OpenNESS: Enabling High Performance Edge for Telco & Enterprise

Prakash Kartha

Amr Mokhtar

OpenNESS
OPEN NETWORK EDGE SERVICES SOFTWARE

certified kubernetes

intel®

Agenda

- Opportunities/Challenges deploying Edge in a Telco environment
- Introduction to OpenNESS
- Architecture Overview
- Deployment Scenarios
- Getting Started
- Wrap up

Delivering Cloud Native Platforms for the Edge

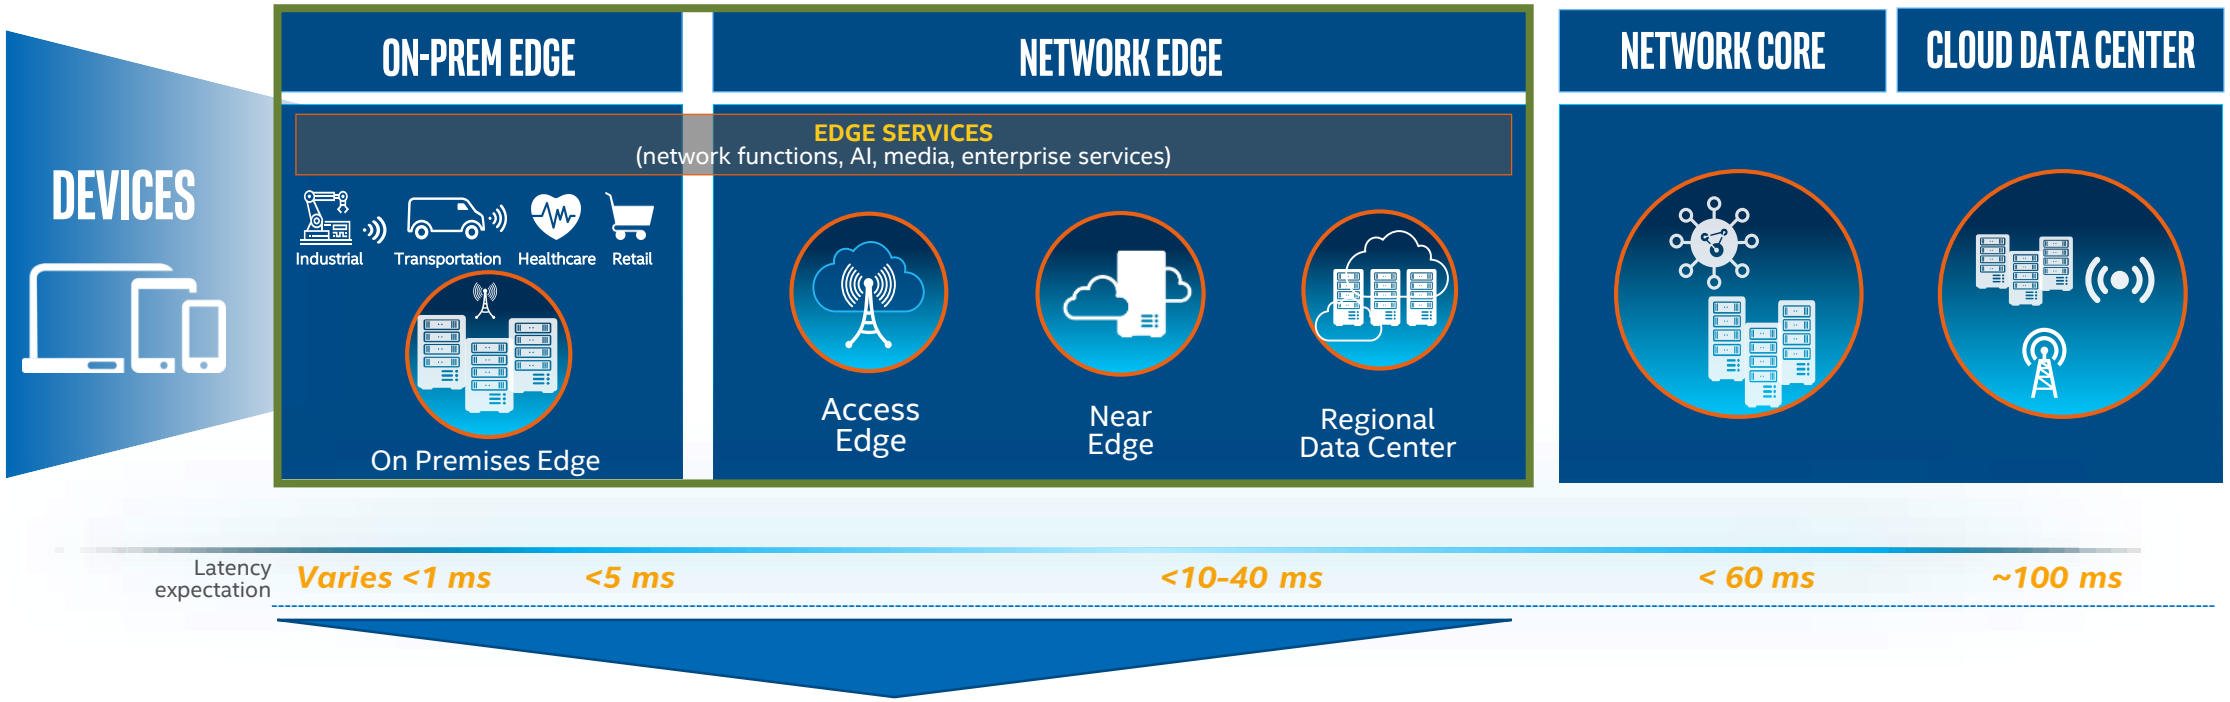

Market Opportunity: Lower TCO with a consistent cloud native platform approach across edge locations

Reference: https://www.openness.org/docs/doc/architecture#deployment-based-on-location

Delivering Cloud Native Platforms for the Edge

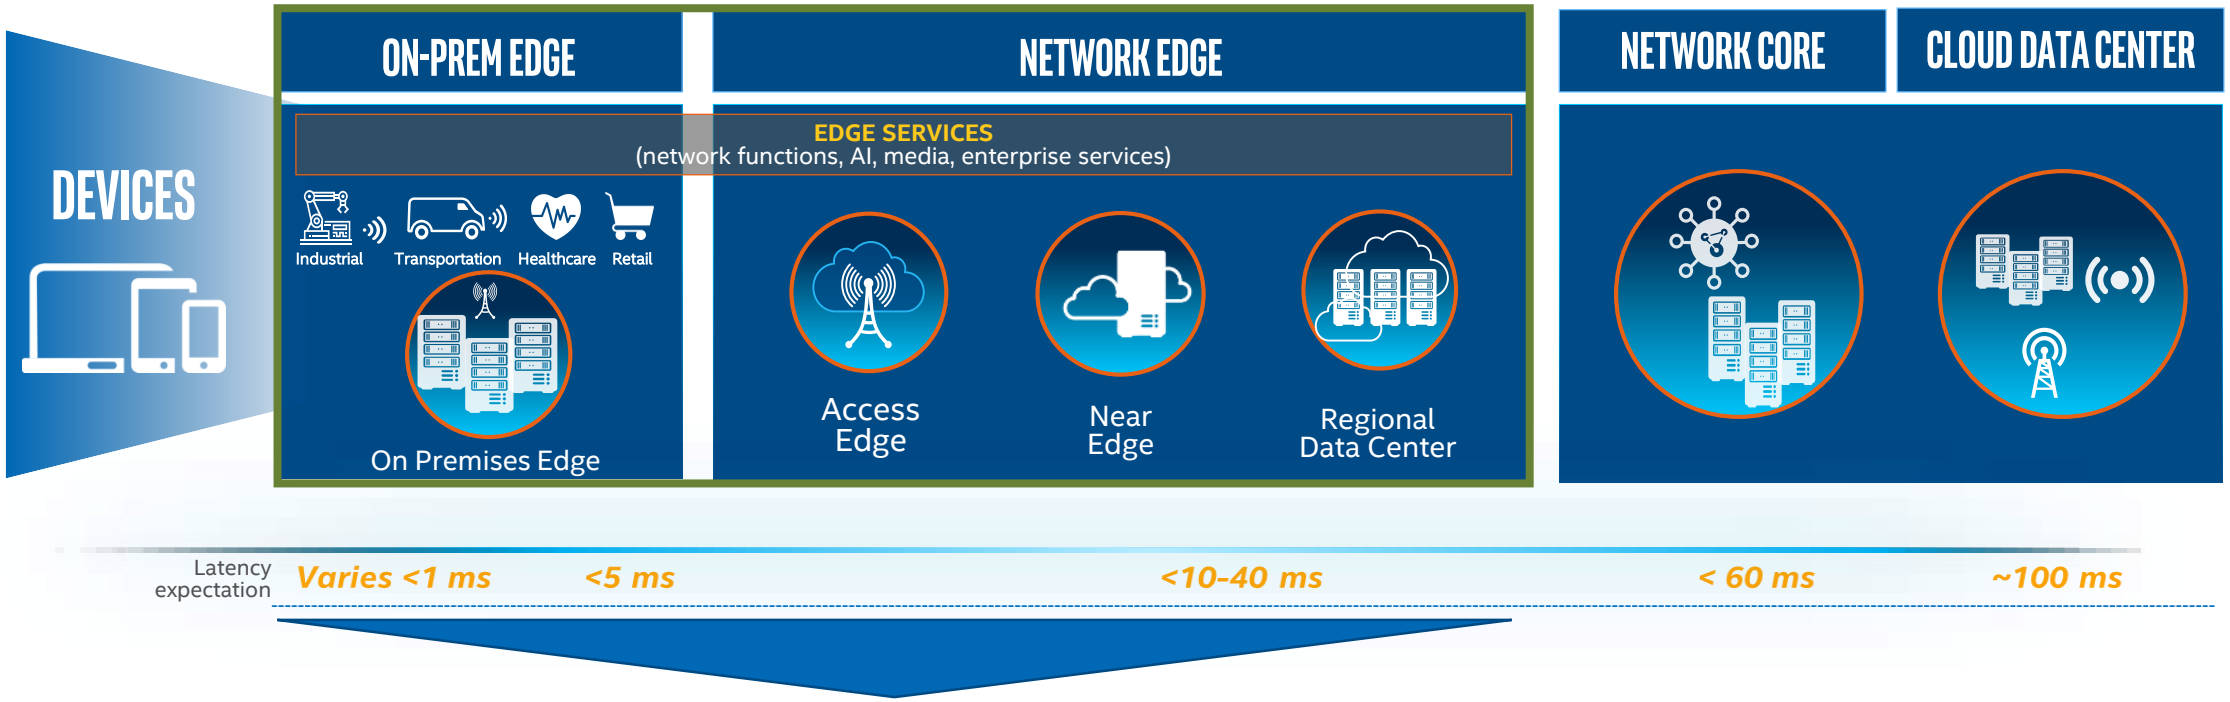

Market Opportunity: Lower TCO with a consistent cloud native platform approach across edge locations

