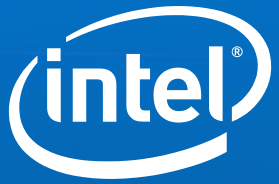

Cloud First

Intel's strategy to accelerate and differentiate cloud growth

Das Kamhout, Sr. Principal Engineer
@dkamhout
March 2018


Autonomous Driving

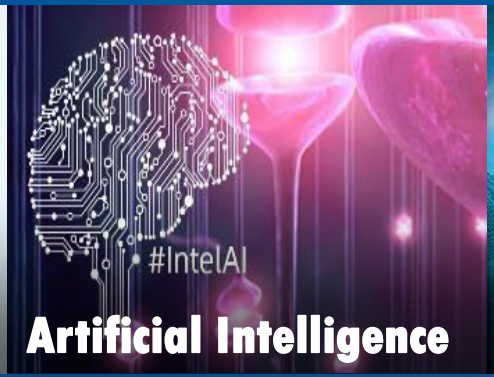
#IntelAI
Artificial Intelligence


Virtual Reality


5G Connectivity

Data Center Computing Progression

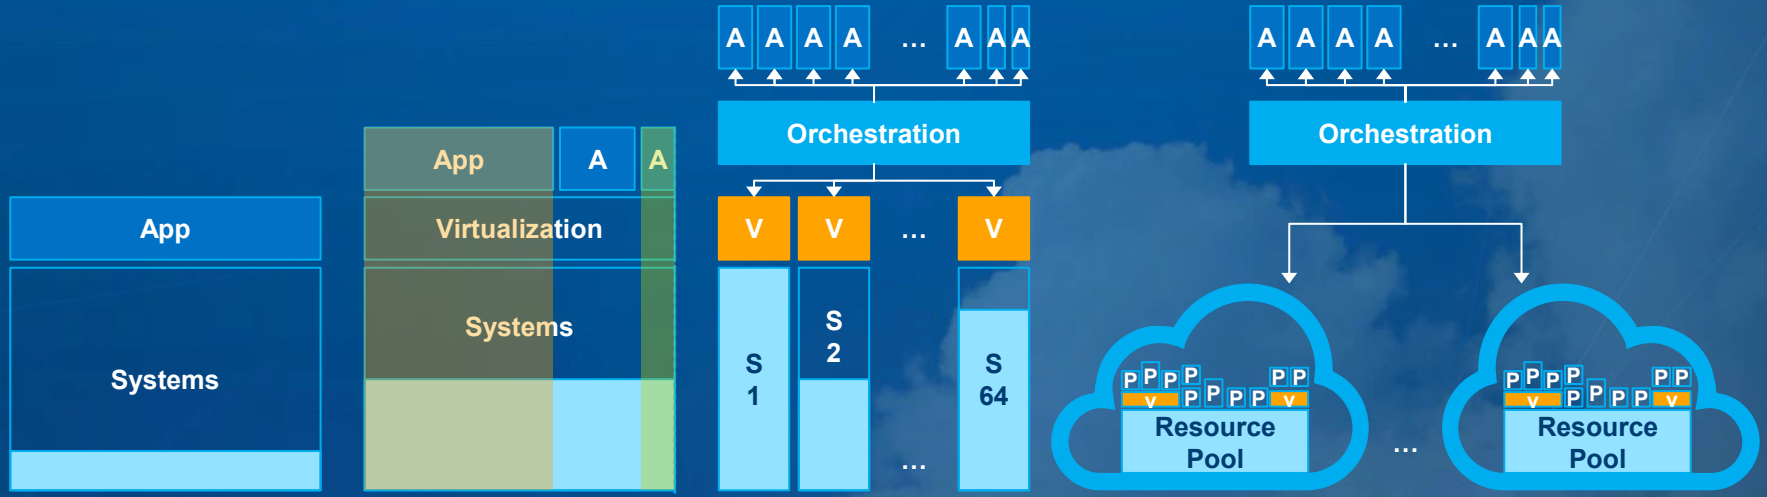

Traditional

Virtualized

Cloud (A)

1. Multi-tenant
2. Rapid elasticity
3. Self-service
4. Measured services
5. Resource pooling

Hyper-Scale Cloud (B)

