AI Bridging Cloud Infrastructure (ABCI) and its I/O Architecture

Algorithm + Big Data + Computing Power + Infrastructure

AI Bridging Cloud

Hirotaka Ogawa

National Institute of Advanced Industrial Science and Technology (AIST), Japan

Frankfurt, June, 2019

ABCI: The World's First Large-Scale Open AI Infrastructure

東京大学 THE UNIVERSITY OF TOKYO AIST


Univ. Tokyo / Kashiwa II Campus

OPERATED SINCE AUGUST 1ST, 2018

ABCI AI Bridging Cloud Infrastructure

- World Top-Level compute and data process capability
- **Open, Public, and Dedicated** infrastructure for Al & Big Data Algorithms, Software, and Applications
- **Open Innovation Platform** to accelerate joint academic-industry R&D for AI

Peak Performance:
550 PFLOPS (FP16)
37 PFLOPS (FP64)
Effective Performance: (as of June 2019)
19.88 PFLOPS (#8 in TOP500)
14.423 GFLOPS/W (#4 in GREEN500)
508.85 TFLOPS (#5 in HPCG)
Power Usage: < 2.3 MW
Average PUE: < 1.1 (Estimated)

AI Infrastructure for Everyone

Expert

- ABCI Grand Challenge: Demonstration of highly challenging academic and/or industrial themes using the whole ABCI resources for 24 hours for free

Advanced & Intermediate

- Up to 512-node computing resource is available for everyone
- Software, datasets, and pre-trained models are ready for use
- High computing capability enables to accelerate AI R&D and promote social implementation

Beginner

- User friendly WebUI based IDE for supporting beginners of deep learning
- PoC platform for B2B2C business model

