CLOUD INFRASTRUCTURE FOR NETWORK FUNCTIONS - REQUIREMENTS AND TESTING

June 24, 2020

Today Discussion

Title:

Cloud Infrastructure for Network Functions - Requirements and testing

Abstract:

Description: Workloads like data plane network-functions require guaranteed, predictable performance, high throughput and low latency. Similar requirements may also reside with other workload categories like the ones serving the Finance sector. To apply cloud principles and satisfy these workloads specific requirements, we have enabled platform technologies and test tools. In this webinar we'll discuss building Kubernetes clusters with appropriate technologies, optimizations, methodology and the usage of tools to characterize performance. We will provide you with examples of how to deliver policy-based orchestration of functions in order to achieve targeted data plane performance.

Speakers

Dana Nehama

Director, Product Management Network Cloud - Intel

Since joining Intel in 2014, Dana has been responsible for driving the adoption of SDN, NFV and Cloud Native in the Communications industry through the introduction of new technologies and marketing initiatives. As a Product Management Director in Intel's Network Platform Group (NPG) she focuses on enabling the Network Transformation and cloudification to support new business and operational models. Throughout her career, Dana has led teams across a spectrum of areas in the telecom and wireless communications industry, bringing to market technologies such as SDN/NFV, Kubernetes, LTE, WiMAX, VoIP and DOCSIS.

Petar Torre

Principal Engineer - Intel

Petar Torre is Principal Engineer in Intel's Communications Service Providers team for data center, network and service transformation. His focus is on NFV and Edge to build consumable Telco Cloud platforms using Cloud Native principles. Petar joined Intel in 1999 and has held various roles from technical presales to strategic alliance management.

The Evolution Toward Cloud Native

SAME GOALS AS WITH NFV – BUILDS ON YEARS OF INDUSTRY EXPERIENCE AND UTILIZES CLOUD NATIVE TECHNOLOGIES TO ACCELERATE ACHIEVING INDUSTRY GOALS

Cloud Native Delivery - What Does it Mean?