1. Highly optimized
2. Highly efficient
3. Managed cross-cloud
4. Heterogeneous

Multiple computing models will persist for foreseeable future

Note: Clouds do not require virtualization

The Progression Merges

Potential Path

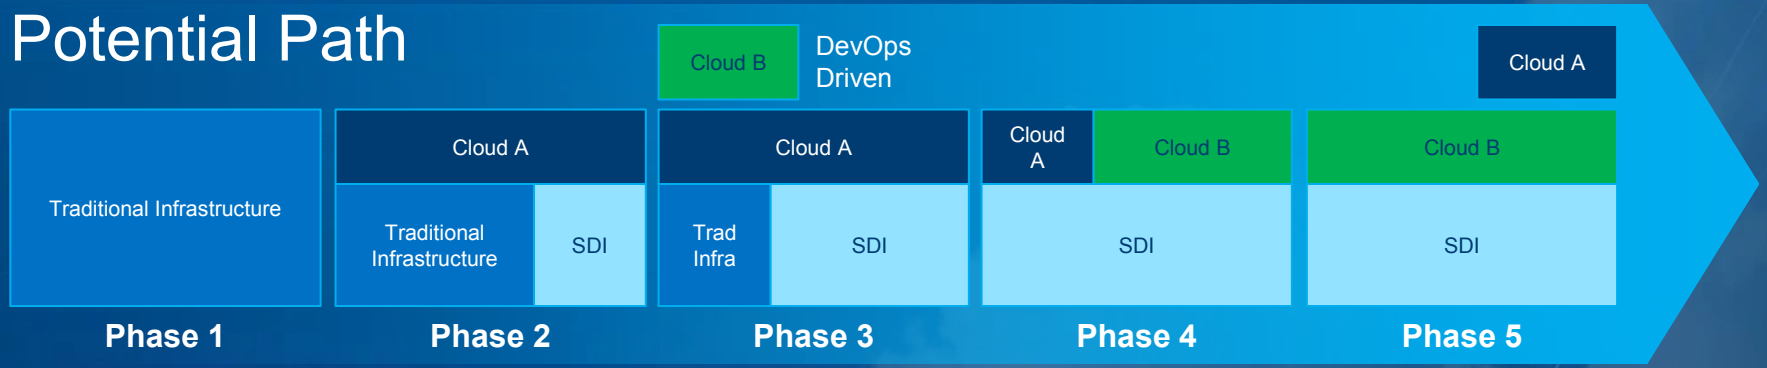

Hyper-Scale Path

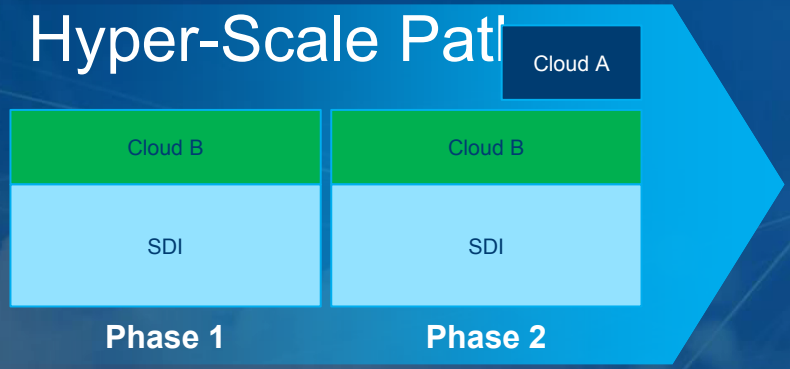

Infrastructure Solution Stack

Powers IaaS, PaaS, and SaaS

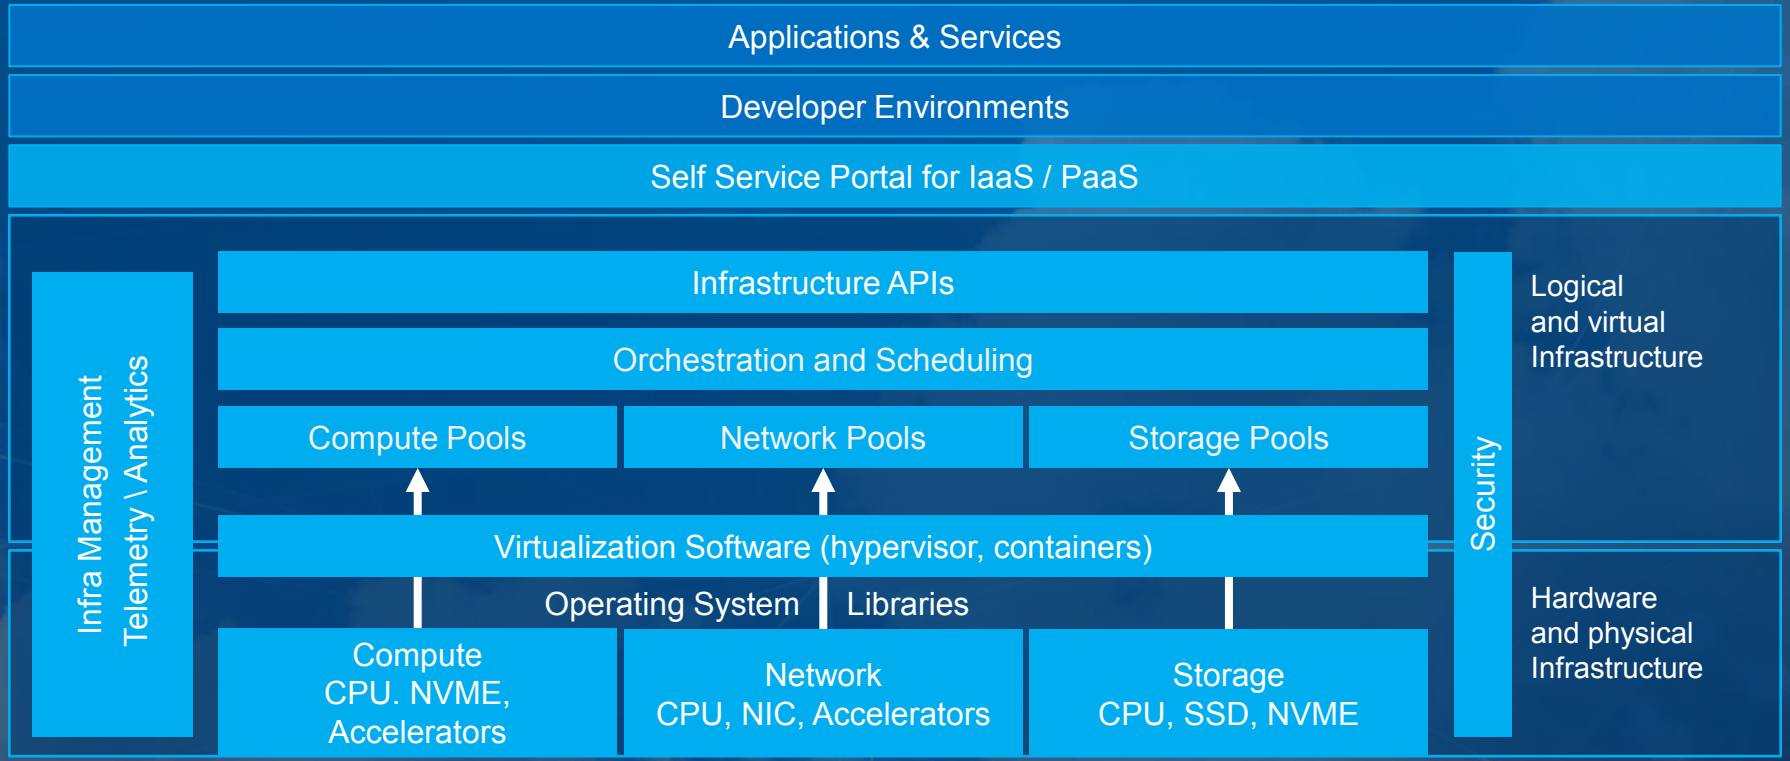

Foundation built on Xeon Scalable Platform

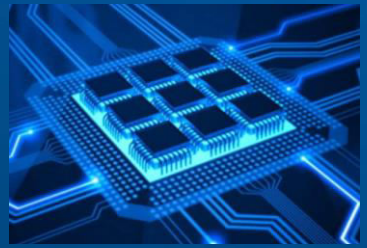

Scalable Configuration of High-Performing Cores

Higher I/O and Memory Bandwidth

Integrated Accelerators

Significant Improvement in Storage Performance

Workload Optimized Frameworks & Telemetry
(e.g. Caffe*, Intel® DAAL, Intel® MKL, DPDK, SNAP*, SPDK)

Advancing virtually every aspect: AI to API

Up to **1.65X**
average generational gains[1]

Up to **2X**
data protection performance gen over gen[2]

4.2X
greater VM capacity vs 4-year-old server[3]

65%
lower total cost of ownership vs 4-year old server[4]

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 www.intel.com/benchmarks. Configuration: Refer to Performance Benchmark Disclosure slide. Results have been estimated or simulated using internal Intel analysis or architecture simulation or modeling, and provided to you for informational purposes. Any differences in your system hardware, software or configuration may affect your actual performance. *Other names and brands may be claimed as the property of others.. Configurations: see Appendix A

