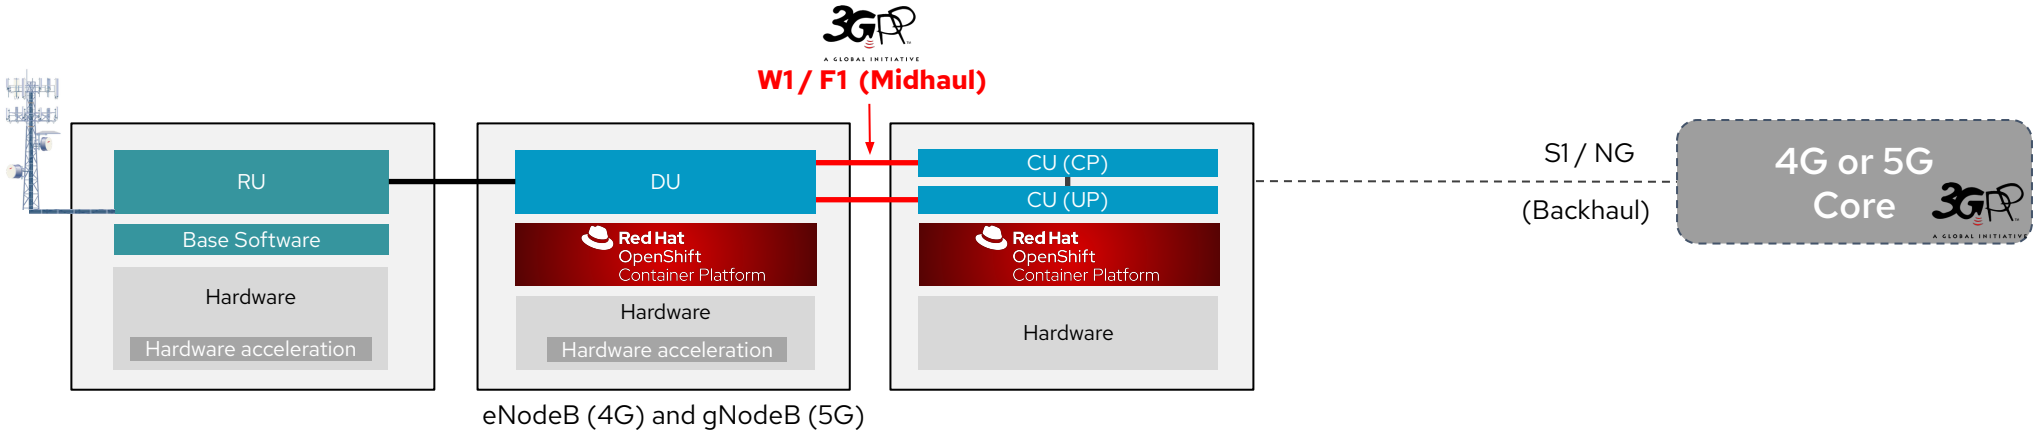

- Functional Split / Disaggregation per 3GPP
- Midhaul interfaces: W1 (Rel 16 for 4G) & F1 (Rel 15 for 5G)
- Standard interfaces (Backhaul) towards Core Network(s)
- Three entity model: Radio Unit (RU), Distributed Unit (DU), Centralized Unit (CU, Control and User Planes)
- Cloud Platform to host DU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

Mobile Network Radio Base Station

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

Open RAN

Open RAN Workload Interfaces

(partially compliant to O-RAN Alliance)

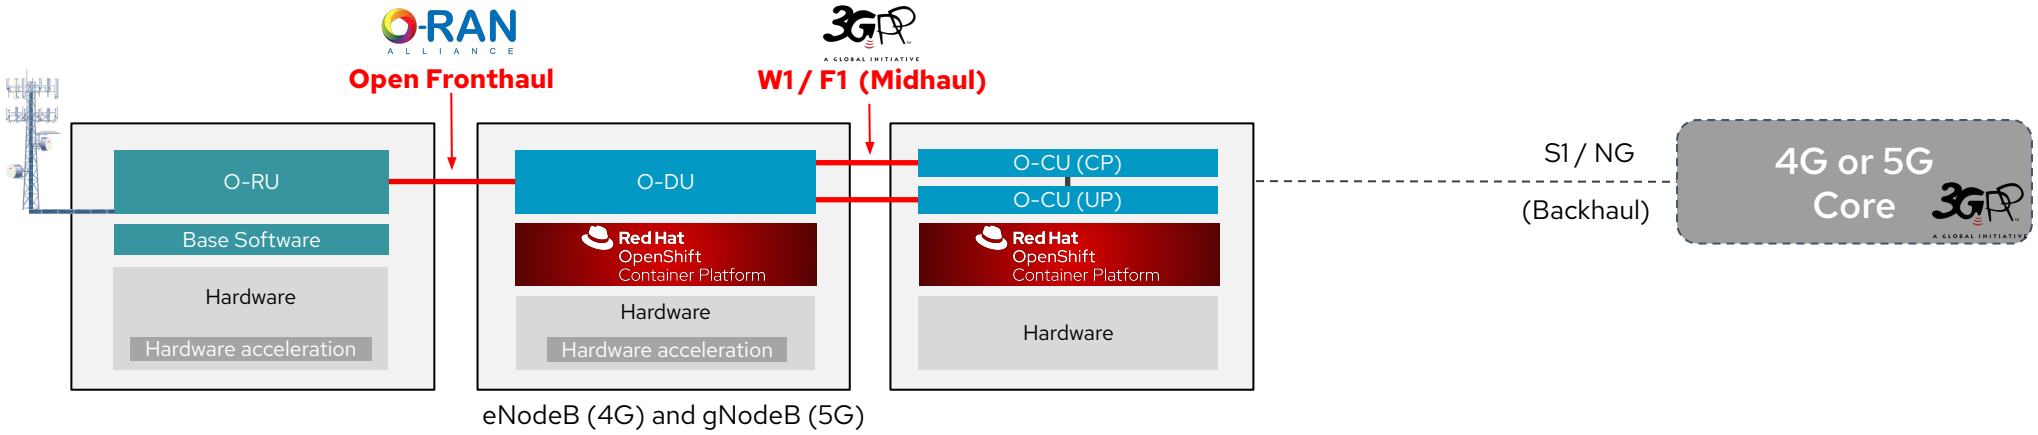

- Added: Open Fronthaul by O-RAN Alliance
- O-RAN Alliance nomenclature: O-RU, O-DU, O-CU
- Potential for multi RAN vendor model