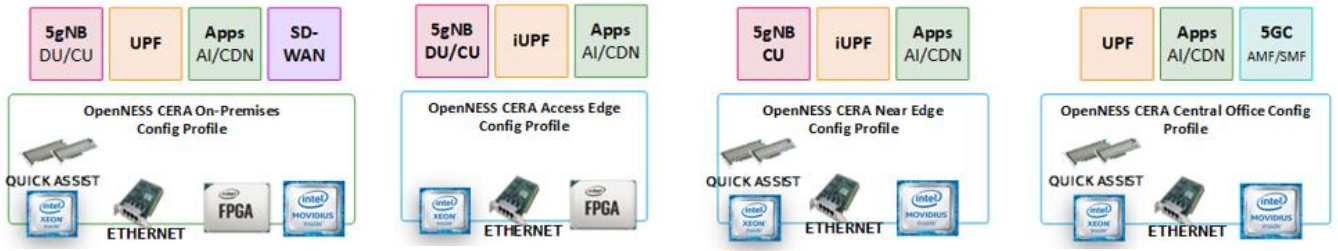

Key challenges to overcome

1. Deliver platform consistency & scalability across diverse edge location requirements
2. Optimize cloud native frameworks to meet stringent edge KPIs and network complexity
3. Leverage a broad ecosystem and evolving standards for edge computing

Reference: https://www.openness.org/docs/doc/architecture#deployment-based-on-location

Open Network Edge Services Software

OpenNESS OPEN NETWORK EDGE SERVICES SOFTWARE

certified kubernetes

OpenNESS is an edge computing software toolkit that enables **highly optimized, secure and performant edge platforms** to on-board and manage applications and network functions with **cloud-like agility across any type of network**

MODULAR

CONSUME AS BUILDING BLOCKS OR REFENCE ARCHITECTURES

MICROSERVICES BASED ARCHITECTURE

Top Use Cases

- Access Edge Aggregation Point (Cloud Native RAN + Apps)
- Near Edge (5G dUPF + Apps)
- uCPE/SD-WAN + Apps
- AI/vision inferencing apps with MEC
- Media apps with MEC

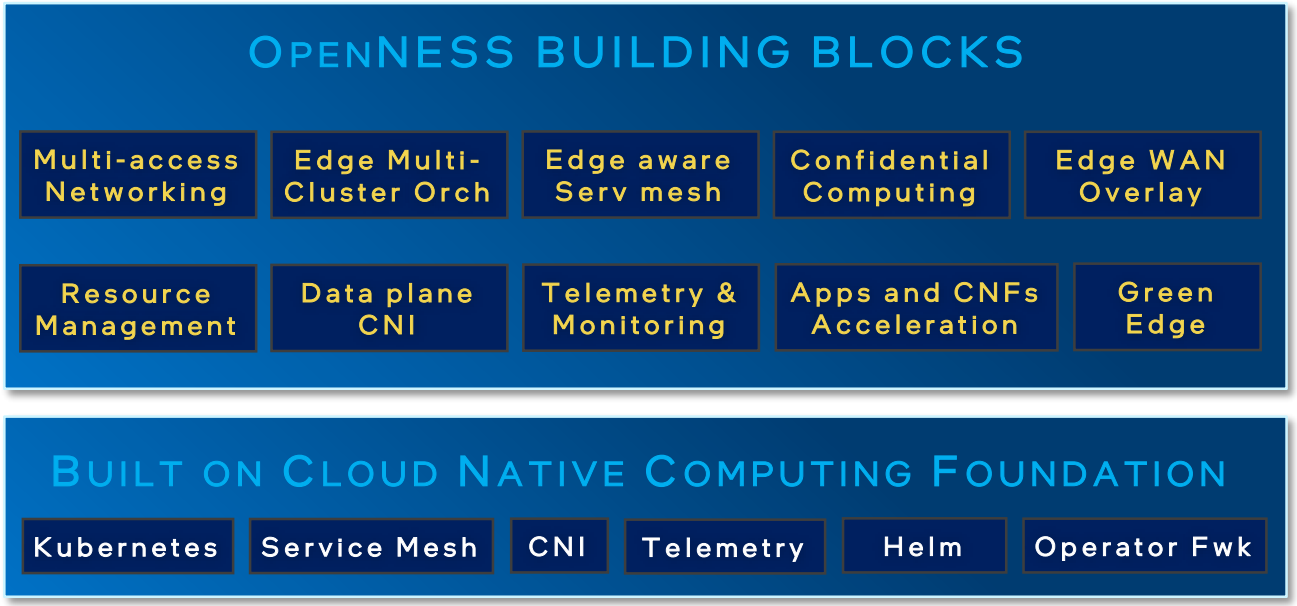

Key Features

- Optimized for Edge KPIs: throughput. determinism, QoS, latency, jitter, security
- Multi-location, Multi-Access, Multi-Cloud
- Delivered via use case specific Reference Architectures for easy of consumption and to accelerate TTM
- Industry Standards (3GPP, CNCF, ORAN, ETSI)