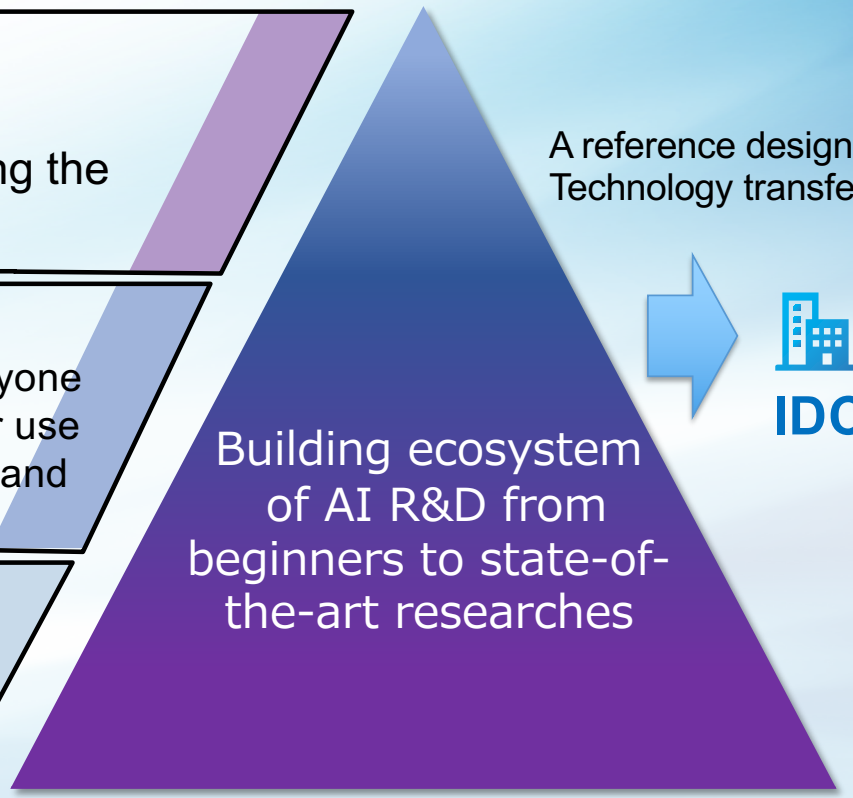

Promoting the use by over 100 institutions and over 1000 researchers/engineers

World's Highest Speed in ImageNet-1k Training

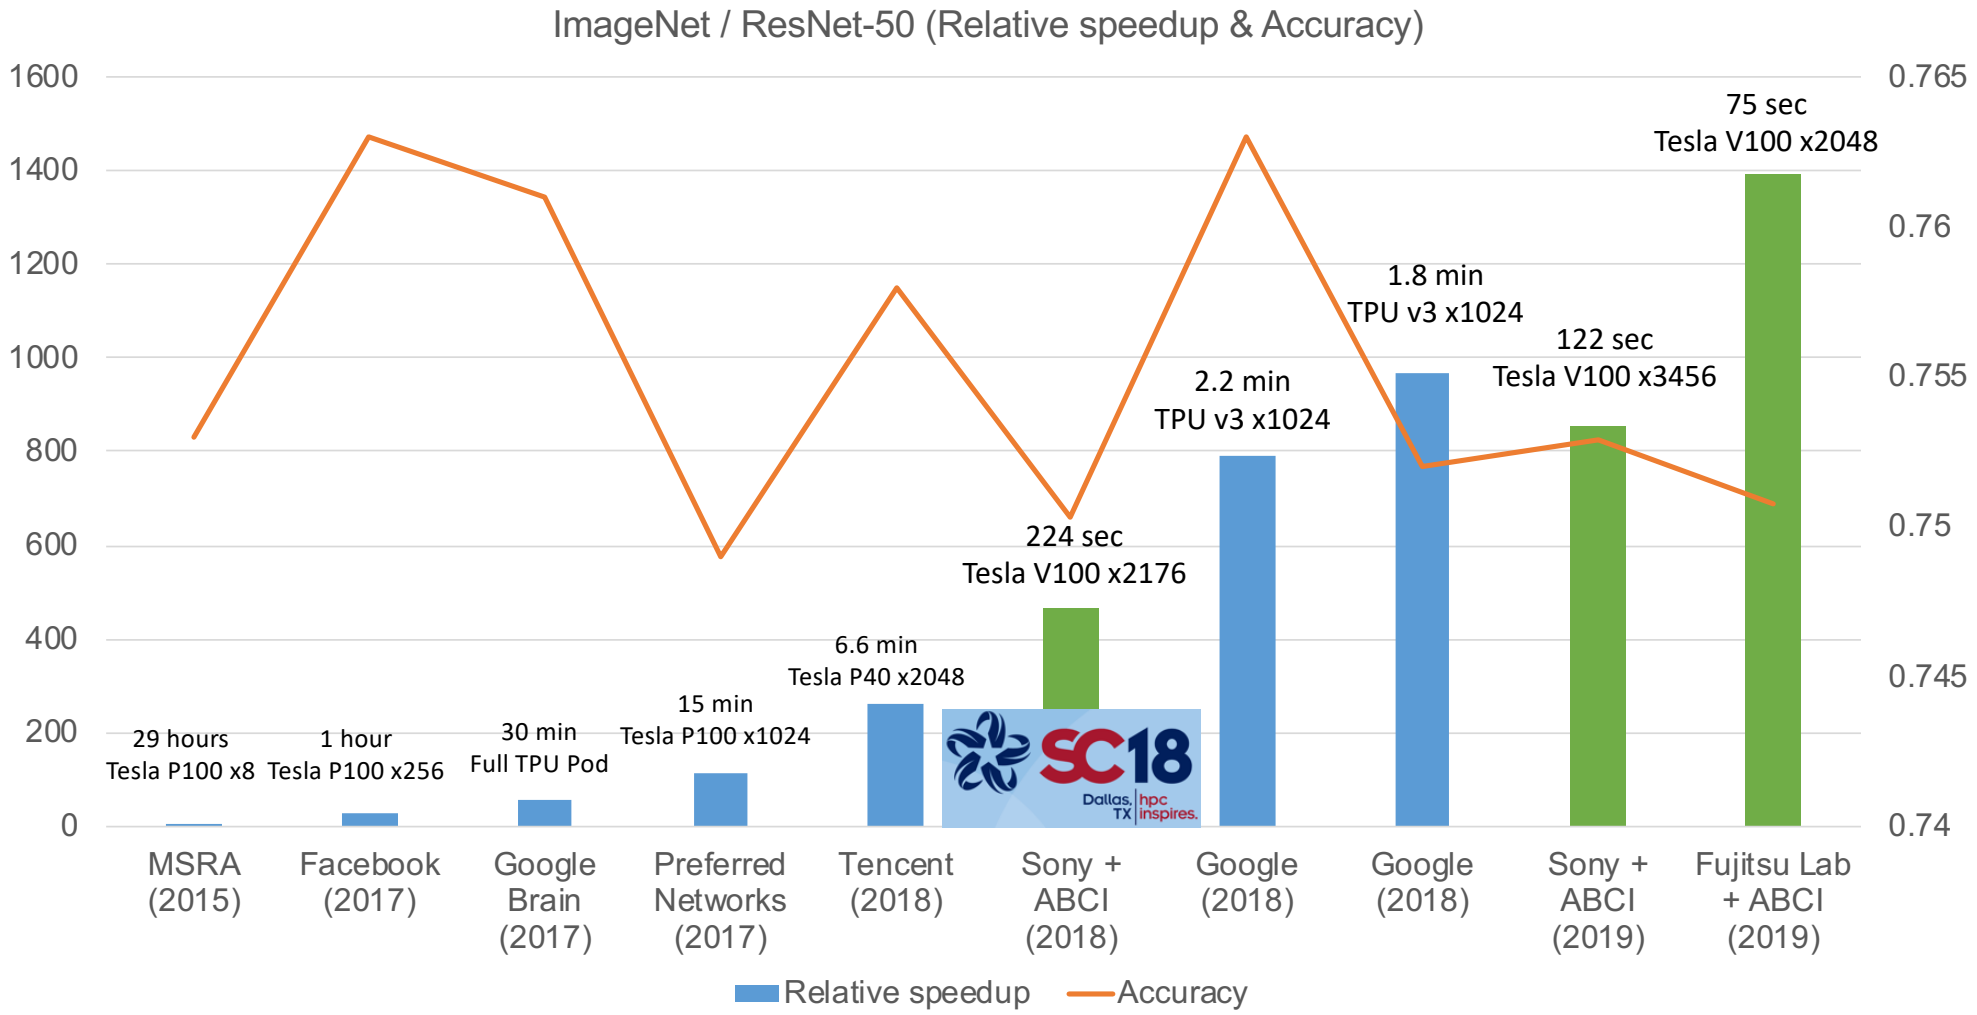

Large-scale Reinforcement Learning for Computer Go

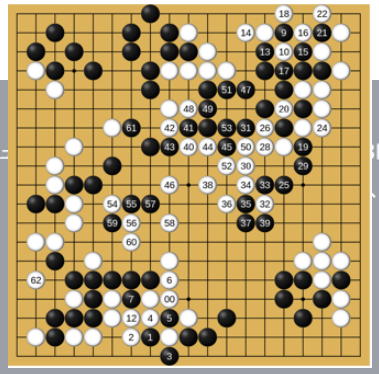

GLOBIS

HOME > ニュ … IS-AQZ」プロジェクトを発表、囲碁AI世界一と若手棋士育成を目指す … 、日本棋院、トリプルアイズ、産総研、東大・松尾研と協働

MENU

JP

EN

2019/04/18

GLOBIS-AQZ

グロービス、「GLOBIS-AQZ」プロジェクトを発表、囲碁AI世界一と若手棋士育成を目指す

囲碁AI開発者・山口祐氏、日本棋院、トリプルアイズ、産総研、東大・松尾研と協働

https://www.globis.co.jp/news/release/20190418_globis.html

Large-scale Video Analysis Demonstration in cooperation between ABCI and SINET5

Large-scale Training on ABCI

3D ResNet [1]

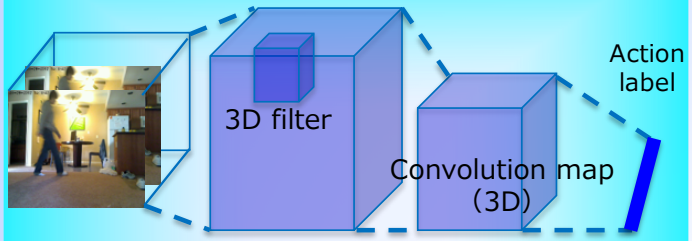

State-of-the-art 3D CNN-based action recognition developed by AIST

Video/Action label

SINET5
https://www.sinet.ad.jp/

High-speed R&E network operated by NII (100Gbps)

Live video from WebCam

AI/SUM in Tokyo

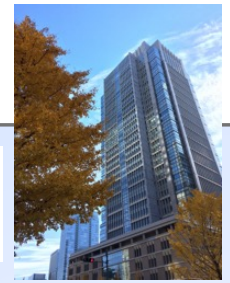

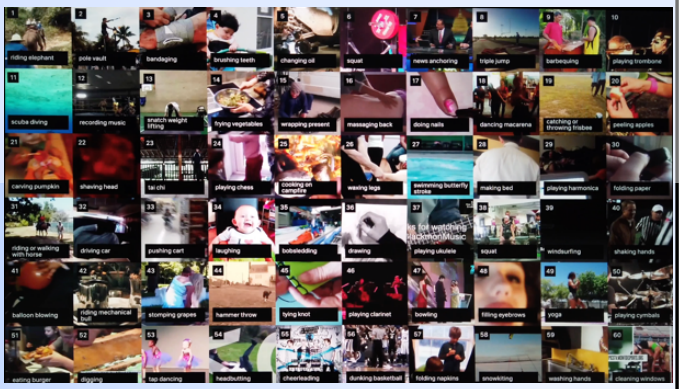

Live demonstration of huge volume video data analysis by ABCI

[1] K. Hara et al., “Can Spatiotemporal 3D CNNs Retrace the History of 2D CNNs and ImageNet?", CVPR18