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

3GPP Interfaces:
- F1
- W1
- NG
- S1

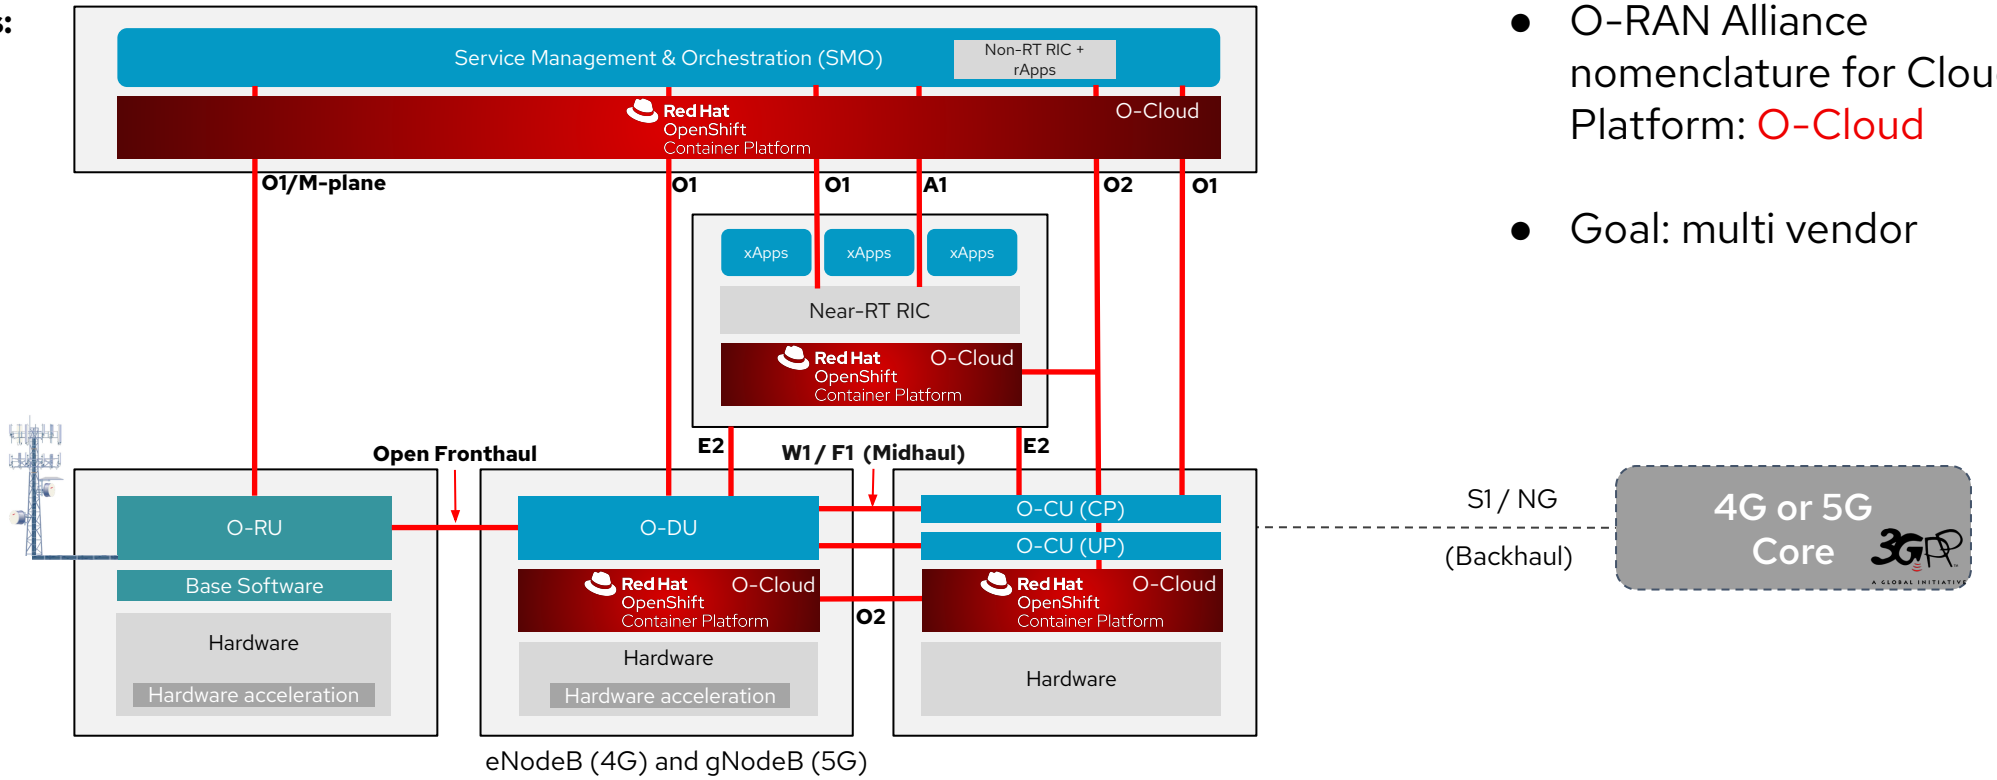

- O-RAN Alliance nomenclature for Cloud Platform: O-Cloud
- Goal: multi vendor

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

Deployment Models for Disaggregated RAN

And Consistent Cloud Platform …

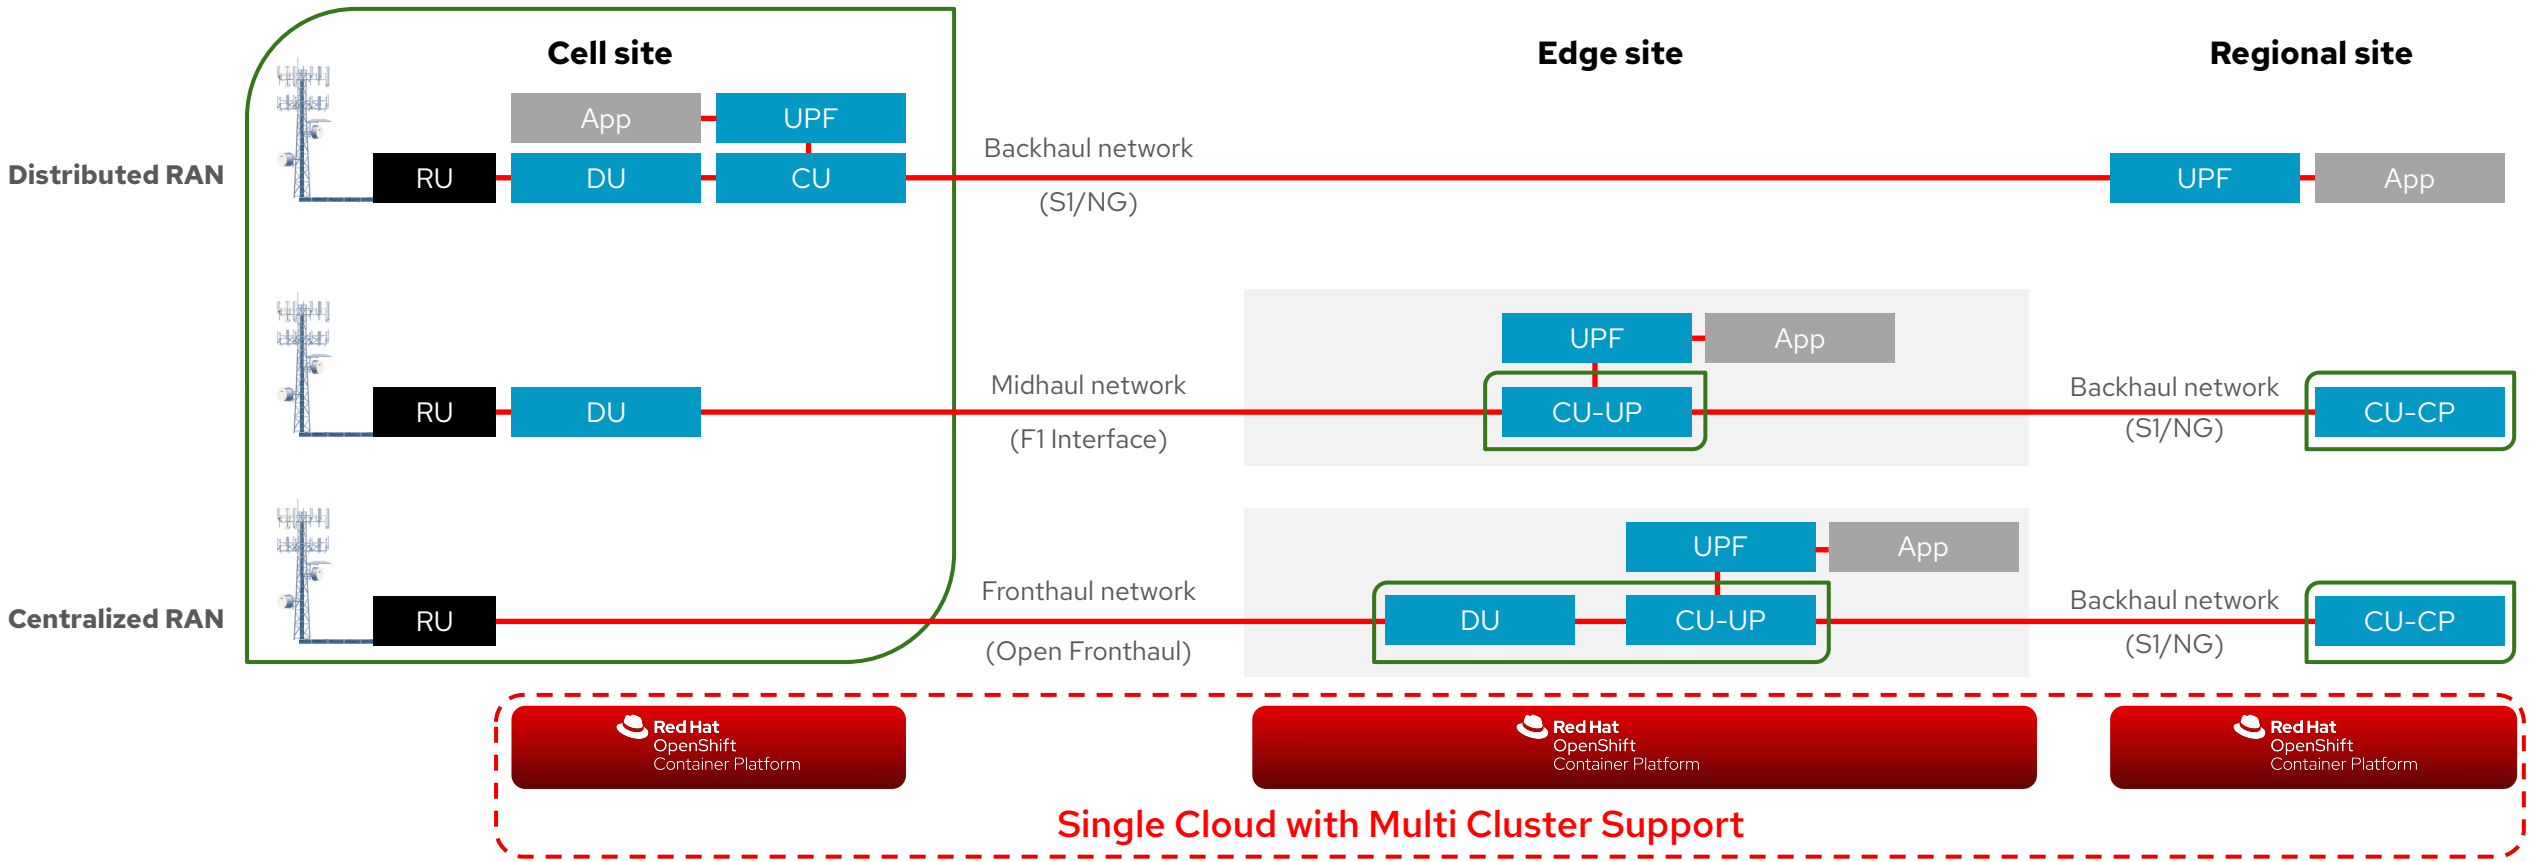

RRH = Remote radio head
BBU = Baseband unit

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

Network Level View

According to TMForum[1], service providers account for about 3% of the total power consumption