OpenNESS Architecture

OpenNESS – Cloud Native Edge Computing Architecture

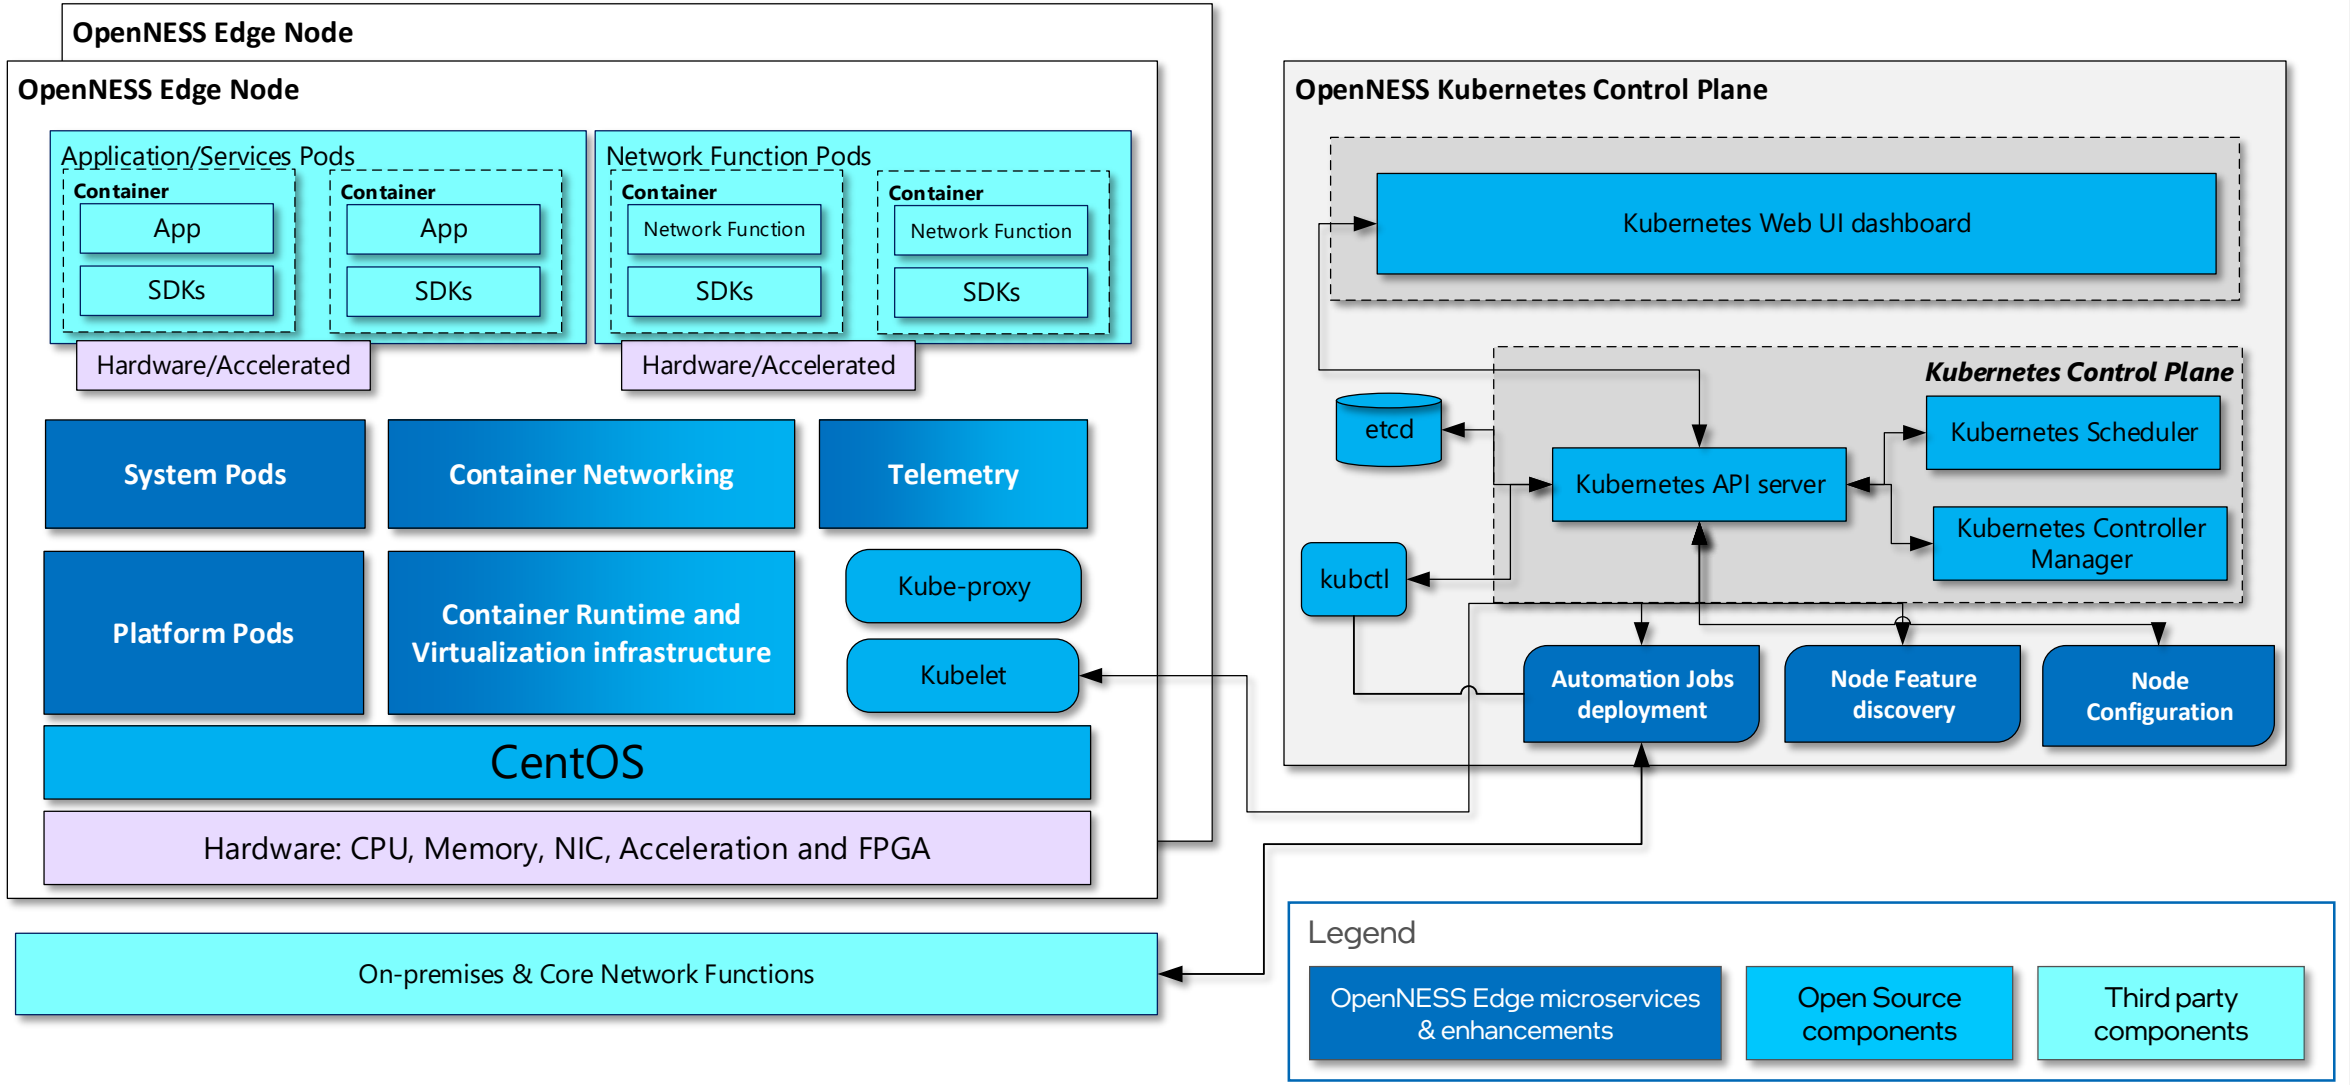

*OpenNESS Edge Node can be deployed on Network Edge or On-Premise Edge

Extending Kubernetes with Edge Capabilities

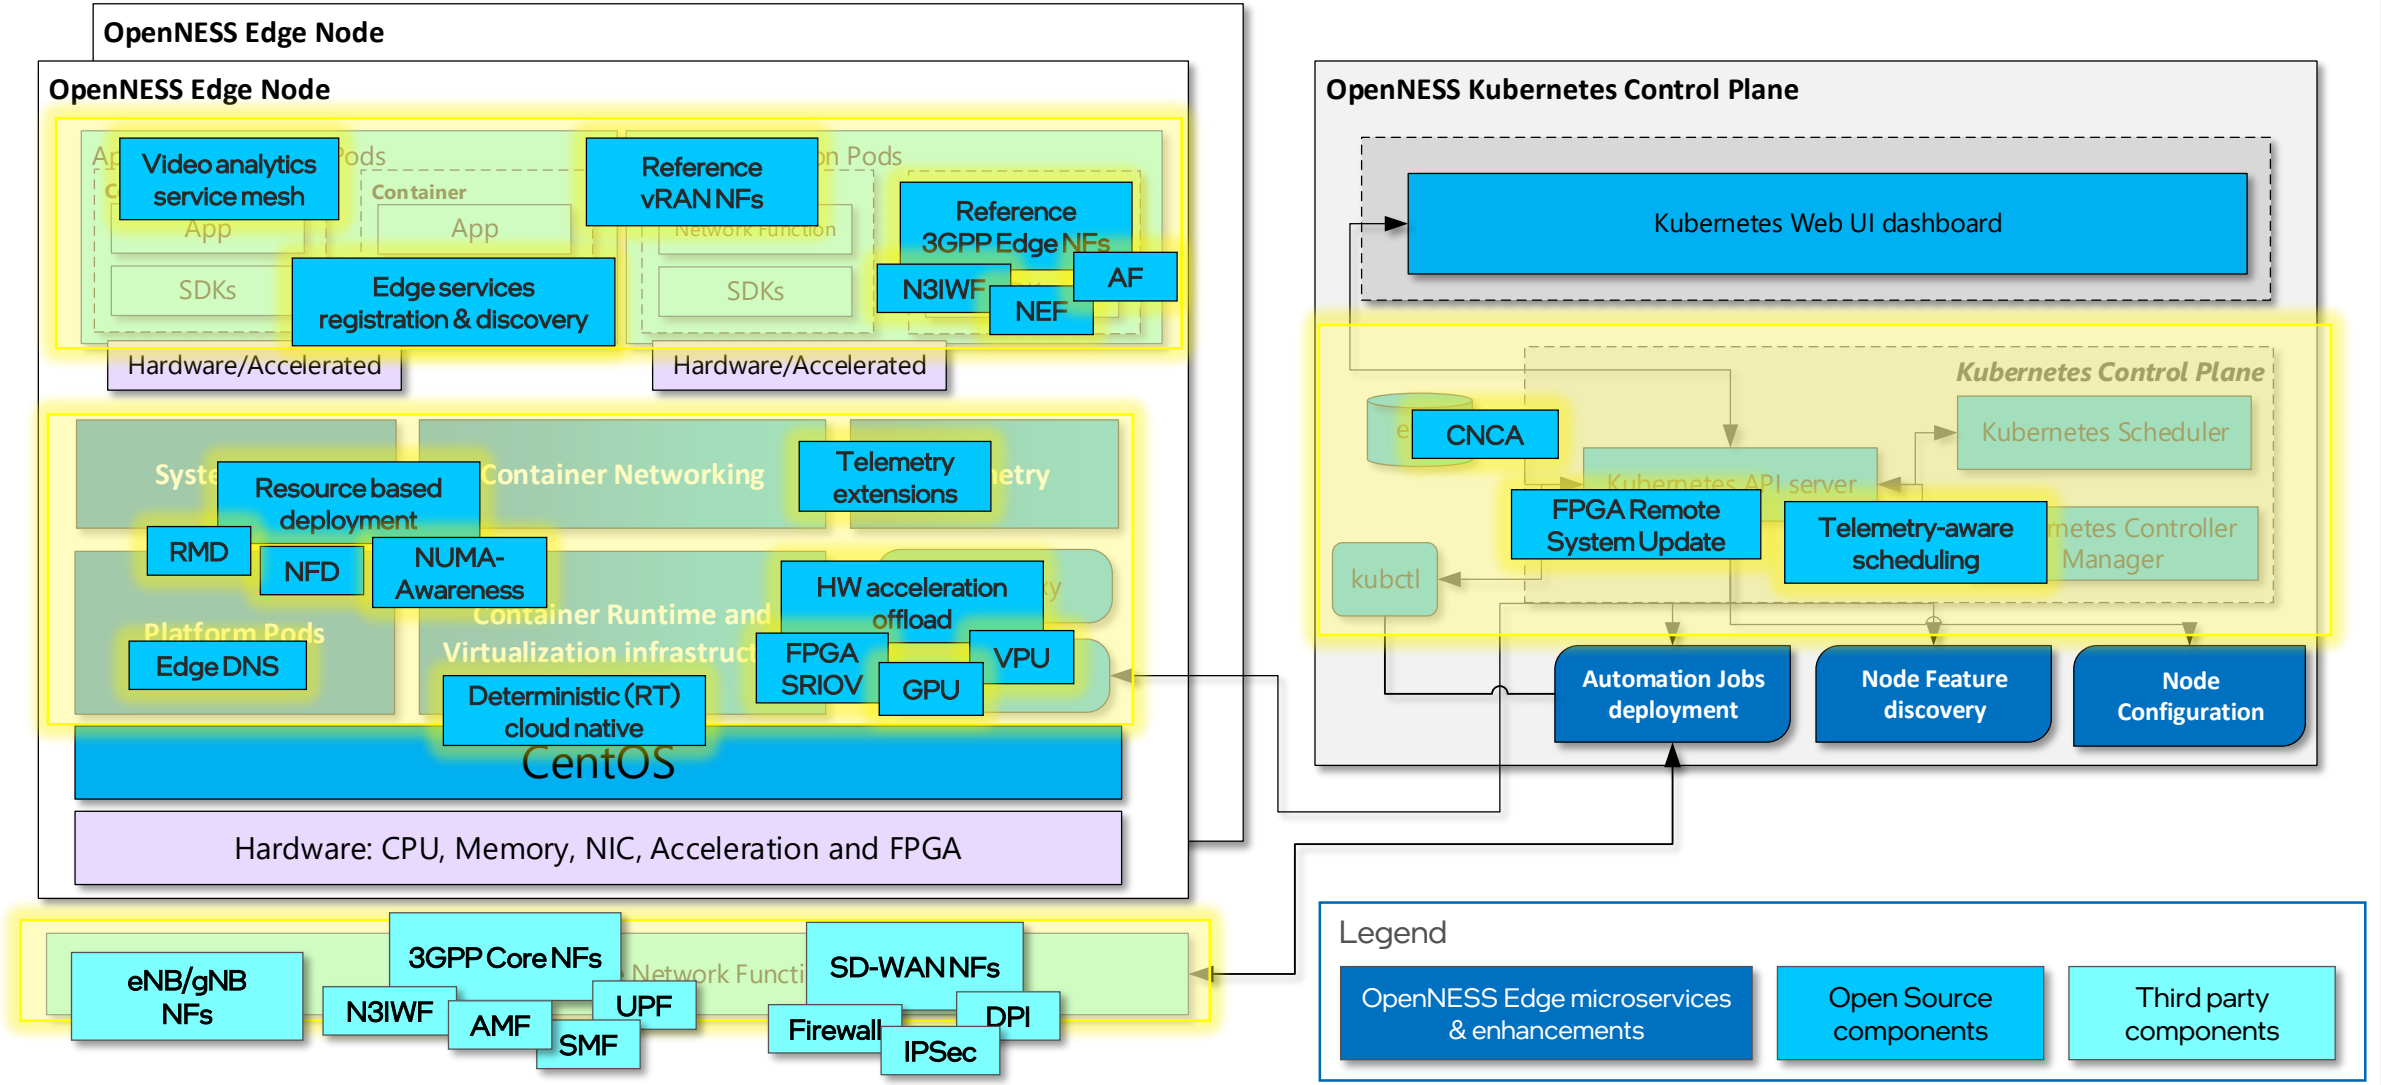

CNCA: Core Network Configuration Agent
NFs: Network Functions

RMD: Resource Management Daemon