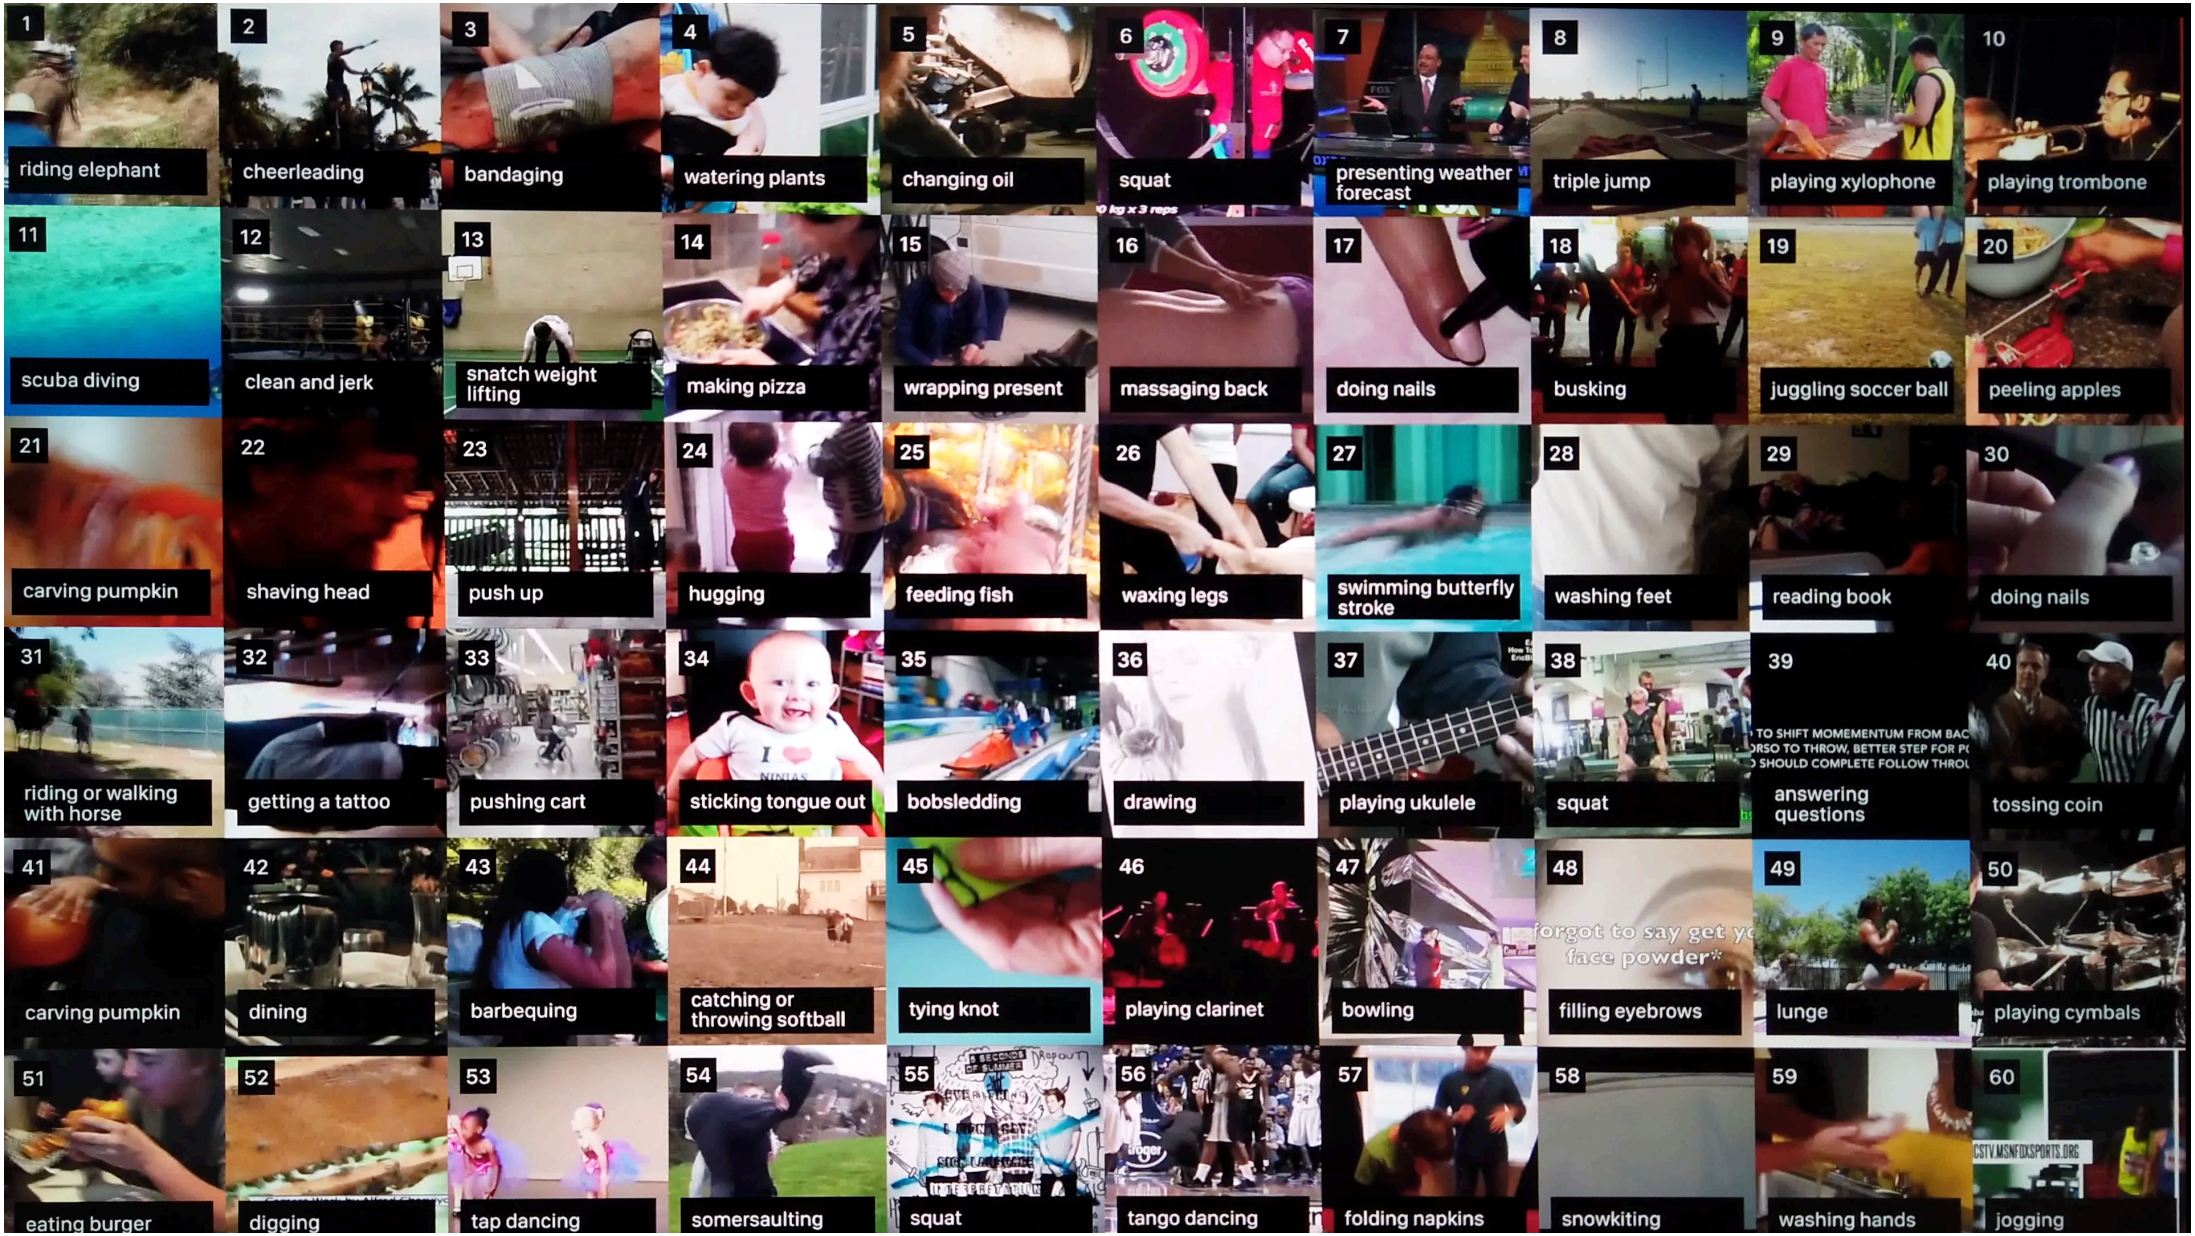
1 riding elephant
2 cheerleading
3 bandaging
4 watering plants
5 changing oil
6 squat
7 presenting weather forecast
8 triple jump
9 playing xylophone
10 playing trombone
11 scuba diving
12 clean and jerk
13 snatch weight lifting
14 making pizza
15 wrapping present
16 massaging back
17 doing nails
18 busking
19 juggling soccer ball
20 peeling apples
21 carving pumpkin
22 shaving head
23 push up
24 hugging
25 feeding fish
26 waxing legs
27 swimming butterfly stroke
28 washing feet
29 reading book
30 doing nails
31 riding or walking with horse
32 getting a tattoo
33 pushing cart
34 sticking tongue out
35 bobsledding
36 drawing
37 playing ukulele
38 squat
39 answering questions
40 tossing coin
41 carving pumpkin
42 dining
43 barbequing
44 catching or throwing softball
45 tying knot
46 playing clarinet
47 bowling
48 filling eyebrows
49 lunge
50 playing cymbals
51 eating burger
52 digging
53 tap dancing
54 somersaulting
55 squat
56 tango dancing
57 folding napkins
58 snowkiting
59 washing hands
60 jogging