of humanity. As shown in figure below, approximately 75%[2] of that is consumed by RAN. This represents a significant portion of the OPEX of service provider networks.

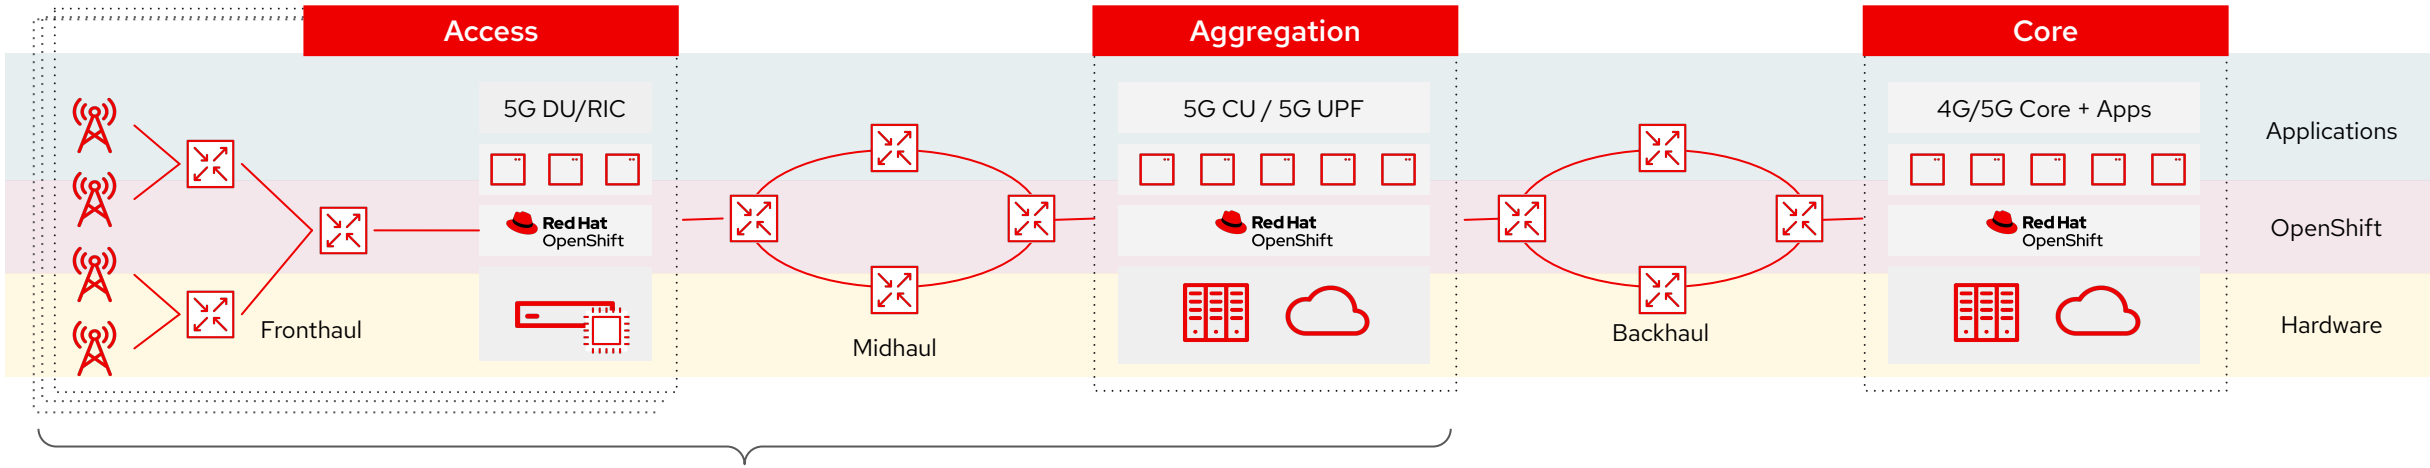

RAN represents more than 75% of the service provider's network power consumption

Legend:
CU: Centralized Unit
DU: Distributed Unit
RIC: RAN Intelligent Controller
UPF: User Plane Function

[1]: "Can the telecoms industry power down its impact on the environment?" Inform TMForum, 19 Aug. 2021.
[2]: "A holistic approach to address RAN energy efficiency." Ericsson, 16 Dec. 2021.