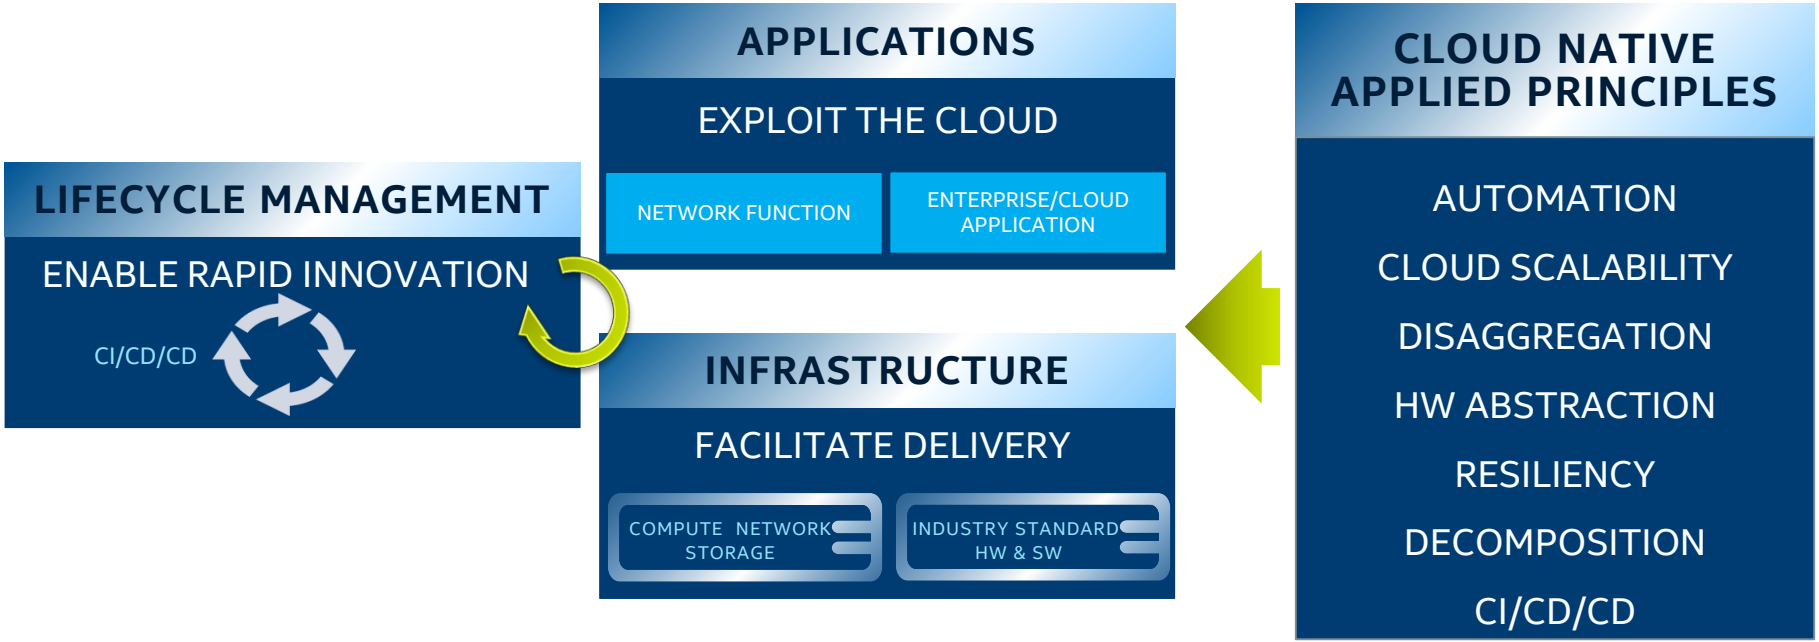

IMPLEMENT PRINCIPLES ALIGNED WITH CNCF CLOUD NATIVE DEFINITION ACROSS NETWORK FOUNDATIONAL ELEMENTS

Driving Rapid Innovation with Cloud Native

APPLICATION DECOMPOSITION

CONTAINERIZED → MICROSERVICES

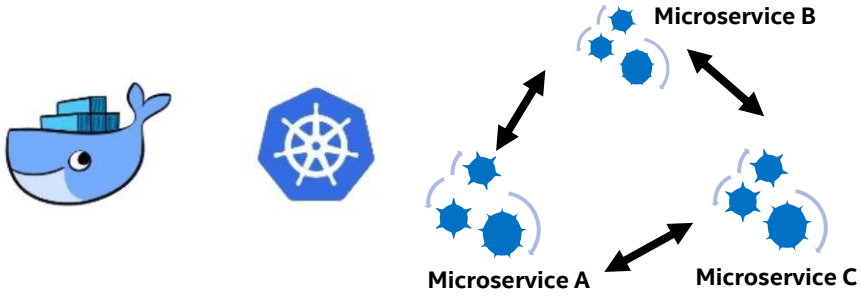

LIFE CYCLE AUTOMATION

CI/CD/CD → CLOUD MANAGED

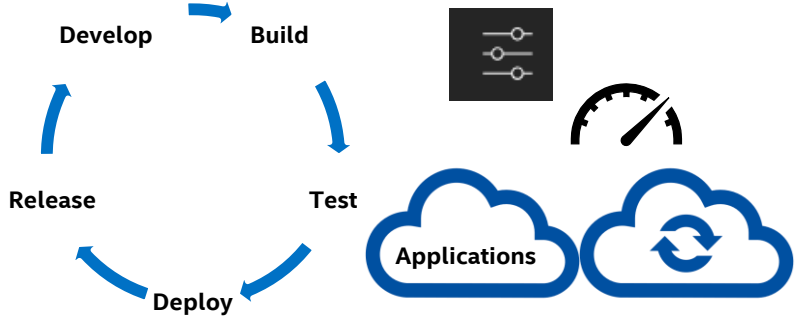

CLOUD NATIVE AGILE APPROACH ENABLES INNOVATION AND SERVICES VELOCITY

Cloud Native Delivery - What Does it Mean?