SONY began providing User friendly WebUI based Deep Learning Toolkit w/ using ABCI (Apr 8th just announced)

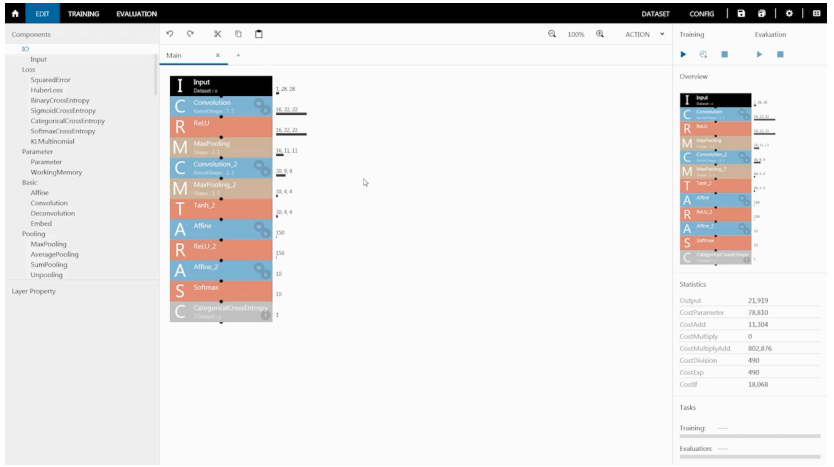

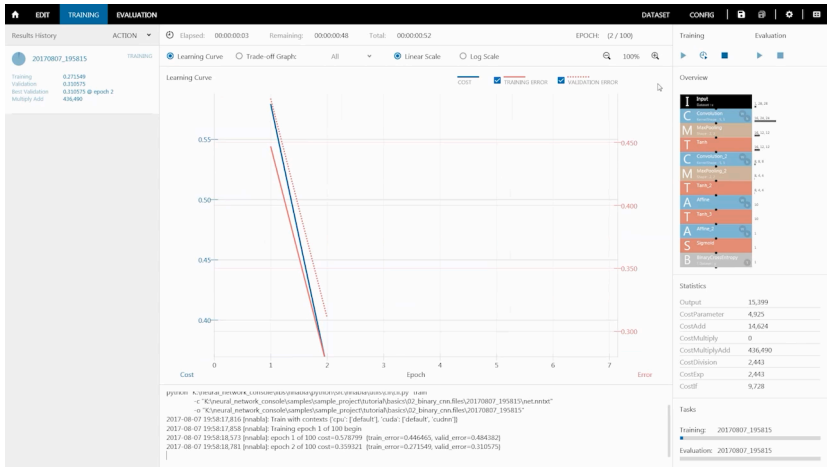

https://dl.sony.com/ja/cloud/abci/index.html

AI Datacenter & Facility

AI Datacenter
"Commoditizing supercomputer cooling technologies to Cloud (70kW/rack)"

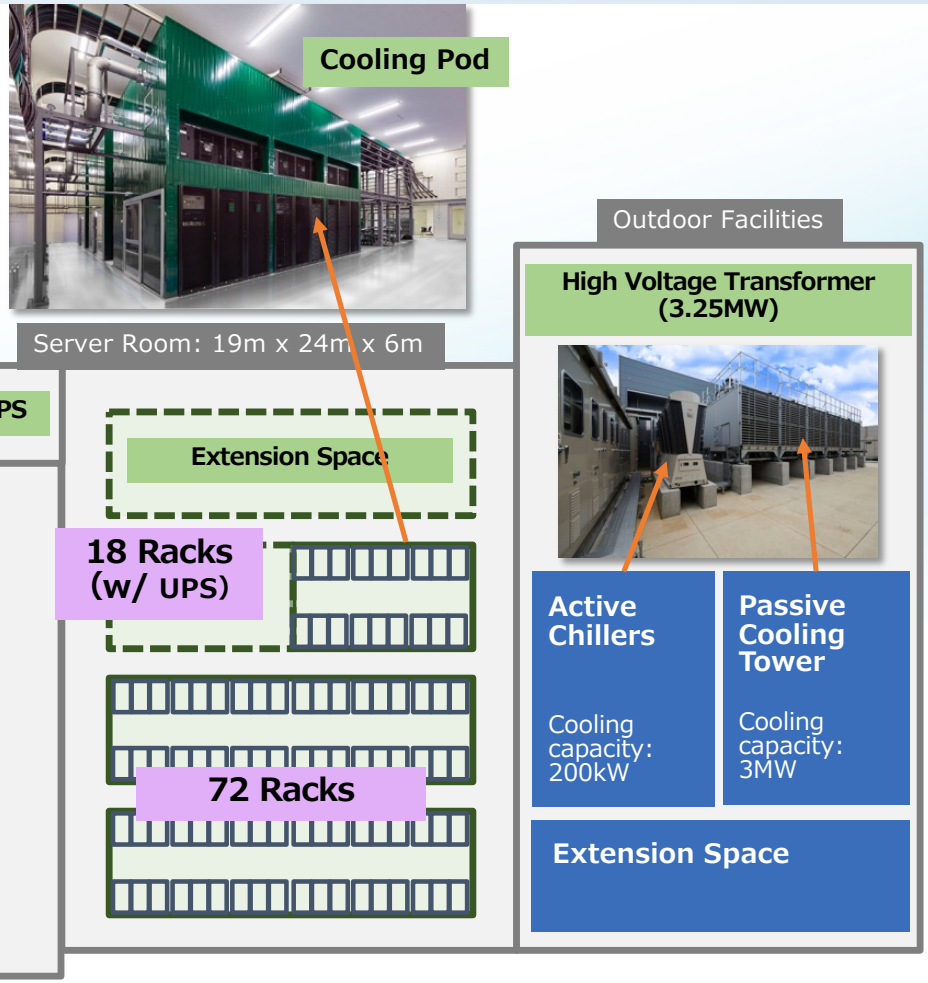

- **Single floor, cost effective building**
- **Hard concrete floor 2t/m² weight tolerance for racks and cooling pods**
- **Number of Racks**
  - Initial: 90 (ABCI uses 41 racks)
  - Max: 144
- **Power capacity: 3.25 MW**
  - ABCI uses 2.3MW max
- **Cooling capacity: 3.2MW**
  - 70kW/rack: 60kW water+10kW air


AIST
May 23rd, 2018
FCU-2R-01

Hybrid Water/Air Cooling System

Fan Coil Unit (FCU)