A Holistic Approach to Sustainability of Mobile Networks

Many dimensions for Power Saving and Energy Efficiency

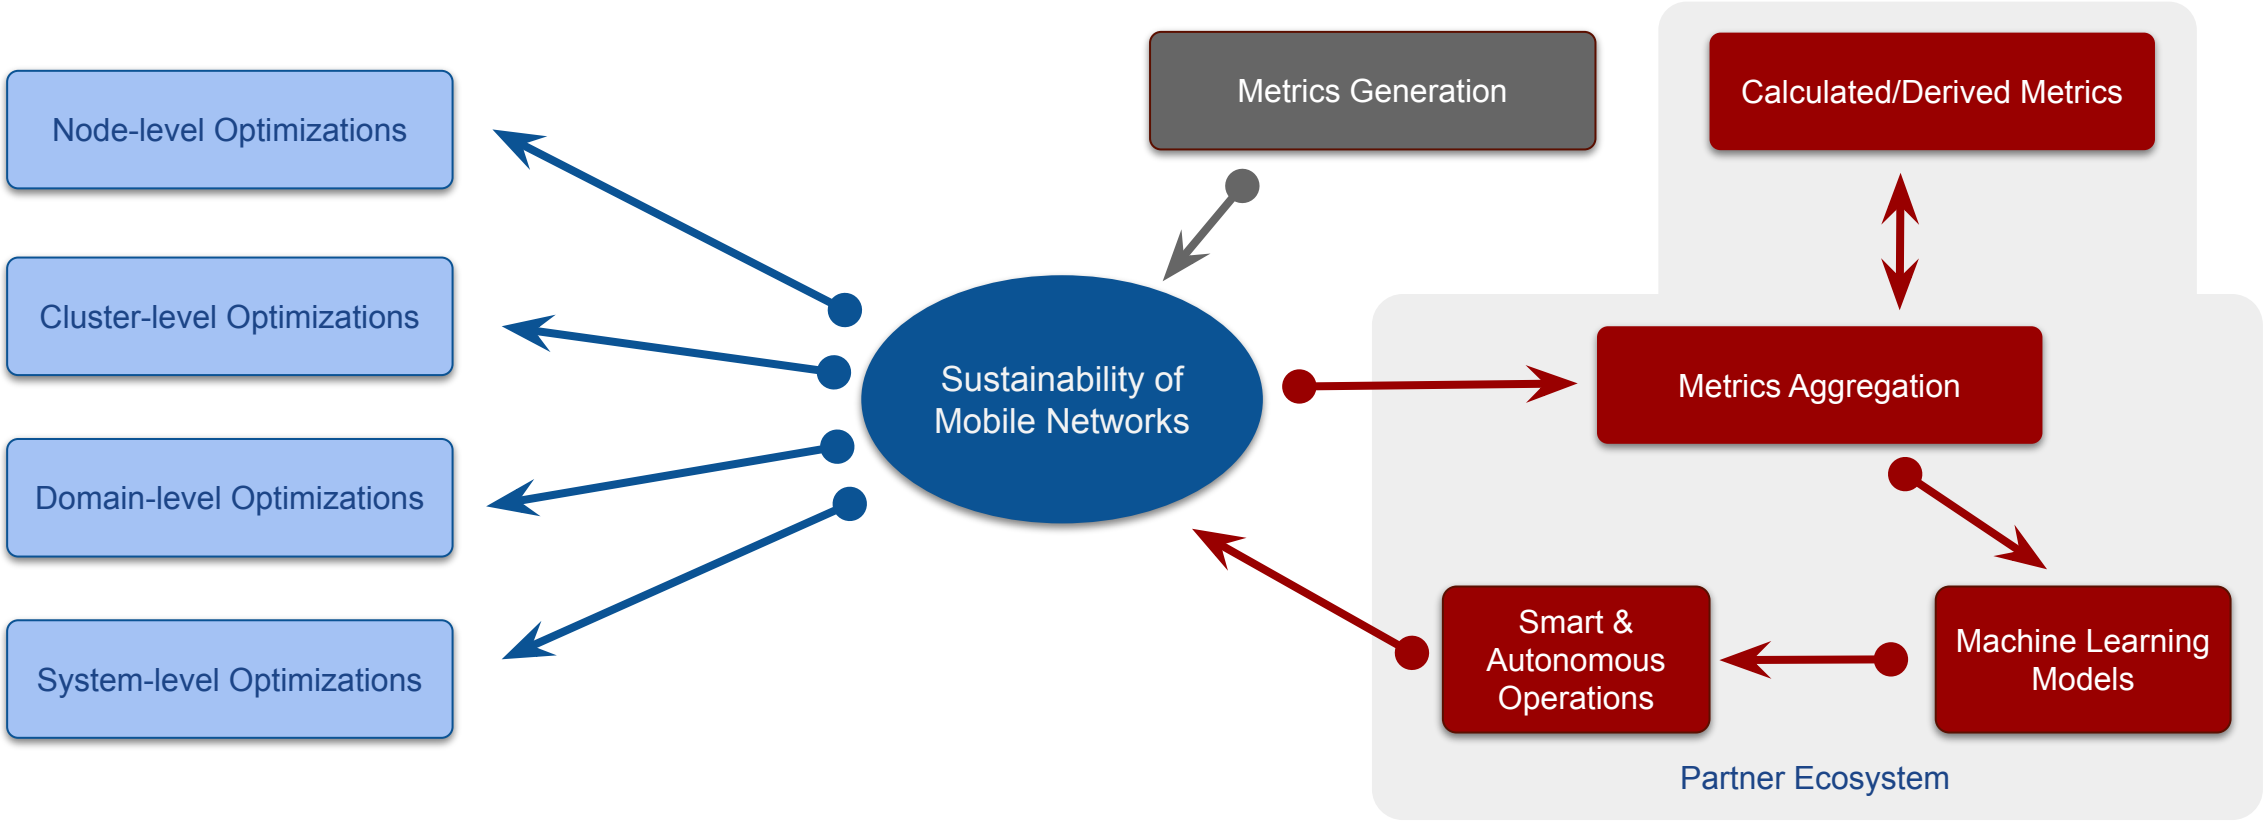

Red Hat

Energy Optimisation is Crucial at Every Layer

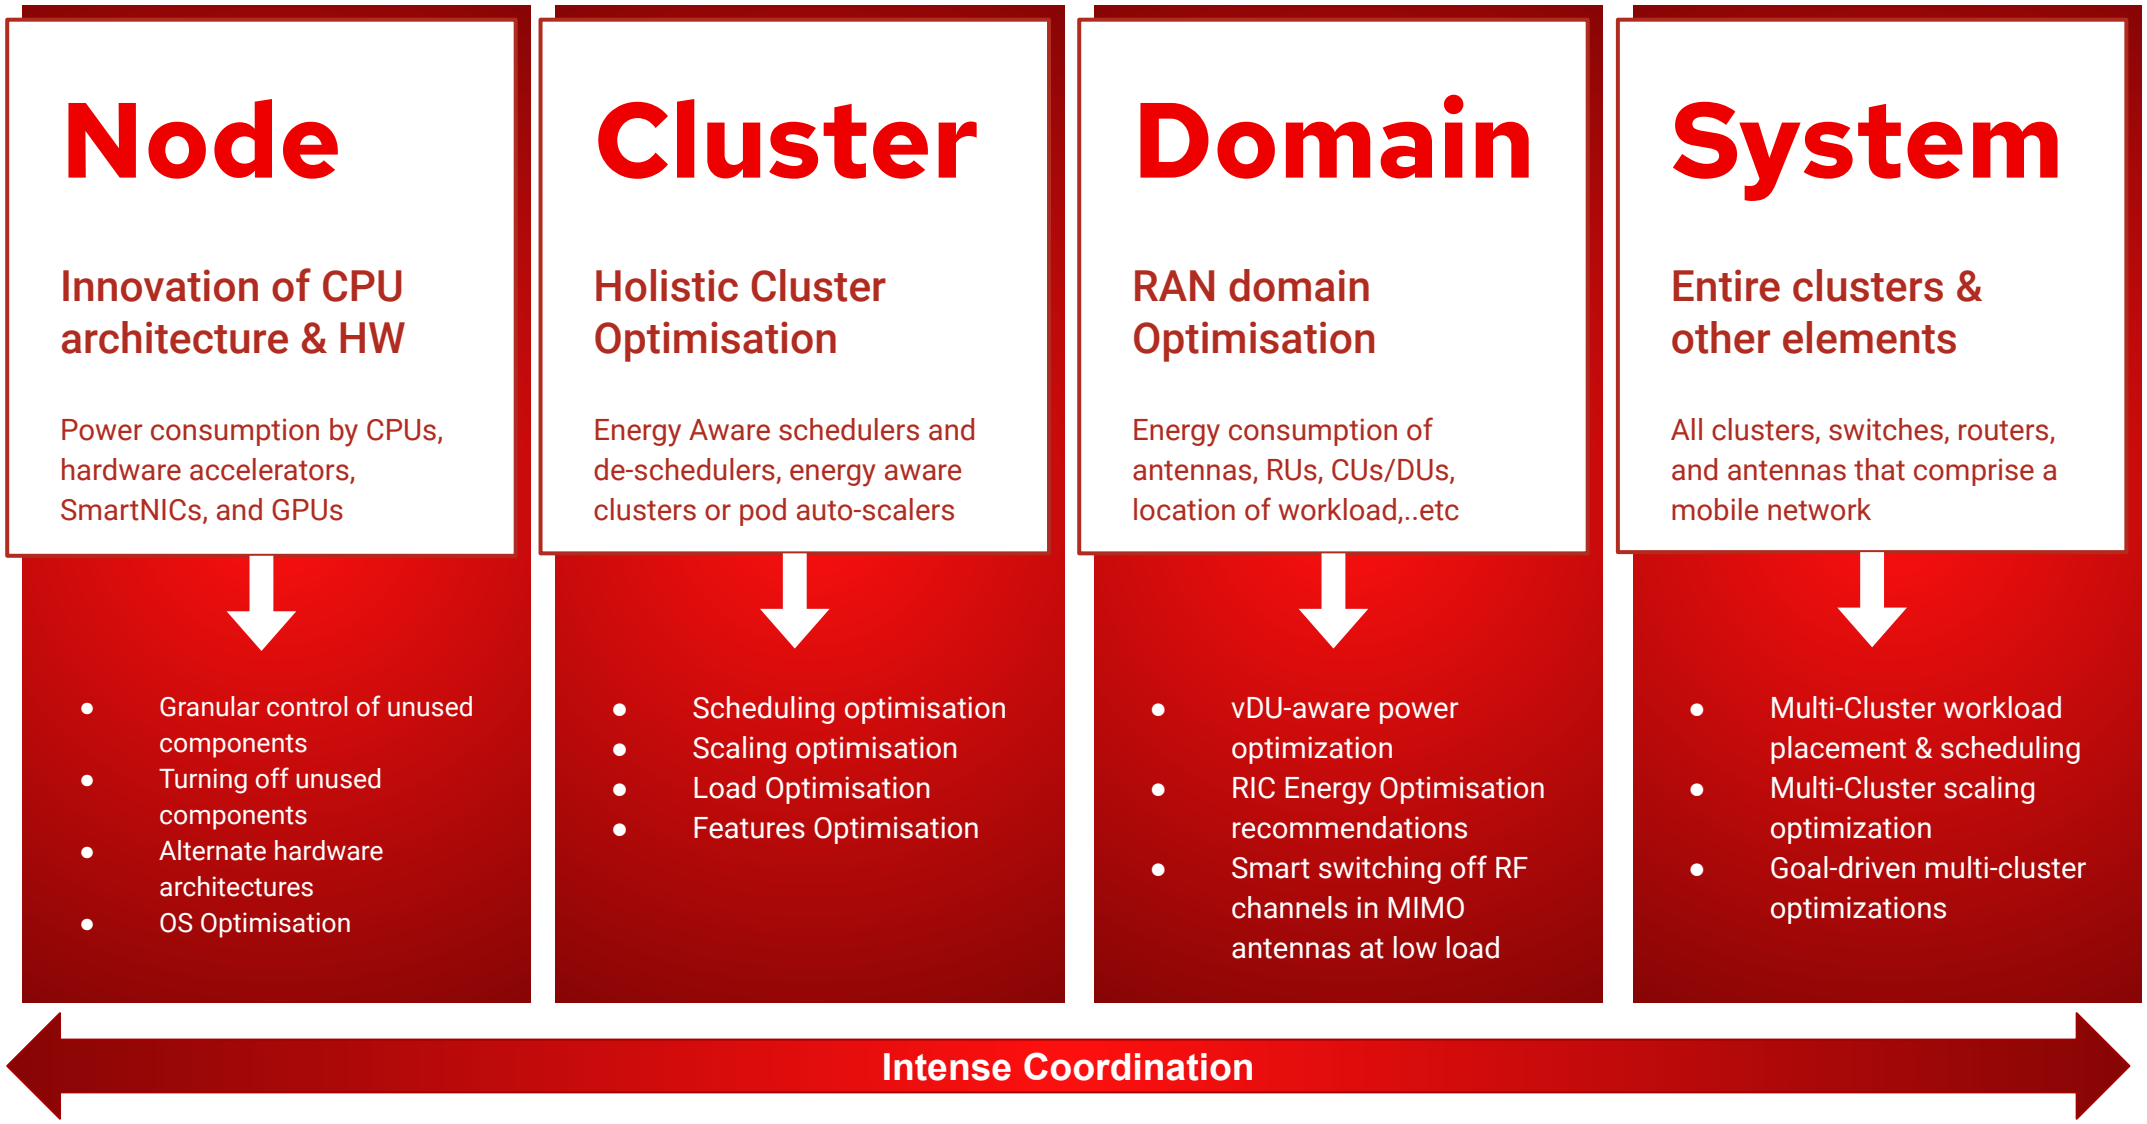

Red Hat

Node Level Optimizations

Red Hat Primary Target for Energy Consumption Management

As shown in the figure below, optimizations range from granular control of hardware components, for example, disks, NICs, and GPUs to identifying and disabling idle components, for example, turning off unused cores or hardware accelerators, to adopting alternate hardware architectures, for example, SmartNICs or ARM CPUs.

Available OS optimizations include power efficient profiles, for example, tuned server-powersave profiles and granular selection of OS settings for energy efficiency, for example, selection of CPU governors, core frequency settings, and adjusting **p/c-state**.

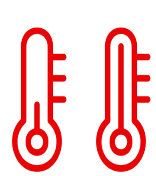

Heat Control

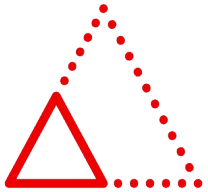

Frequency Scaling

Voltage Scaling

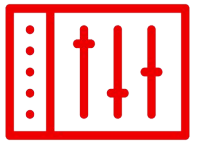

Clock Gating

Sleep States

Disabling or turning off unused hardware components or having the ability to disable unused CPU cores or vGPUs adds up to a noticeable reduction in power consumption. Taking into account hundreds of thousands of nodes that are part of a service provider network, the power reduction achieved by the combined nodes could be significant.