Intel® Xeon® Scalable Processors

The Foundation for agile, secure, workload-optimized clouds

Best

UP TO 28 CORES

UP TO 2, 4 & 8 SOCKET SUPPORT

DDR4 2666 MHZ WITH UP TO 1.5 TB

HIGHEST

great

UP TO 22 CORES

2 & 4 SOCKET SUPPORT

UP TO 3 UPI LINKS

ADVANCED

Good

SCALABLE PERFORMANCE AT LOWER POWER

STANDARD RAS

MODERATE TASKS

ENTRY

SCALABLE PERFORMANCE

STANDARD RAS

Light TASKS

TOPLINE MEMORY CHANNEL BANDWIDTH

UP TO 6 UPI LINKS

ACCELERATOR THROUGHPUT

RELIABILITY, AVAILABILITY AND SERVICEABILITY

INTEL® TURBO BOOST TECHNOLOGY AND INTEL® HYPER-THREADING TECHNOLOGY FOR MODERATE WORKLOADS

ENTRY PERFORMANCE FOR LIGHT WORKL

LATEST

MAINSTREAM

Efficient

ENTRY

Intel's Contributions Optimize Features & Platforms

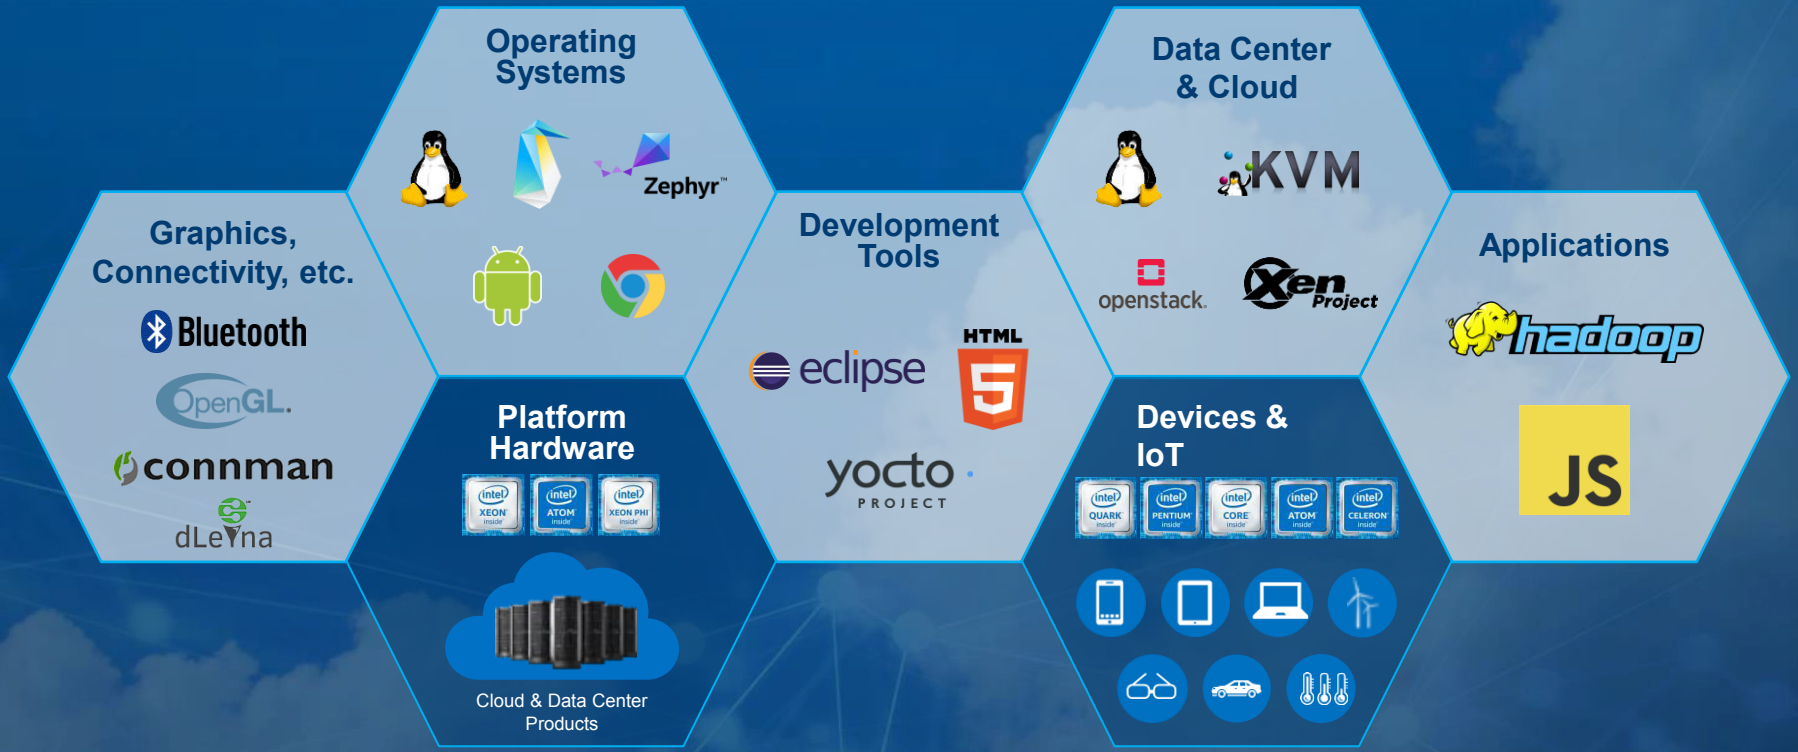

Linux Kernel Contributions

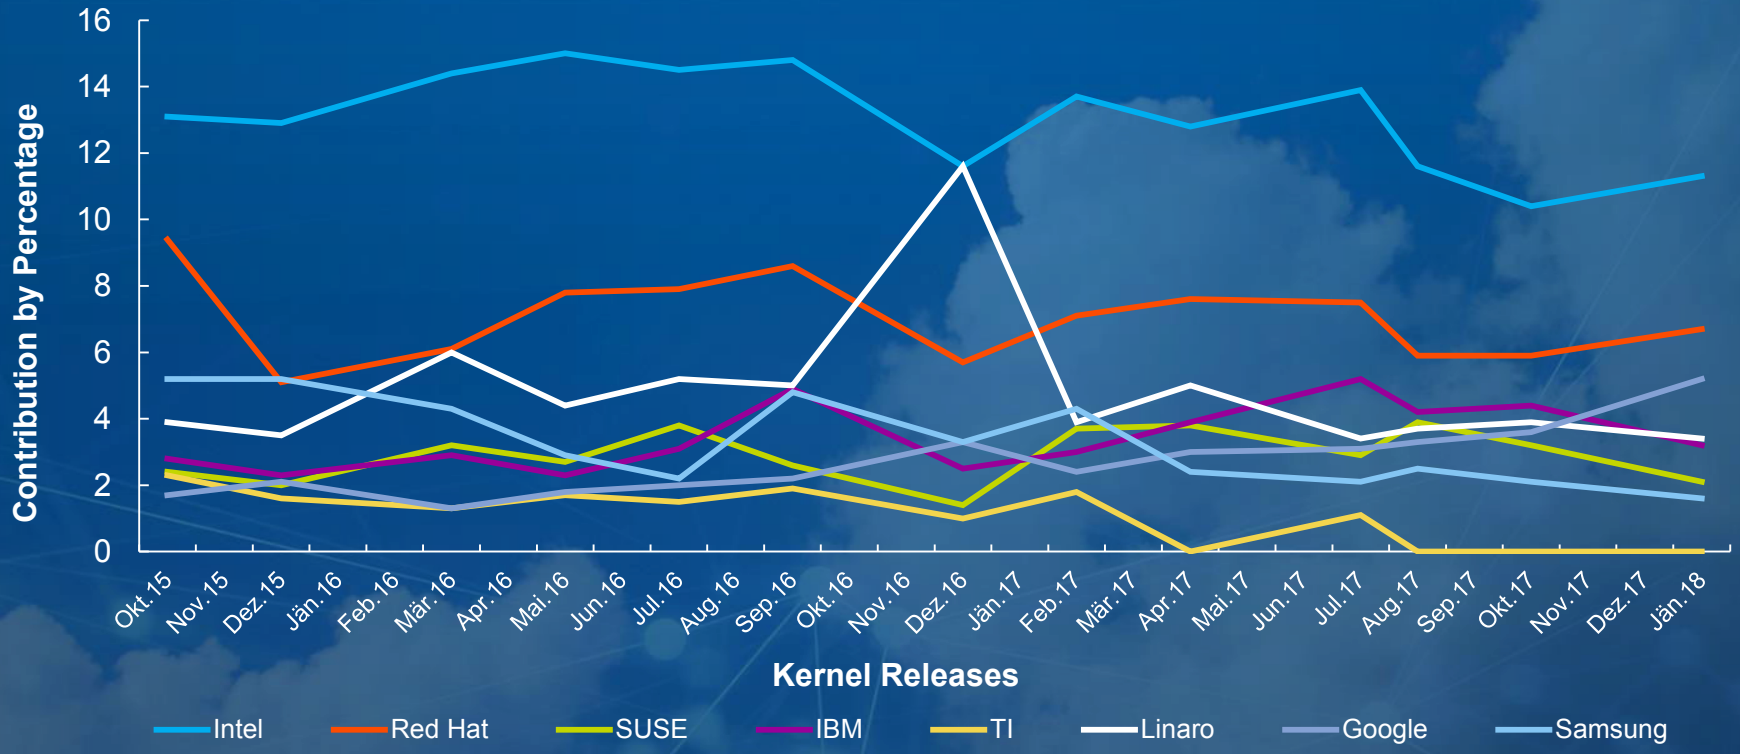

Source: http://lwn.net

Intel's Leadership in OpenStack

Foundation Board

Setting strategy

Platinum (permanent) board seat (Imad Sousou)

Individual (elected) board seat (Shane Wang)

Maximum representation allowed for a single company

Working Groups

Prioritizing development

Product WG

Enterprise WG

Telco WG

App Ecosystem Development WG