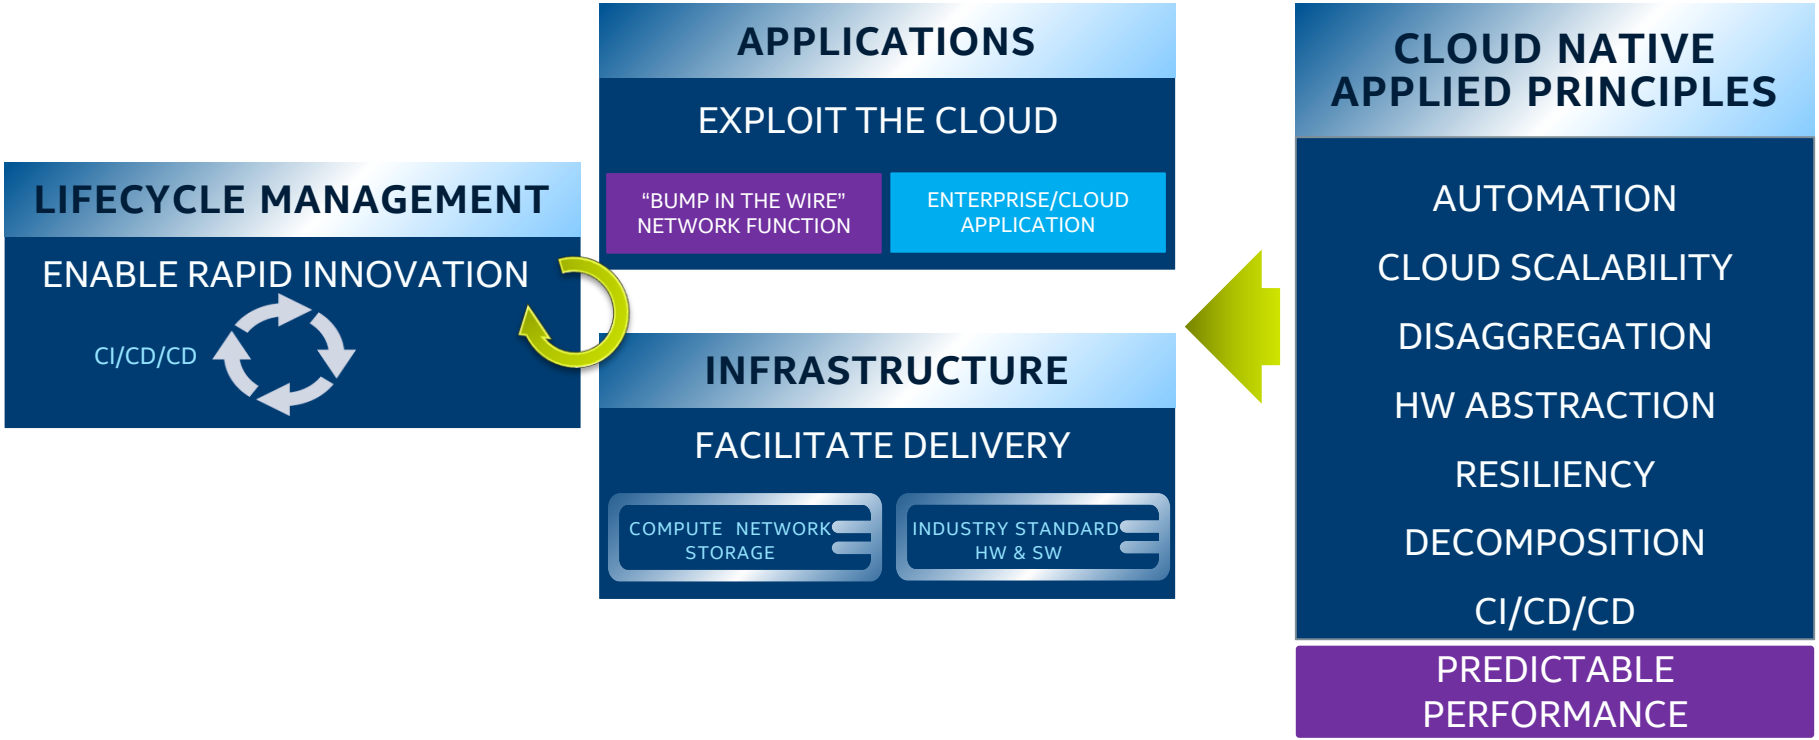

NO ONE SIZE THAT FITS ALL - NETWORKING CLOUD NATIVE APPLICATIONS DEMAND PERFORMANCE AND HIGH RELIABILITY

Requirements per Network Location

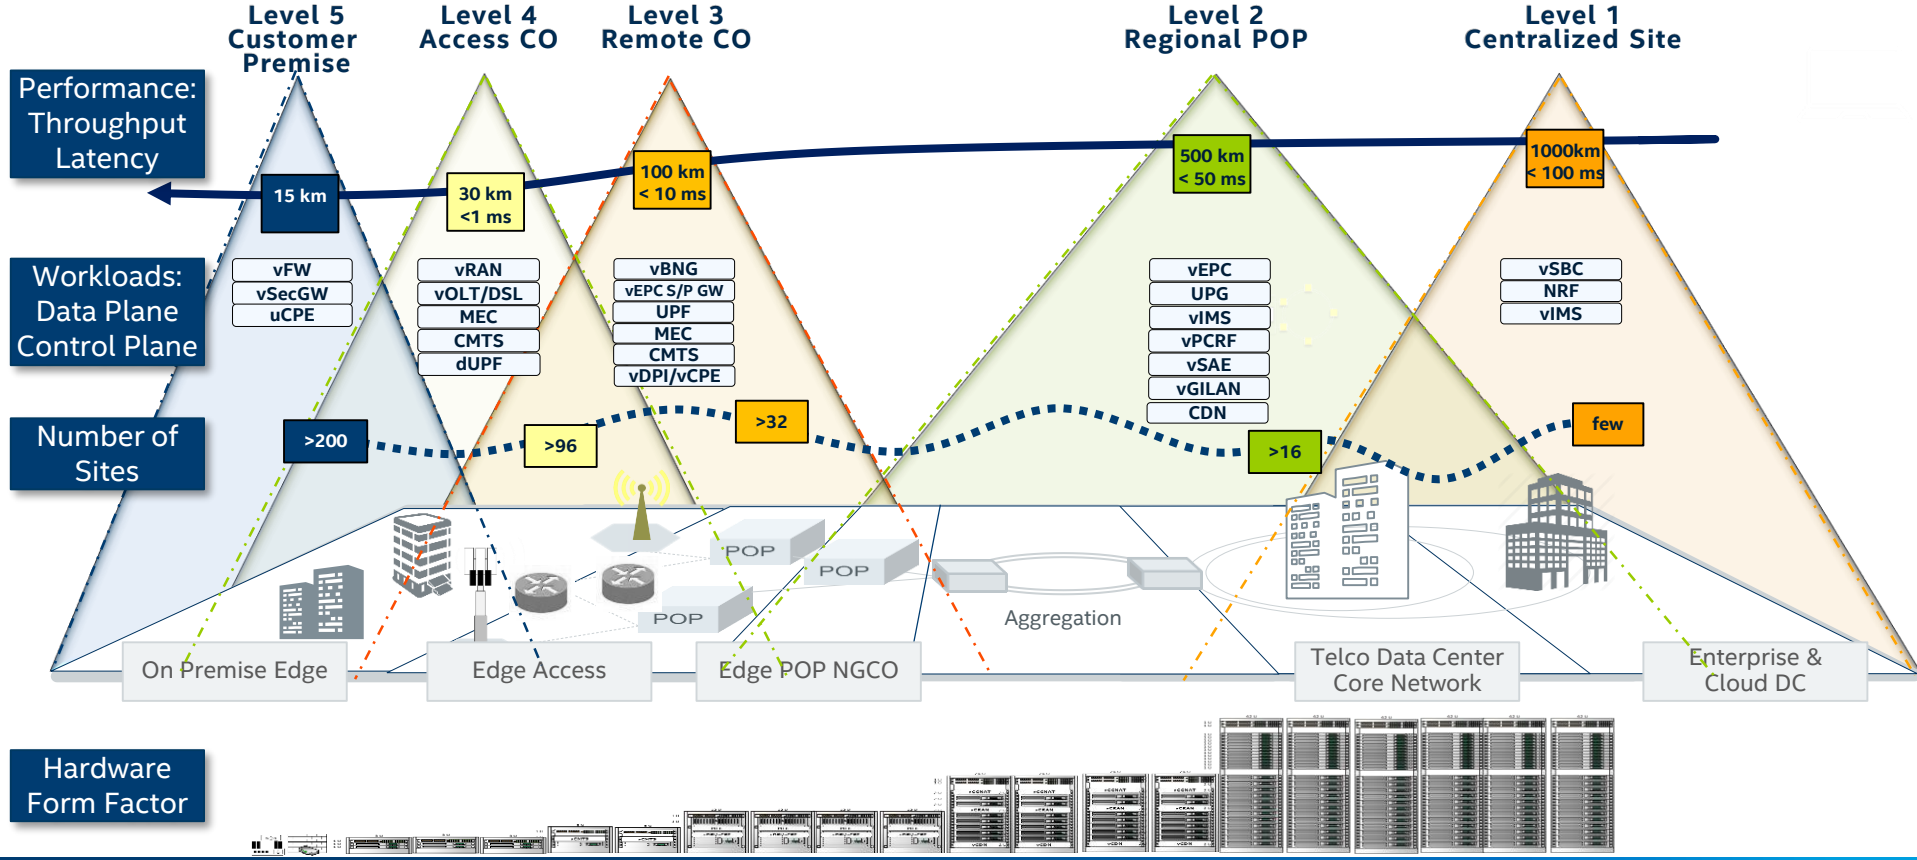

Address Gaps in Data Plane Application

REQUIREMENT	ADDRESSED BY
I/O interfaces to efficiently get packets to and from the application.	Balance flexibility vs performance