CPU Power Modes Feature

- OpenShift supports configuring CPU Power states p-states & c-states per pod
- It is possible to enable power savings for a node that has low priority workloads that are collocated with high priority workloads without impacting the latency or throughput of the high priority workloads.
- This allows us to enable CPU power saving features on all CPUs by default, and CPU Power saving features are then disabled for isolated pods that require it via the two annotations that you can see on in the illustration.
- This means also that we can now have a single CPU power efficient configuration for a whole cluster, as high power demanding pods and regular pods can now run on the same nodes.

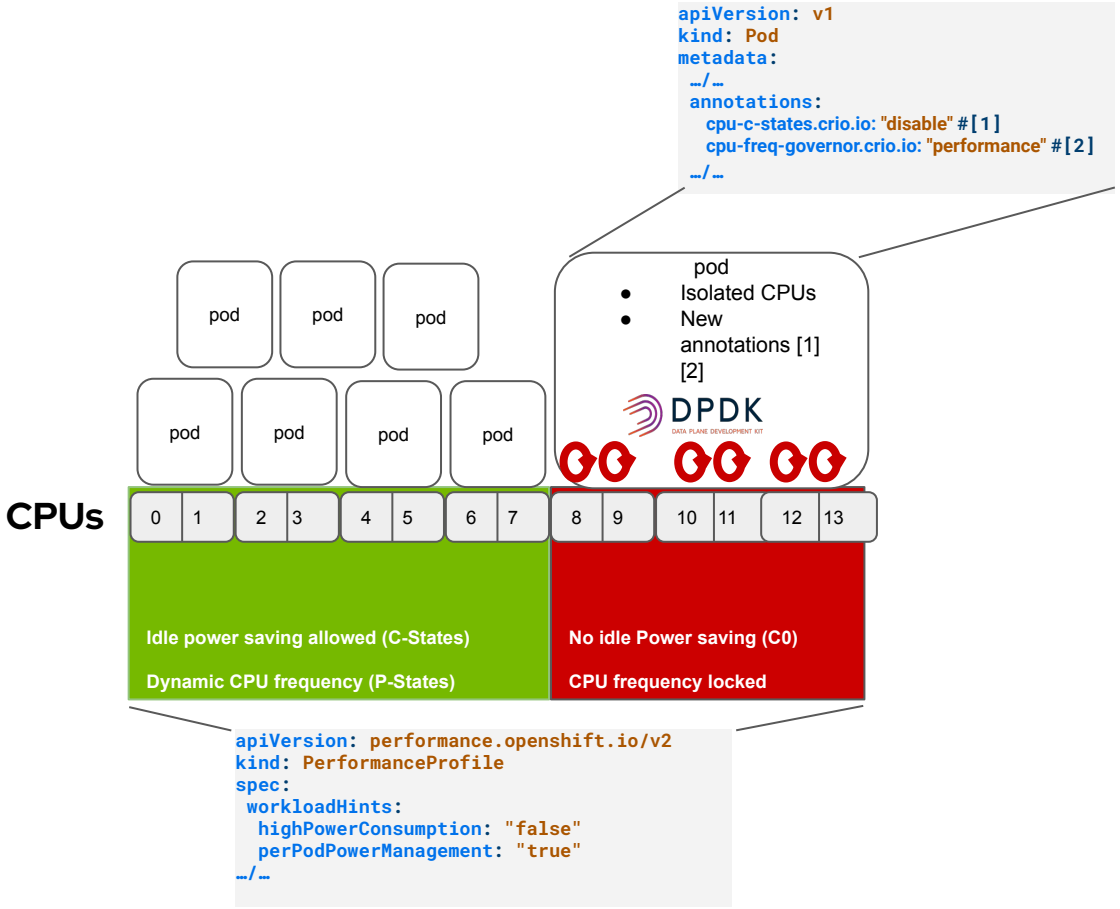

Red Hat

OpenShift Observability
Kubernetes-based Efficient Power Level Exporter

Drive energy cost down for IT operations and contribute to achieve sustainability goals

KEPLER

Power monitoring with Kepler

- Uses eBPF* to probe energy related system stats and exports as Prometheus metrics that can be leveraged for workload scheduler and auto-scaling and drive CI/CD pipelines
- Power monitoring with Kepler will be Dev Preview with OpenShift 4.14