Diversity WG

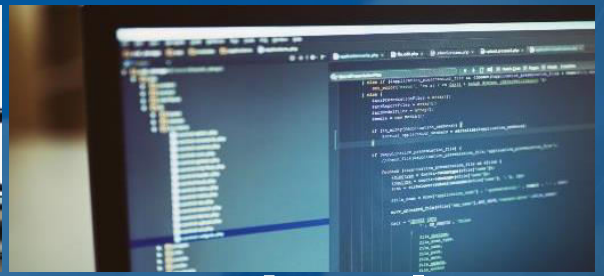

Technical Leadership

Code talks

Project Technical Leads (PTLs) (*related to usability*)

Core reviewers

Top 10 contributor

CLOUD NATIVE COMPUTING FOUNDATION

kubernetes

CNI

Prometheus


#IntelAI
Artificial Intelligence

High Risk Inspection by Drones: 1 CPU Node

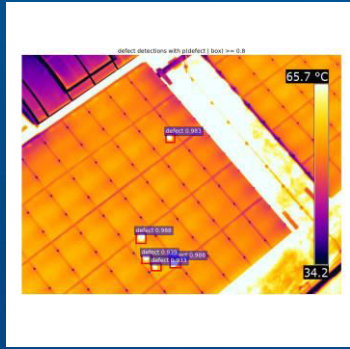

Caffe

intel XEON inside

FRAMEWORK HARDWARE

Time to train: 6 hours

Chong Y., Yiqiang Z and Jiong G., "Automatic Defect Inspection Using Deep Learning for Solar Farm" Dec. 2017. https://software.intel.com/en-us/articles/automatic-defect-inspection-using-deep-learning-for-solar-farm

*Optimization notice slide 24

Silicon Package Defect Detection: 8 cpu nodes

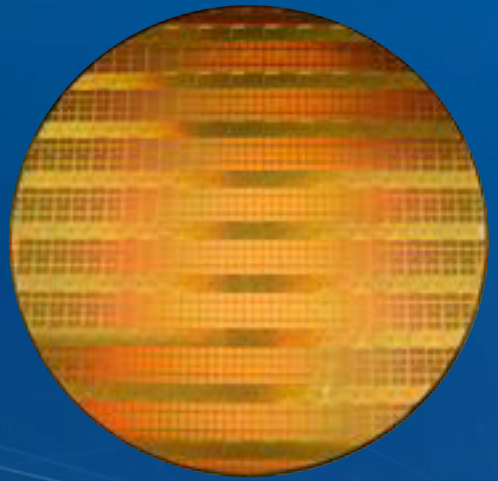

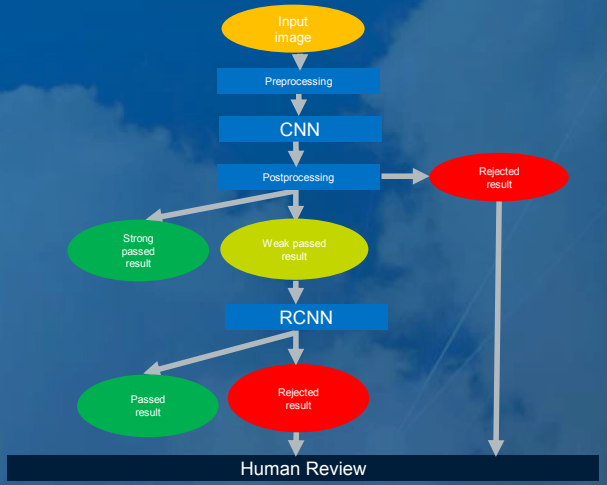

Training within one hour on 8 CPU nodes.

Caffe

FRAMEWORK

intel XEON inside

HARDWARE

Z. Yiqiang and J. Gong, "Manufacturing package fault detection using deep learning." Aug. 2017. https://software.intel.com/en-us/articles/manufacturing-package-fault-detection-using-deep-learning

intel

Scaling analysis : 32/64 CPU nodes

Throughput Scaling (1 node → 32/64 nodes)