VPU: Intel® Movidius™ Vision Processing Unit

Reference Deployment Scenarios with OpenNESS

Public Wireless Deployment

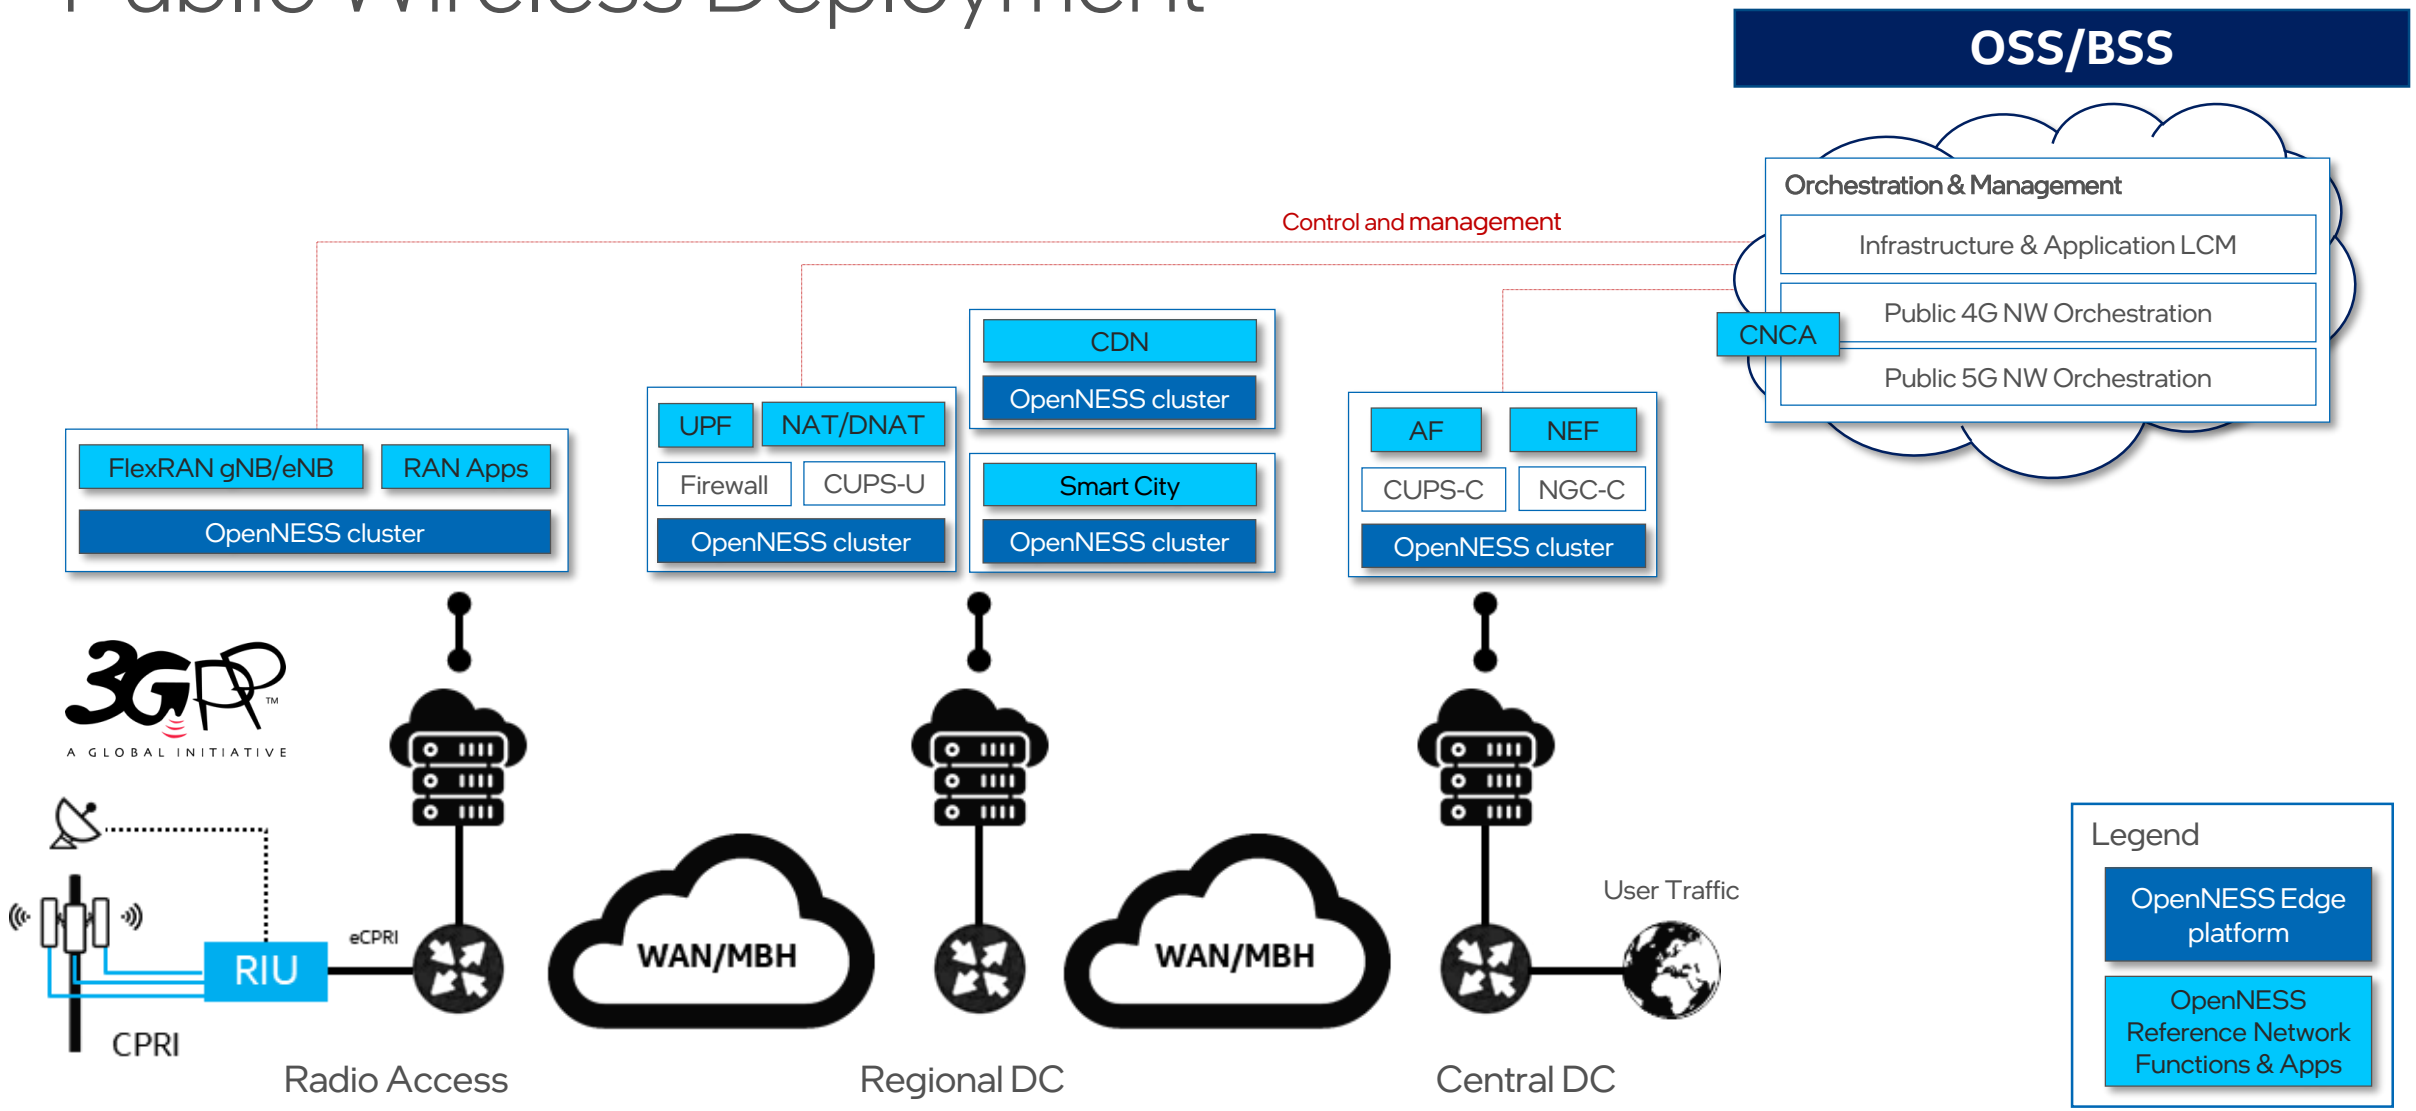

LCM: Life Cycle Management
CUPS: Control and User Plane Separation
RIU: Radio Interface Unit

NGC: Next-Gen Core (5G)
CNCA: Core Network Configuration Agent
MBH: Microwave Backhaul