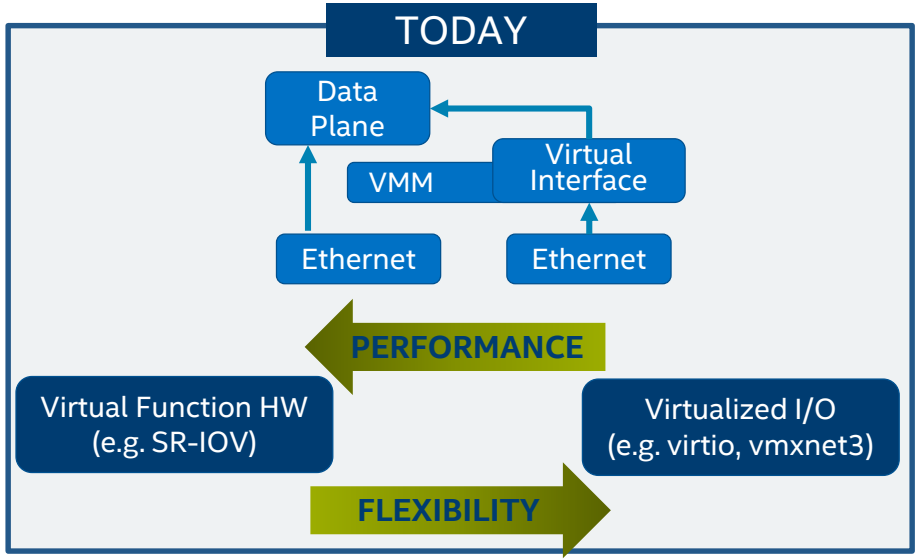

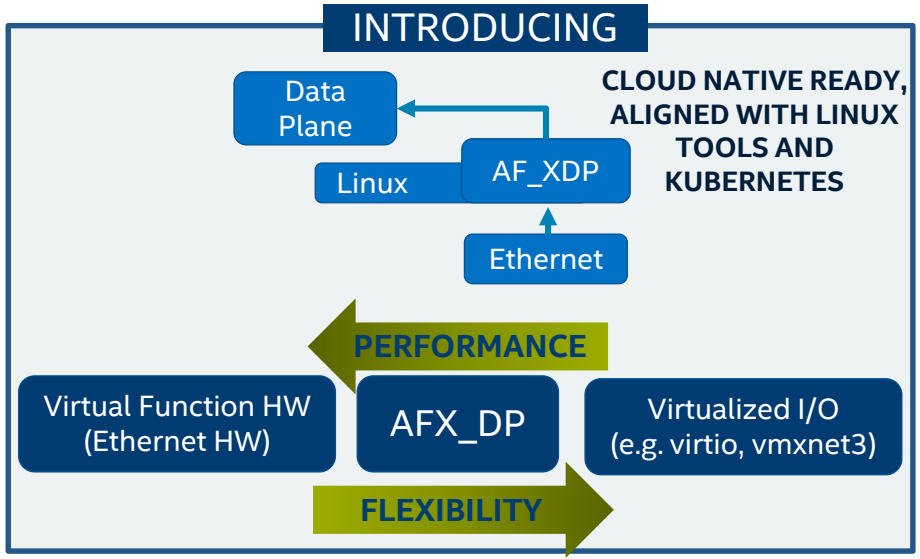

Address Gaps in Cloud Native Kubernetes Networking

TECHNOLOGY GAPS	ADDRESSED BY
K8s NETWORKING	Multiple Network for CNF — MULTUS
PACKET PROCESSING	High Performance E-W — CNI USERSPACE CNI DPDK
	High Performance N-S — CNI SR-IOV DPDK
	OVS-DPDK SDN Control — KUBE-OVN
	Portable Data Plane N-S — CNI AFX-DP CNI
	HA Networking — CNI BOND-CNI
RESOURCE MANAGEMENT	Platform discovery — Node Feature Discovery (Intel® AVX; SR-IOV etc.)
	CPU pinning/isolation #1 — CMK - CPU Mgr. for Kubernetes
	CPU pinning/isolation #2 — Native CPU Manager for k8s
	Dynamic Huge Page — Huge page support for K8s
	Manage Devices — Device Plugin (SR-IOV, Intel® QAT, GPU, user spc)
	Set NUMA Alignment — Topology Manager (NUMA)
TELEMETRY	Scheduling per Telemetry — Telemetry Aware Scheduler
	Platform Telemetry — collectd
DEPLOYMENT	Ref. Architecture Playbook — HELM

LIFECYCLE MGMT.

Metal^3 : Bare Metal Provisioning

Operators: Day1-Day2 Deployment

EXPERIENCE KITS

https://networkbuilders.intel.com/network-technologies/container-experience-kits

*Other names and brands may be claimed as the property of others.

Kubernetes* Networking – Recent Developments Status

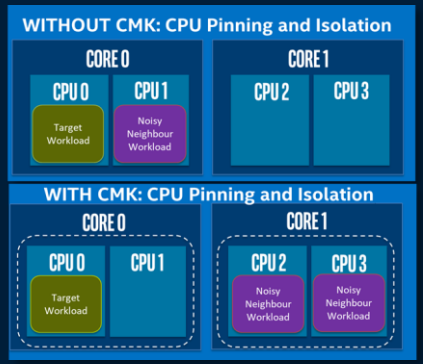

CMK
NATIVE CPU MANAGER

CPU pinning and Isolation

CMK - Open Source
https://github.com/intel/CPU-Manager-for-Kubernetes

Native CPU Manager
in K8s 1.16

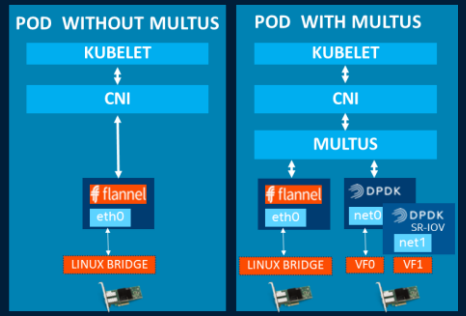

MULTUS

Multi Network Interfaces

A “meta – plugin”

Open Source

K8s Network Plumbing WG

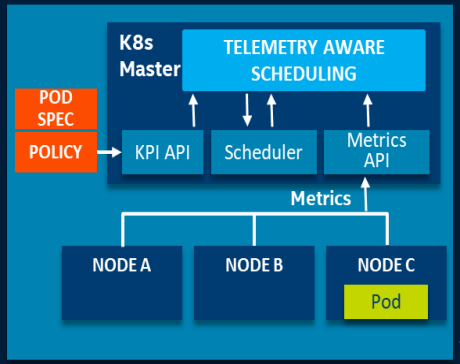

TELEMETRY AWARE SCHEDULING

Consider telemetry to schedule workloads

Apply a policy to make decisions

Open Source
https://networkbuilders.intel.com/network-technologies/container-experience-kits

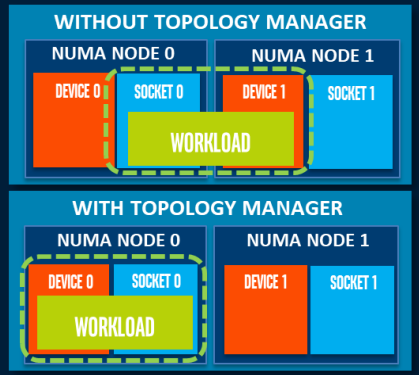

TOPOLOGY MANAGER

Address performance needs

NUMA resource allocation

Start with CPU Manager and Device Manager

In K8s 1.16

*Other names and brands may be claimed as the property of others

Accelerate Network Solutions' Development & Adoption

HARDWARE & OPEN SOURCE SOFTWARE

- K8s NETWORKING
- PACKET PROCESSING
- RESOURCE MANAGEMENT
- TELEMETRY

Intel FPGA

Intel® Intel® SSD NVMe

Intel® Optane™ Persistent Memory

Intel® Ethernet Controller

Intel® QuickAssist Technology

Intel® Visual Cloud Accelerator

INSTALLATION PLAYBOOK

- ANSIBLE SCRIPTS
- OPERATORS
- HELM CHARTS
- KUBESPRAY

<30 min

INTEL CONTAINER BARE METAL K8s REFERENCE ARCHITECTURE

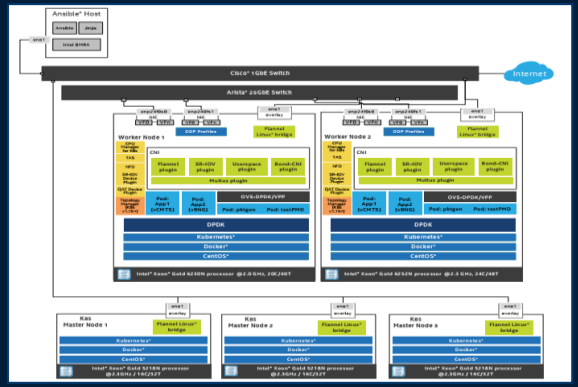

ARCHITECTURE PROFILES

ON PREM	REMOTE USER PLANE	OVC
WEB APPS; VISUAL CLOUD: ADI; SMTC	5G-UPF; CMTS; BNG; CDN	VISUAL CLOUD TRANSCODING

. . .

