

5G 云网融合论坛

基于 CERA 参考架构 构建 5G 现场边缘

高纪明

英特尔 网络平台事业部 网络平台解决方案架构师



intel®

免责声明

- 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.
- © Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.

大量新兴业务的涌现对业务系统提出更高要求



80%

数据流量是视频相关的流量, by 2021³

75%

企业数据是在数据中心外产生并完成处理, by 2025¹

43%

AI任务发生在边缘设备上(相对于云), in 2023

▼ Require New KPIs ▼



多种因素驱动算力在边缘快速发展



边缘应用要求业务能力融合

跨现场边缘和网络边缘的IT, OT & 网络能力融合

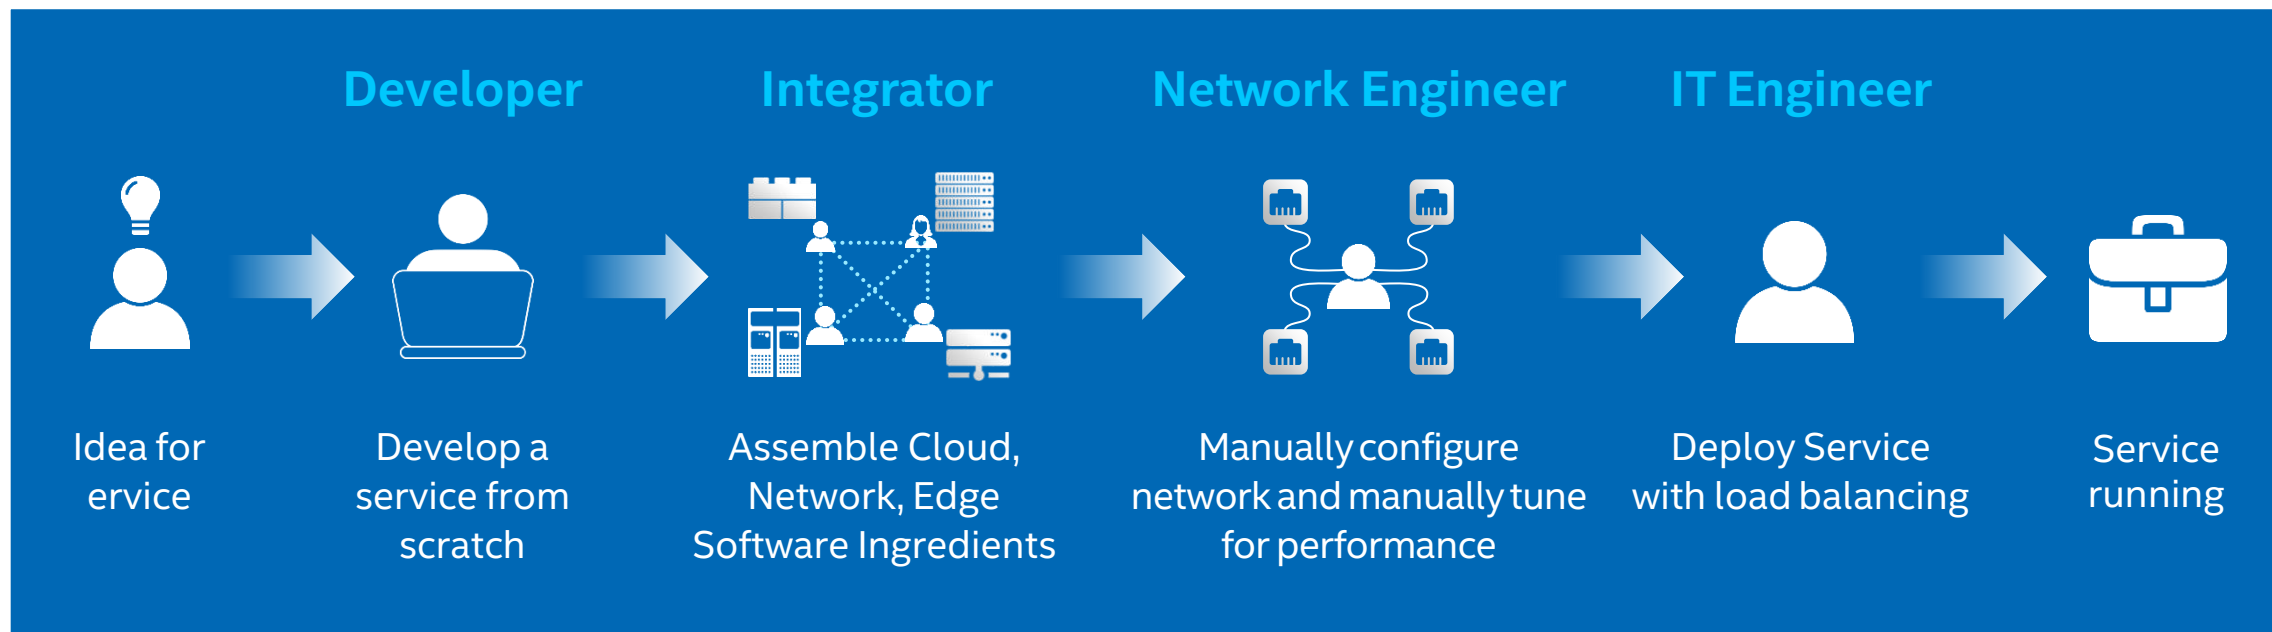
一个优化的业务平台：融合各类能力到统一的平台



边缘业务的部署复杂, 耗时且成本高

Devops新业务部署时间: 几个月到一年

TCO: \$ ↑



企业需要灵活, 敏捷和经济高效的方法来部署边缘服务

企业业务创新痛点 需要开源云原生的边缘计算平台



原有基础设施业务
能力差，需要拆分



巨型实体难以支持
业务多样性



网络复杂性



IOT领域创新难

如何帮助这些企业？

为业务创新而进行的基础设施变革要满足以下几个条件：

- ✓ 降低Total Cost of Ownership
- ✓ 加快TIME TO MARKET
- ✓ 减低网络与生态系统整合的复杂度

英特尔: 开源云原生融合边缘参考架构

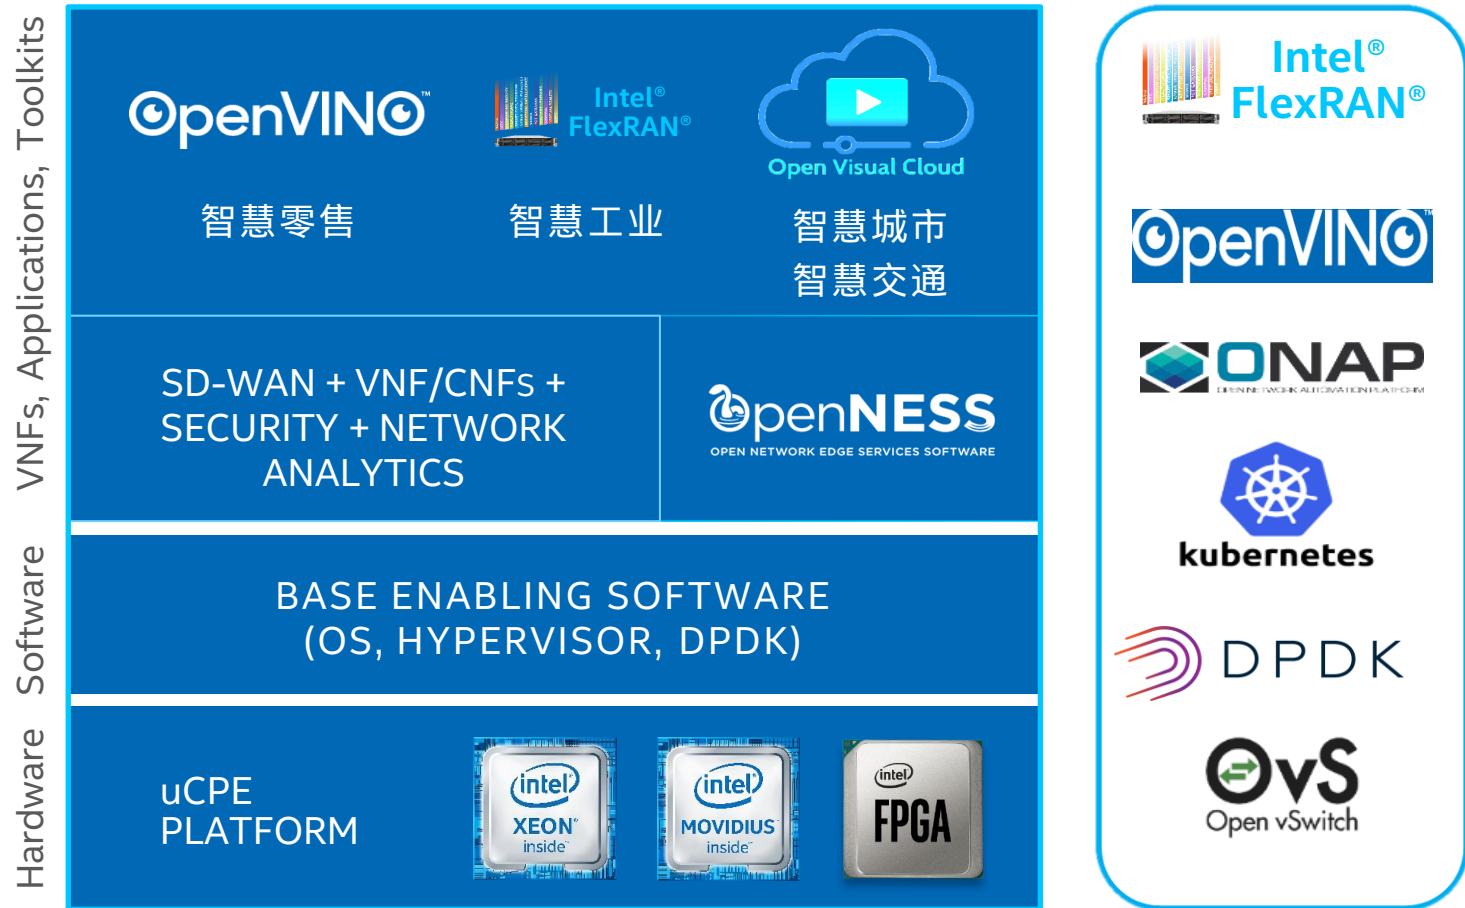
Converged Edge Reference Architecture (CERA)

适用于现场边缘的参考架构CREA

Converged Edge Reference Architecture

基于优化平台的业务创新：

- 在统一的云原生基础设施上支持业务的多样性
- 提供丰富的网络 / IoT / AI / Media 功能，简化网络和应用部署复杂度
- 与实际业务对齐，提供预集成一体化方案




参考母版: 硬件规格/外观形态

Common Mother Board

1S Intel® Xeon® D Processor

- 64 GB RAM across 4x memory channels
- 4x 10GbE plus \geq 4x GbE
- 2x SSD
- USB ports + VGA/HDMI port
- 2x PCIe FH,³/₄L with x16

Available Form Factors




Rack-mount Wall-mount Pole-mount


1S 2nd Intel® Xeon® or Intel® Xeon® Scalable Processor

- 96 GB RAM across 6x memory channels
- 2x 10GbE
- 2x SSD
- USB ports + VGA/HDMI port
- 2x PCIe FH, ³/₄L with x16


