Introduction to KISTI 5th supercomputer

Division of Supercomputing, KISTI

Ji-Hoon Kang
jhkang@kisti.re.kr

22 Feb. 2018, Japan-Korea HPC winter school, Center for Computational Science, Univ. of Tsukuba

Background of 5th system procurement(I)

Current KISTI-4 system, Tachyon II, was procured in 2009 and served for 8 years.

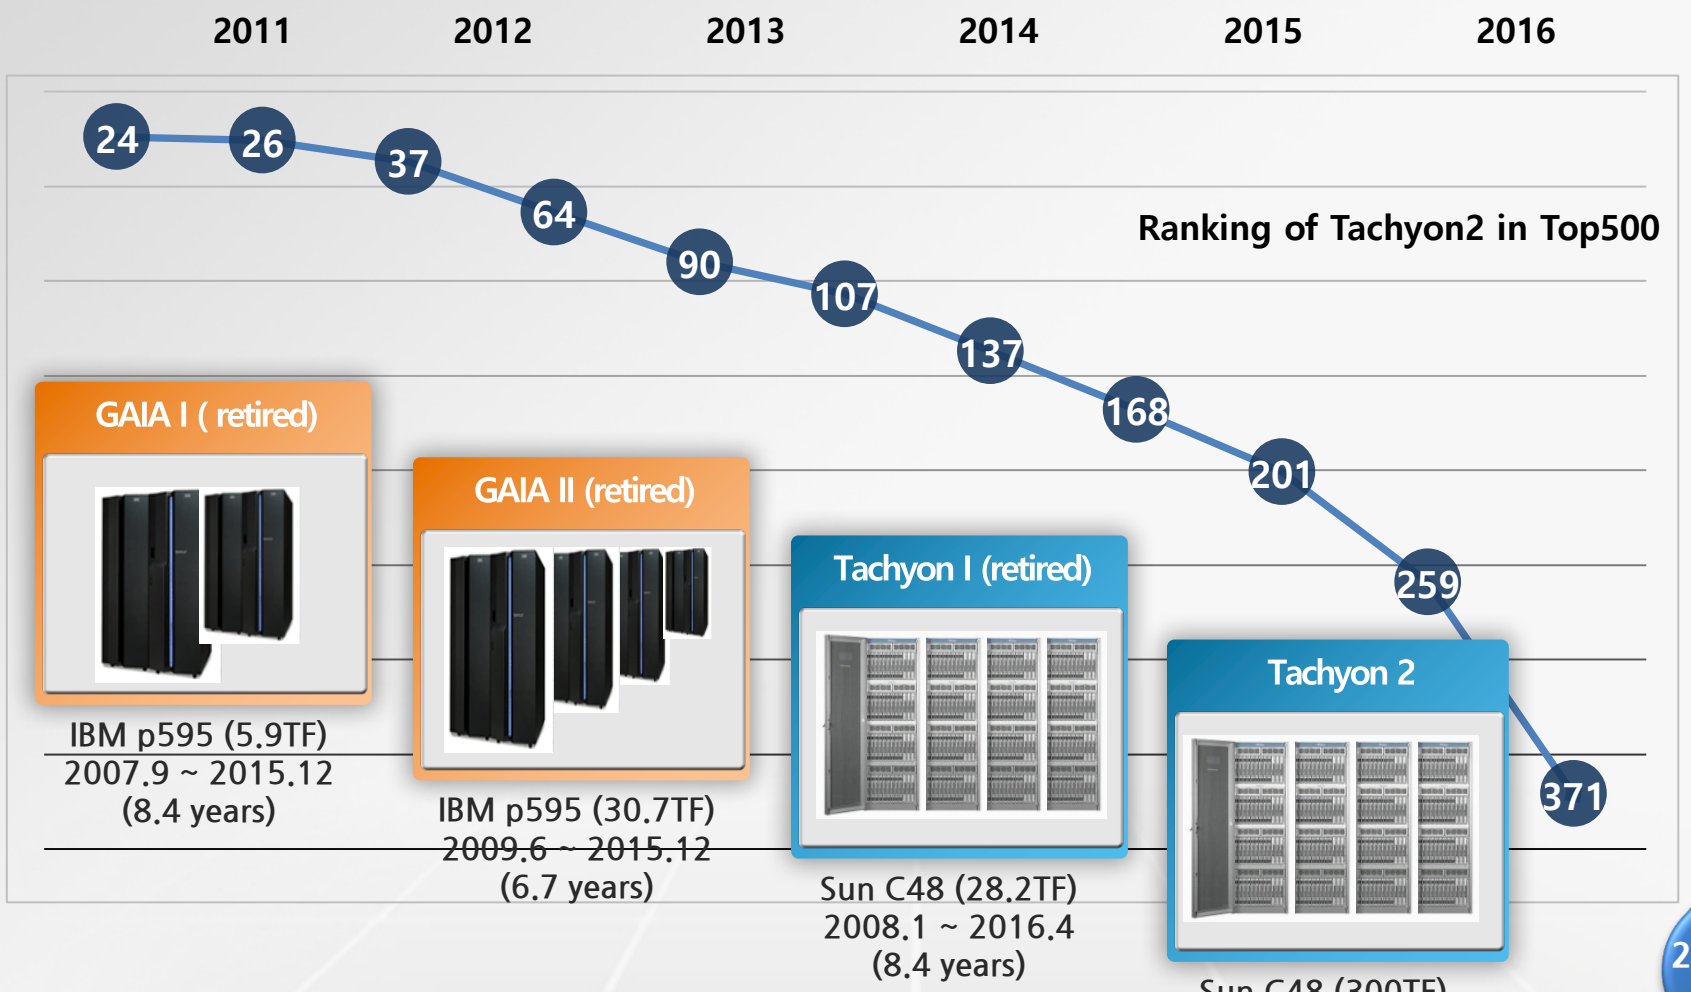

Tachyon2 had been ranked 14th in Top 500 of Nov. 2009, but was out of Top 500 rankings in 2016.

Background of 5th system procurement(II)

