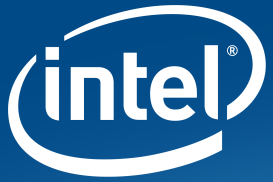

Red Hat® Ceph Storage and Network Solutions for Software Defined Infrastructure

Frank Ober – Intel Non-Volatile Memory Solutions Group
Tony Dempsey – Intel Network Platforms Group
Dr. John Fitzpatrick – Openet Labs

Data Center needs are **Changing**

- From collecting to analyzing data
- To virtualizing networks and storage
- To delivering secure & compliant cloud like services

Traditional Data Center Resources are Limited

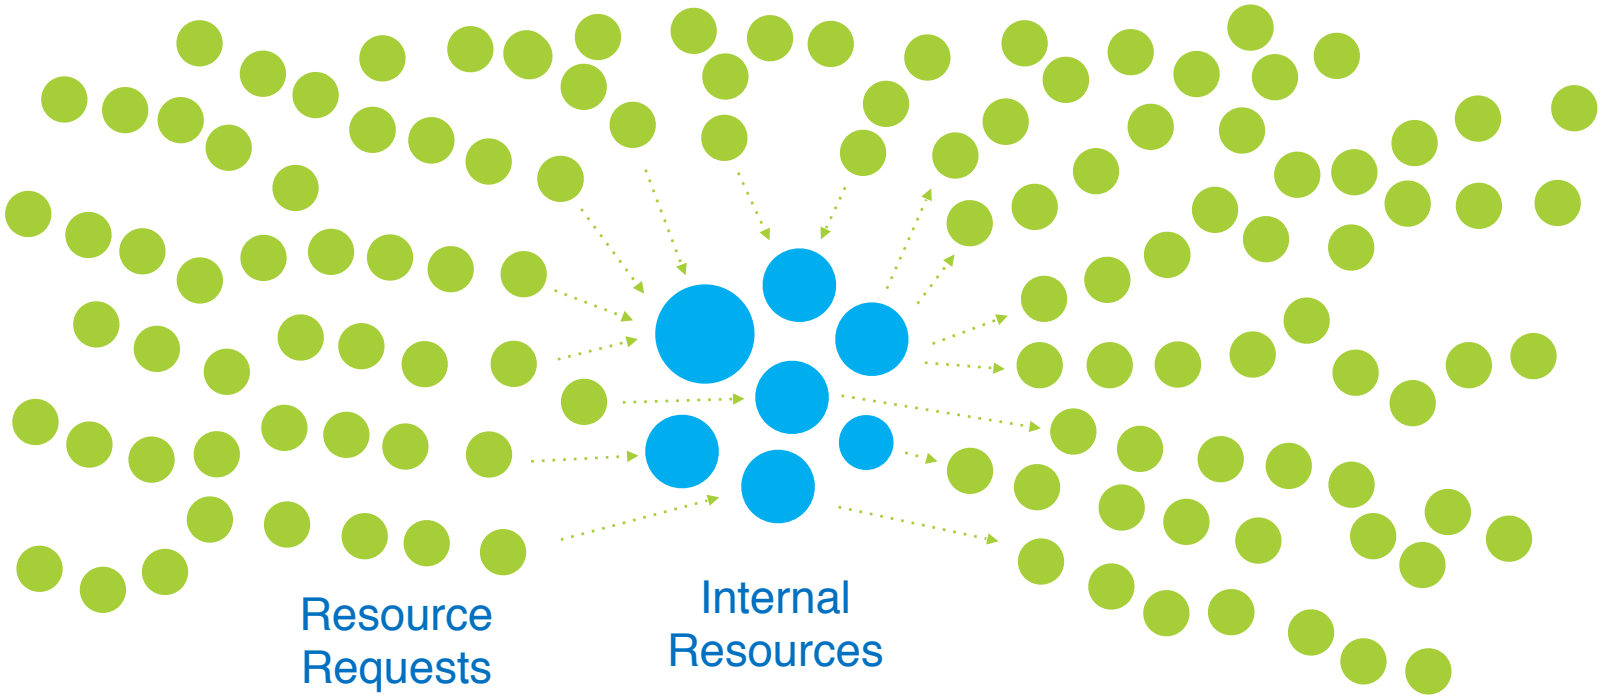

Intel Software Defined Infrastructure (SDI) Vision

Dynamic, policy-driven resource management

Abstraction of SW from HW,
provides flexibility and scalability

Provisioning of resources
dynamically (pay-as-you-grow)

Orchestration of diverse systems
through application SLOs
to enable seamless access

SDI is *not*:

A single product or solution

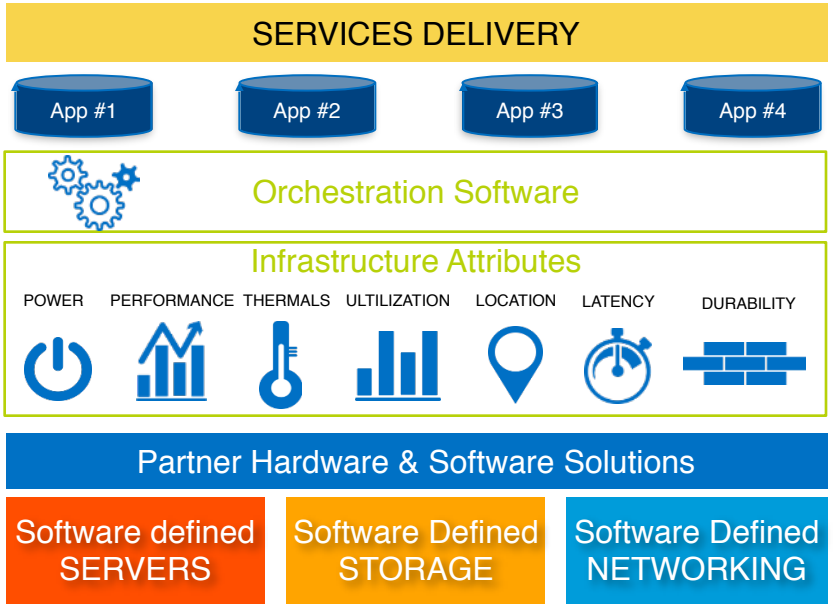

SDI is a framework: providing efficient, flexible,
scalable standards-based compute, networking, and storage resources

Intel SDI Assets

Intel® QuickAssist Technology

Intel® Storage Acceleration Library (ISA-L) & Storage DPDK

PROCESSOR PLATFORMS

NETWORKING AND FABRIC

SSDs

Intel® Cache Acceleration Software

SERVER BOARDS & SYSTEMS

Software Defined Infrastructure

SOFTWARE ECOSYSTEM

+

SOLUTION REFERENCE ARCHITECTURES

Software Defined Storage

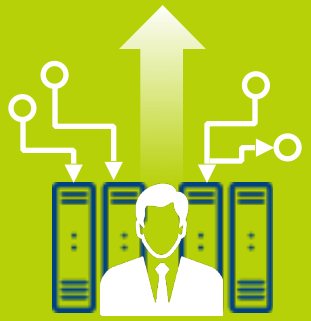

Software Defined Storage (SDS) Architecture

DATA SERVICES

- Part of the storage SW that runs in data plane to optimize storage
- For example: Data Deduplication, Erasure Code, Tiering