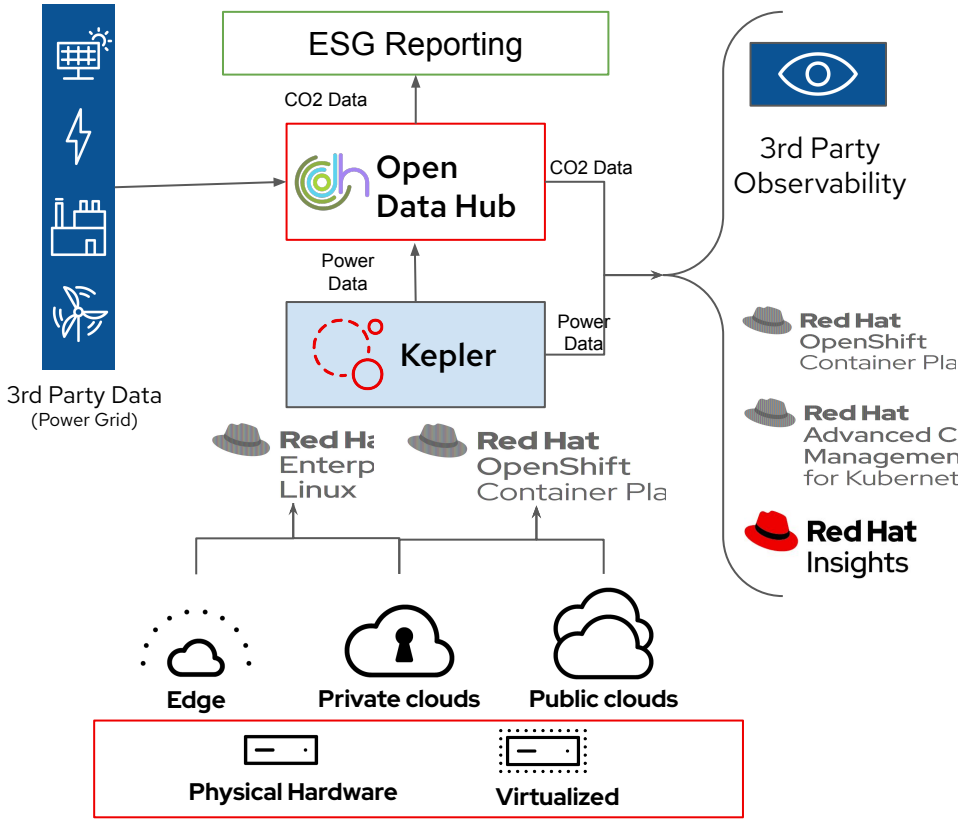

Project Scope

- Monitor/Report Energy Costs and CO2 Emissions
- Hybrid Cloud Energy and CO2 Monitoring and Reporting
- Data/Analytics for Energy Optimization
- Data/Analytics for Green IT and Green (Re)-Engineering
- Data for ESG reporting

Kepler was accepted to CNCF on May 17, 2023 and is at the Sandbox project maturity level: https://www.cncf.io/projects/kepler/ https://github.com/sustainable-computing-io/kepler

* eBPF is part of the Linux kernel. eBPF was added to the Linux kernel to enable the small sandboxed programs. the eBPF program runs as native instructions, which is a faster and more efficient method for running the underlying bytecode

Red Hat

Partner: Intracom-Telecom

NFV-RI: Optimizations after a period of time

Average Power Savings with AI from NFV-RI

Side-by-side comparison of power consumption **before and after** optimizations by NFV-RI

UPF Bandwidth Throughput

AI-based frequency tuning of cores used by UPF instances

OpenShift: Cloud RAN app pod Power Management

ERICSSON

Performance profile allows per-core runtime tuning of CPU power states

```
apiVersion:
performance.openshift.io/v2
kind: PerformanceProfile
spec:
 workloadHints:
  highPowerConsumption: "false"
  perPodPowerManagement: "true"
…/…
```

Allowed to control C-States. Dynamic CPU freq (P-States)

```
apiVersion: v1
kind: Pod
metadata:
…/…
annotations:
  cpu-c-states.crio.io: ”enable" #[1]
  cpu-freq-governor.crio.io: "performance" #[2]
…/…
```