Cooling Pod / Rack

Water 32℃

Water 40℃

Air Cooling

Air 35℃

Air 40℃

Outdoor Facilities

Computing Node

Cooling Tower

CDU

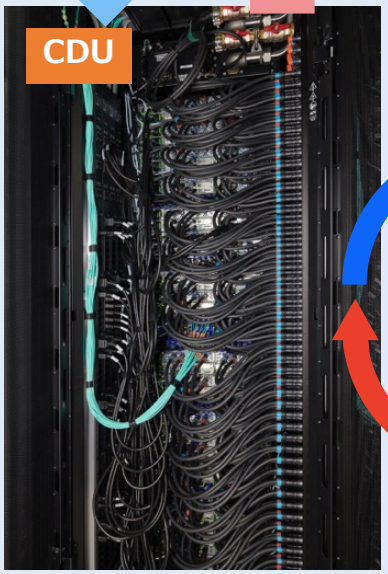

Water 32℃

Water 40℃

2nd closed circuit

1st closed circuit

Water Cooling

Water Block

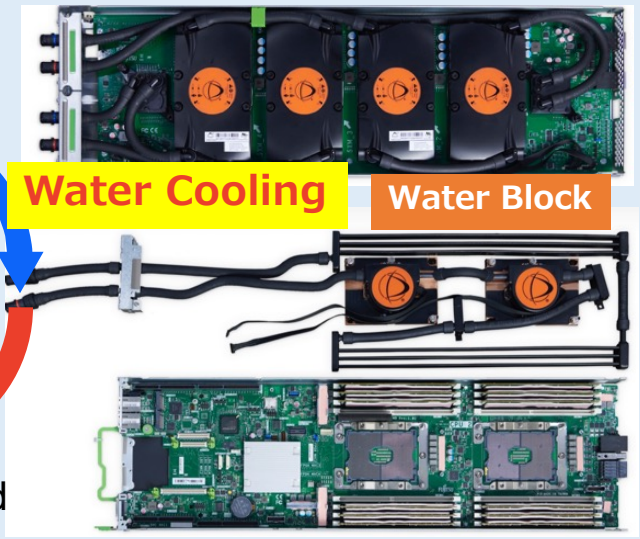

ABCI HW/SW Overview

ABCI Hardware Overview

High-Performance Computing System

550 PFlops(FP16), **37.2** PFlops(FP64)
476 TiB Memory, **1.74** PB NVMe SSD

Computing Nodes (w/ GPU) x 1088

GPU	NVIDIA Tesla V100 SXM2 x 4
CPU	Intel Xeon Gold 6148 (2.4GHz/20cores) x 2
Memory	384GiB
Local Storage	Intel SSD DC P4600 (NVMe) 1.6TB x 1
Interconnect	InfiniBand EDR x 2

Multi-platform Nodes (w/o GPU) x 10

- Intel Xeon Gold 6132 (2.6GHz/14cores) x 2
- 768GiB Memory, 3.8TB NVMe SSD

Interactive Nodes x 4

Management and Gateway Nodes x 15

Large-scale Storage System

22 PB GPFS + **1** PB Lustre +
17 PB Scality RING Object Storage

DDN SFA14K (w/ SS8462 Enclosure x 10) x 3

- 12TB 7.2Krpm NL-SAS HDD x 2400
- 3.84TB SAS SSD x 216
- NSD Server x 12

DDN SFA14KX (w/ SS9012 Enclosure x 5) x 1

- 7.68TB SAS SSD x 185, 900GB x 13
- MDS Server x 2, OSS Server x 4

HPE Apollo 4510 Gen10 x 24

- 12TB SATA HDD x 1440
- 3.2TB SSD x 24

Interconnect (Infiniband EDR)

- Mellanox CS7500 x 2
- Mellanox SB7890 x 229

Service Network (10GbE)

Gateway and Firewall

- Nexsus 3232C x2
- FortiGate 1500D x2
- FortiAnalyzer 400E x1

100Gbps

SINET5

ABCI Computing Node

FUJITSU PRIMERGY Server (2 servers in 2U)	
CPU	Xeon Gold 6148 (27.5M Cache, 2.40 GHz, 20 Core) x2
GPU	NVIDIA Tesla V100 (SXM2) x4
Memory	384GiB DDR4 2666MHz RDIMM
Local Storage	1.6TB NVMe SSD (Intel SSD DC P4600 u.2) x1
Interconnect	InfiniBand EDR x2

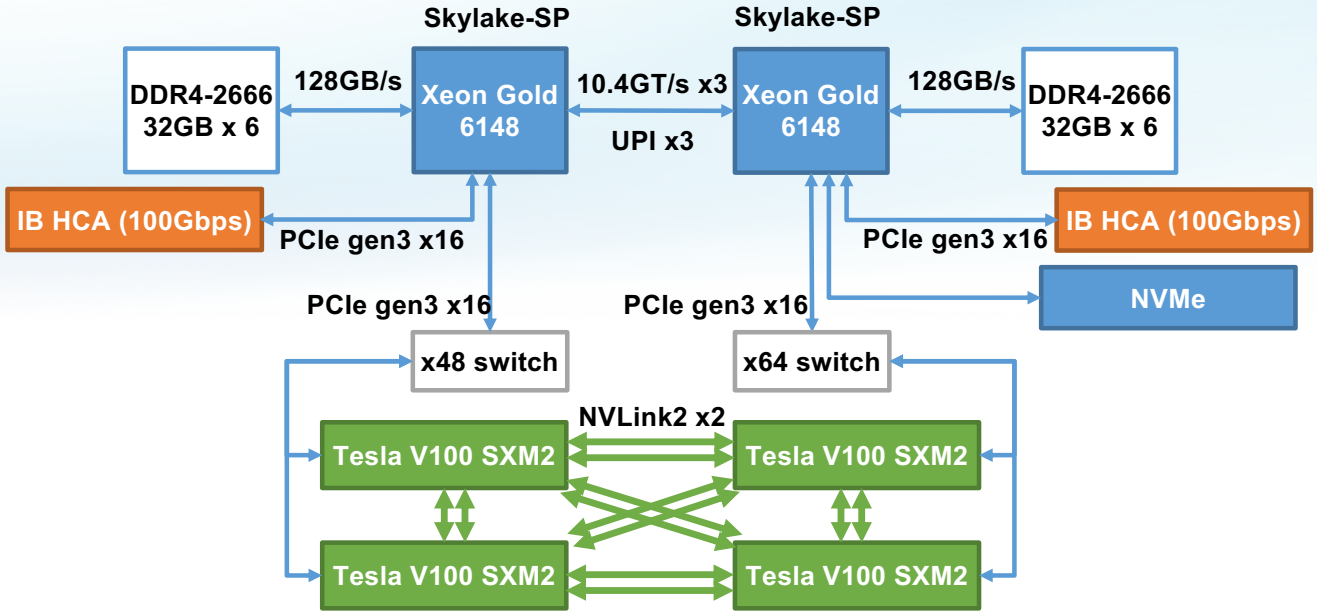

CPU blade

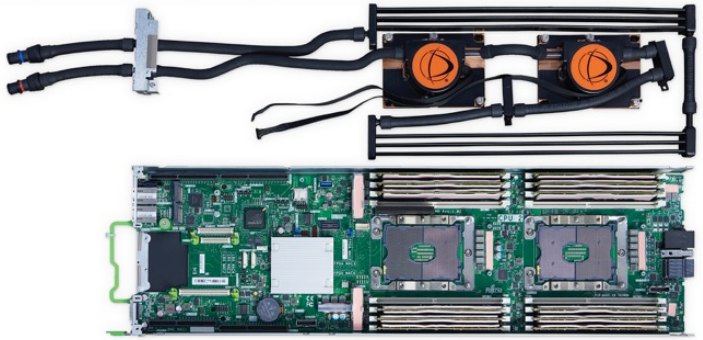

GPU blade

ABCI High-Performance Computing System

Chips (GPU, CPU)

Tesla V100

Xeon Skylake-SP

Compute Node (4GPUs, 2CPUs)

PRIMERGY CX2570 M4

Node Chassis (2 Compute Nodes)

PRIMERGY CX400 M4

Rack (17 Chassis)

System (32 Racks)

1088 Compute Nodes
4352 GPUs

GPU:	CPU:				
7.8 TFlops(FP64)	**1.53 TFlops(FP64)**	**34.2 TFlops(FP64)**	**68.5 PFlops(FP64)**	**1.16 PFlops(FP64)**	**37.2 PFlops(FP64)**
125 TFlops(FP16)	**3.07 TFlops(FP32)**	**506 TFlops(FP16)**	**1.01 PFlops(FP16)**	**17.2 PFlops(FP16)**	**0.55 EFlops(FP16)**
		～3.72 TB/s MEM BW		～131TB/s MEM BW	～4.19 PB/s MEM BW
NVIDIA Tesla V100 (16GB SMX2)		384 GiB MEM 200 Gbps NW BW 1.6TB NVMe SSD		Full Bisection BW within Rack 70kW Max	1/3 of Oversubscription BW 2.3MW