APPLICATIONS

ORCHESTRATOR

Northbound API

DATA SERVICES

SDS CONTROLLER

Southbound API

SDS CONTROLLER

- Has visibility into and control of ALL storage resources
- Coordinates between apps, orchestrator, and storage systems
- Allocates storage resources to meet SLAs

STORAGE SYSTEM [SAN]

STORAGE SYSTEM [Capacity]

STORAGE SYSTEM [DAS]

STORAGE SYSTEM [Performance]

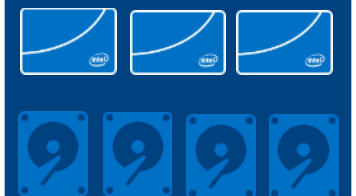

STORAGE SYSTEM [NAS]

SDS Vision to Action

- SDS will provide a convenient way to manage all storage in the datacenter
- Framework is still in development
- So, how to get started with SDS?
  - Focus on the storage system layer
  - Implement scale-out storage systems from OEMs and ISVs built on standard high-volume servers

Deploying storage systems on standard, high-volume servers today supports a seamless transition to SDI tomorrow

Advent of Solid-State Drives (SSDs)

- CPU = 175x vs. HDD IO = 1.3x
- IOs reach the spindles in a random fashion
- Gets worse with higher # of apps or VMs per LUN

- SSDs $30/GB -> <$1/GB (‘08-’15)
- Lower TCO than HDD ($, Watts, Space)
- $ / IOPS better for SSDs ($0.01 vs $0.80)
- SSDs are reliable (2MHr MTBF)

SSDs are a cost effective, reliable means to remove the storage bottleneck

* Performance and cost figures are estimates for informational purposes only and may vary

Red Hat® Ceph Storage

- Ceph is an open-source, massively scalable, software-defined storage system which provides object, block and file system storage in a single platform. It runs on commodity hardware—saving you costs, giving you flexibility—and because it's in the Linux kernel, it's easy to consume.
- Object Store (RADOSGW)
  - A bucket based REST gateway compatible with S3 and swift
- File System (Ceph FS)
  - A POSIX-compliant distributed file system
- Block device service (RBD)
  - OpenStack native support
    - Cinder-backend
    - Glance-backend
  - Kernel client and QEMU/KVM driver

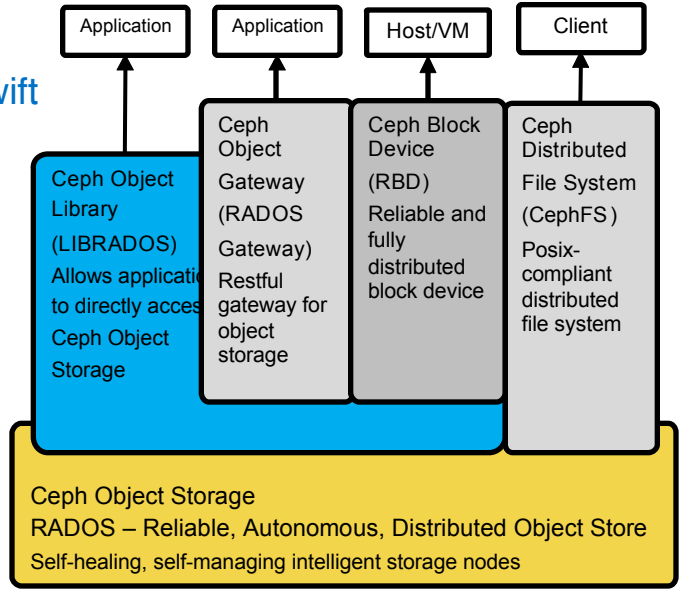

Ceph Design Considerations

- Consistency of an SSD Caching Drive within Ceph is key to overall consistent performance
- The developers (technical documentation), warn about using the most consistent drives possible. Drives that protect the data.
- Relatively inexpensive SSDs are often hiding issues that you need to find.

Important: We recommend exploring the use of SSDs to improve performance. However, before making a significant investment in SSDs, we **strongly recommend** both reviewing the performance metrics of an SSD and testing the SSD in a test configuration to gauge performance.

Since SSDs have no moving mechanical parts, it makes sense to use them in the areas of Ceph that do not use a lot of storage space (e.g., journals). Relatively inexpensive SSDs may appeal to your sense of economy. Use caution. Acceptable IOPS are not enough when selecting an SSD for use with Ceph. There are a few important performance considerations for journals and SSDs:

http://ceph.com/docs/master/start/hardware-recommendations/#data-storage

Open Source Virtual Storage Manager for Ceph

VSM is designed for a bundled ceph storage appliance, it creates a ceph cluster and does management plus monitoring. See https://github.com/01org/virtual-storage-manager