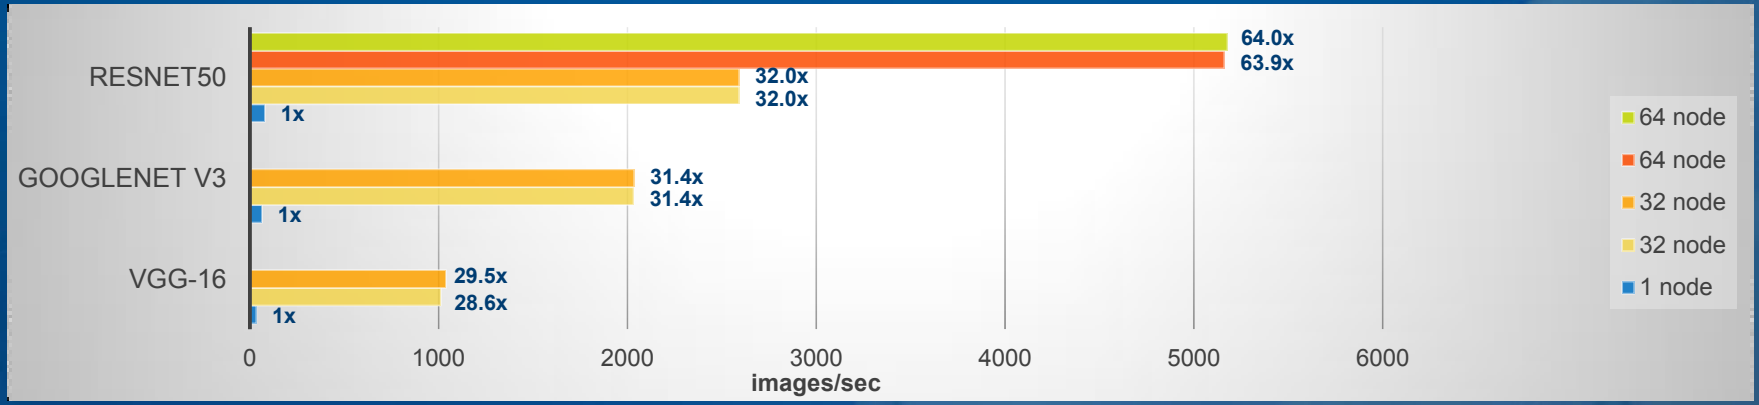

Resnet-50 Time to Train Performance

System Configuration	Network Fabric	Minibatch Size	Top-1 Accuracy	Measured TTT
64-node Intel® Xeon® Scalable Processor system Intel® Xeon® Gold 6148 Processor based *	10Gb Ethernet	8192	75.9%	7.3 hours

* Intel Internal measured data on system configuration noted above. Configuration slides attached

All Intel data measured with Intel® Distribution of Caffe* and Intel® Machine Learning Scaling Library (Intel® MLSL)

Scaling Efficiency (ResNet-50): 256 CPU NODES

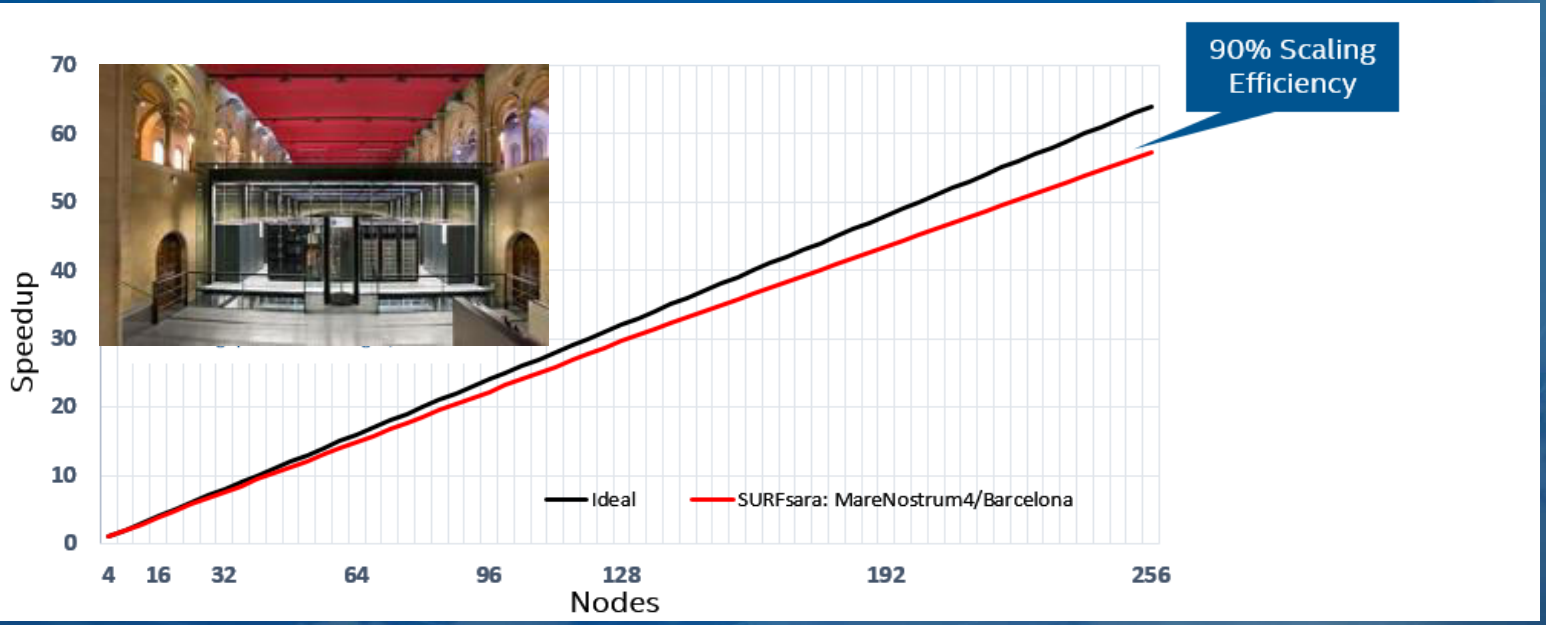

V. Codreanu et al, "Achieving Deep Learning Training in less than 40 Minutes"
https://blog.surf.nl/en/imagenet-1k-training-on-intel-xeon-phi-in-less-than-40-minutes/

Intel® - SURFsara* Research Collaboration

HARDWARE

90% scaling efficiency with up to 74% Top-1 accuracy on 256 nodes

Time to Train: 1600 cpu nodes – UC Berkeley

ResNet-50 Time to Train

31
minutes

AlexNet Time to Train

11
minutes

HARDWARE

Large Batch Size method with Layer-wise Scaling Layer-wise Adaptive Rate Scaling (LARS) algorithm

- Technical Report by Y. You, Z. Zhang, C-J. Hsieh, J. Demmel, K. Keutzer: https://people.eecs.berkeley.edu/~youyang/publications/imagenet_minutes.pdf

Cloud Service Progression

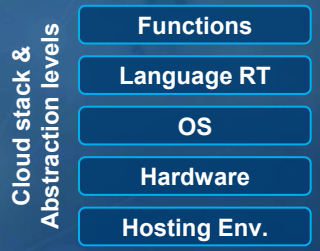

Decreasing concern (& control) over cloud stack implementation

Rainer Sträter

VP of Data Center Infrastructure
1&1 Internet Group

The “beautiness” of CaaS (Container as a Service)

- SharedWebHosting eco-system
- Fast growing number of apps, tools and variants of dev stacks
- Managed and unmanaged
- Seamless transfer of apps/data