Available Form Factors



Rack-mount Wall-mount



Wall-mount server
(<300 mm wide)
Indoor, wall or shelf



1U, 19" Rack-mount server
(<450 mm wide)
Indoor IT Rack



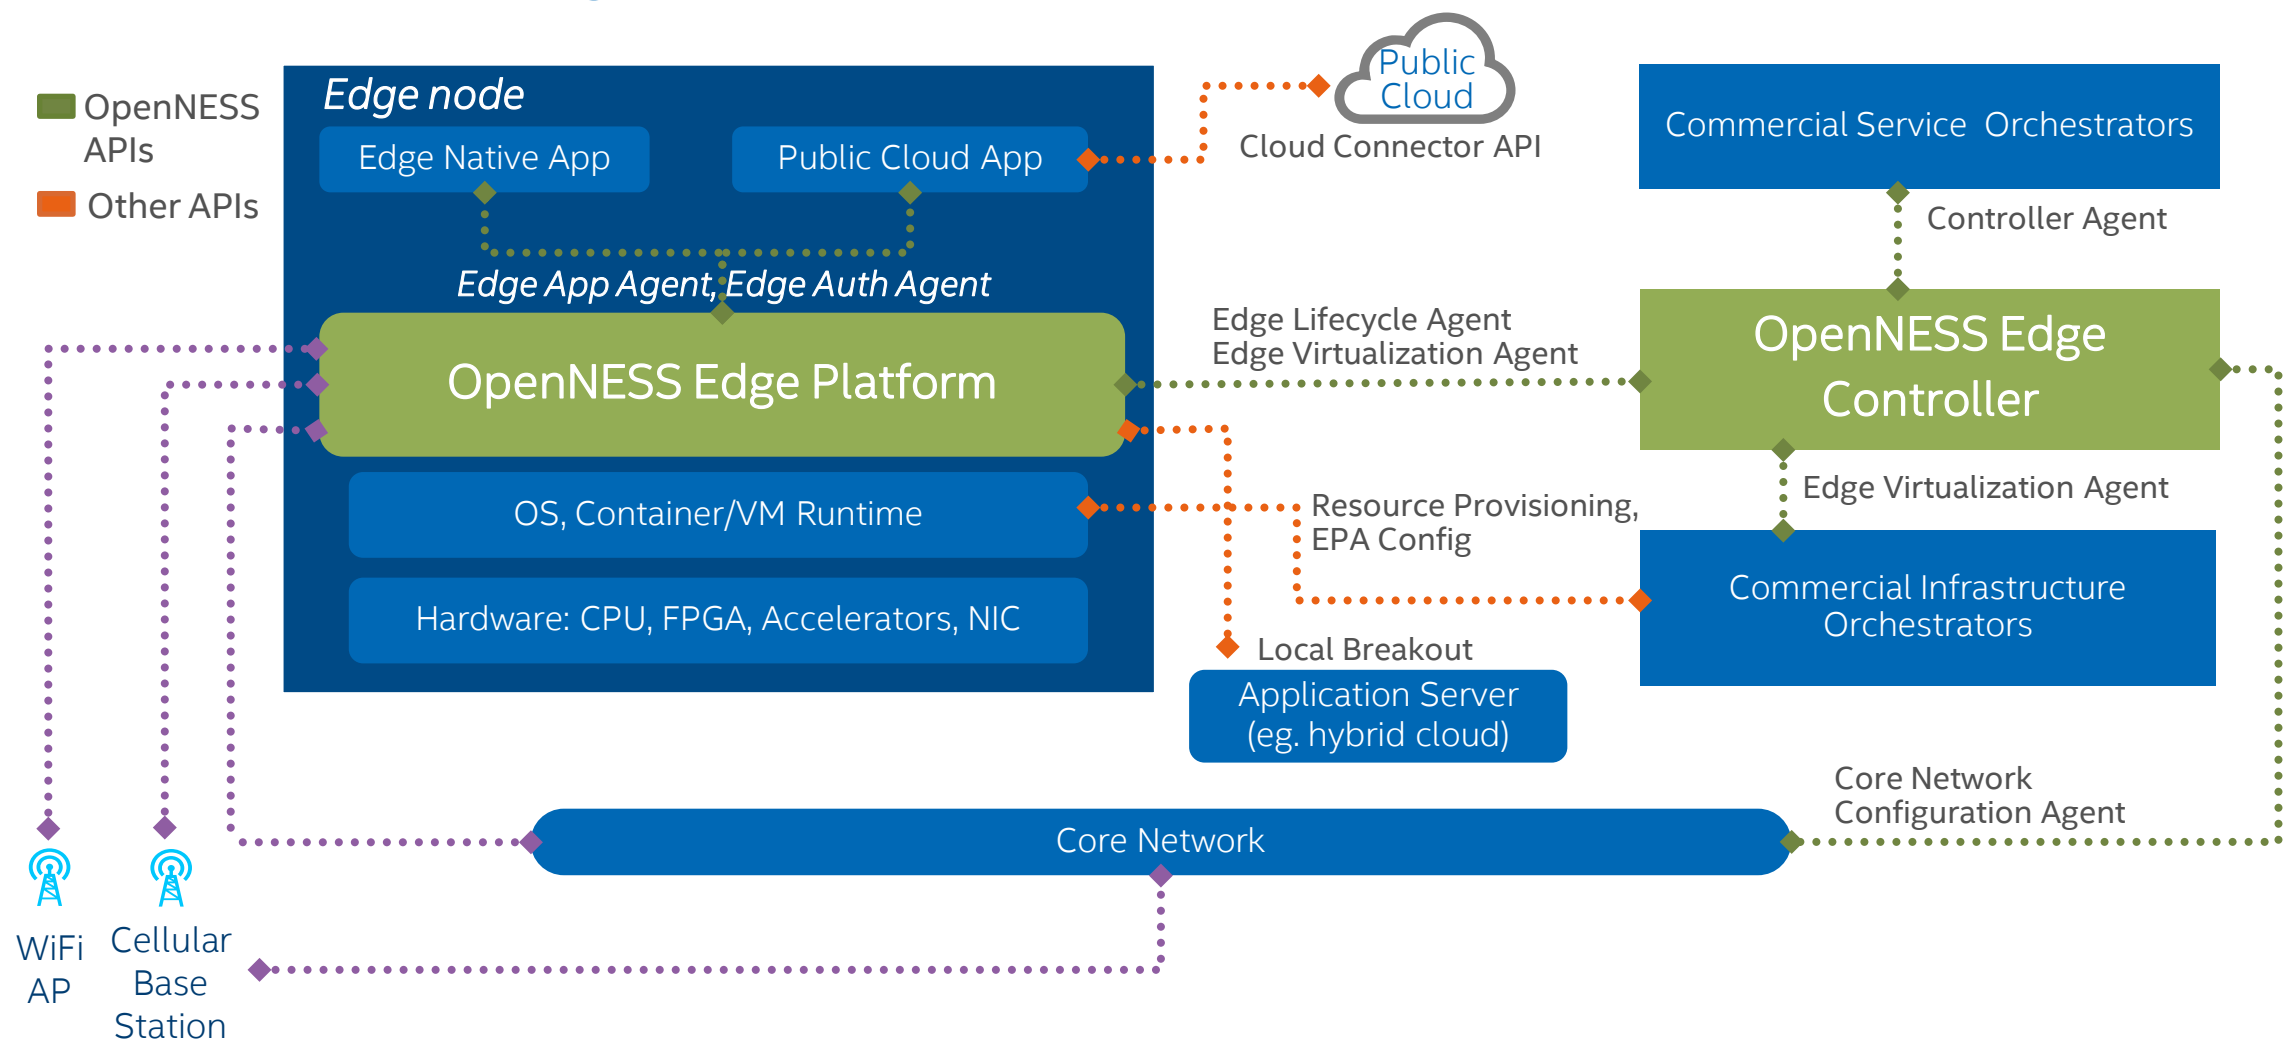
Pole-mount server for Intel® Xeon® D Processor (<350 mm wide)



Intel® AVX-512 for enabling 4G and 5G RAN. RAN require a min of 16c CPU configuration