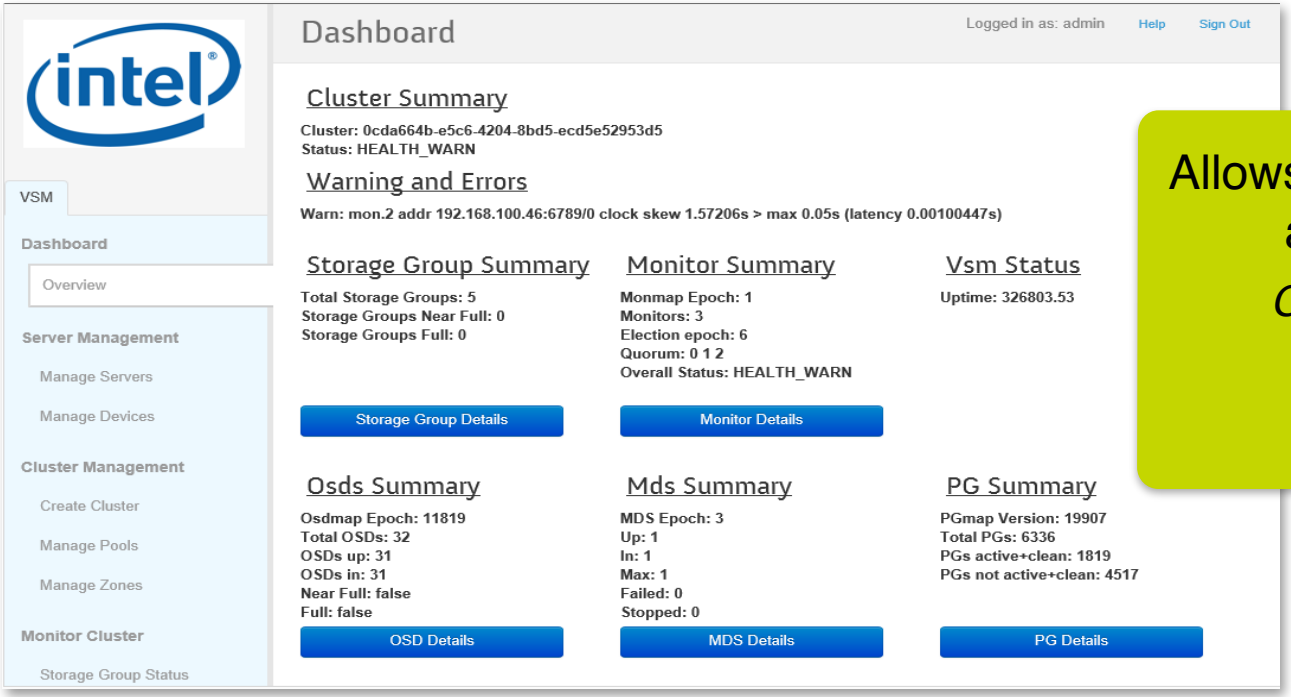

Allows easier Ceph Creation and Management

Creates Ceph cluster

Manages Ceph

Monitors Ceph

Software Defined Networking

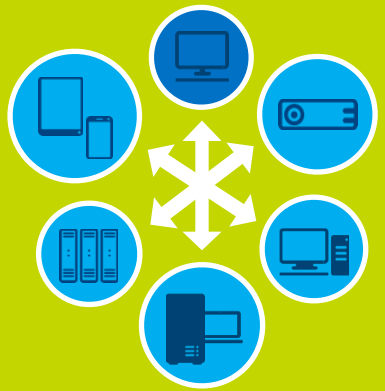

Intel Innovations in Networking

ACCESS NETWORKS EDGE/CORE NETWORKS ENTERPRISE

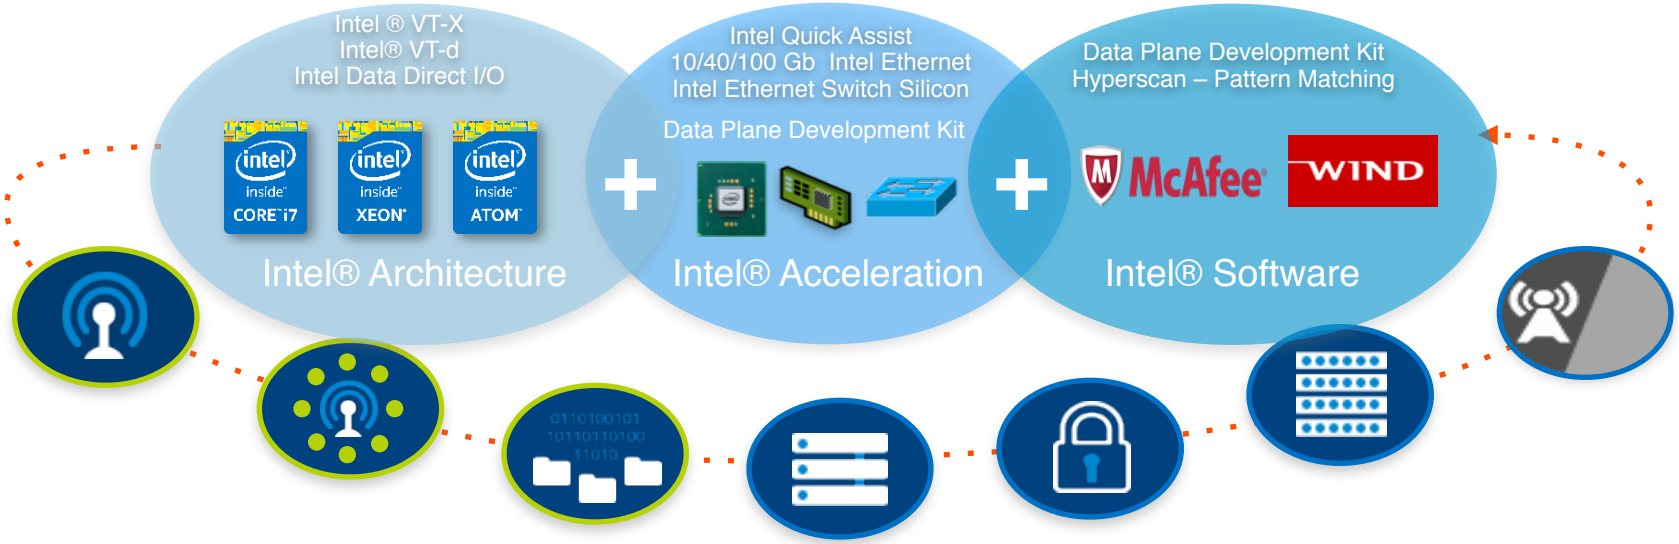

Moore's Law: The number of transistors incorporated in a chip will approximately double every 24 months