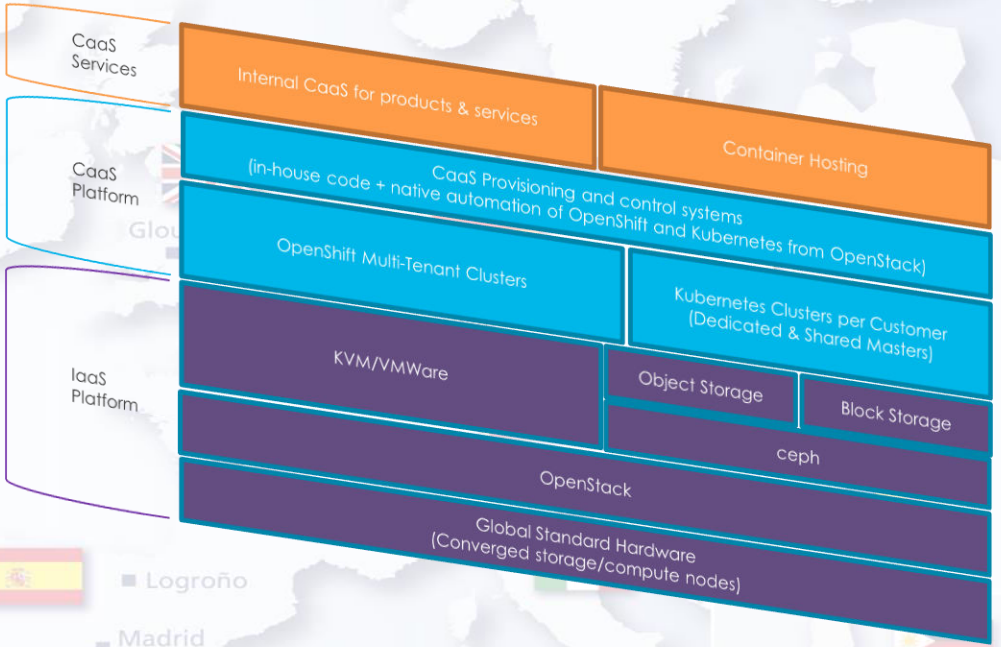

Inside the machine room

- High container density driven by Intel® Xeon ® SP performance
- High speed east-west traffic
- Container relocation driven by AI based on telemetry