Operating cost beyond the residual value

Increase of replacement due to hardware failure

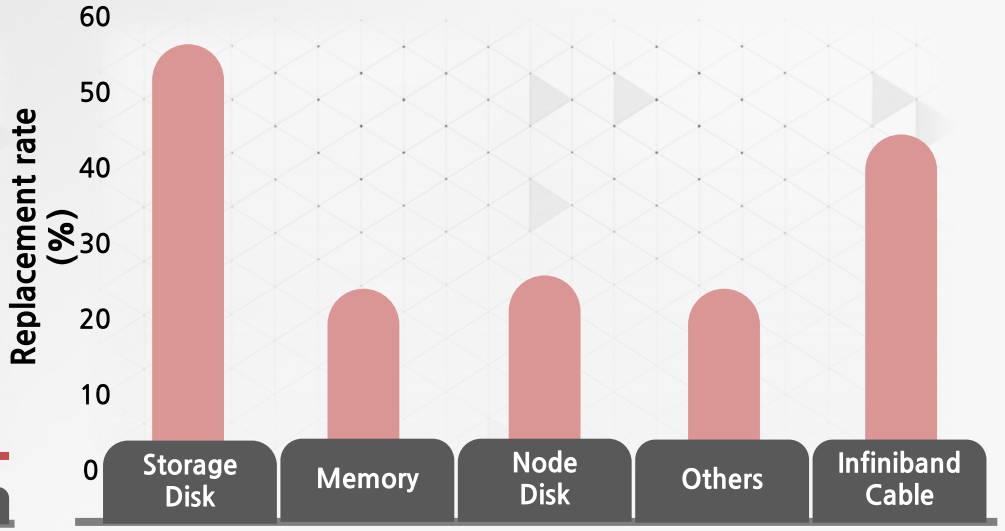

Increased waiting time due to limited resources

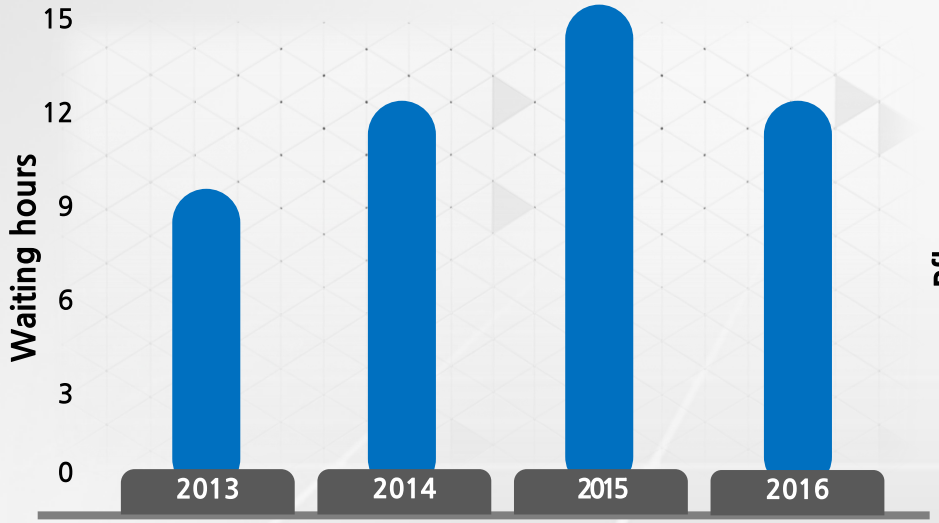

Insufficient infrastructure supply to demand

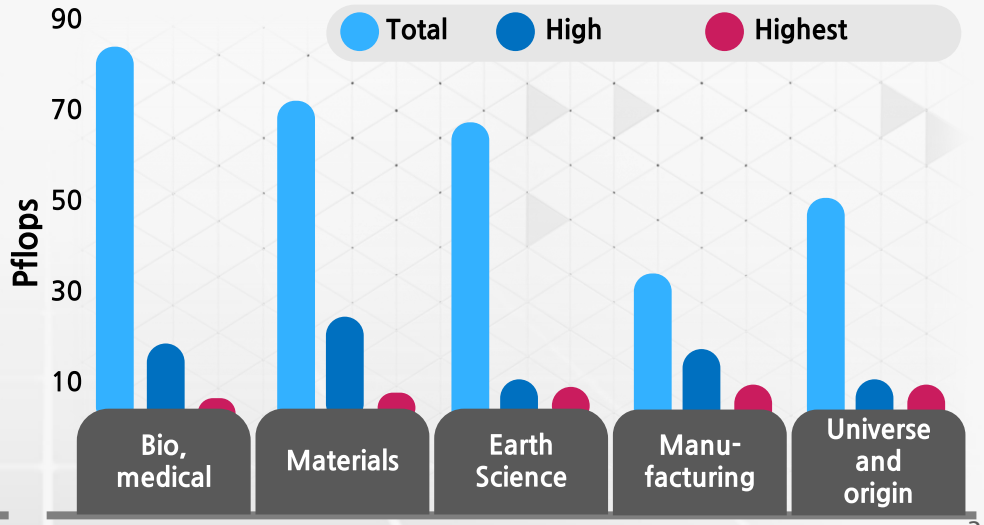

Overview of KISTI-5 system requirement