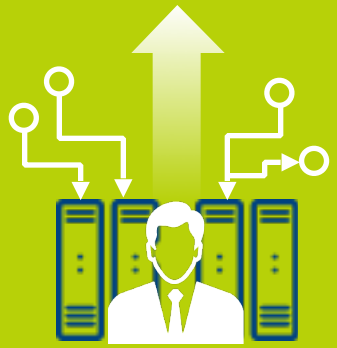

INDUSTRY-LEADING
INNOVATION AND
TECHNOLOGY
LEADERSHIP

COMMITMENT TO
OPEN STANDARDS
AND **PLATFORMS**

Intel Delivers **SOFTWARE-DEFINED INFRASTRUCTURE**

INVESTMENT
IN BUILDING
STRONG
ECOSYSTEMS

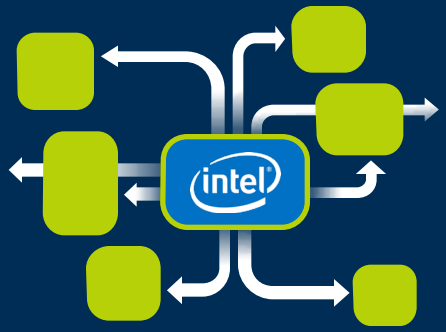

ACCELERATING
THE MARKET

Workload Consolidation to IA

A Decade of Investment

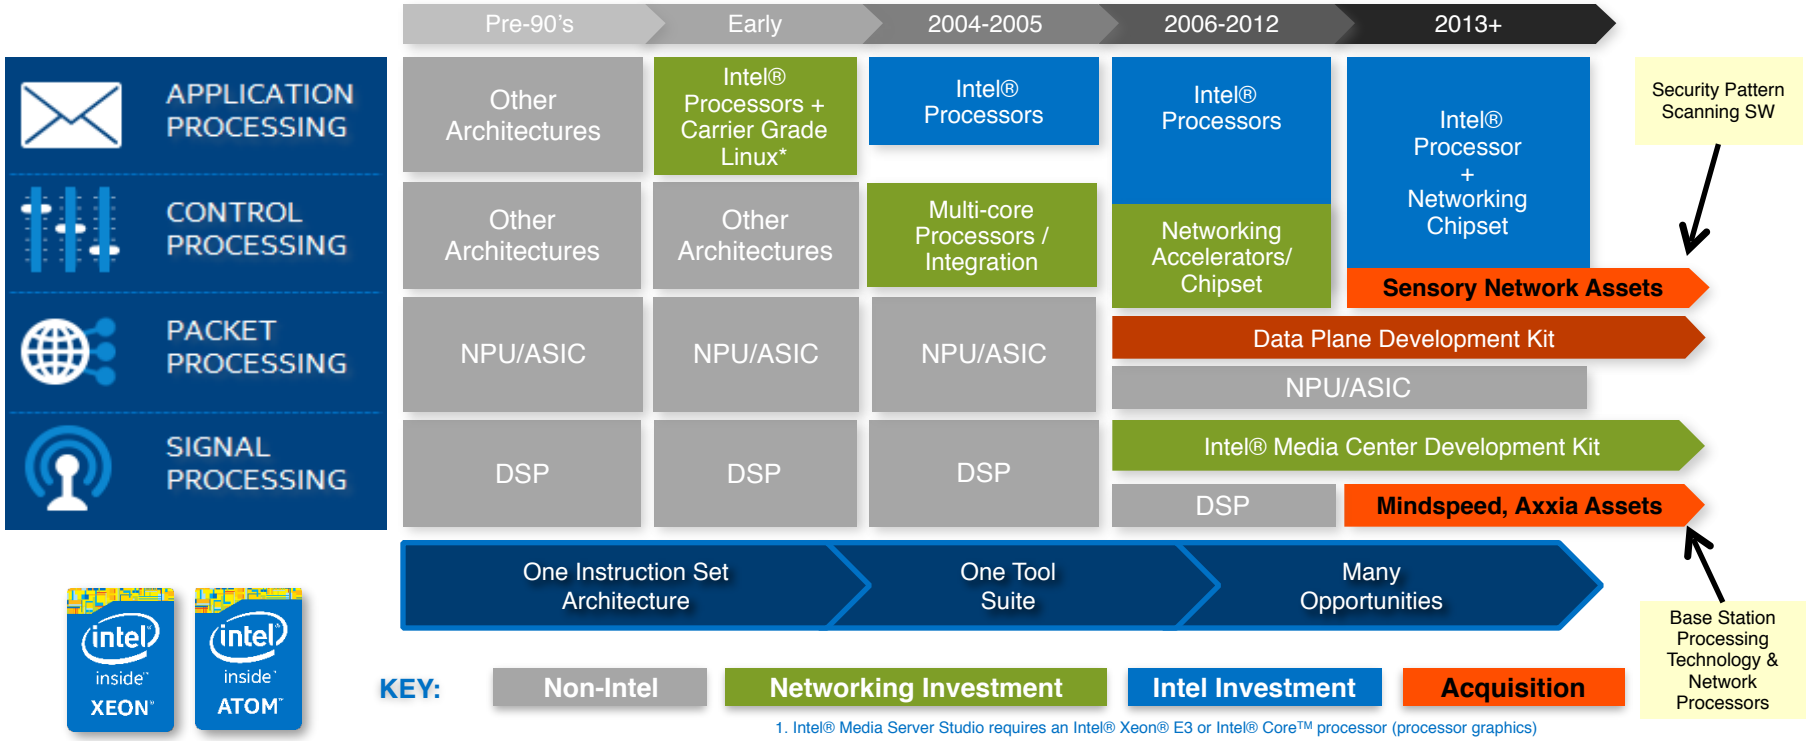

1. Intel® Media Server Studio requires an Intel® Xeon® E3 or Intel® Core™ processor (processor graphics)

4:1 Strategy = Single architecture that consolidates the workloads into scalable and simplified solution

Open Source/Standards Fuel Transformation: Intel is a Leading Contributor

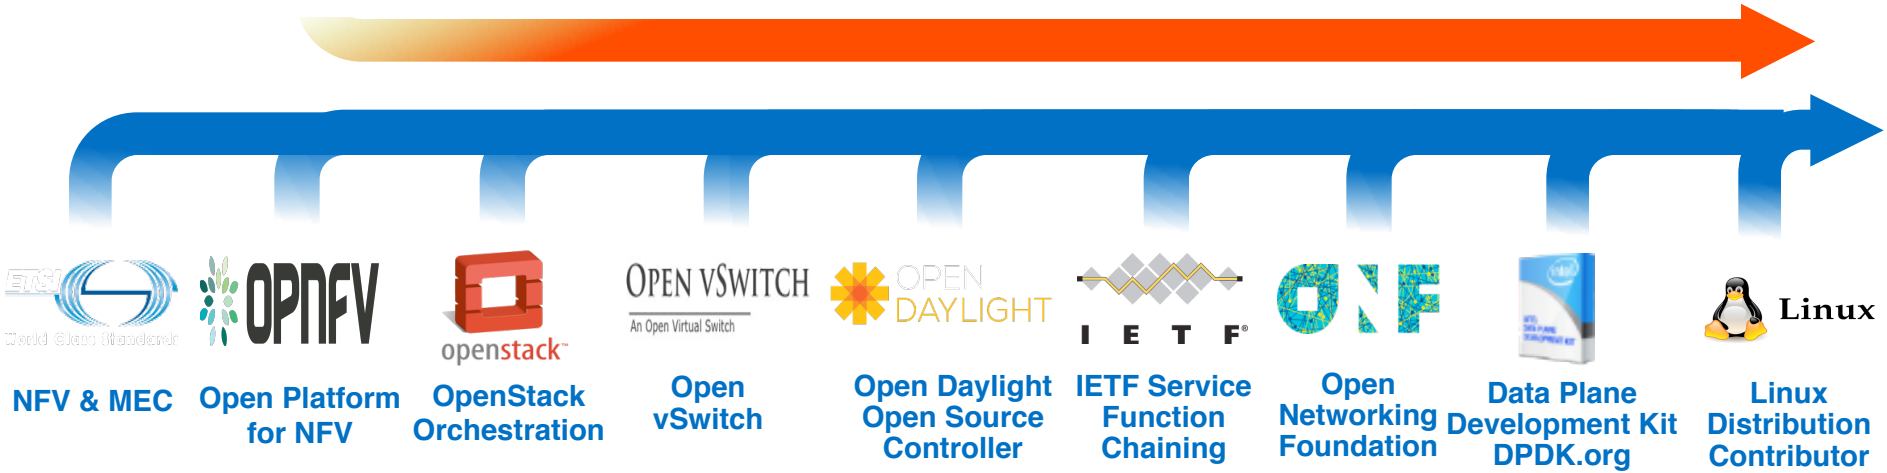

Investing in contributions to SDN/NFV initiatives
Specifications -> Code -> Test -> Reference Platforms