DC: Data Center

Private Wireless Deployment

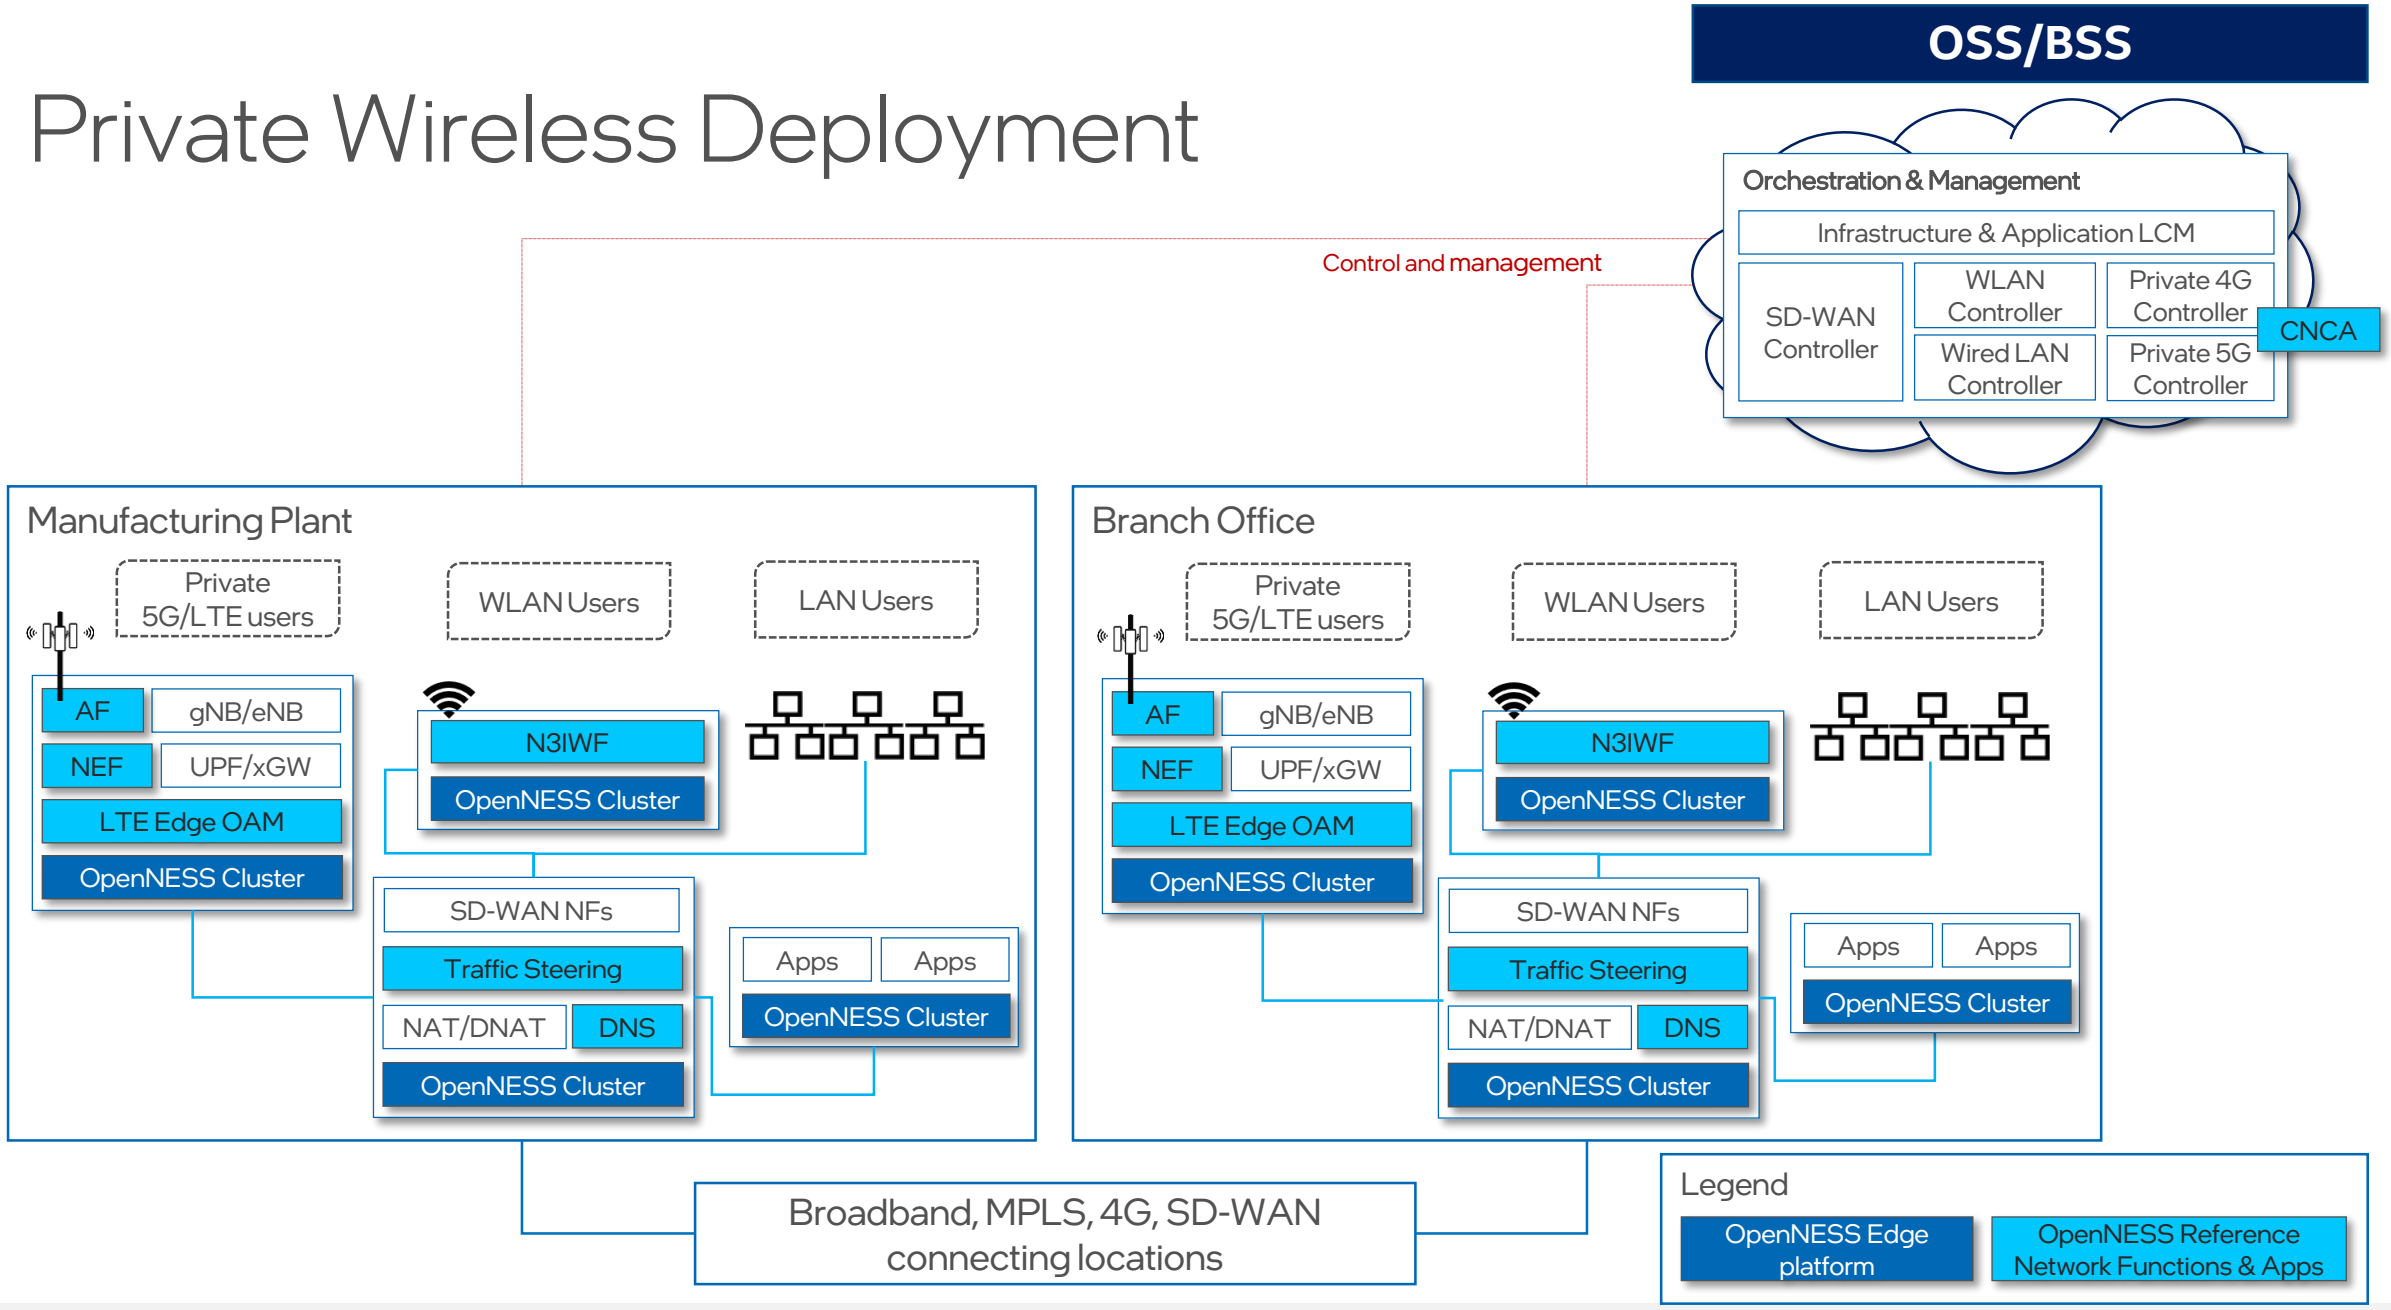

N3IWF - Non-3GPP Interworking Function
OAM: Operations And Management
CNCA: Core Network Configuration Agent

Radio Access Edge (vRAN DU)