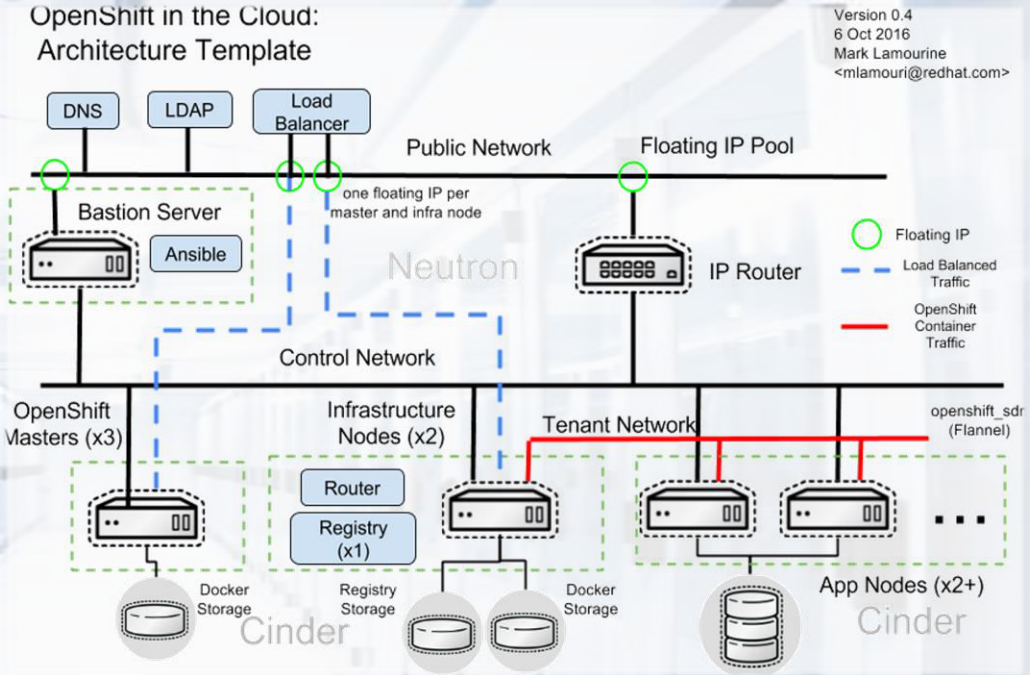

Architecture

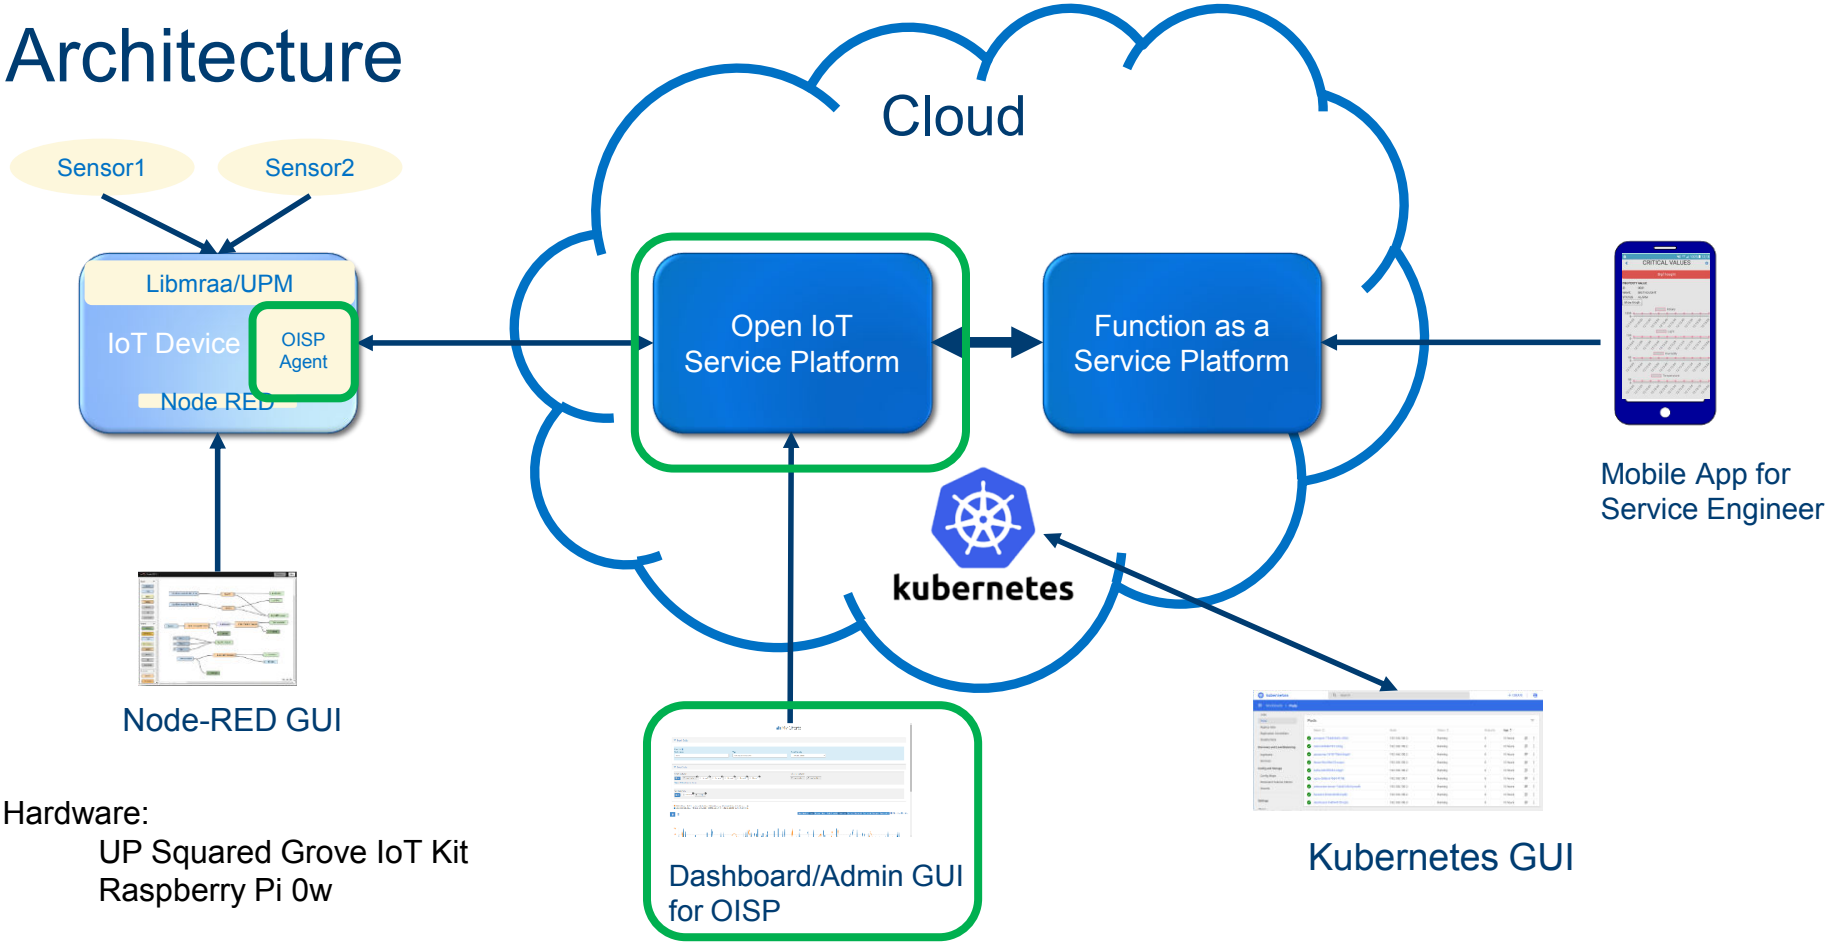

Hardware:

- UP Squared Grove IoT Kit
- Raspberry Pi 0w

It’s all real!

Screenshots

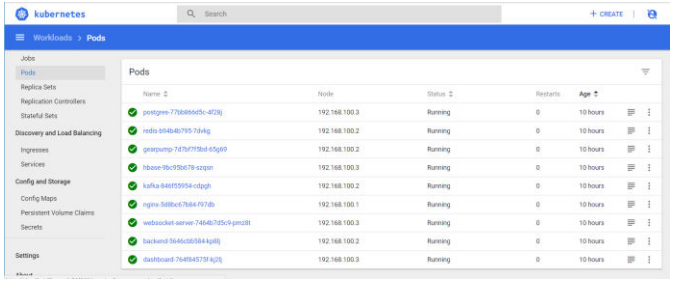

Kubernetes UI for OISP deployment

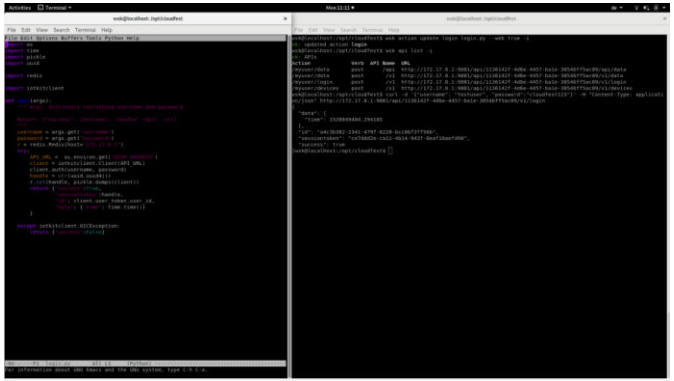

FaaS console to submit function

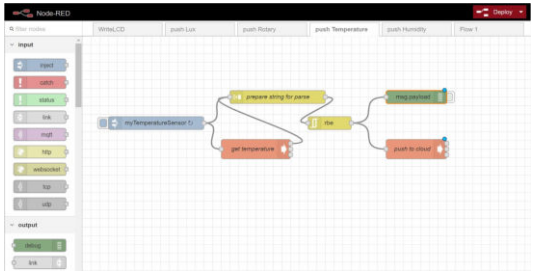

Node RED IoT configuration

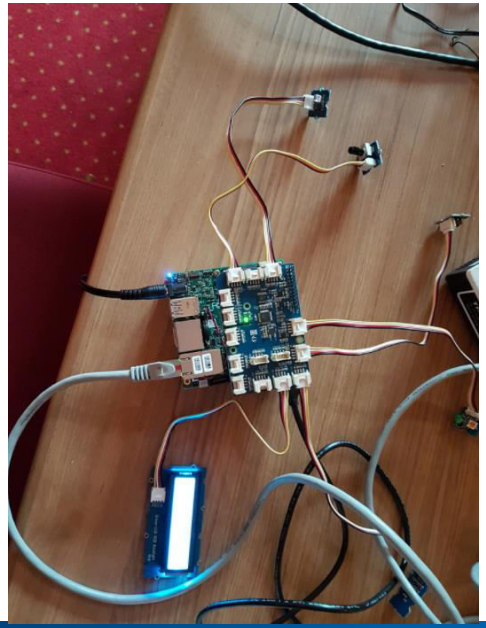

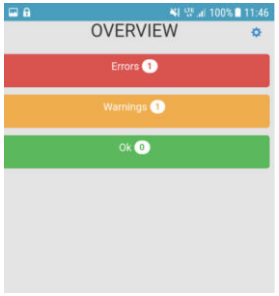

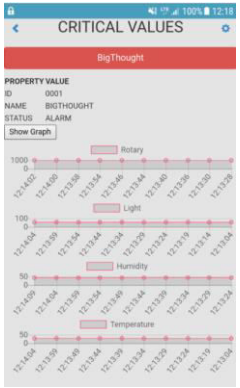

Mobile App for Service Engineer

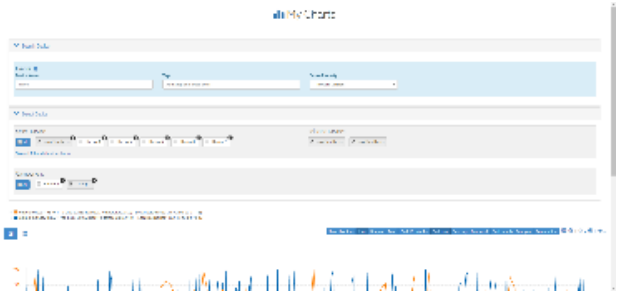

Service/Admin GUI

Get Started with Intel's Cloud Solutions Today

3 **Joint GTM** — Work with Intel to create, market and sell new cloud services.

2 **Try it yourself** — Test drive new Intel® technologies through Cloud Insider Sandbox.