Intel® Open Network Platform (ONP) Server

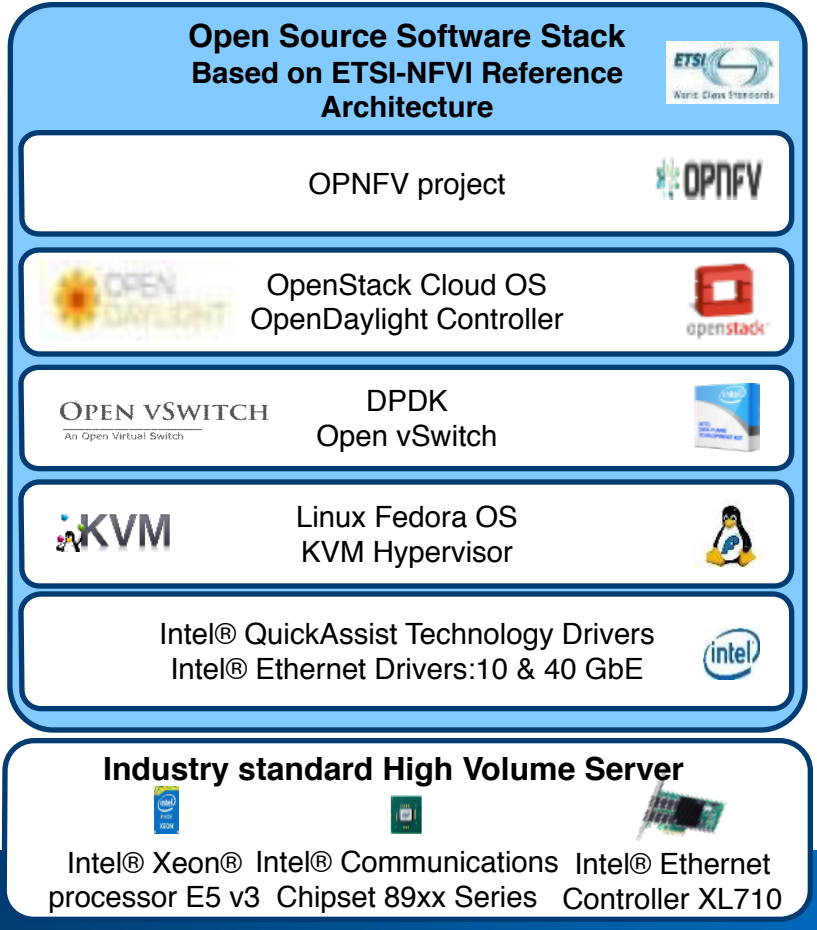

What is it?

- Server Hardware & Software reference design integrating Intel HW optimizations with Open Source ingredients used in SDN/ NFV

Who is it targeted at?

- Directly to NFVI+VIM vendors: OEMs, TEMs, SIs, ODMs
- Indirectly to: ISVs, CommSP
- Adjacent Markets: Enterprise, Cloud SP

Where does my customer get it?

- ONP Software is released on Intel's 01.org on quarterly basis in the form of: Reference Architecture Document; Benchmark Test Report; Collateral ; Application Demo Setup*

How does it relate to Linux Foundation Open Platform for NFV (OPNFV)?

- ONP Server is available today and will stay as Intel's progr to contribute covered ingredients into OPNFV

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

Intel® Network Builders Program - Partners

Customer Success Stories: SDN/NFV

25+ PILOTS IN PROCESS
WORLDWIDE

JUST A FEW
EXAMPLES

Network Operators & Enterprises **Embracing SDN/NFV** Using **General-Purpose Processing Technology** in Their Networks

Intel is Investing in SDN/NFV Transformation

Advance Open Source and Standards

Deliver Open Reference Designs

Collaborate on Trials and Deployments

Enable Open Ecosystem on IA

Accelerating Network Transformation Through Industry Collaboration

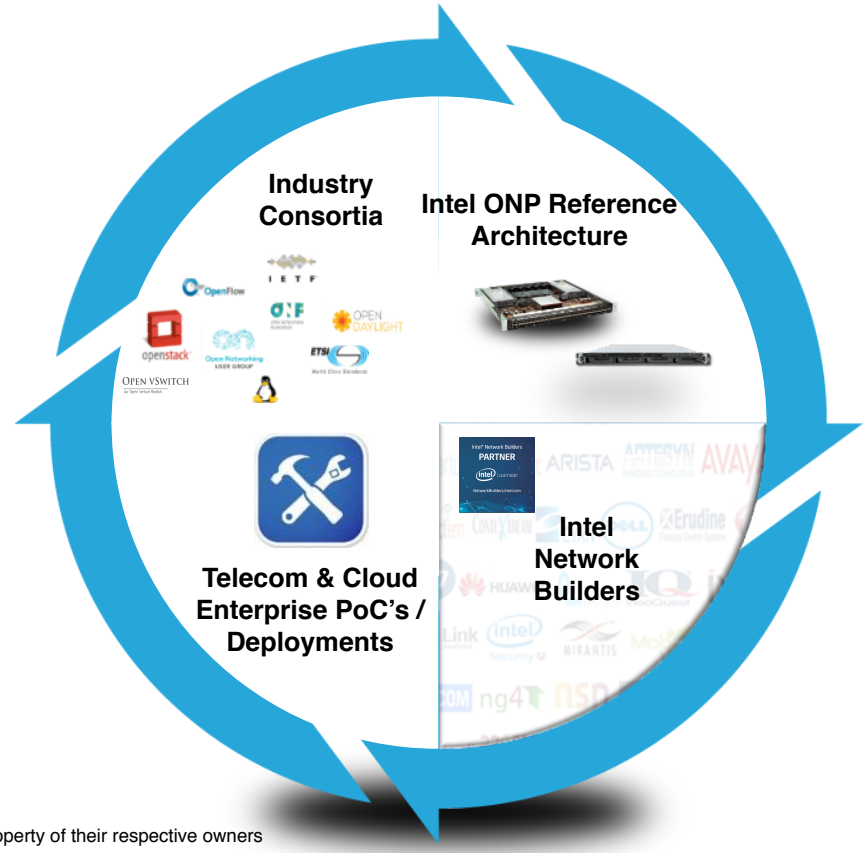

*Other brands and names are the property of their respective owners

Openet ETSI NFV Use Case

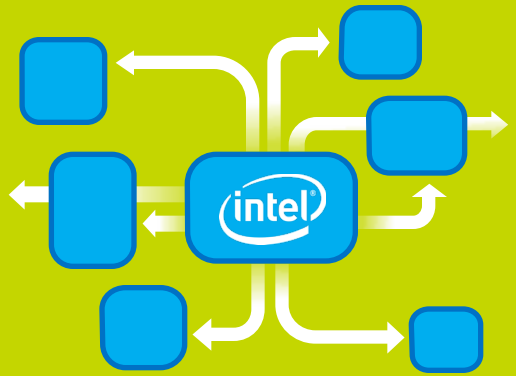


he world of computing has gone through
apid changes in recent years

This has lead to great improvements in how software is created, deployed, managed and used but **how does this relate to mobile networks?**

ur networks have not
dapted to the same degree
s many other computing
riven industries