Reference Architecture Guide including playbook : Container Bare Metal for 2nd Generation Intel® Xeon® Scalable Processor Reference Architecture

*Other names and brands may be claimed as the property of others.

Working With The Community

CLOUD NATIVE COMPUTING FOUNDATION

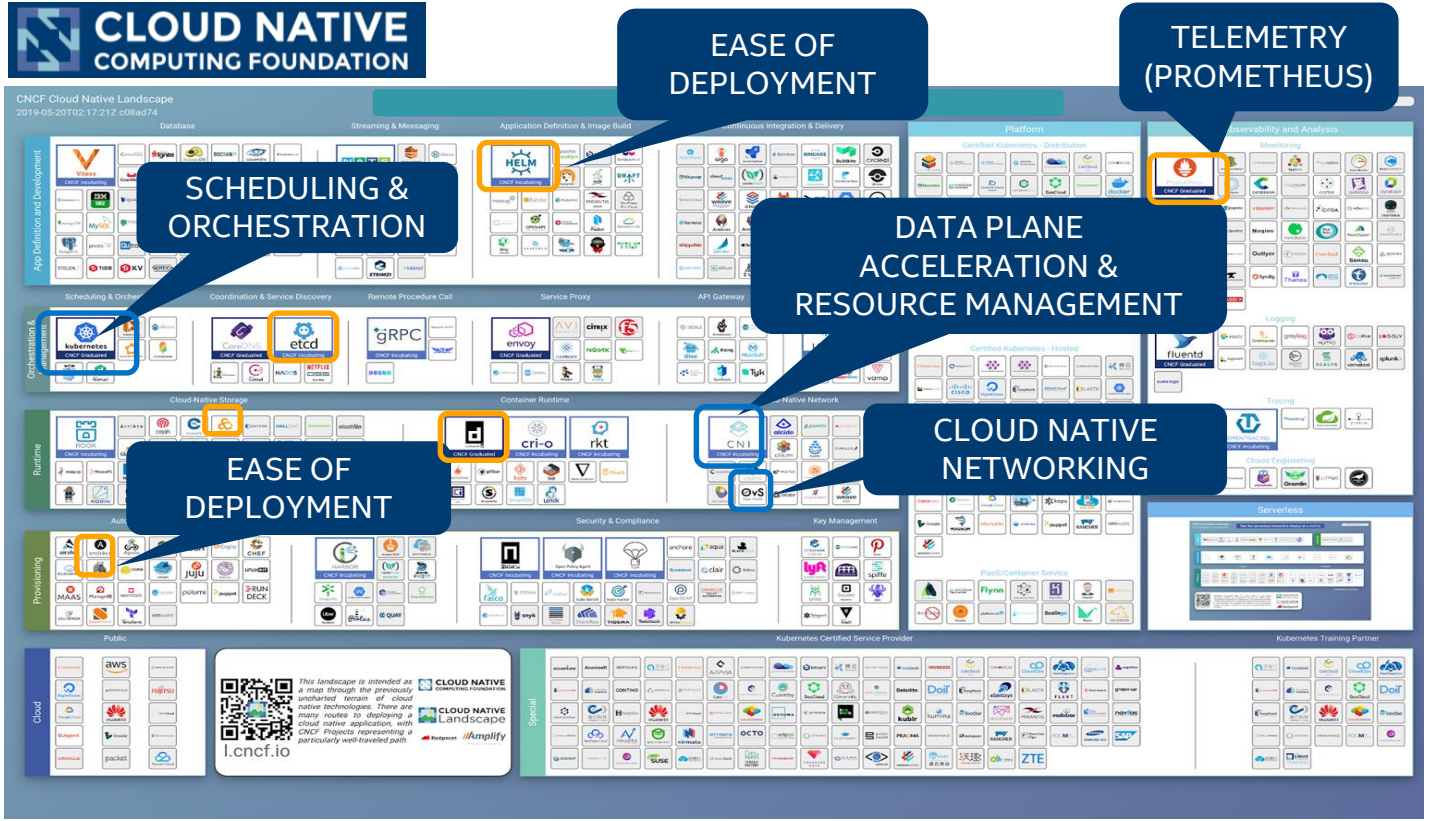

Source: https://github.com/cncf/landscape

CNCF

CNCF Telco User Group

CNCF CNF Test Bed and CNF Conformance

Kubernetes* Special Interest Groups (SIGs) and Work Groups

- WG: Resource Management
- SIGs: Network, Node, Scheduling & Instrumentation

Projects contributing to
Tools used

CNTT | Cloud iNfrastructure Telco Task Force

SPONSORS GSMA THE LINUX FOUNDATION

*Other names and brands may be claimed as the property of others.

BEST PRACTICES

Services from Cross-Domain Apps. on Shared Infrastructure

CLEAR LAYERS OF ABSTRACTION	CROSS-DOMAIN MODULAR ORCHESTRATION/ AUTOMATION	CROSS-DOMAIN MODULAR SERVICES
SERVICES	Service orchestration	Services
APPLICATIONS	Application managers	app1 app2 ...
INFRASTUCTURE	Virtualized infra. mgr. Physical infra. mgr.	Virtual infrastructure for POP, on-prem, MEC, IOT, IT... Public Cloud ...