ABCI Computing Rack / Interconnect

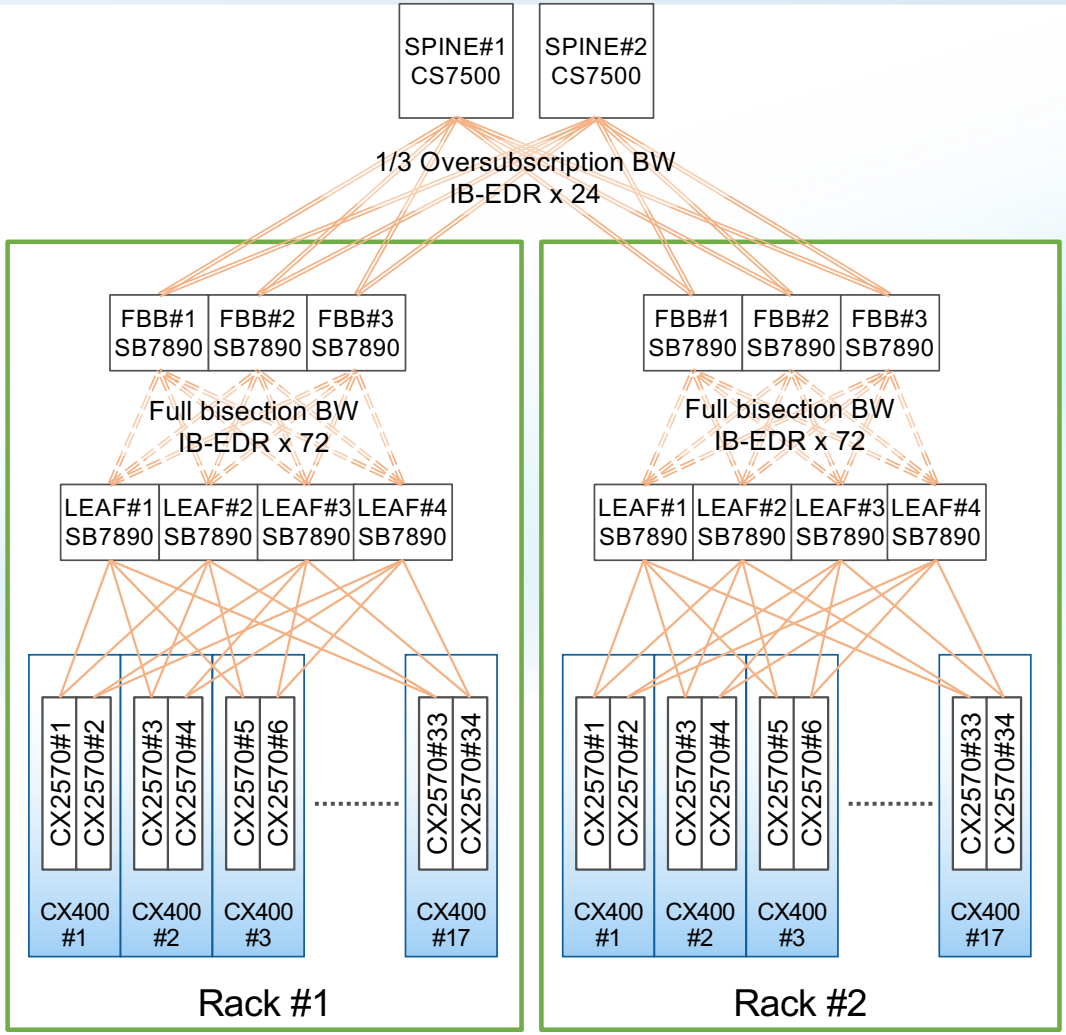

- **Dense-packaged rack: 34 nodes, 136 Tesla V100**
  - Theoretical peak performance per rack : 1.16 PFlops (FP64), 17 PFlops (FP16) c.f. Google TPU 3.0 Pod (>100PFlops/8racks)
  - Power consumption per rack: 67.33 kW
- **Interconnect**
  - Fat-tree topology
  - Intra-rack: full bisection BW
  - Inter-rack : 1/3 over-subscription (2400/6800)
  - Large-scale storage system: full bi-section BW

.....