Virtualization to the Rescue

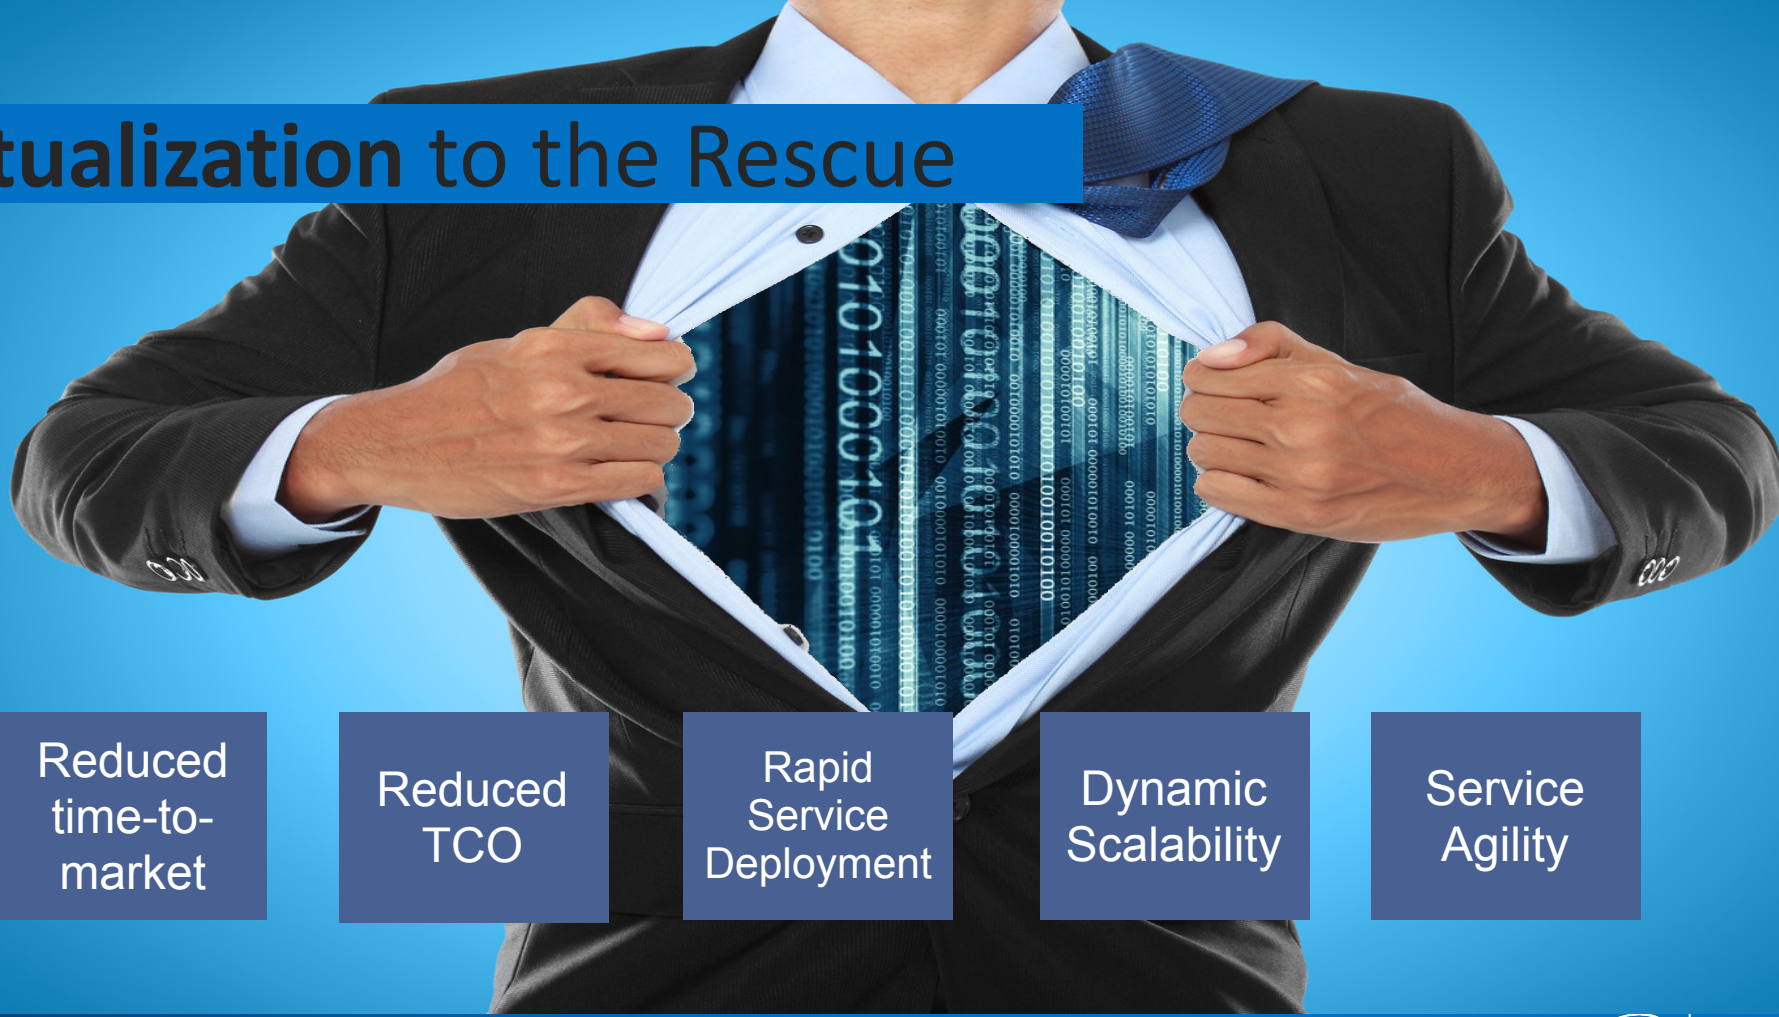

- Reduced time-to-market
- Reduced TCO
- Rapid Service Deployment
- Dynamic Scalability
- Service Agility