More in white paper
https://www.intel.com/content/dam/www/public/us/en/documents/white-papers/hr-**disaggregation-and-the-virtual-network**-brief.pdf

Testing CaaS and CNF - Examples

STEP	LAYER IN STACK, TYPE OF TESTING	GOAL OF TESTING	EXAMPLE TOOLS
STEP 0 Platform functionality and performance	CaaS functionality	Validate platform capabilities	To test Huge Pages, DPDK interfaces, Multus…
	NFVI performance characterization	Ensure that underlying platform performs well	OPNFV* PROX as load generator and swap function with automation scripts
STEP 1 CNF Cloud Ready/Friendly/Native and Performance	5GC functionality	5GC element re-architected towards Cloud Native	CNCF CNF Conformance
	5G UPF performance	Performance characterization of UPF	To create sessions and load
LATER STEP (advanced) CNF Cloud Resilient	VNF resilience	Understand resilience to different impairments	PROX as Impair Network GW

Testing with OPNFV PROX

1. For NFVI Performance Characterization:
   - Chaos SW engine in the same virtualized environment
   - Is not HW switch outside
2. For VNF Dataplane Resilience Testing:
   - Measure from “the point of view of CNF”, which is what other CNFs will see
   - Not from point of view of added generator which is external to the virtualized environment

Same image/engine as generator, swap (and impair gateway); can be packaged as container bare metal, container in VM or typical guest VM

TEST FRAMEWORK SOFTWARE

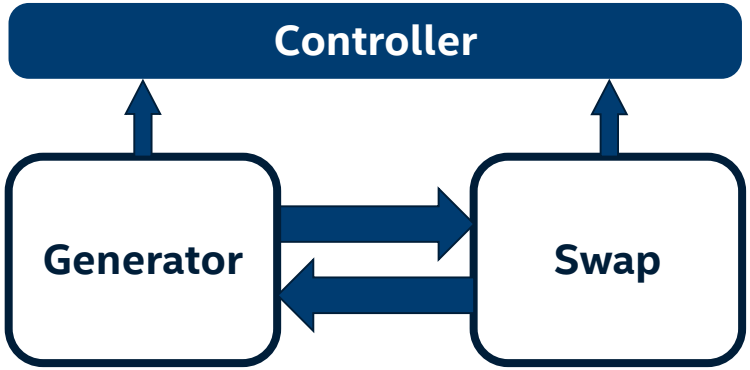

runrapid.py

Uses configuration file defining:
- where to place generator and swap
- max. target throughput
- packet sizes and number of flows
- thresholds for acceptable latency and packet loss

Prometheus

Grafana

User Guide: https://builders.intel.com/docs/networkbuilders/packet-processing-execution-engine-prox-performance-characterization-for-nfvi-user-guide.pdf

Defragmenting NFVI options, Decoupling VNFs/CNFs

Cloud iNfrastructure Telco Task Force (CNTT) Deliverables:

FRAMEWORK FOR RA:
- Design principles
- HW profiles

NFVI SPECIFICATIONS:
- #1 OpenStack*
- #2 Kubernetes*

BASED ON RA:
- OPNFV*

DELIVER VERIFIED NFVI, VNF AND CNF TO MARKET

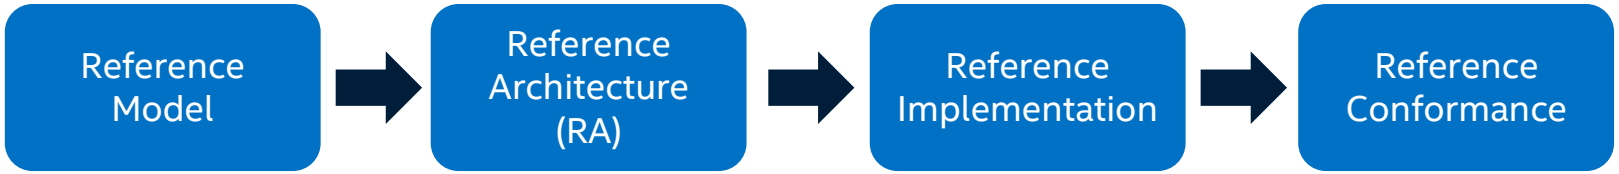

Source: https://github.com/cntt-n/CNTT

Intel Contribution to CNTT – Intel Containers Bare Metal K8s RA

Verify Intel Containers Bare Metal K8s Ref. Architecture alignment with RA2 Specifications

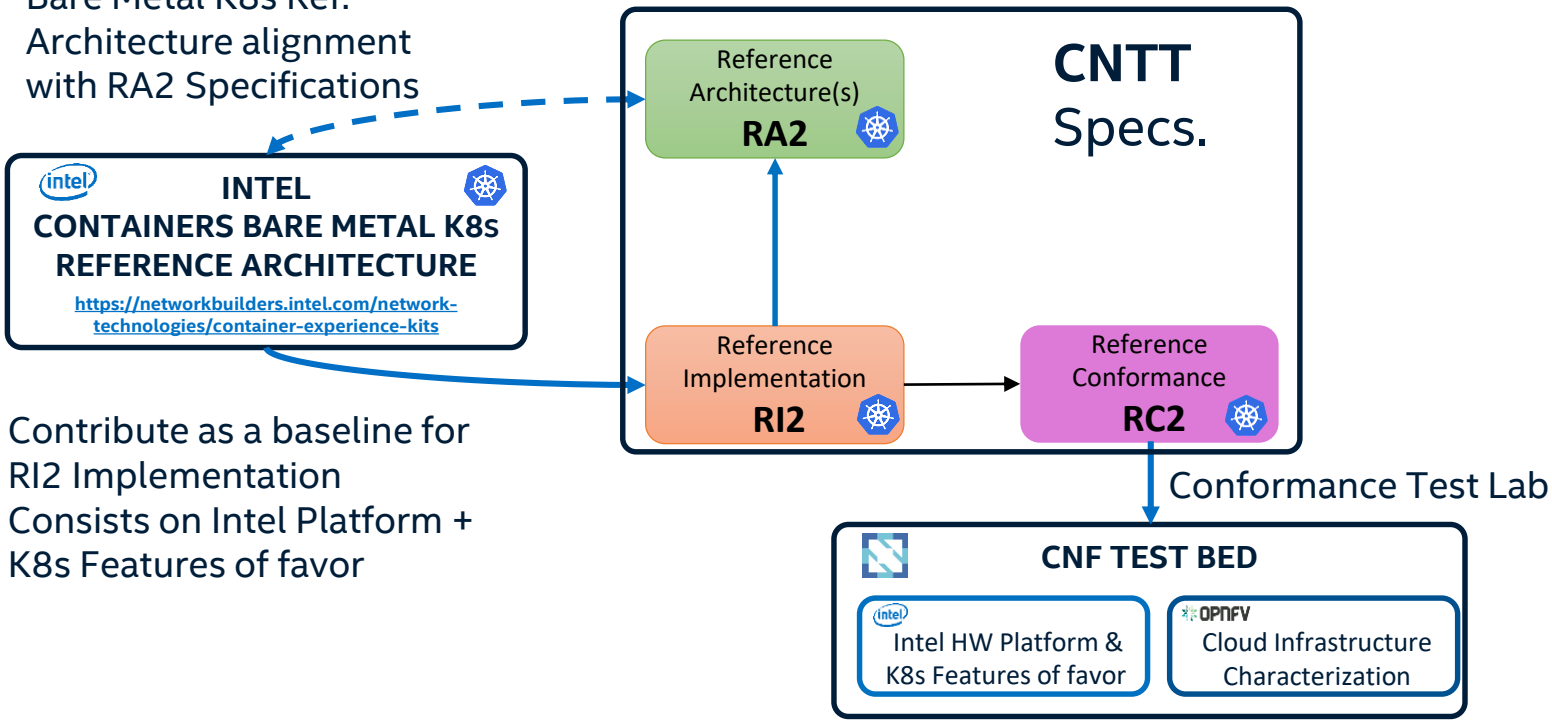

Contribute as a baseline for RI2 Implementation Consists on Intel Platform + K8s Features of favor