1 **Learn** — Go to www.intel.com/CSP, or join Intel® Cloud Insider Program.

in the hyper evolution of new apps and services

intel HW+SW investments are driving tech forward

Let's partner together to create the new normal

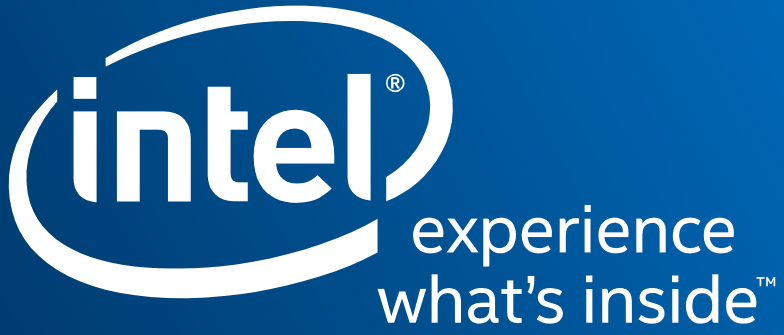
intel®
experience
what's inside™

Notices & 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 computer system 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 .

Benchmark results were obtained prior to implementation of recent software patches and firmware updates intended to address exploits referred to as "Spectre" and "Meltdown." Implementation of these updates may make these results inapplicable to your device or system.

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.

Configuration details
32/64-node CPU system Intel® Xeon® 6148 Gold processor with 10GB Ethernet / OPA

Benchmark Segment	AI/ML
Benchmark type	Training
Benchmark Metric	Images/Sec or Time to train in seconds
Framework	Caffe
Topology	Resnet-50, VGG-16, GoogleNet V3
# of Nodes	32/64
Platform	Wolfpass (Skylake)
Sockets	2S
Processor	Xeon Processor code named Skylake, B0, ES2*, 24c, 2.4GHz, 145W, 2666MT/s, QL1K CPUID=0x50652
BIOS	SE5C620.86B.01.00.0412.020920172159
Enabled Cores	24 cores / socket
Platform	Wolfpass (Skylake)
Slots	12
Total Memory	192GB
Memory Configuration	12x16GB DDR4 2R, 1.2V, RDIMM, 2666MT/s
Memory Comments	Micron MTA 18ASF2G72PDZ-2G6B1
SSD	800GB Model: ATA INTEL SSDSC2BA80 (scsi)
OS	Oracle Linux Server 7.3, Linux kernel 3.10.0-514.6.2.0.1.el7.x86_64.knl1

Ethernet Configurations	Intel Corporation Ethernet Connection X722 for 10GBASE-T (rev 03)
Omni-Path Configurations	Intel Omni-Path HFI Silicon PCIe Adapter 100 Series [discrete]. OFED Version 10.2.0.0.158_72. 48 port OPA switch, with dual leaf switches per rack 48 nodes per rack, 24 spine switches
HT	ON
Turbo	ON
Computer Type	Server
Framework Version	Internal Caffe version
Topology Version	Internal ResNet-50 topology Internal VGG-16 topology Internal GoogleNet V3 topology
Batch size	ResNet-50 : 128 x # of node VGG-16 : 64 x # of node GoogleNet V3 : 64 x # of node
Dataset, version	Imagenet, ILSVRC 2012 (Endeavor location), JPEG resized 256x256
MKLDNN	aab753280e83137ba955f8f19d72cb6aaba545ef
MKL	mklml_lnx_2018.0.1.20171007
MLSL	2017.2.018
Compiler	Intel compiler 2017.4.196

Configuration details of Amazon* EC2 C5.18xlarge 1/32/64/128 node systems

Benchmark Segment	AI/ML
Benchmark type	Training
Benchmark Metric	Images/Sec
Framework	Intel Caffe
Topology	Resnet-50
# of Nodes	1/32/64/128
Platform	Amazon EC2 C5.18xlarge instance
Sockets	2S
Processor	Intel® Xeon® Platinum 8124M CPU @ 3.00GHz (Skylake)
BIOS	N/A
Enabled Cores	18 cores / socket
Platform	N/A
Slots	N/A
Total Memory	144GB
Memory Configuration	N/A
SSD	EBS Optimized 200GB, Provisioned IOPS SSD
OS	Centos 7.4 (HVM)

Network Configurations	Amazon Elastic Network Adapter (ENA) 25 Gbps of aggregate network bandwidth Installed Enhanced Networking with ENA on Centos Placed the all instances in the same placement group
HT	ON
Turbo	ON
Computer Type	Server
Framework Version	Intel Caffe version 1.0.6
Topology Version	Intel Caffe ResNet-50 and GoogleNet V3 internal version available from https://github.com/intel/caffe/tree/master/models/intel_optimized_models
Batch size	GoogleNet V3 : 64 x # of node ResNet-50 : 128 x # of node
Dataset, version	Imagenet, ILSVRC 2012, JPEG resized 256x256
MKLDNN	c7ed32772affaf1d9951e2a93d986d22a8d14b88
MKL	mklml_lnx_2018.0.20170908
MLSL	ecc6db2a133bab3894993baac54a01334c12b95a with internal patch
Compiler	gcc/g++: 4.8.5 Icc/icpc: 17.0.5

Backup

Ceph Community

Contributors by Company†

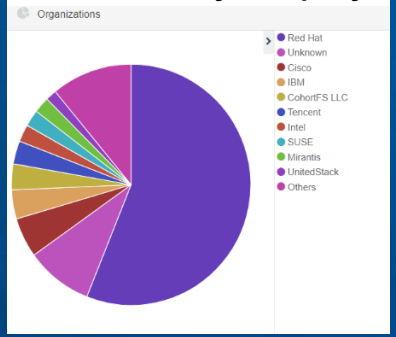

Large # of contributors

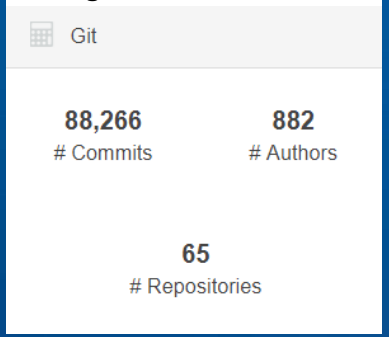

2018 Focus areas

- NFVi optimizations for hyperconverged Ceph
- Containerized control plane for Ceph (ongoing)
- DPDK/SPDK-based Ceph OSD – Support AT&T to deploy Ceph multi-OSD in OpenStack Helm
- Rack Scale Design

Community Advisory Board

- Red Hat (chair)
- **Intel (Anjaneya Chagam)**
- Canonical
- CERN
- Cisco
- Fujitsu
- 42on
- SUSE
- SanDisk

Technical Leadership and Contributions

Upstreamed Features and Enhancements

- Intel Storage Acceleration Library (ISA-L) Integration
- Storage Performance Development Kit (SPDK) Integration for NVMe drivers
- BlueStore as new ObjectStore and RocksDB enhancements
- Cache Tiering
- CeTune tool

Available for Ceph but not upstream... yet

- Quick Assist Technology (QAT) Integration for Ceph Encryption and Compression
- Ceph Performance Tracing and Profiling
- Remote Direct Memory Access (RDMA) Enabling
- Persistent client-side cache

† Intel in Top 5 for the latest Luminous release
References: https://metrics.ceph.com