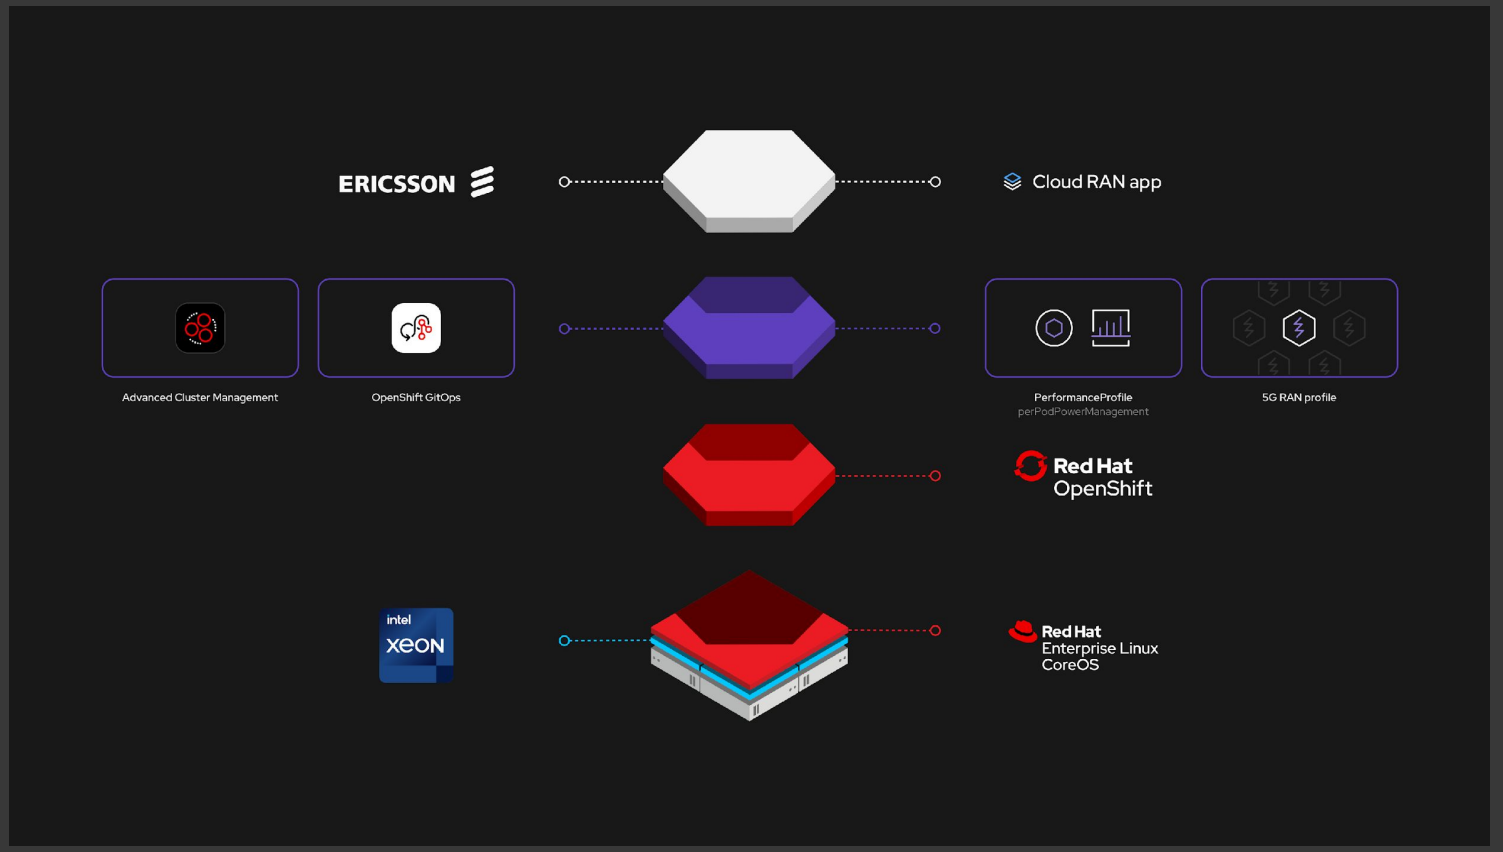

Red Hat Telco Lab front end:
https://open-telco.vercel.app/dashboard

In partnership with

intel Red Hat

IPM Demo Architectural Overview

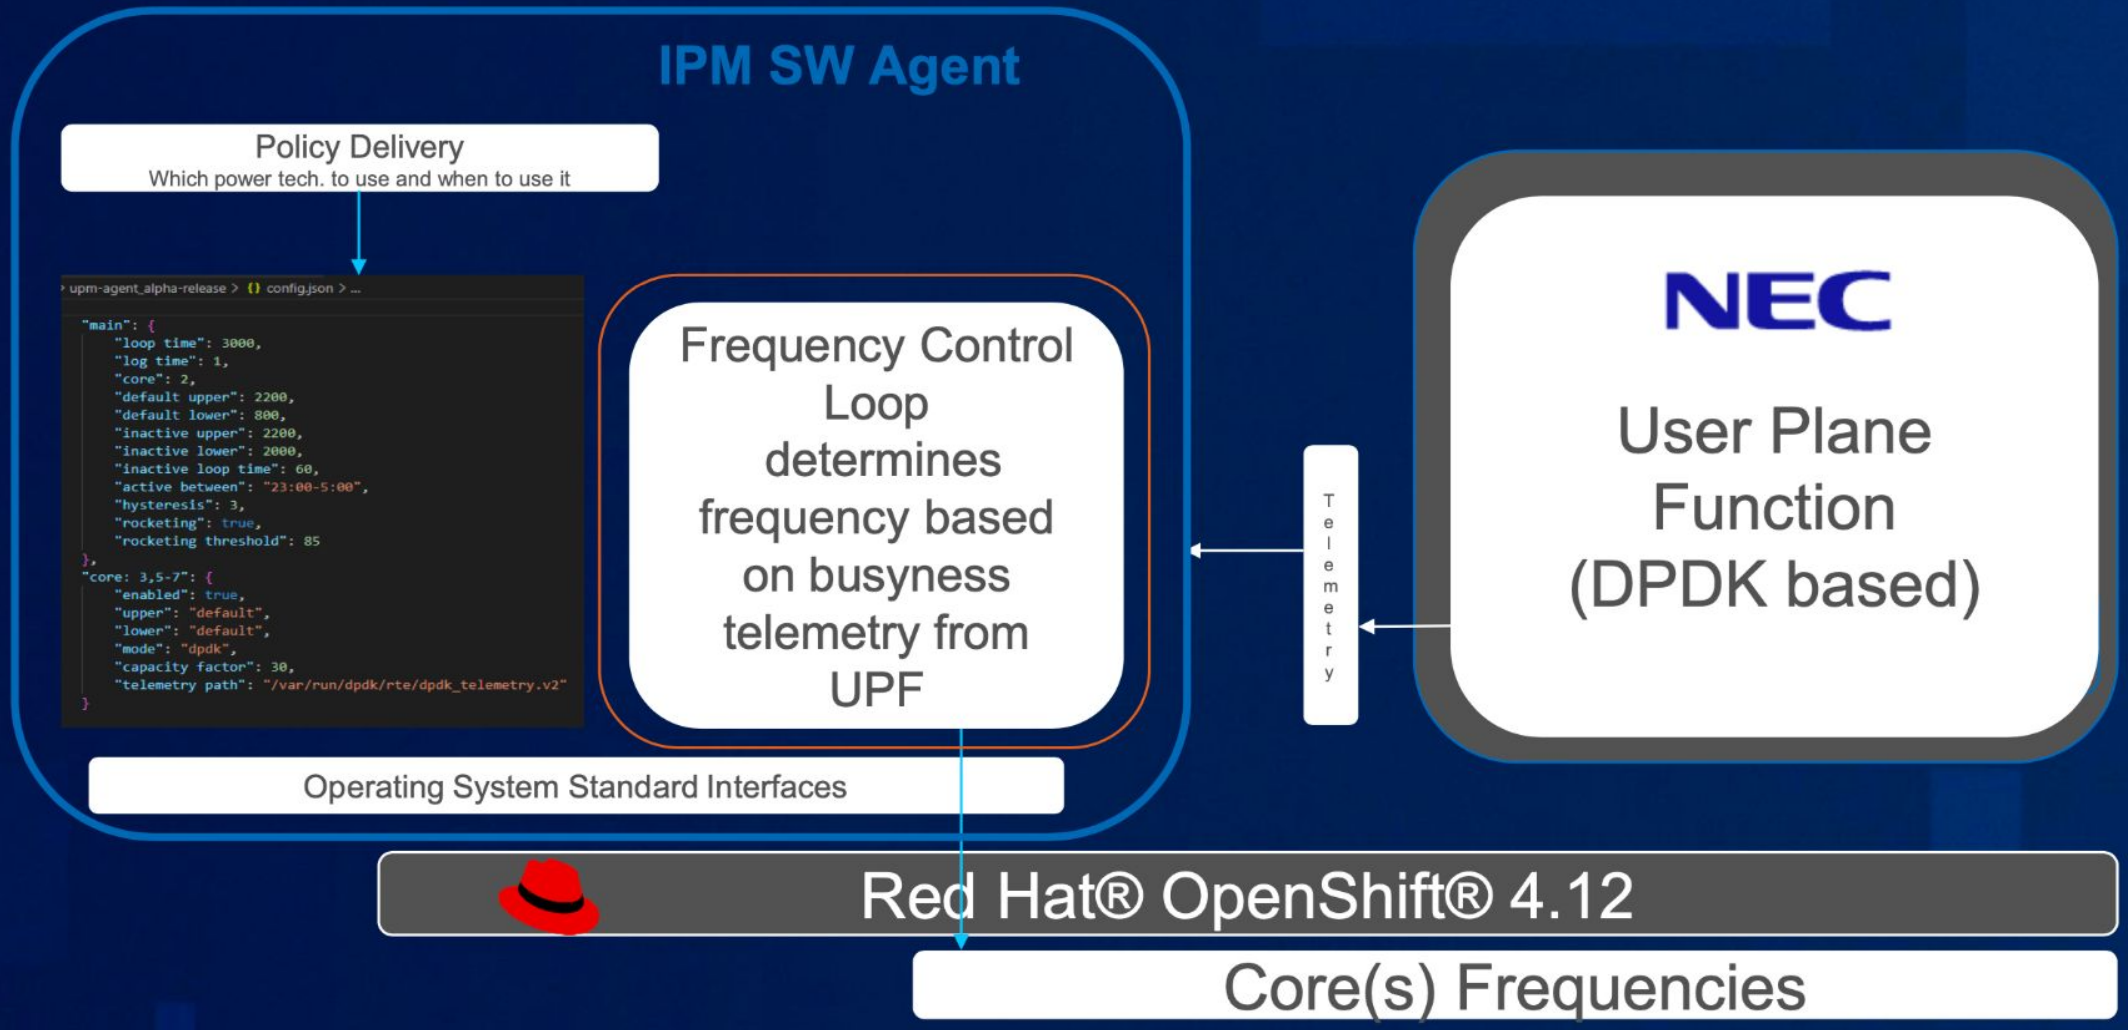

Some Additional Upstream Projects (subset of ongoing projects)

KEDA[1]	Community to bring native CO2 and energy aware auto scaling capabilities to Kubernetes event-driven autoscaling (KEDA)
CLEVER[2]	Container Level Energy-efficient VPA recommender to enable energy consumption metrics to be used for the vertical pod autoscaling for an application.
PEAKS[3]	Power Efficiency Aware Kubernetes Scheduler (PEAKS) on a Kubernetes scheduler that will take power metrics into consideration for the scheduling of an application.

[1] https://www.cncf.io/projects/keda/
[2] https://github.com/sustainable-computing-io/clever
[3] https://github.com/sustainable-computing-io/peaks