InfiniBand EDR x4
InfiniBand EDR x6
InfiniBand EDR x1

Hierarchical Storage Layer of ABCI

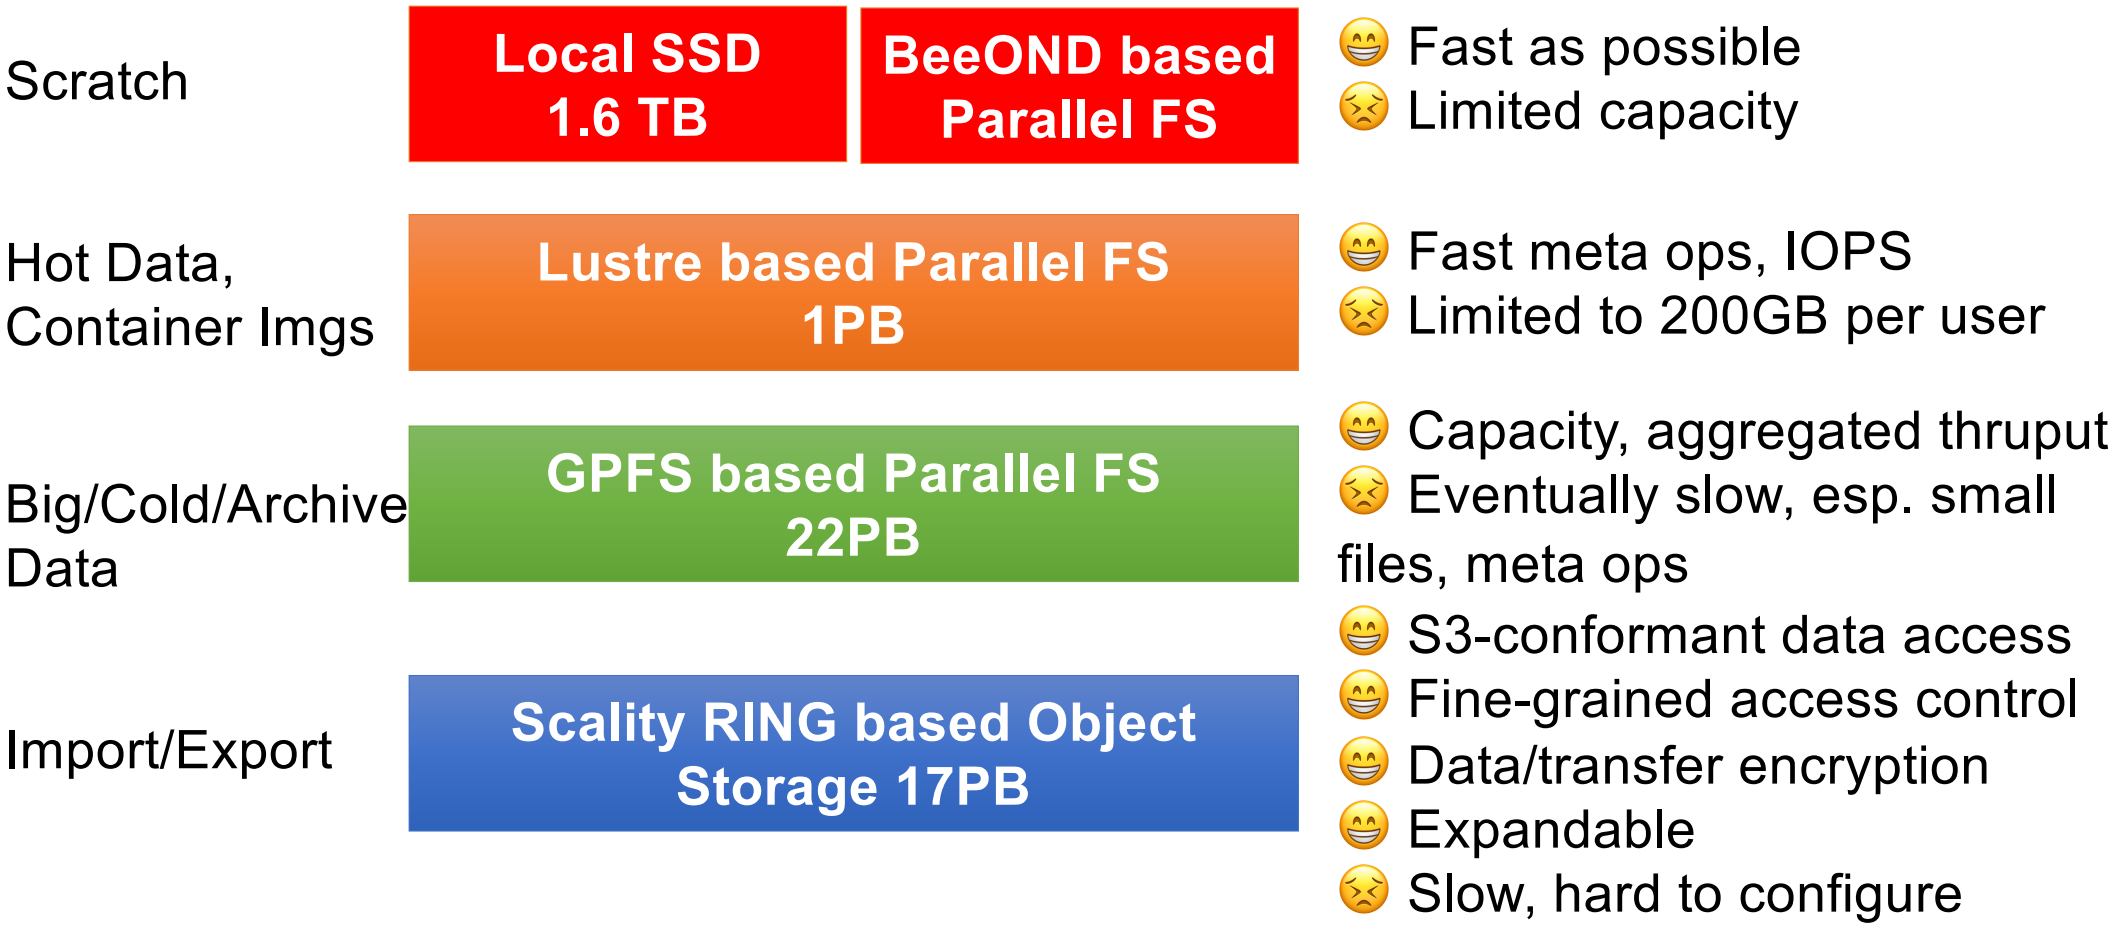

ABCI Software Stack

Software

Operating System	CentOS / RHEL 7.6
Job Scheduler	Univa Grid Engine 8.6.3
Container Engine	Docker 17.12.0 Singularity 2.6.1
MPI	Intel MPI 2018.2.199 MVAPICH2 2.3rc2, 2.3 / MVAPICH2-GDR 2.3a, 2.3rc1, 2.3, 2.3.1 OpenMPI 1.10.7, 2.1.3, 2.1.5, 2.1.6, 3.0.3, 3.1.0, 3.1.2, 3.1.3
Development tools	Intel Parallel Studio XE Cluster Edition 2017.8, 2018.2, 2018.3, 2019.3 PGI Professional Edition 17.10, 18.5, 18.10, 19.3 NVIDIA CUDA SDK 8.0, 9.0, 9.1, 9.2, 10.0 cuDNN 5.1, 6.0, 7.0, 7.1, 7.2, 7.3, 7.4, 7.5 NCCL 1.3.5, 2.1, 2.2, 2.3, 2.4 Intel MKL 2017.8, 2018.2, 2018.3, 2019.3 GCC, Python, Ruby, R, OpenJDK, Go, Perl
Deep Learning	Caffe, Caffe2, TensorFlow, Theano, Torch, PyTorch, CNTK, MXnet, Chainer, Keras, etc. *(Frameworks provided by NVIDIA GPU Cloud can also be deployed)*
Big Data Processing	Hadoop, Spark

ABCI Software Stack

Container support

- Containers enable users to instantly try the state-of-the-art software developed in AI community
- ABCI supports two container technologies
  - **Docker**, having a large user community
  - **Singularity**, recently accepted HPC community
- ABCI provides various single-node/distributed deep learning framework container images optimized to achieve high performance on ABCI
  - **NGC (NVIDIA GPU Cloud)** plays the key role for realizing this

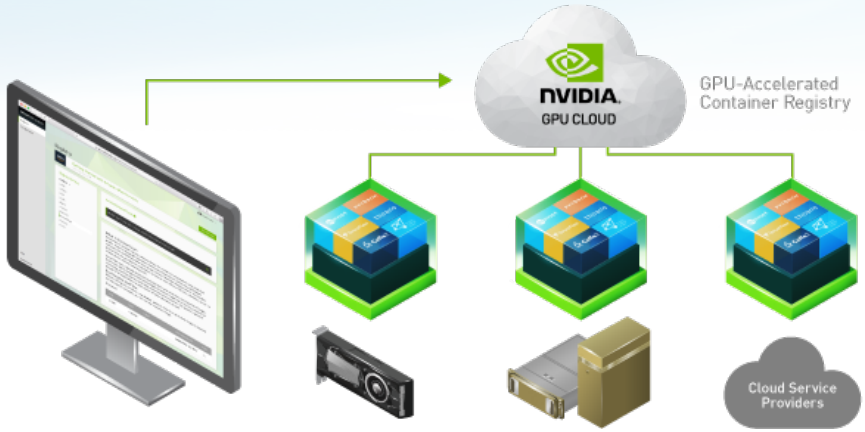

BeOND on ABCI

BeeOND on ABCI

- Build a temporal BeeGFS (v7.1.2) partition for each job
  - Users have to conduct explicit data staging between global filesystems and BeeOND filesystem
- The feature is enabled by a directive of job script

```
#!/bin/sh
#$-l rt_F=2
#$-l USE_BEEOND=1
#$-l h_rt=1:00:00
#$-cwd

Job contents
```

- Only Full node instance is supported
- Specify USE_BEEOND => mounted on /beeond

BeeGFS Default Parameters

Parameter	Value
#MDS *	1
#SS *	#Hosts
Filesystem Size	#Hosts * 1.6TB
Stripe Size *	512KB
Stripe Count *	4

* : Parameters can be changed by users

BeeOND Performance on ABCI

- Use IOR and MDTEST to measure weak scaling performance
  - 2 nodes to 32 nodes (within a rack)
- Sequential Read/Write
  - Use 40 process/node, 16 GiB/process, File/proc => 640 GiB/node
  - Transfer size is 4KB
- Random Read/Write
  - Same as Sequential Read/Write
  - 300 seconds IO
- IO Petterns
  - Each process access to different files (FPP: File Per Process)

Single SSD IOR Performance

Sequential Read	2959 MB/s
Sequential Write	1452 MB/s
Random Read	1519 MB/s
Random Write	302 MB/s