Decoupling Applications

Decoupling from infrastructure and PaaS, other application components, and application management/control

CLOUD DEPLOYMENT MODEL:

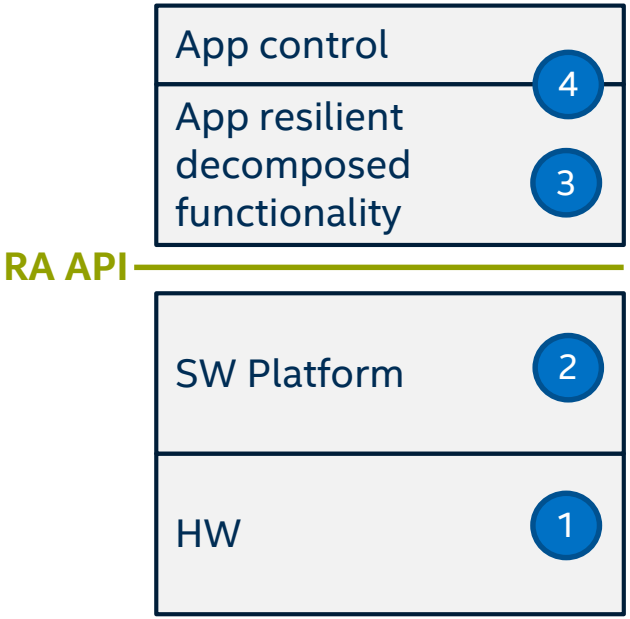

Service agility and easier operations will depend on levels of decoupling between application and:

1. Infrastructure
2. Platform Services
3. Application Resiliency

Relevant for sizing infrastructure and operations:

4. Other app functionality (decomposition, manageability)

Example of Decoupled Stack

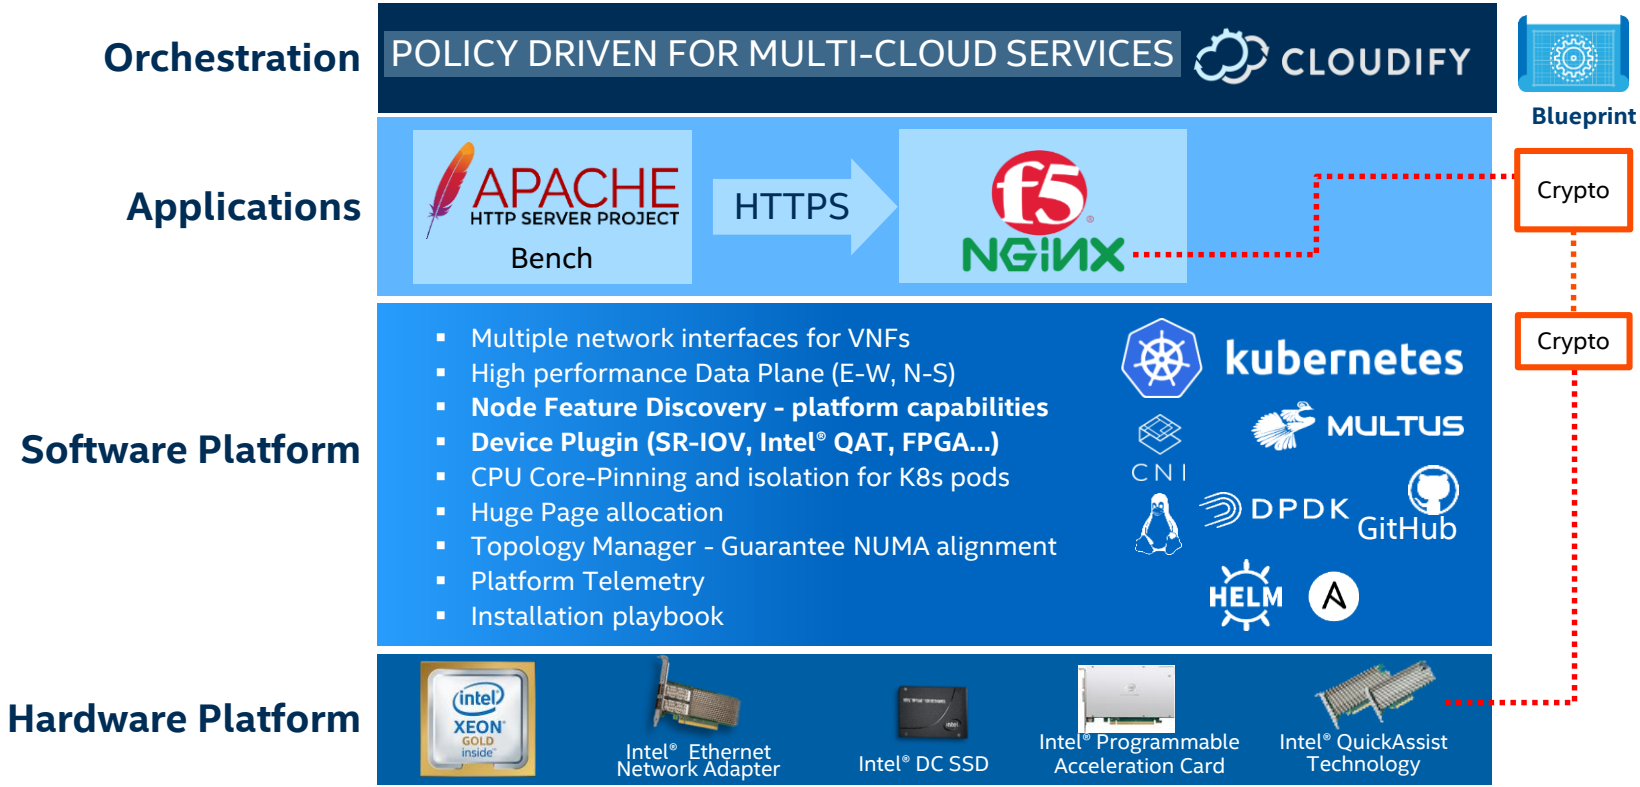

White Paper: https://builders.intel.com/docs/networkbuilders/multi-cloud-services-on-kubernetes-with-cloudify-orchestration-and-f5-networks-functions.pdf

*Other names and brands may be claimed as the property of others.