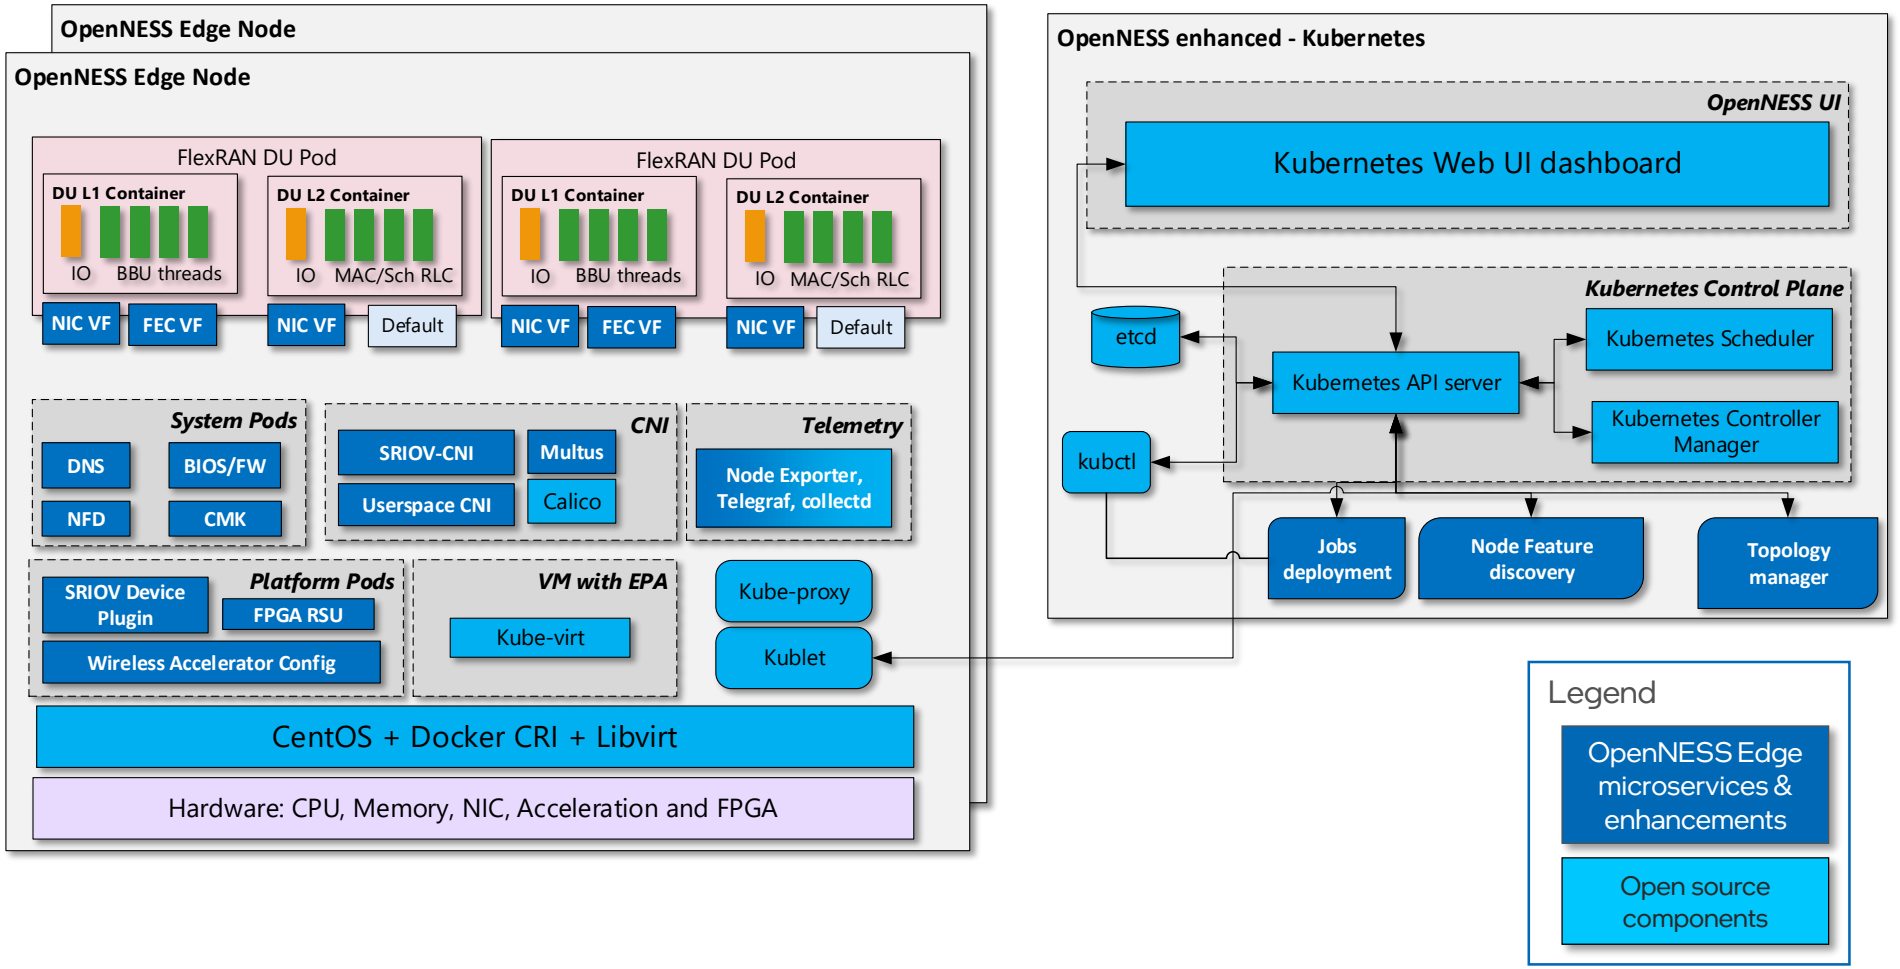

Deterministic IO

- DPDK
- NIC Offload (DDP)
- Kernel Driver
- Multiplex NW

Deterministic Acceleration

- FEC
- BBDev API (SW & HW Offload)
- FPGA & eASIC Support
- Arbitration, UL/DL Ratios

Deterministic Orchestration

- NUMA Awareness (Topology Manager)
- CMK & Native CPU Manager
- Node Feature Discovery

Deterministic Platform & Environment

- Real Time Kernel
- Complete Core Isolation
- Frequency & Power Management

References:
- https://www.openness.org/docs/doc/architecture#ran-node-flavor
- https://www.openness.org/docs/doc/ran/openness_ran

Accelerators – Intel® FPGA PAC

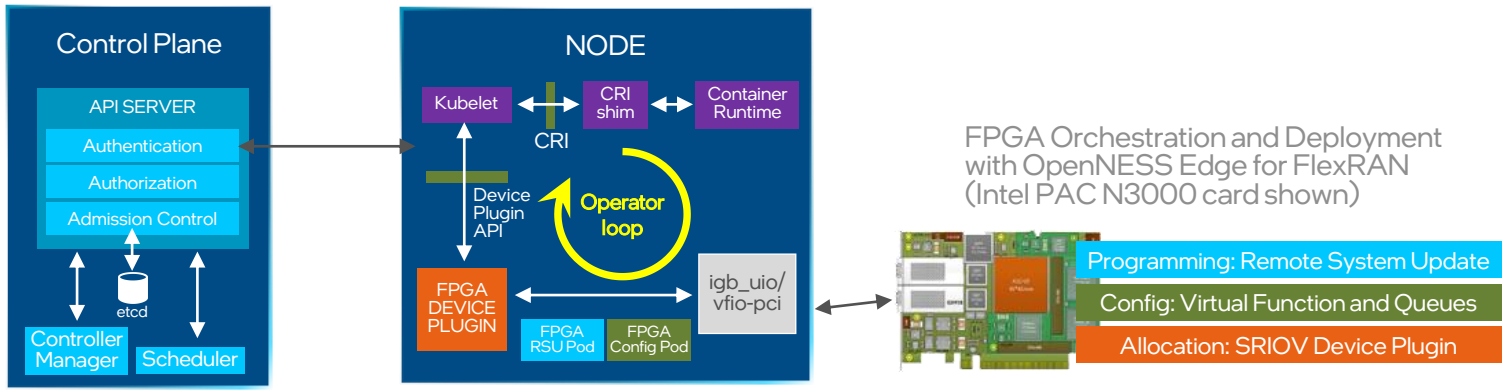

FPGA Orchestration and Deployment with OpenNESS Edge for FlexRAN (Intel PAC N3000 card shown)

Programming: Remote System Update

Config: Virtual Function and Queues

Allocation: SRIOV Device Plugin

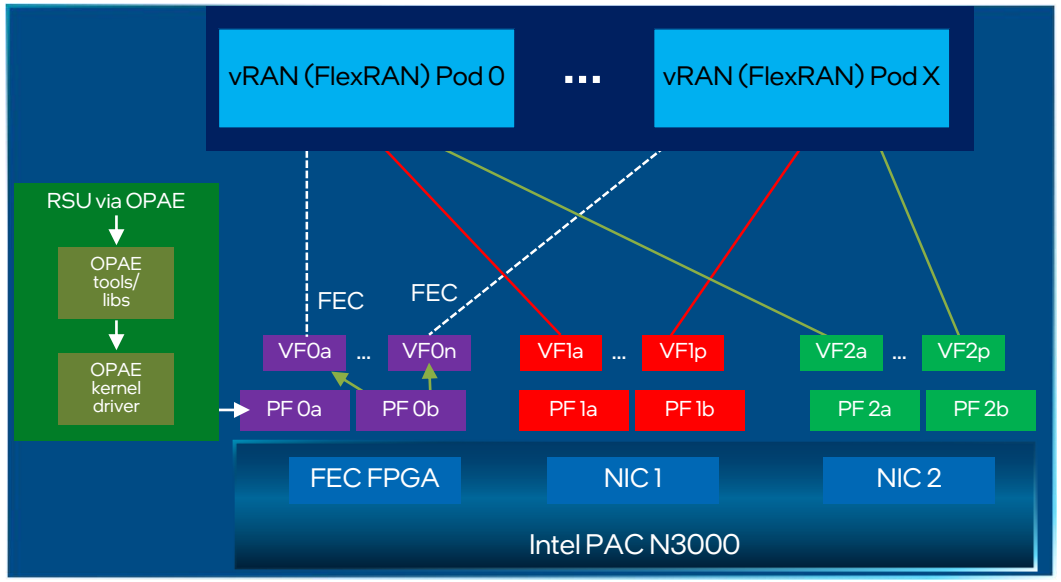

OpenNESS Network Edge Intel PAC N3000 RSU and resource allocation

What does it do?