BeeOND Performance on ABCI Sequential Read/Write

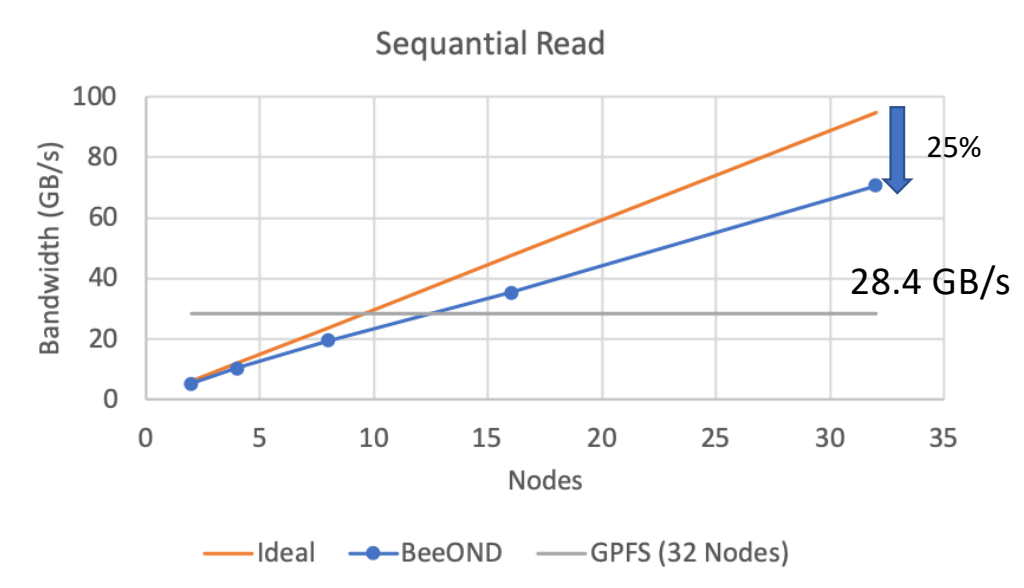

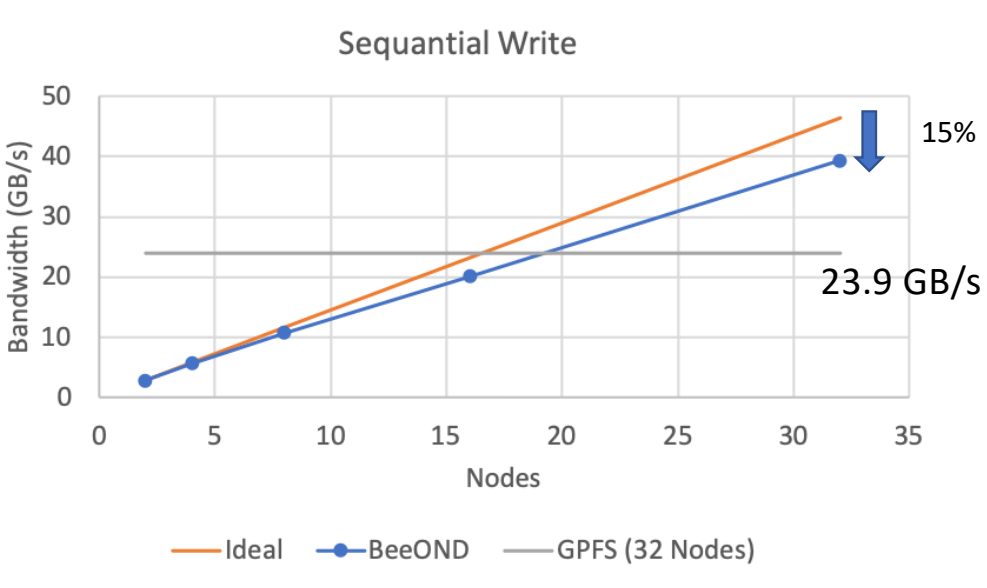

- Good weak scalability but some performance degradation against Ideal
- BeeOND outperforms GPFS when more than 32 nodes (within a rack) are used

BeeOND Performance on ABCI Random Read/Write

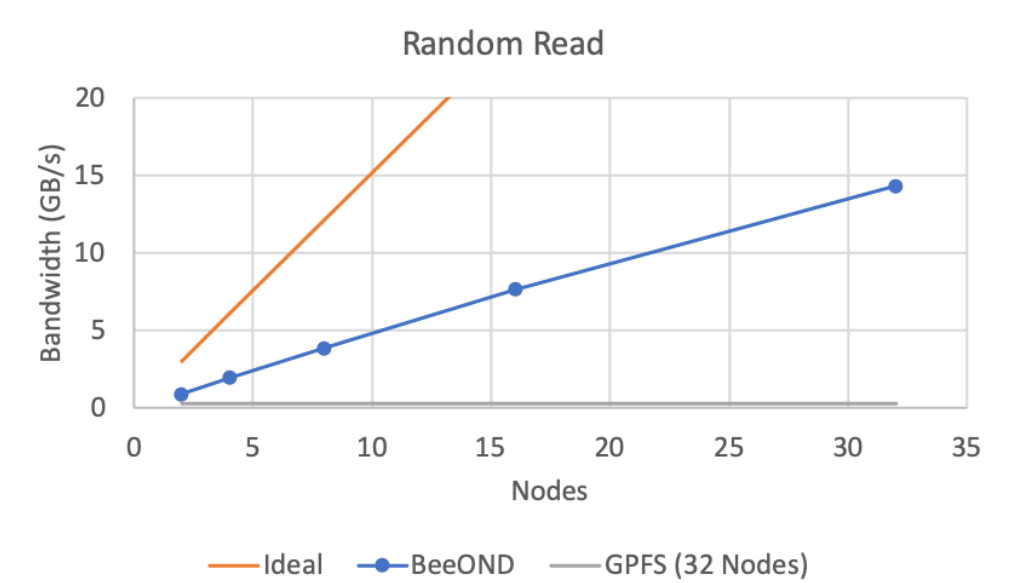

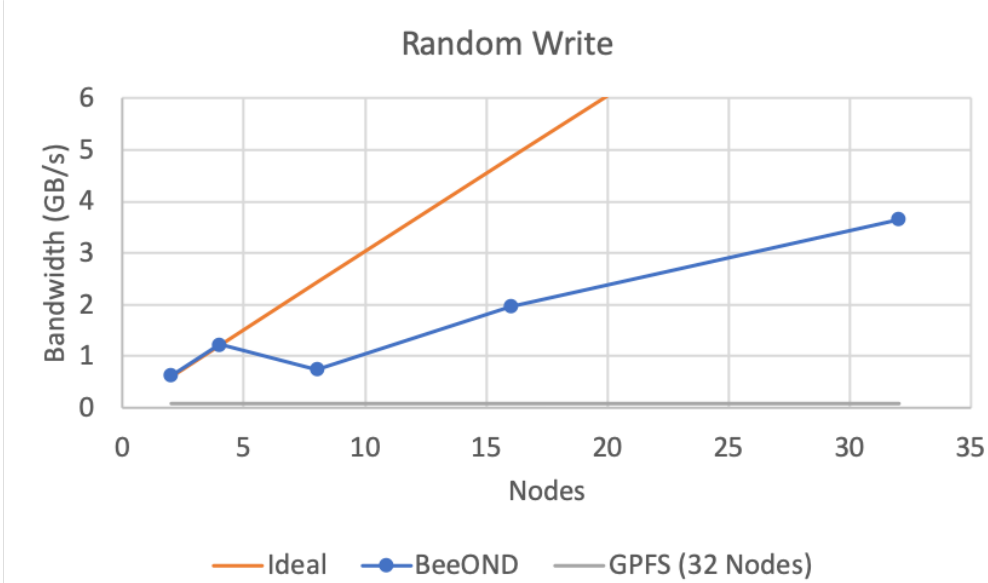

- Scalability of BeeOND is good although there are large gap between Ideal scalability
  - Performances of read and write against Ideal are 30% and 40%, respectively
- Performance of GPFS was very bad because of mismatch of transfer size (4K) and GPFS block size (4M)


ABCI
https://abci.ai/
FCU-2R-01