Computing Tiers with Network Level View

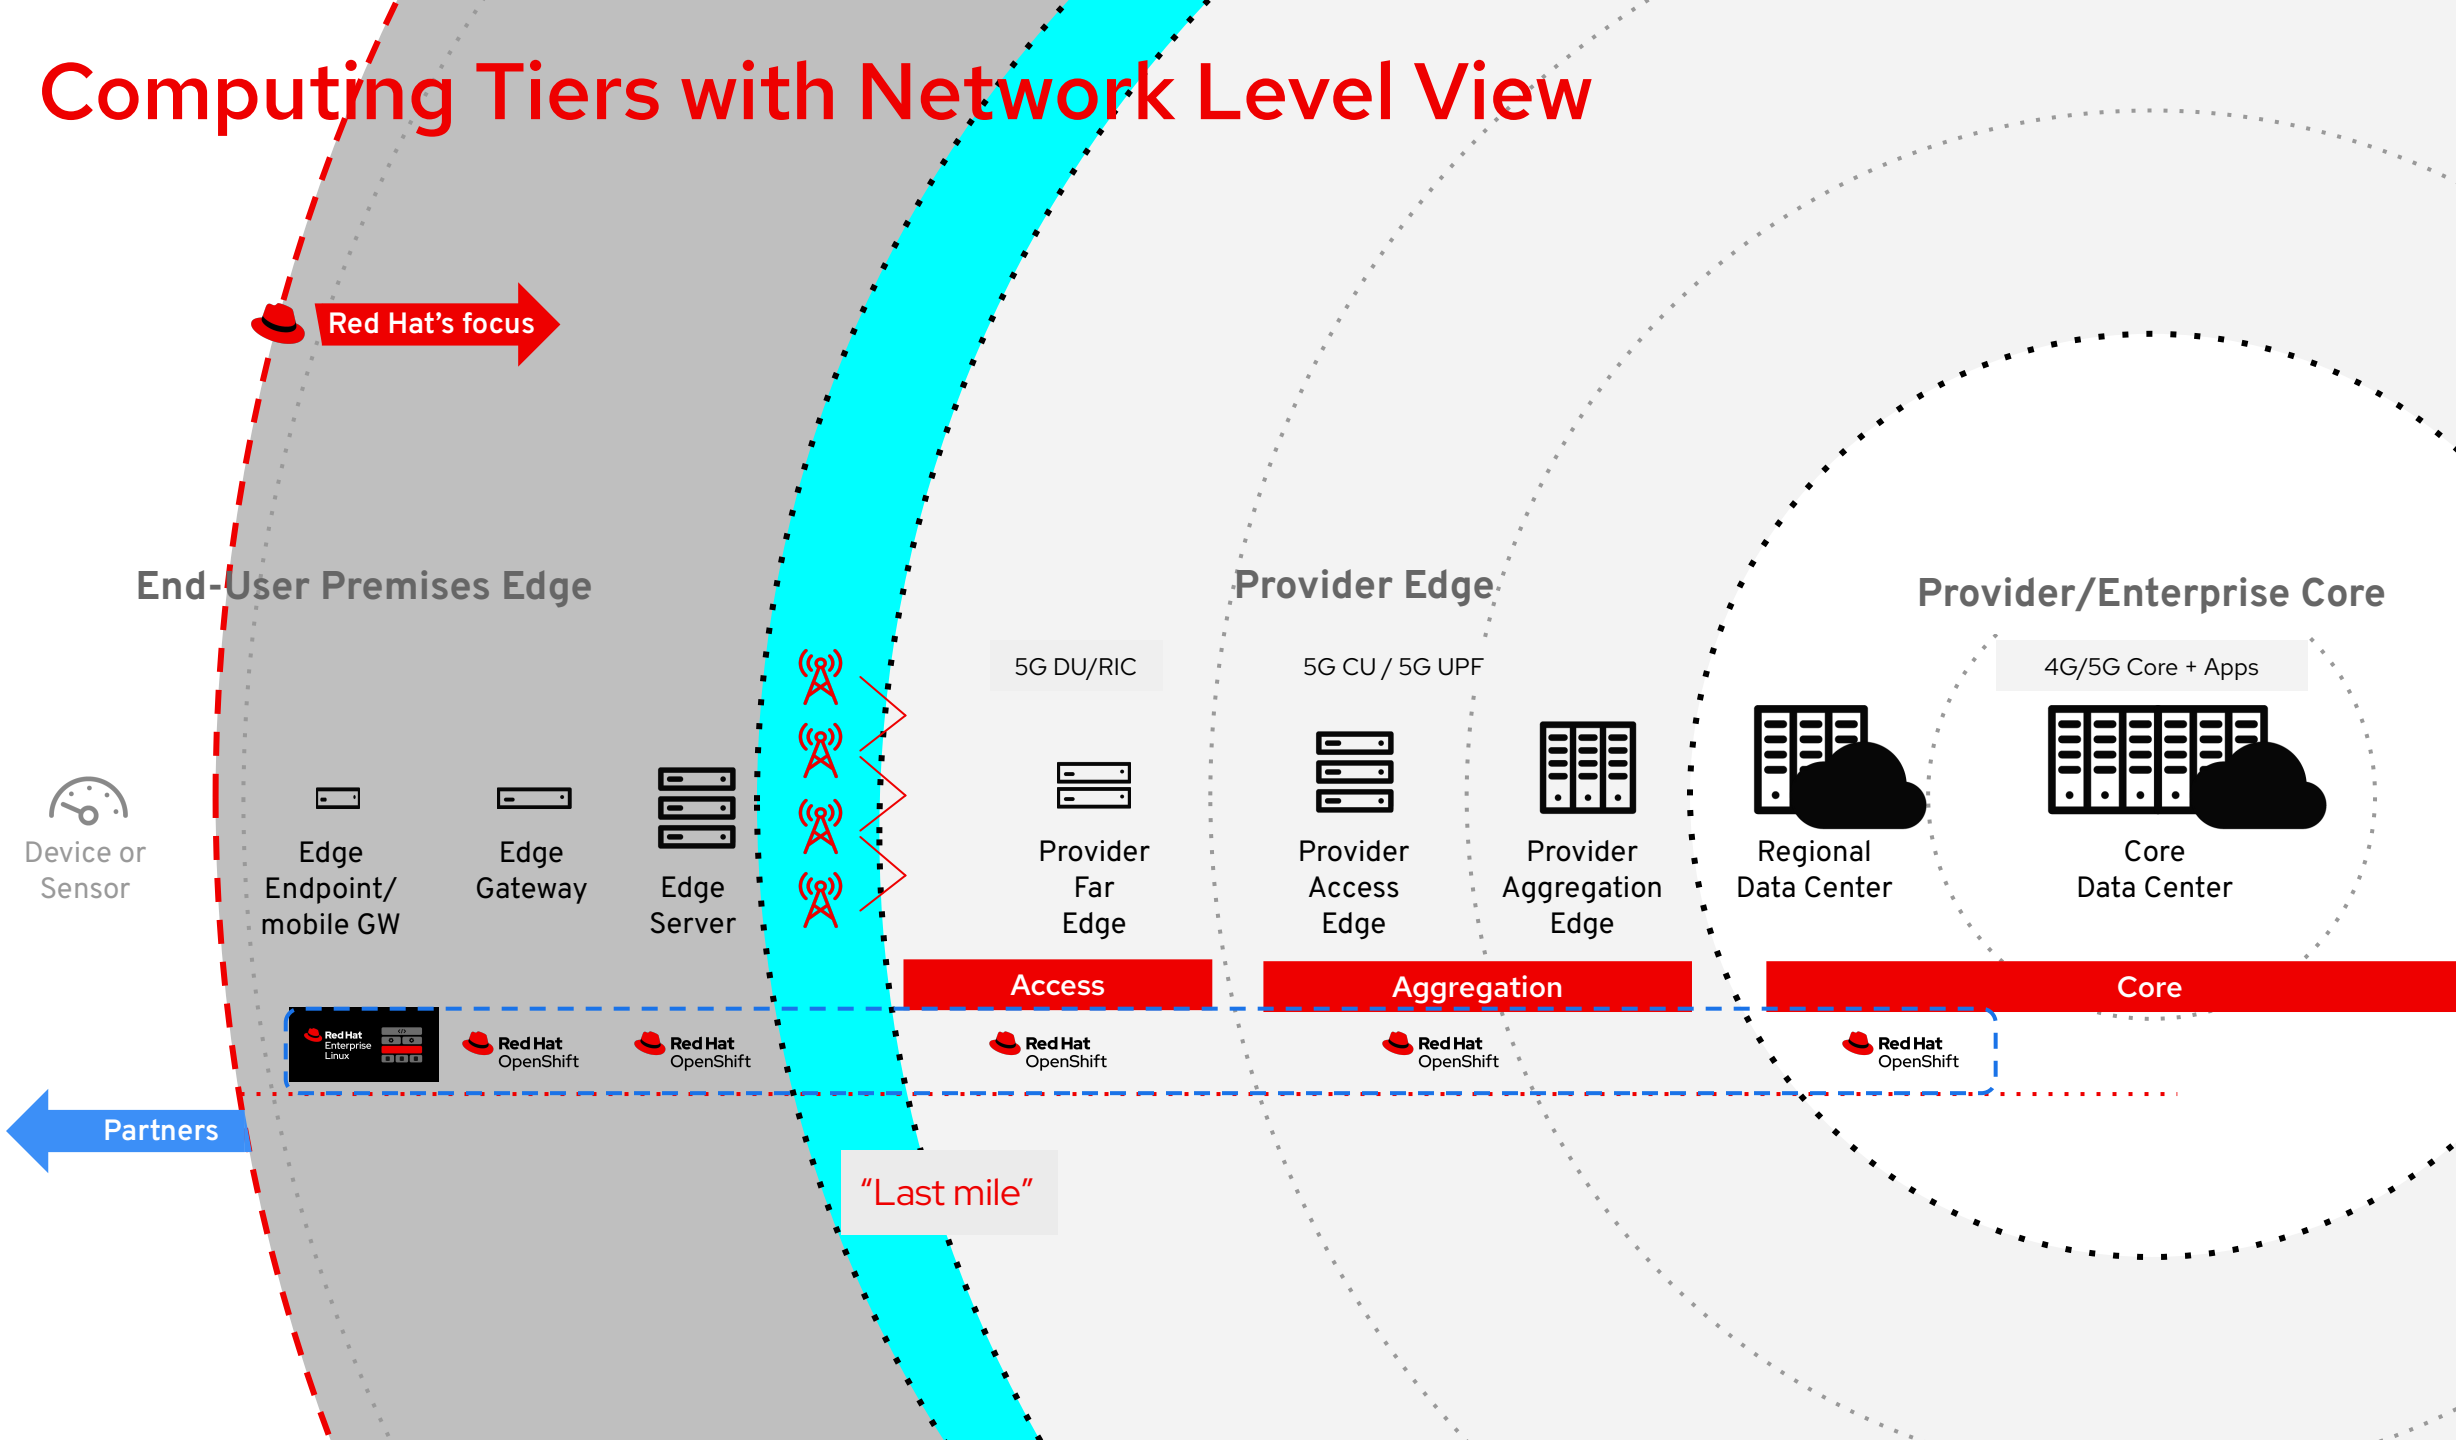

redhat.com/telco

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