OPERATOR FRAMEWORK

- SRIOV NIC, OPAE & SRIOV FEC operators are deployed for Intel ® FPGA PAC device
- Introduces an extended control loop to deploy device plugin & monitor the FPGA device resources
- User-defined config CRD instructs the control loop to configure the SRIOV virtual functions (VFs) with the FEC encode/decode queues and the encoding technology (LTE Turbo or 5G LDPC) and allocate to vRAN (FlexRAN) pods
- FPGA firmware image upgrades are autonomously executed through OPAE operator to the FPGA devices installed in remote sites
- In-field configuration/programming capabilities are currently not available in off-the-shelf k8s

Alternative Solutions:

- Physically program/config the FPGA (truck roll)
- Implement custom/proprietary solutions

Enables automated remote programing/configuration of the FPGA and resource allocation to it

FEC: Forward-Error Correction
SRO: Special Resource Operator
PAC: Programmable Acceleration Card

NIC: Network Interface Card
OPAE: Open Programmable Acceleration Engine

Open Visual Cloud – Video Analytics Service Mesh

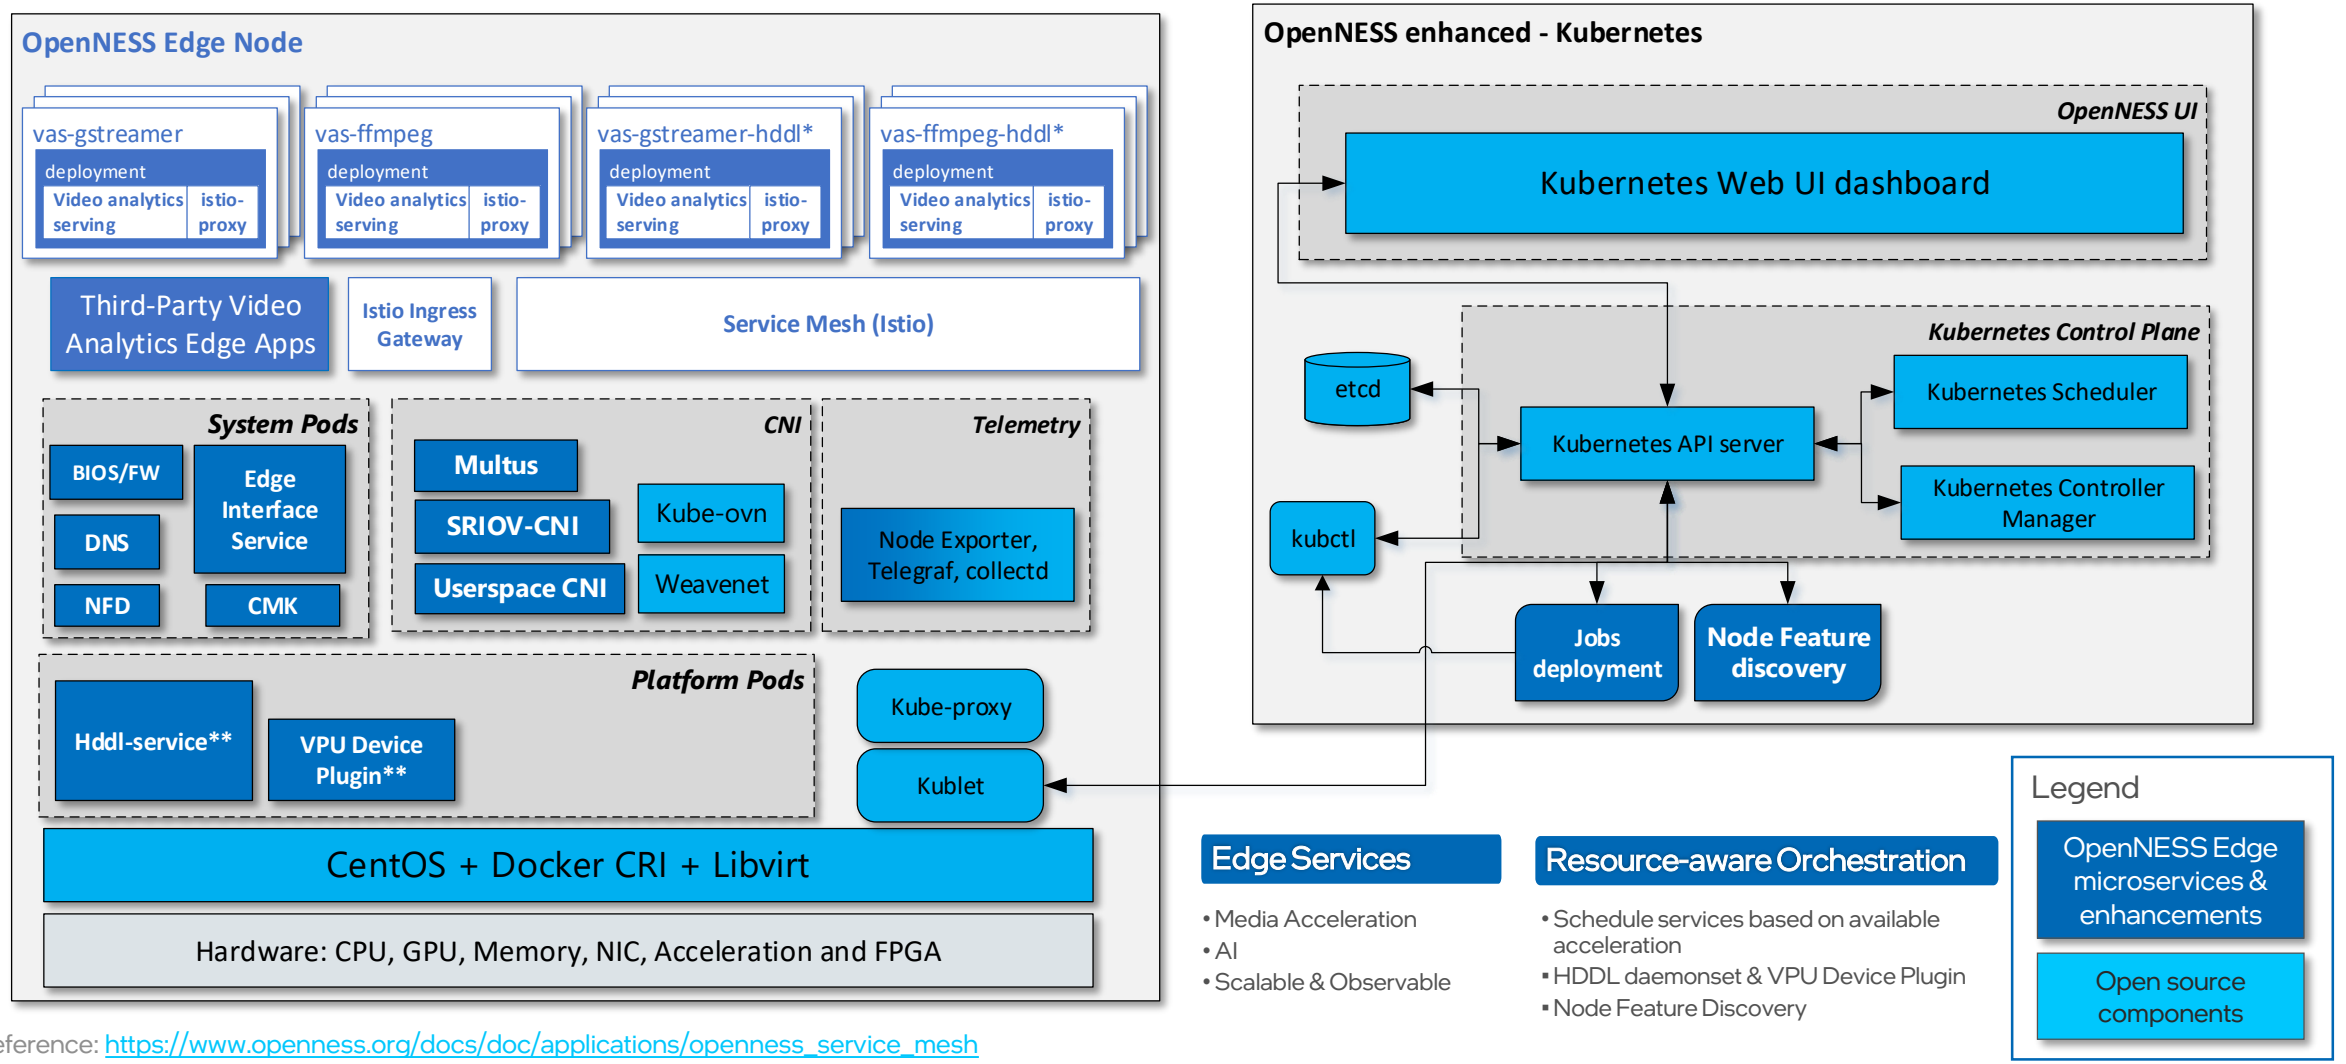

Reference: https://www.openness.org/docs/doc/applications/openness_service_mesh

*In-planning

**Provisioned when Intel® Movidius Myriad X High Density Deep Learning (HDDL) acceleration is available in the cluster

Accelerators – Media Analytics & Transcoding

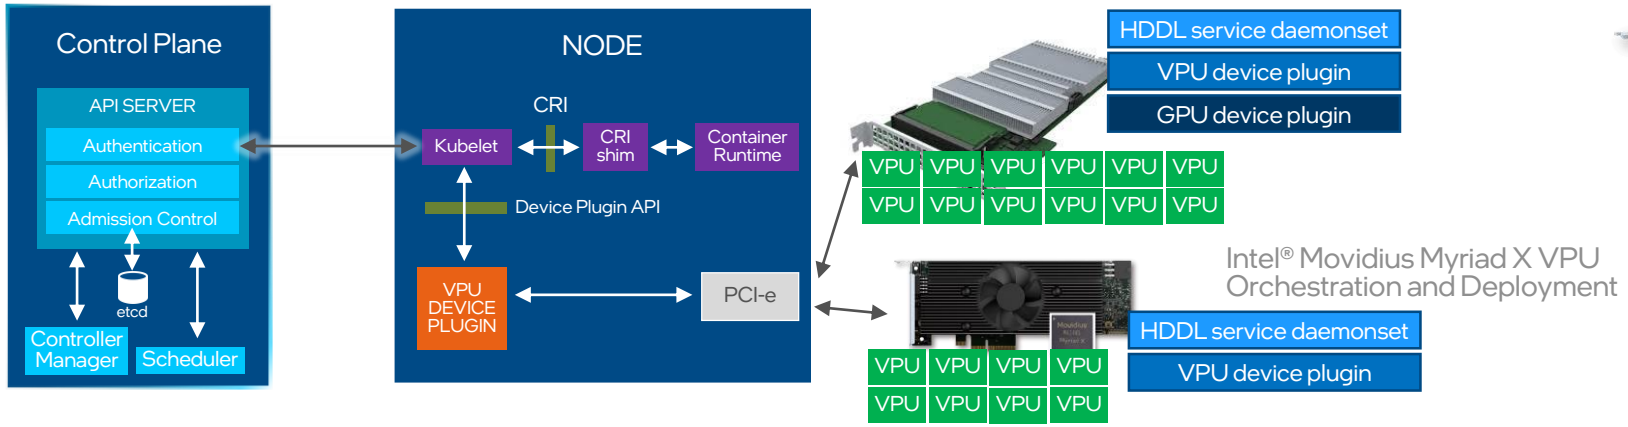

Intel® Movidius Myriad X VPU Orchestration and Deployment

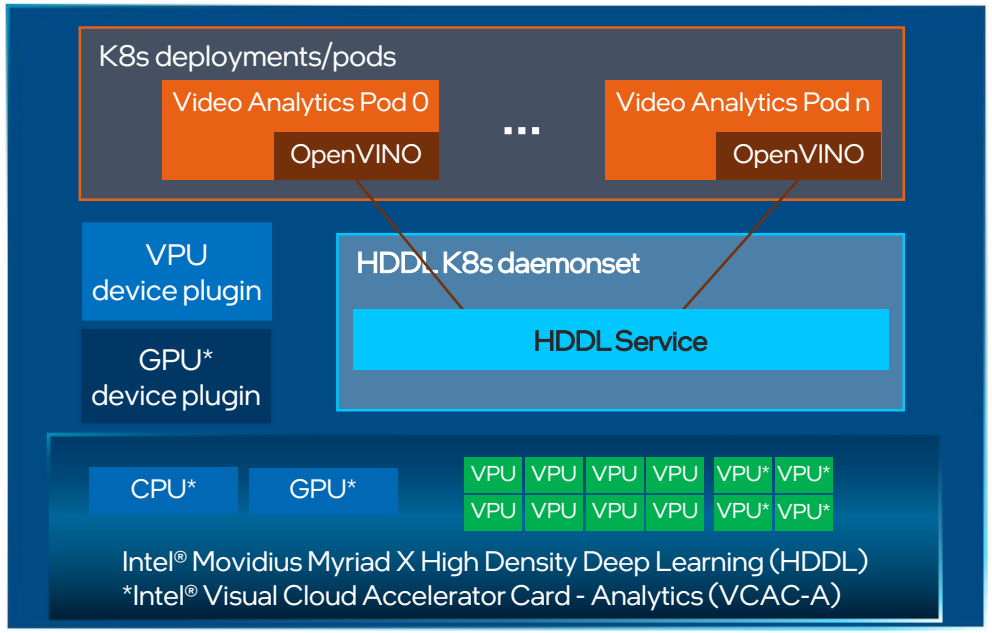

Intel® OpenVINO based seamlessly consume VPU resources through HDDL service daemonset

What does it do?

- Intel® OpenVINO based applications & services are accelerated using Intel® Movidius Myriad X VPU engines
- VPU device plugin & HDDL service is deployed as daemonset on all edge nodes that has HDDL cards installed
- GPU device plugin is deployed on edge nodes with VCAC-A cards for enhanced media transcoding using the on-board Iris® Pro Graphics GPU

VPU: Intel® Movidius™ Vision Processing Unit
OpenVINO: Open Visual Inference and Neural network Optimization (https://software.intel.com/content/www/us/en/develop/tools/openvino-toolkit.html)

Getting Started

Have an edge application idea?

App vendor participate in Intel® Network Builders → App vendor contribute app or C/VNF through GitHub → OpenNESS maintainers review and accept → Partner is added to the portal logo wall

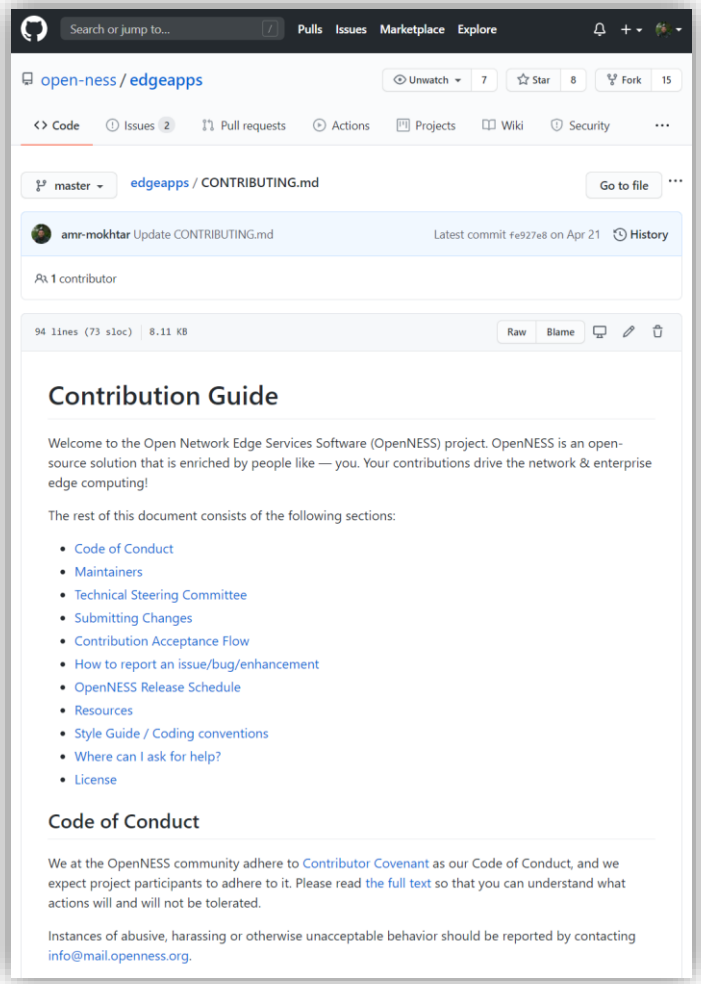

App Contribution

Enrolment Form

Edge Applications Hub

https://networkbuilders.intel.com/commercial-applications/participate

Get in touch

OpenNESS
OPEN NETWORK EDGE SERVICES SOFTWARE

certified kubernetes

www.openness.org

github.com/open-ness

developer@mail.openness.org

Edge Applications

github.com/open-ness/edgeapps

networkbuilders.intel.com/commercial-applications

Summary

OpenNESS is an edge computing software toolkit that enables highly optimized, secure and performant edge platforms to on-board and manage applications and network functions with cloud-like agility across any type of network

OPENNESS

ACCELERATING SERVICE INNOVATION AT THE EDGE

Notices and Disclaimers

- Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration.
- No product or component can be absolutely secure.
- Tests document performance of components on a particular test, in specific systems. Differences in hardware, software, or configuration will affect actual performance. For more complete information about performance and benchmark results, visit http://www.intel.com/benchmarks .
- Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit http://www.intel.com/benchmarks .
- Intel Advanced Vector Extensions (Intel AVX) provides higher throughput to certain processor operations. Due to varying processor power characteristics, utilizing AVX instructions may cause a) some parts to operate at less than the rated frequency and b) some parts with Intel® Turbo Boost Technology 2.0 to not achieve any or maximum turbo frequencies. Performance varies depending on hardware, software, and system configuration and you can learn more at http://www.intel.com/go/turbo.
- Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.
- Cost reduction scenarios described are intended as examples of how a given Intel-based product, in the specified circumstances and configurations, may affect future costs and provide cost savings. Circumstances will vary. Intel does not guarantee any costs or cost reduction.