OpenNESS: 英特尔网络边缘业务软件套件

Open Network Edge Services Software



OpenNESS 微服务架构

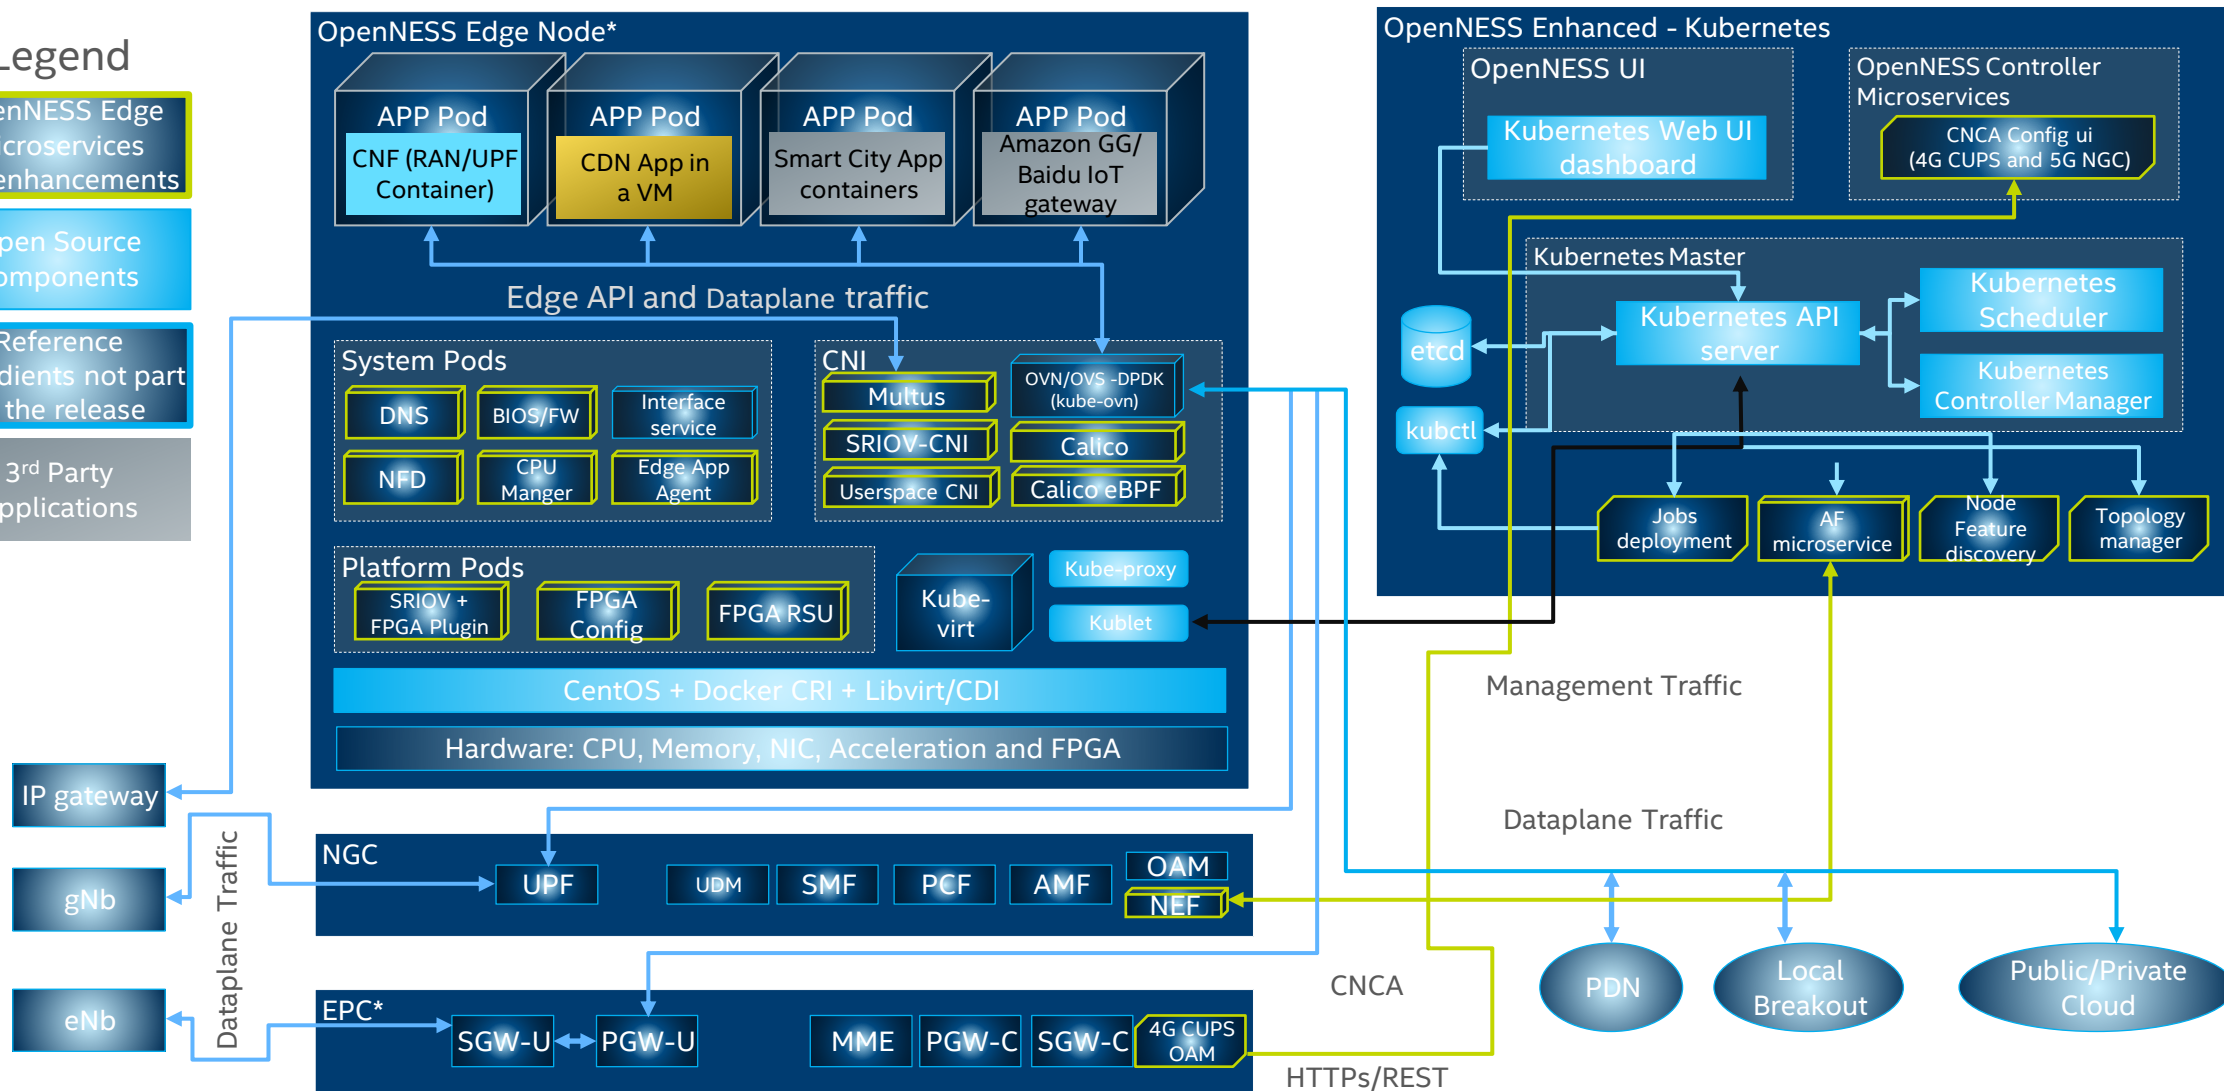
Legend

OpenNESS Edge
Microservices
and enhancements

Open Source
components

Reference
ingredients not part