Openet's Virtualization Architecture

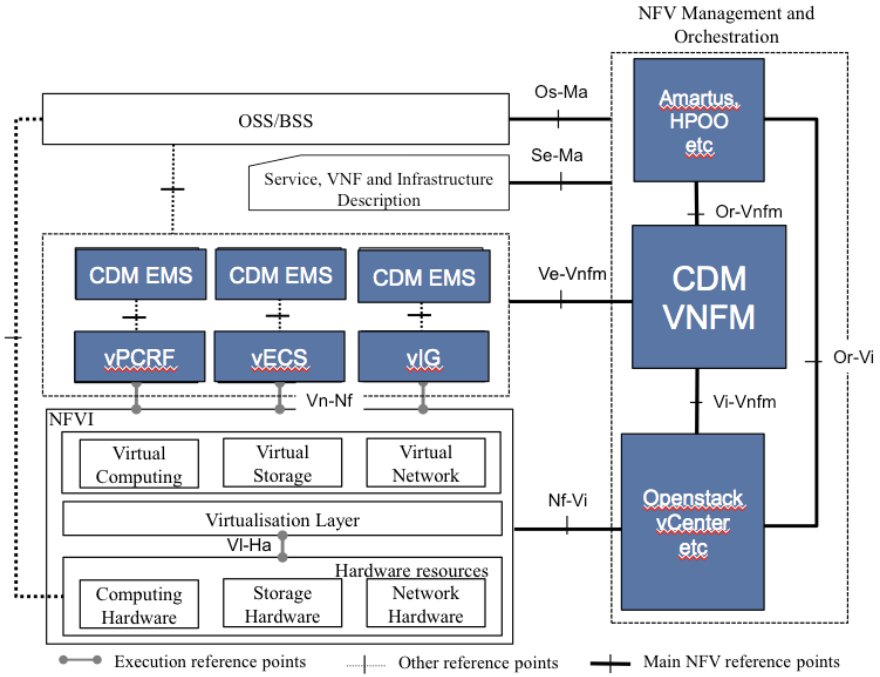

- Openet have being deploying solutions in fully virtualised environments for a number of years
- Openet Configuration and Deployment Manager (CDM) was developed to provide VNF orcehstration and management prior to NFV
- CDM was deployed in the worlds first Telco-grade virtualised PCC orchestration solution (pre NFV)
- Currently deployed in 5 operator deployments
- Supporting two large scale tier one deployments
  - 600 Virtual Machines in its first deployment
  - 400 Virtual Machines in its second deployment
- Reduced upgrade time from weeks to a couple of hours

ETSI NFV PoC #32 - Virtualized vEPC

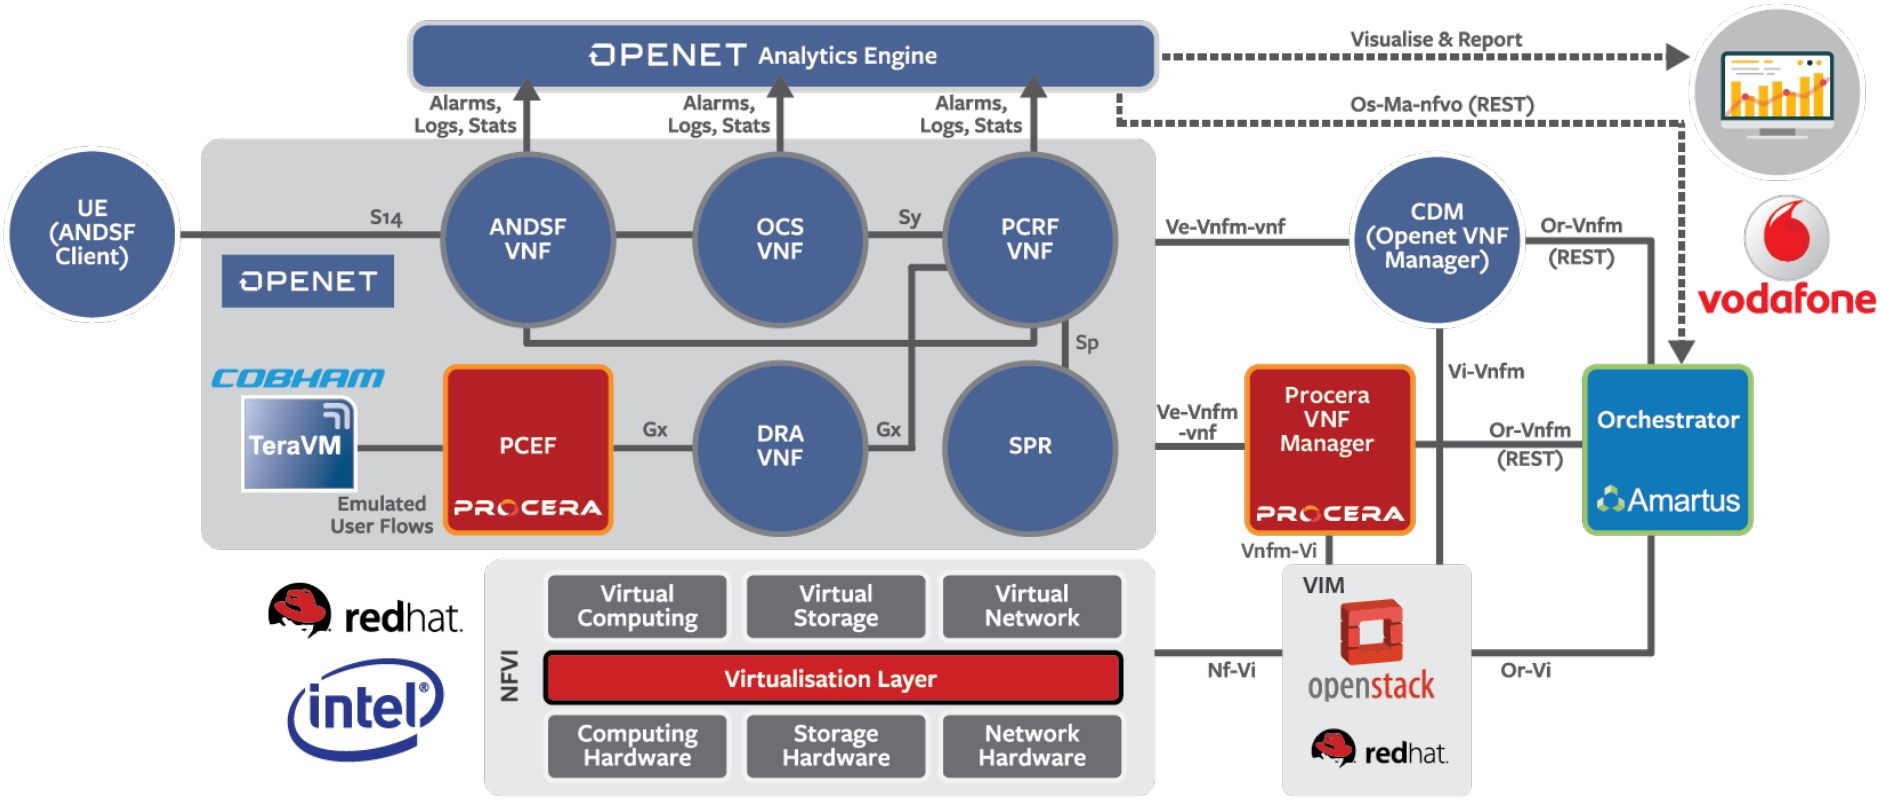

ETSI NFV PoC#32

- Fully integrated & orchestrated vEPC enabling the rapid creation and deployment of new services using common service templates (Using VNFD)
- Multiple independent PCC service deployments within single compute infrastructure
- Common Or-Vnfm REST based interface (in line with NFV MANO) enabling:
  - Multi-Vendor VNF Management and full lifecycle management (Instantiate, Upgrade, Modify, Terminate, Monitor, Check feasibility)
- Automated analytics based scaling
- Monitor revenue generated per service
- VNF based traffic emulation
- SR-IOV based hardware acceleration