5G User Plane Function Performance

Terabit 5G UPF

Live Throughput
1,343.09 Gbps
Delay < 200ms

Cloud Native UPF
Using Zeon 6230N in 4U

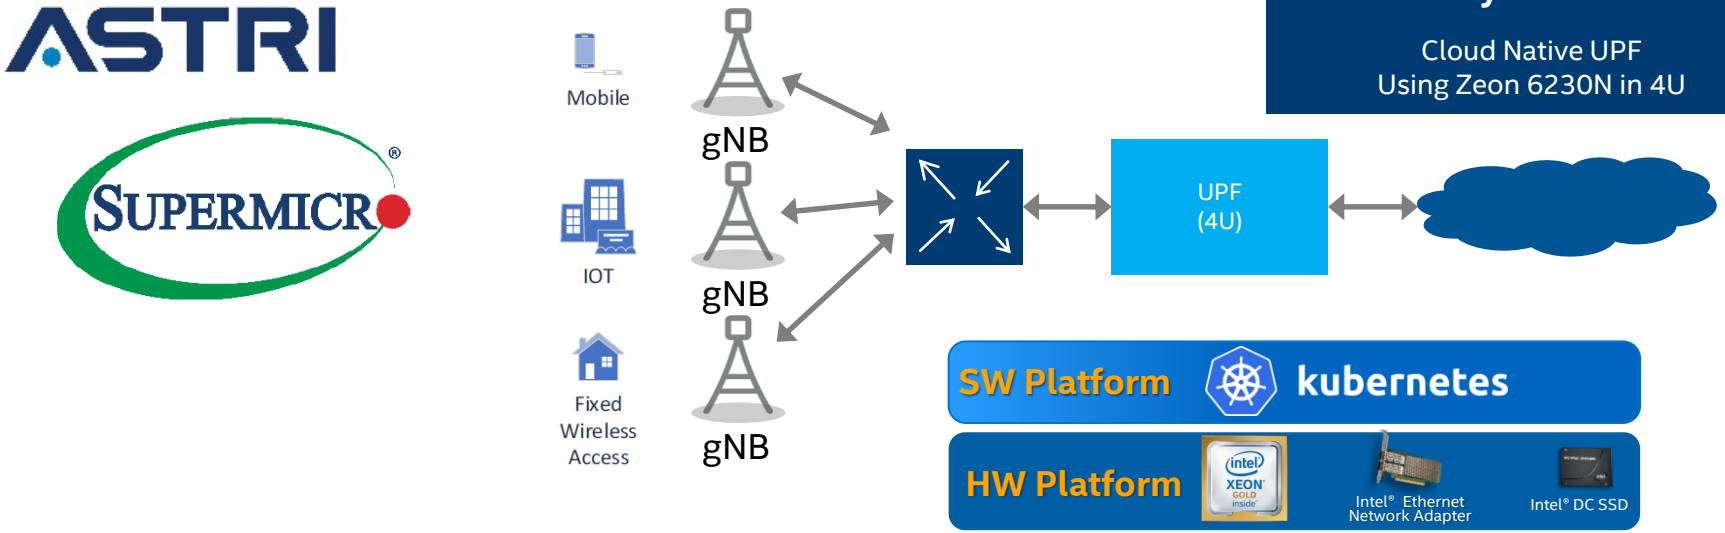

White paper: https://builders.intel.com/docs/networkbuilders/5g-user-plane-function-upf-performance-with-astri-solution-brief.pdf

*Other names and brands may be claimed as the property of others.

Abstract Summary

Depth and speed of transformation is defined more by widely accepted design principles than by size, current skills, etc.

WITH MULTI-VENDOR NFV PROCESSES, ORGANIZATIONS CHANGE FROM MANAGEMENT OF SILOS TO **MANAGEMENT OF LAYERS**

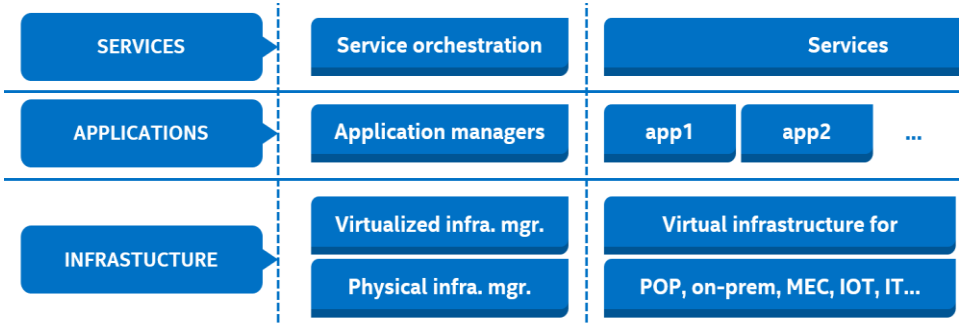

LAYERS OF COMPLEXITY:

- External
- Processes
- Orgs
- Reward system
- Culture/mindset

HOW HARD IS EACH TO CHANGE?

Summary

Cloud native is an application development approach that uses cloud computing delivery models

Service providers want to embrace cloud native technologies within their operational ecosystems, mainly to achieve service velocity

Intel supports closing the gaps in Cloud Native for service providers and delivers platforms that enable flexibility, agility and performance optimization that are foundational in modern network deployment

Find Out More

LEARN

About Kubernetes Networking Technologies, please click on:
https://networkbuilders.intel.com/network-technologies/container-experience-kits

About Network Transformation Solutions, please click on:
https://networkbuilders.intel.com/network-technologies/network-transformation-exp-kits

About the Containerized 5G Core, please click on:
https://www.intel.com/content/www/us/en/communications/why-containers-and-cloud-native-functions-paper.html

EXPLORE

Free foundational 5G training from Intel Network Academy *
https://www.coursera.org/learn/network-transformation-101

Intel Containers Bare Metal Kubernetes Reference Architecture
Container Bare Metal for 2nd Generation Intel® Xeon® Scalable Processor Reference Architecture

ENGAGE

Please contact your Intel representative

* Registration required.

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.
- 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'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.
- Your costs and results may vary.
- Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.
-

Legal notices & disclaimers

- This document contains information on products, services and/or processes in development. All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest forecast, schedule, specifications and roadmaps.
- The products described may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.
- No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.


intel®
experience
what's inside™