of the release

3rd Party
Applications



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

OpenNESS微服务: Node Feature Discovery (NFD)

NFD是EPA(Enhanced Platform Awareness)组件的一部分:

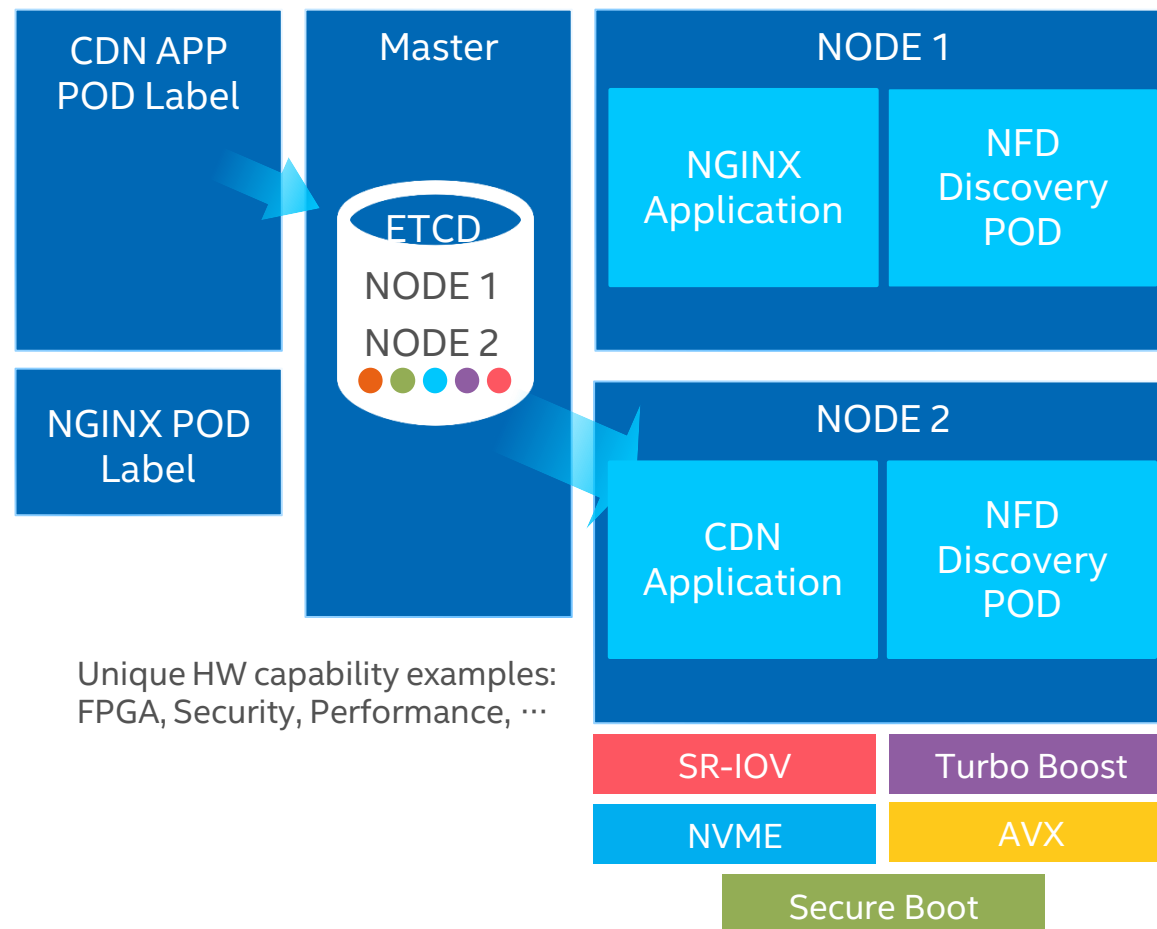
主要功能:

- NFD将边缘节点的能力通告给协调器
- 协调器将APP部署到最优的边缘节点上, 实现最佳的边缘KPI

其他方案:

- 如果没有NFD, 部署的应用程序性能会降低
- 或者实现某种自定义解决方案

通过提前收集节点的硬件资源功能特性, 实现工作负载与特殊硬件资源的智能匹配, 以优化应用程序的性能和可管理性。



OpenNESS微服务: NUMA Topology Manager

TM是EPA组件的一部分:

主要功能:

非平衡NUMA节点，网卡仅连接到一个socket上:

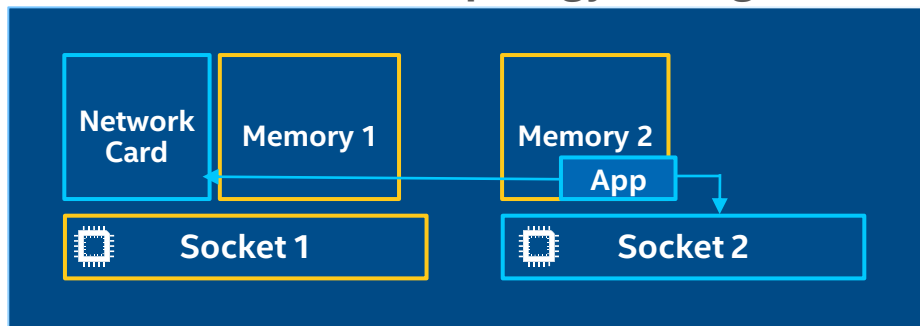
- TM微服务可以暴露网卡连到哪个socket上
- 协调器能够将对性能敏感的边缘应用程序部署到合适的socket上，满足吞吐量和延迟的性能要求。

其他方案:

资源分配不理想，导致性能下降

编排器具有NUMA意识，可以优化性能敏感边缘应用程序的部署

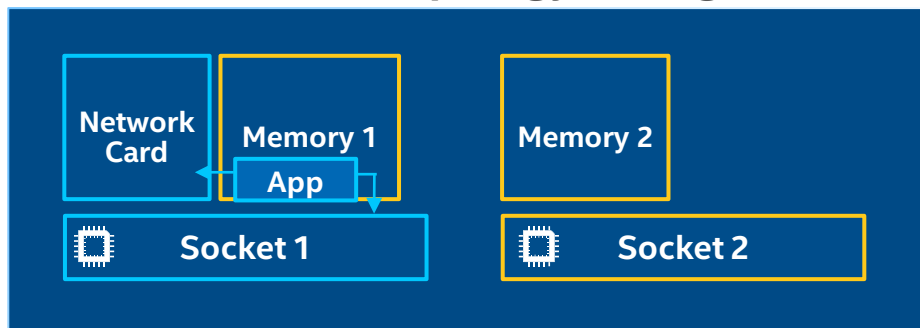
Node without Topology Manager



Not NUMA Aware

Application may be deployed to the wrong socket leading to performance degradation

Node with Topology Manager



NUMA Aware

Application always deployed on the socket with the desired resources for deterministic & reliable performance

NUMA: 非一致内存访问

OpenNESS微服务: Telemetry Aware Scheduler

TAS是EPA组件的一部分:

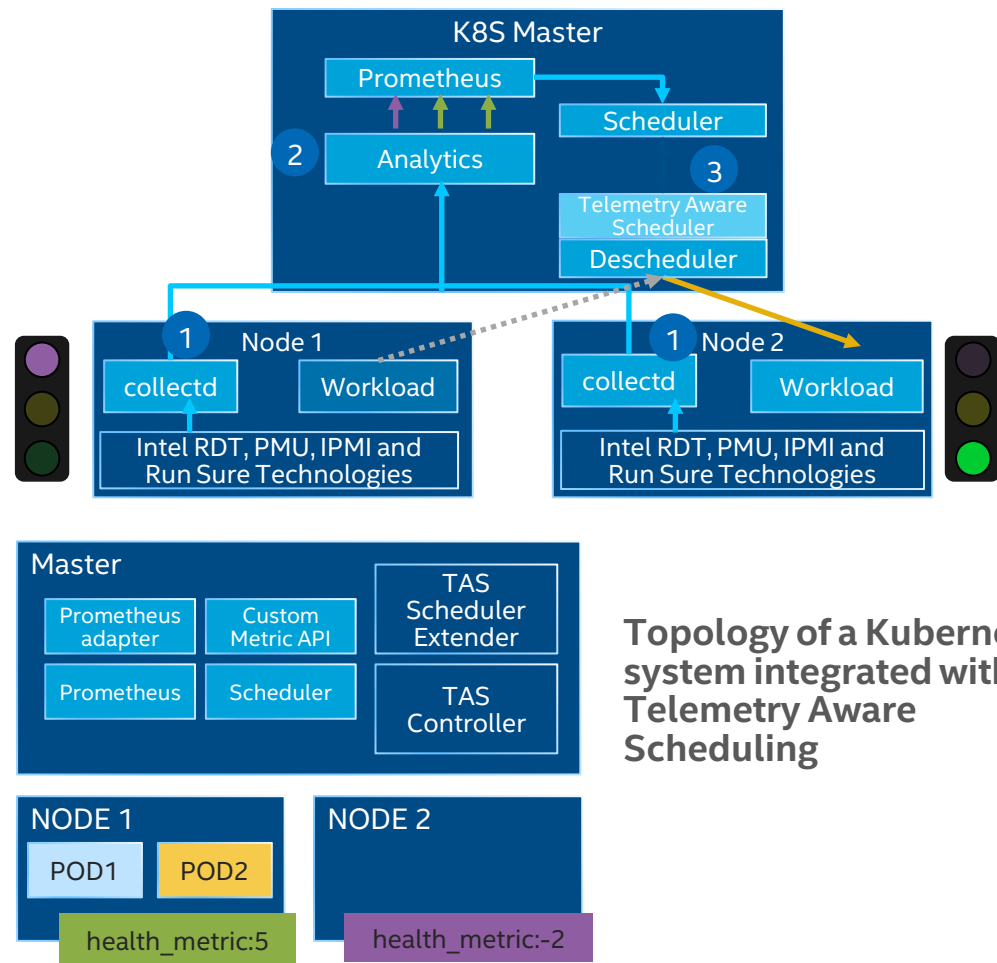
主要功能:

1. TAS收集整个平台telemetry数据给Master
2. master主机能够监视各个节点的性能，并动态部署/迁移工作负载以获得最佳性能
3. 示例：位于地点A的CDN为一个服务提供商支持高收入的流媒体应用以及语音应用。
 - Loc A由于本地紧急情况而导致语音应用程序过载。
 - 使用来自TAS的Telemetry，K8s Master可以观察到位置A的瓶颈并确定位置B能够增加工作量。
 - 服务提供商快速无缝地将流媒体应用程序移动到Loc B，而对客户体验的影响最小。

其他方案:

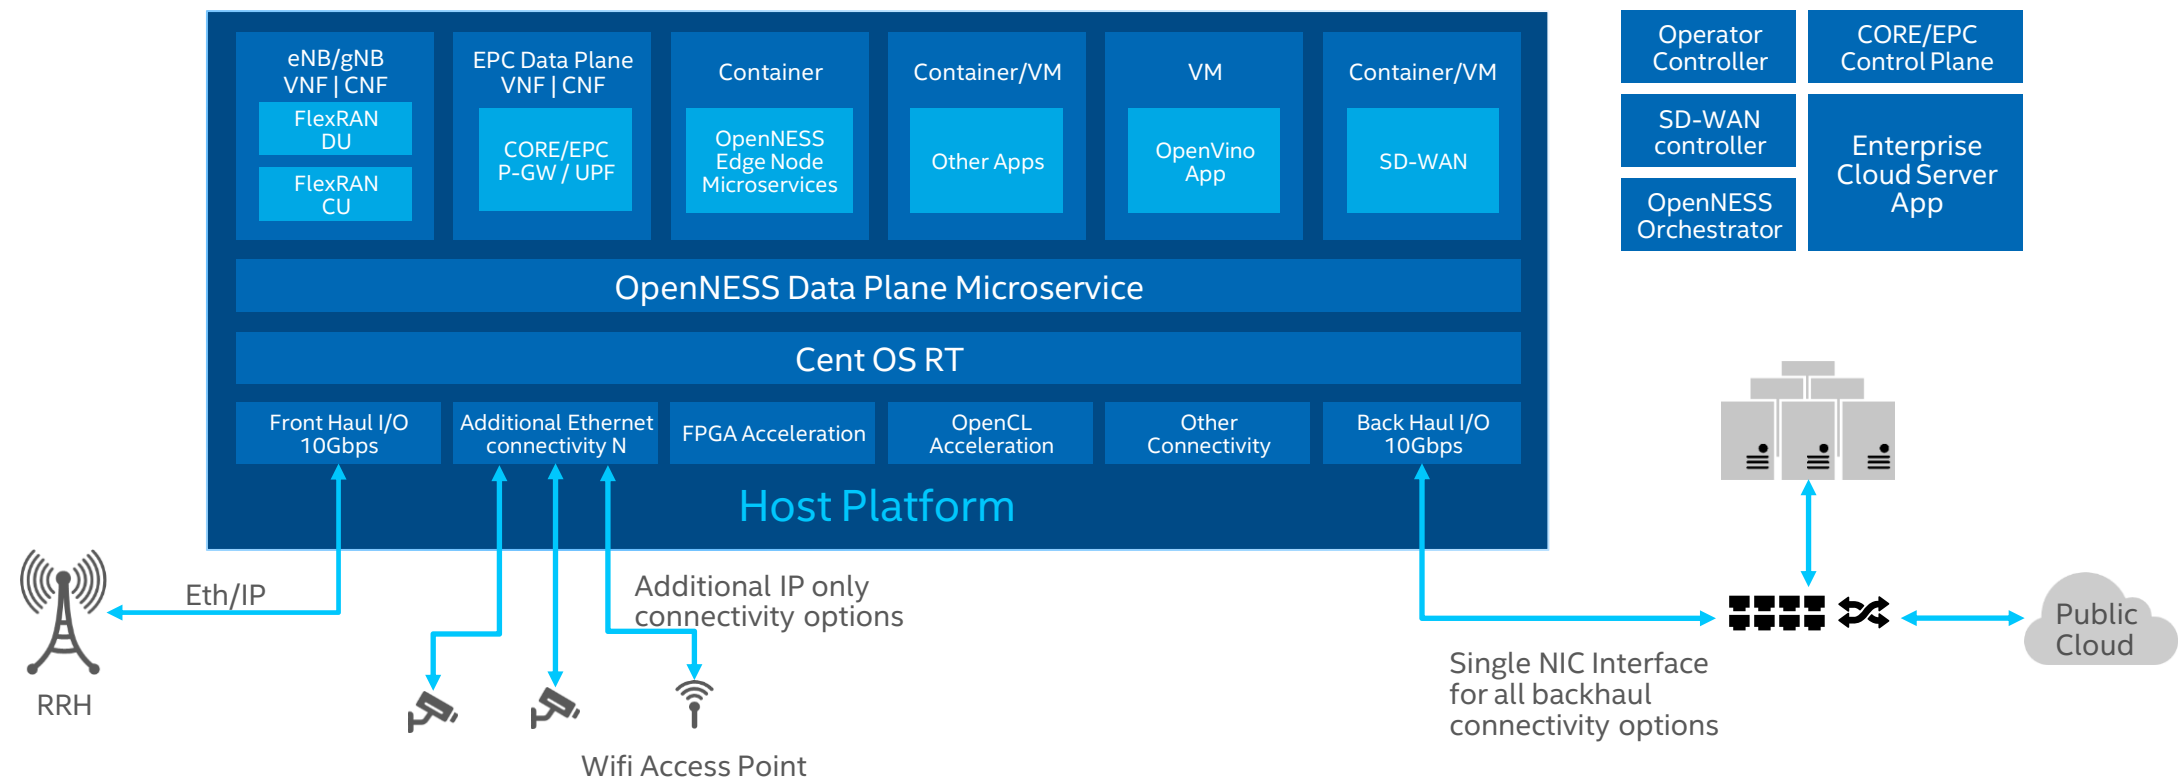
缺少这种telemetry集成，就无法获取识别瓶颈和适合动态卸载的节点信息的实时数据

通过边缘节点telemetry指标，使服务提供商能够实施基于规则的工作负载布置，以实现业务最佳性能和弹性部署。



All-in-One的一体化参考解决方案

Workloads: RAN, CUPs/UPF, IOT, SD-WAN, Edge Services



基于现场边缘CERA的垂直行业解决方案

基于CERA平台，为垂直行业定义个性化平台

垂直行业

智慧工业

Networking, AI, edge services,
industrial insights & private
wireless

- ⊕ Fault Detection Apps
- ⊕ Predictive Maintenance Apps
- ⊕ Edge Insights Software

智慧交通

Networking, AI, edge
services & private wireless

- ⊕ Traffic Monitoring Apps
- ⊕ Pedestrian Detection Apps
- ⊕ Near Miss/Collision Detection Apps