OPENET

Rapid Service Deployment & Monetization through NFV

Service Monetization

- The ability to rapidly deploy and scale new services will lead to new revenue opportunities and business models

- It is essential that operators can easily monetise and monitor these new service innovations

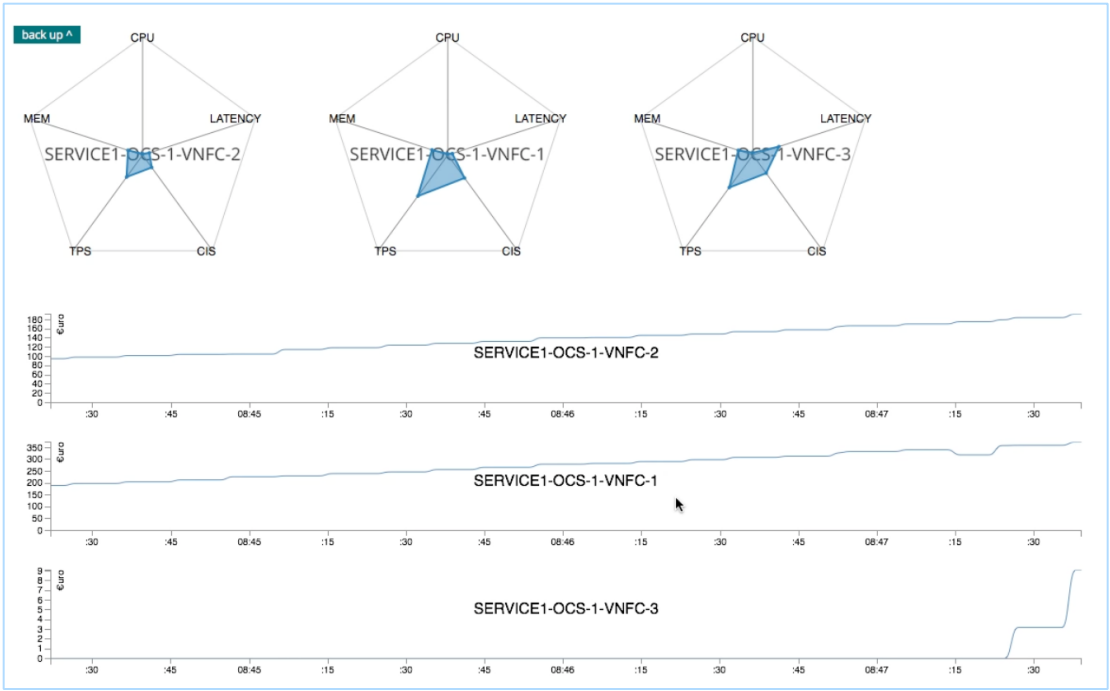

VNF Monitoring and Automated Scaling

- Incorporates Openet Streaming Analytics into a Network Function Virtualisation (NFV) architecture
- Enables real-time monitoring of Virtualised Network Function (VNF) telemetry
- Just in time delivery of network resources:
  - Predict demand and mitigate underutilisation of resources through automated feedback into orchestration layer to predict when to scale out

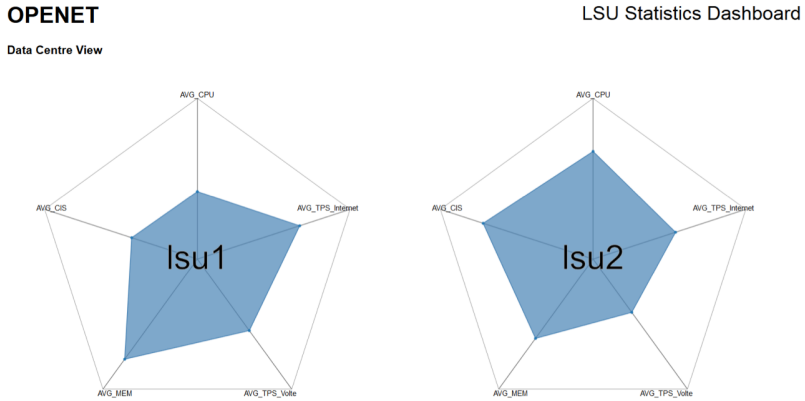

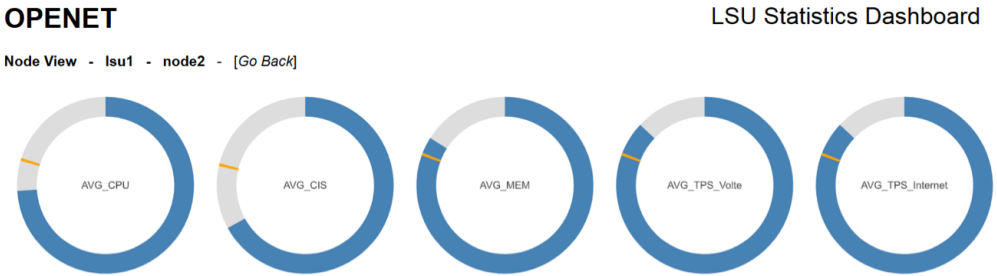

SR-IOV and CPU Pinning

Openet's PCRF, OCS and DRA have been optimized for performance using SR-IOV and CPU Pinning

- Virtualized instances of Openet's PCRF/OCS/DRA with these performance optimizations have been deployed in live networks for a number of years
- Reduce CPU overhead and latencies compared to using virtual bridges
- Provides performance values within 10% of bare metal performance

With the latest release of OpenStack now providing SR-IOV and CPU Pinning support, Openet's VNFs can now utilize these in an OpenStack enabled NFV environment

Summary

Data center infrastructure evolving to keep pace

Open standards accelerating SDI

Intel® Network Builders is enabling a broad ecosystem on open source/standards

Intel & Red Hat optimizing Ceph for Software Defined Infrastructure

Learn more at the Intel booth #905

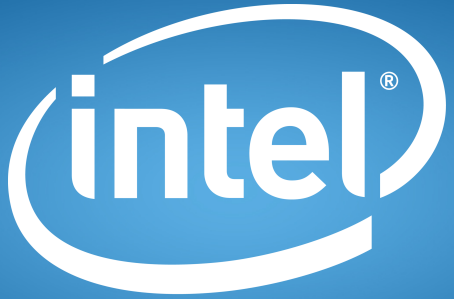

experience
what's inside™