- Intel does not control or audit third-party benchmark data or the web sites referenced in this document. You should visit the referenced web site and confirm whether referenced data are accurate.
-

LEGAL Disclaimers

Performance results are based on testing as of dates shown in configuration and may not reflect all publicly available security updates. See configuration disclosure for details. No product or component can be absolutely secure. Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit www.intel.com/benchmarks.

Optimization Notice: Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. Check with your system manufacturer or retailer or learn more at intel.com.

Up to 2.5x Accelerated AI Performance: Workload: images per second using AIXPRT Community Preview 2 with Int8 precision on ResNet-50 and SSD-Mobilenet-v1 models. Intel preproduction system, ICL-U, PL1 15w, 4C/8T, Turbo TBD, Intel Gen11 Graphics, GFX driver preproduction, Memory 8GB LPDDR4X-3733, Storage Intel SSD Pro 760P 256GB, OS Microsoft Windows 10, RS5 Build 475, preprod bios. Vs. Config – HP spectre x360 13t 13-ap0038nr, Intel® Core™ i7-8565U, PL1 20w, 4C/8T, Turbo up to 4.6Ghz, Intel UHD Graphics 620, Gfx driver 26.20.100.6709, Memory 16GB DDR4-2400, Storage Intel SSD 760p 512GB, OS – Microsoft Windows 10 RS5 Build 475 Bios F.26.

Approximately 2x Graphics Improvements: Workload: 3DMark11 v 1.0.132. Intel PreProduction ICL U4+2 15W Configuration (Assumptions):, Processor: Intel® Core™ i7 (ICL-U 4+2) PL1=15W TDP, 4C8T, Memory: 2x8GB LPDDR4-3733 2Rx8, Storage: Intel® 760p m.2 PCIe NVMe SSD with AHCI Microsoft driver, Display Resolution: 3840×2160 eDP Panel 12.5", OS: Windows* 10 RS5-17763.316, Graphics driver: PROD-H-RELEASES_ICL-PV-2019-04-09-1006832. Vs config – Intel PreProduction WHL U4+2 15W Configuration (Measured), Processor: Intel® Core™ i7-8565U (WHL-U4+2) PL1=15W TDP, 4C8T, Turbo up to 4.6Ghz, Memory: 2x8GB DDR4-2400 2Rx8, Storage: Intel® 760p m.2 PCIe NVMe SSD with AHCI Microsoft driver, Display Resolution: 3840×2160 eDP Panel 12.5", OS: Windows* 10 RS4-17134.112., Graphics driver: 100.6195

Nearly 3x Faster Ultra-High-Speed Wireless Connectivity: 802.11ax 2×2 160MHz enables 2402Mbps maximum theoretical data rates, ~3X (2.8X) faster than standard 802.11ac 2×2 80MHz (867Mbps) as documented in IEEE 802.11 wireless standard specifications and require the use of similarly configured 802.11ax wireless network routers.

Forward-Looking Statements: Statements in this release that refer to future plans and expectations, including with respect to Intel's future technologies and the expected benefits of such technologies, are forward-looking statements that involve a number of risks and uncertainties. Words such as "anticipates," "expects," "intends," "goals," "plans," "believes," "seeks," "estimates," "continues," "may," "will," "would," "should," "could," and variations of such words and similar expressions are intended to identify such forward-looking statements. Statements that refer to or are based on estimates, forecasts, projections, uncertain events or assumptions, including statements relating to total addressable market (TAM) or market opportunity, future products and the expected availability and benefits of such products, and anticipated trends in our businesses or the markets relevant to them, also identify forward-looking statements. Such statements are based on current expectations and involve many risks and uncertainties that could cause actual results to differ materially from those expressed or implied in these forward-looking statements. Important factors that could cause actual results to differ materially from the company's expectations are set forth in Intel's most recent earnings release dated April 25, 2019, which is included as an exhibit to Intel's Form 8-K furnished to the SEC on such date. Additional information regarding these and other factors that could affect Intel's results is included in Intel's SEC filings, including the company's most recent reports on Forms 10-K and 10-Q. Copies of Intel's Form 10-K, 10-Q and 8-K reports may be obtained by visiting our Investor Relations website at www.intc.com or the SEC's website at www.sec.gov.