零售, CDN, DSS

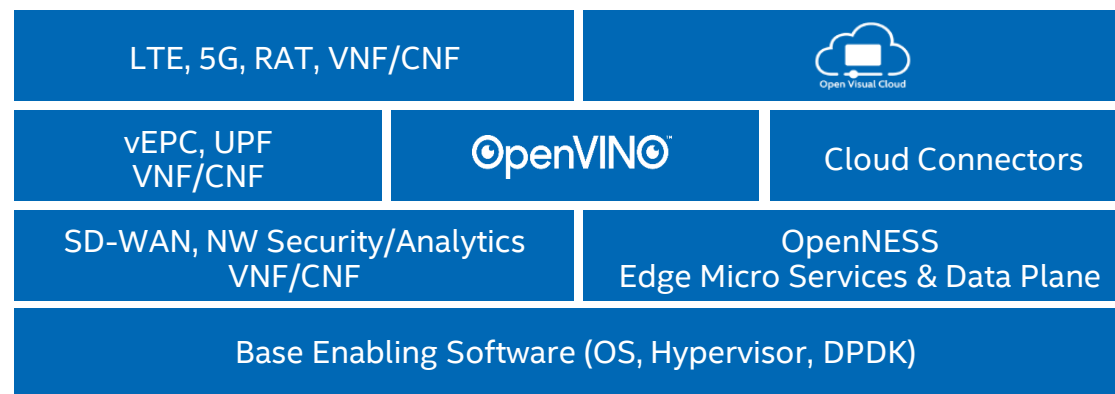
Networking, AI, edge services,
visual cloud & private wireless

- ⊕ Customer Hotspot Apps
- ⊕ POS Apps
- ⊕ Traffic Flow Analysis Apps

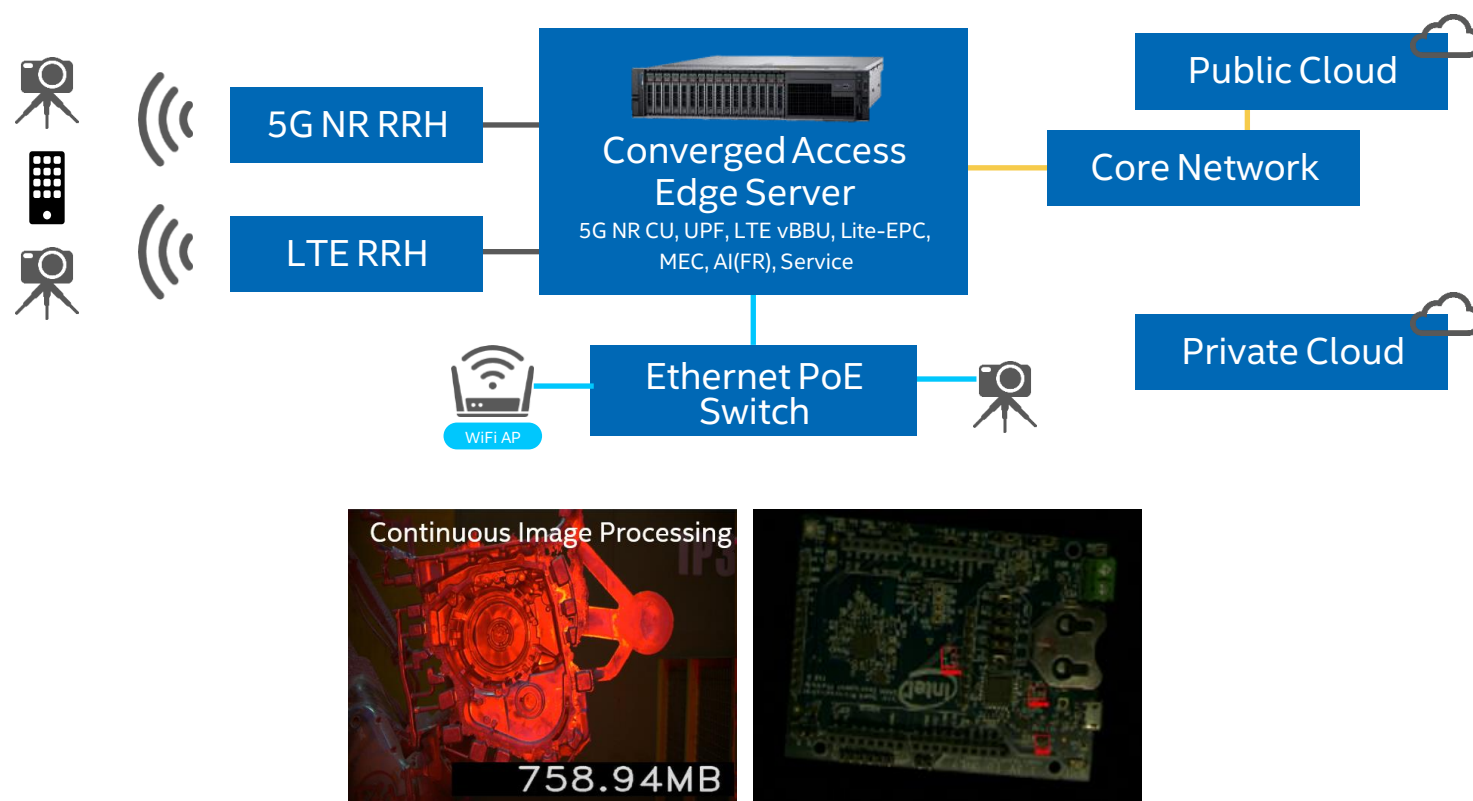
CERA
On-Premise

NETWORKING, AI, MEDIA AND EDGE SERVICES

(Intel® Xeon® D & Intel® Xeon® Scalable Platforms)



现场边缘应用案例 – 工业场景



What is the use case?

Video Analytics and 5G enabled AGV use cases for the factory floor.

Business value that CERA drives:

Friendly, simple, easy to use and maintain, run IT/OT Apps and services. Increase operational efficiency and factory floor insight with AI.

CERA architecture benefits:

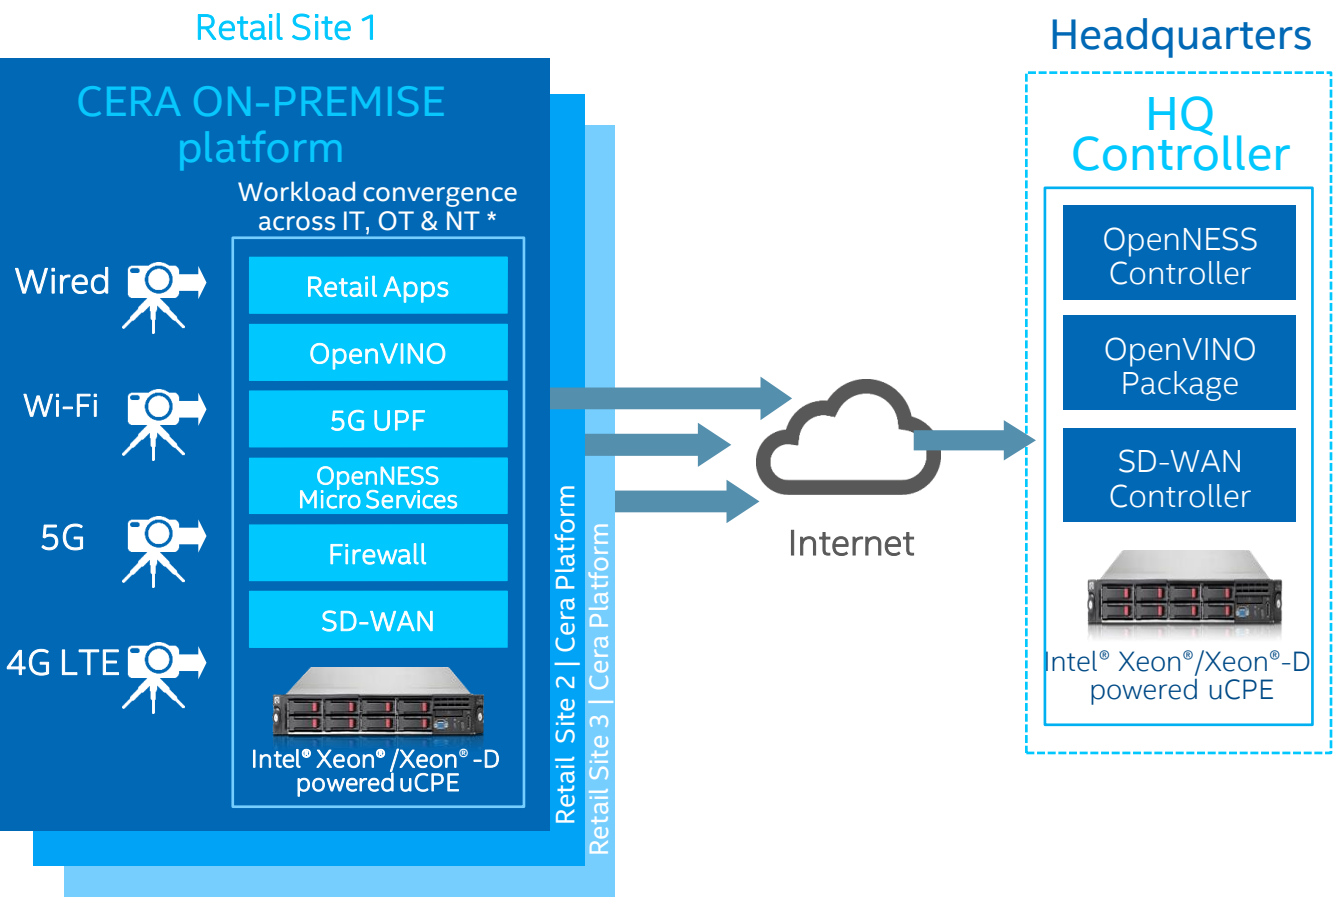
Cloud native CERA architecture supporting convergence of 5G CU/DU, 5G UPF, OpenNESS, OpenVINO on Intel Xeon based platform.

How did Intel help the ecosystem?

Intel is working closely with Foxconn in developing 4G and 5G private wireless and MEC solutions based on CERA. Solution is in trial now.

现场边缘应用案例 – 零售场景

- Digital Advertising
- Personalized Shopping
- Inventory Management
- POS Detection
- Customer Behavior Analysis



What is the use case?

Smart branch networking solution combined with artificial intelligence (AI) to drive advanced smart check out and facial recognition applications.

Business value that CERA drives

Build up a stable, easy to scale IT. Friendly, simple, easy to use and maintain IT systems that do not require dedicated IT team in each shop.

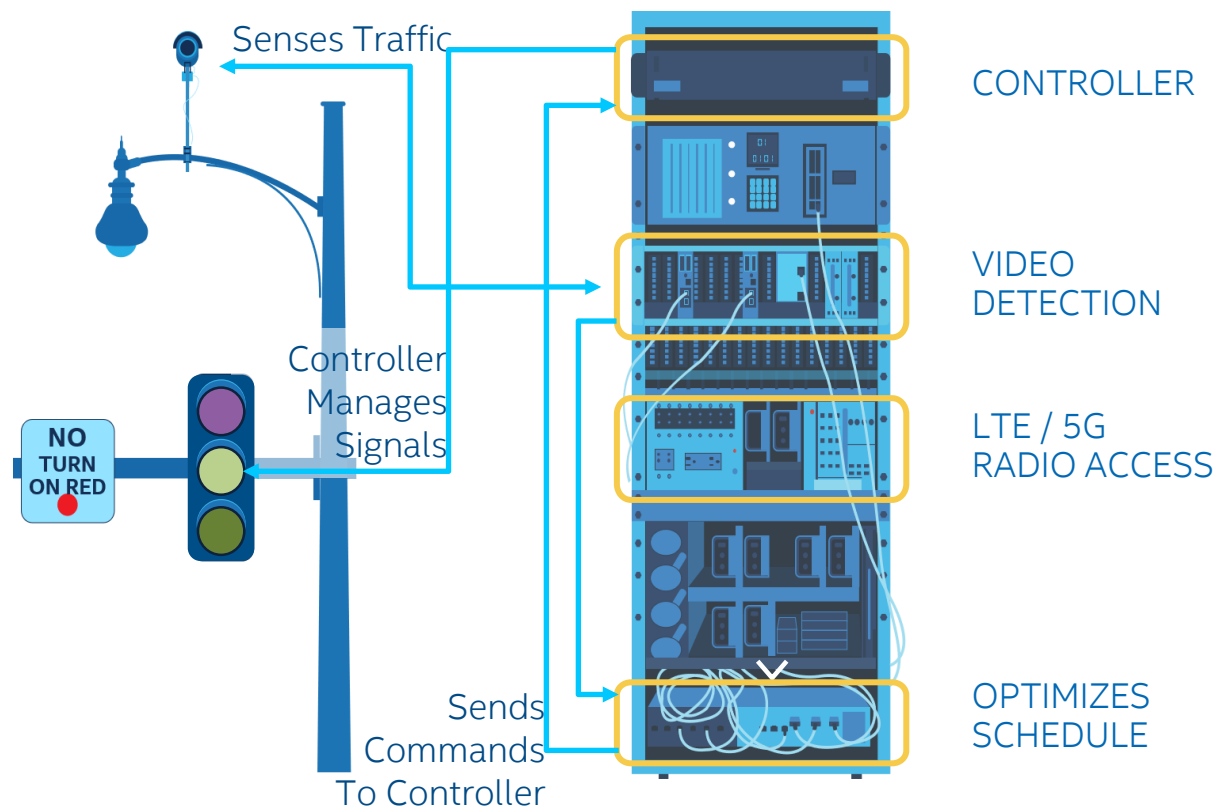
CERA architecture benefits:

Cloud native CERA architecture supporting convergence of 5G UPF, OpenNESS, OpenVINO, SD-WAN on Intel Xeon based platform.

How did Intel help the ecosystem?

QNAP and Intel closely collaborated in developing the uCPE platform for Retail Office. Solution has completed trial and RFP Ready Kit (RRK) created via Intel's market scale program is available now.

现场边缘应用案例 — 智慧城市/路边站



Adaptive traffic control is the future of mobility

Smart Intersections leverages AI, Edge Compute, Networking to improve city mobility

What is the use case?

Smart Traffic Management & Advanced Roadside Infrastructure.

Business value that CERA drives:

In order to support the growing traffic and urban mobility needs, traffic management is leveraging AI to optimize the lights and traffic flow. This workload convergence hence with better TCO, less complex edge nodes, and ability to scale the solutions based on demand.

CERA architecture benefits:

Moving from individual/purpose built boxes into a foundational edge compute node that can service various functions: traffic controller, tolling, video analytics, V2X, sensor fusion, MEC, 5G etc..

How did Intel help the ecosystem?

Intel is enabling ITS Traffic Management OEM partners leveraging outdoor edge solutions from SuperMicro and other Intel ODM partners for converging WLs (video analytics/V2X/MEC focus).