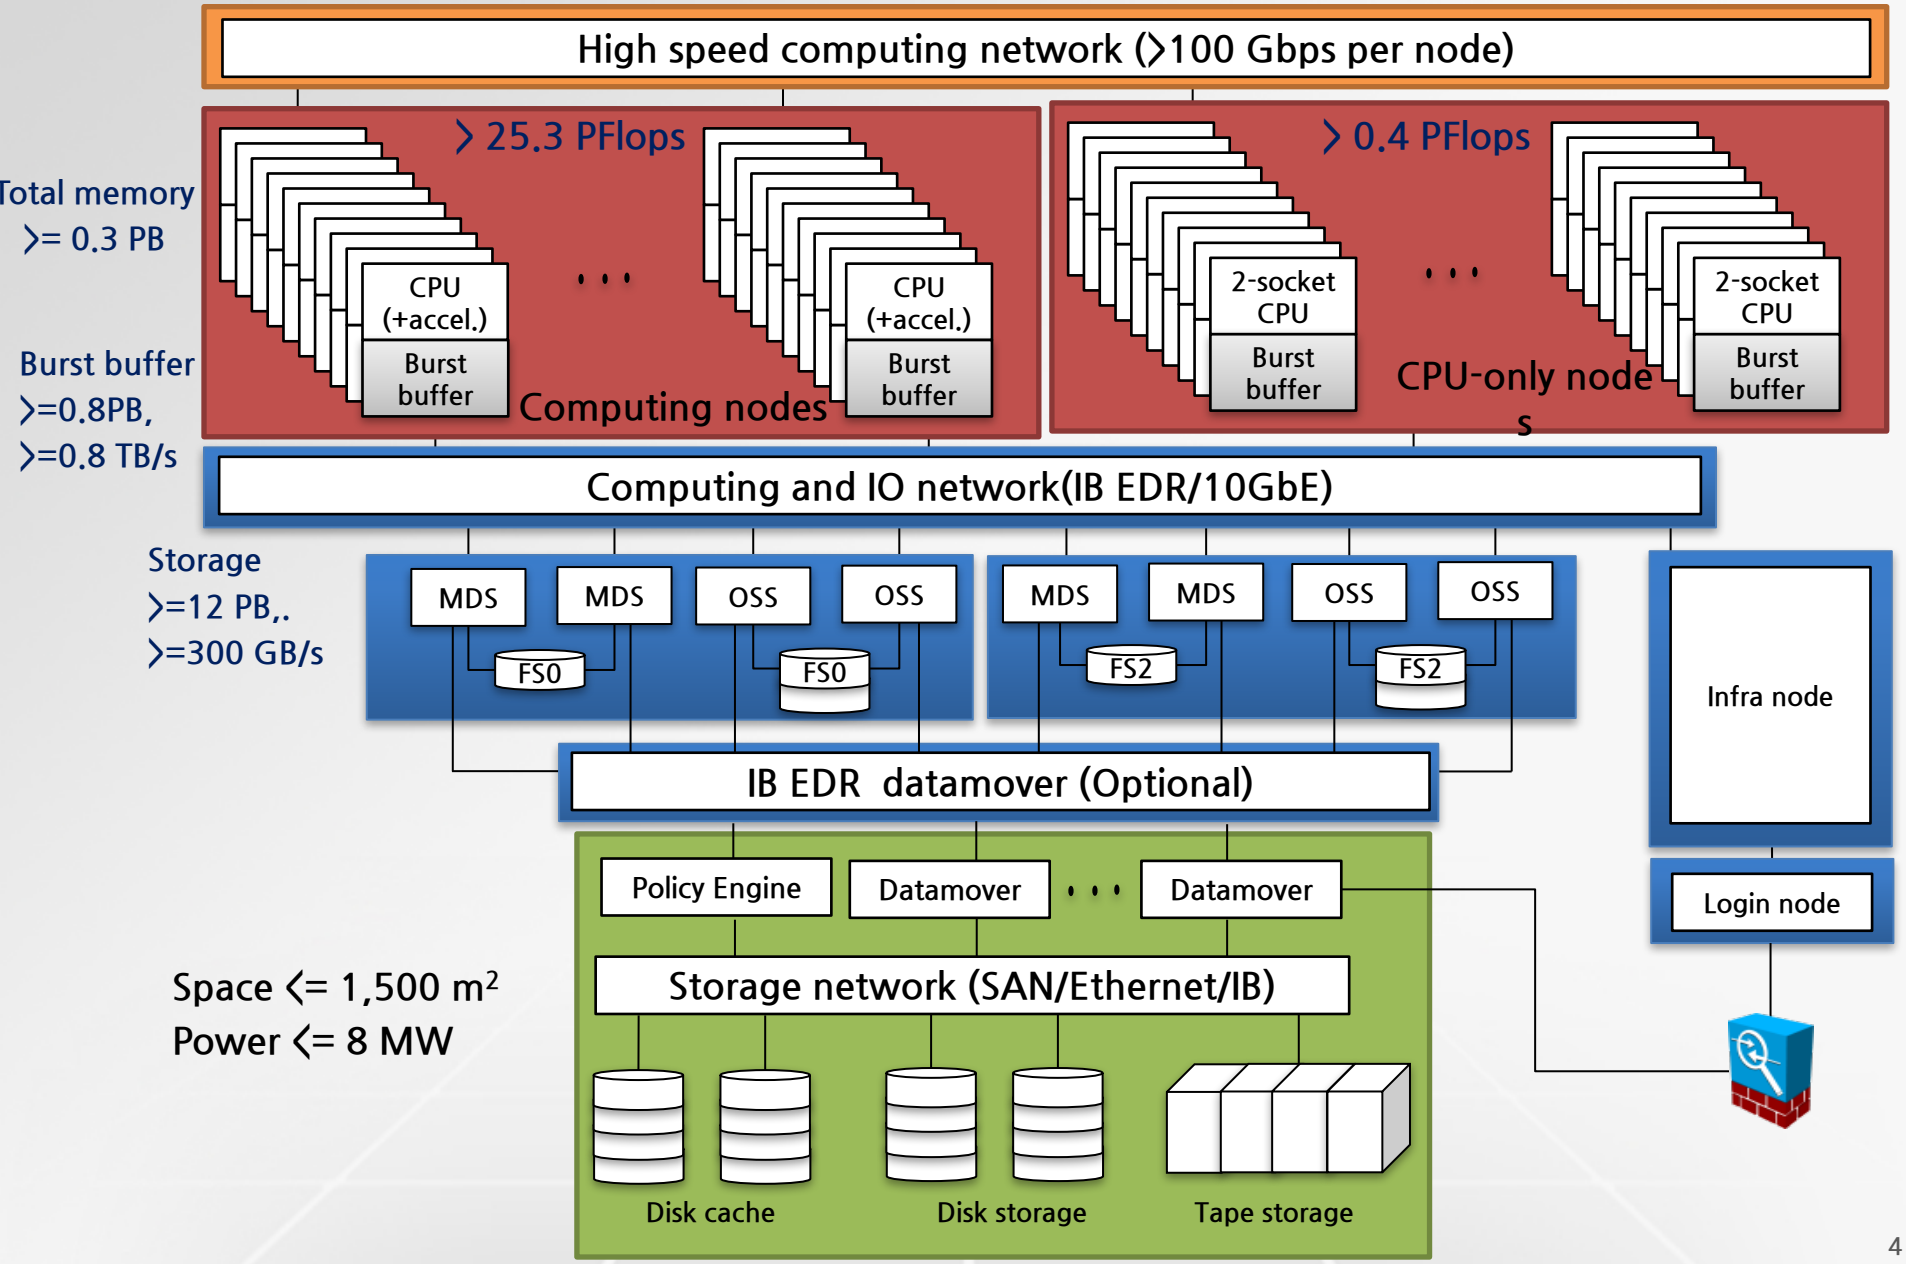

KISTI-5 procurement history

'15.06	▶ Building construction completed
'15.07.07	▶ Preliminary feasibility study approved
'16.03	▶ RFI and BMT annoucement
'16.07	▶ RFP and transfer to PPS(調達廳)
'16.10~12	▶ Bidding
'17.01	▶ Failure in bidding and RFP modification
'17.02~05	▶ Bidding
'17.06~07	▶ Cray winning bid and negotiation
'17.08	▶ Contract (49M USD)
'17.11	▶ Pilot system(16nodes) delivery completed
'17.12	▶ Main system delivery started
'18.01~02	▶ Main system construction
'18.03~04	▶ BMT verification and system inspection
'18.05~06	▶ NOA and early access

Efficiency: PUE < 1.3

Performance: 25.7PFlops

Pictures

- Chilling doors (currently under construction)
- System power-on (in Feb.)

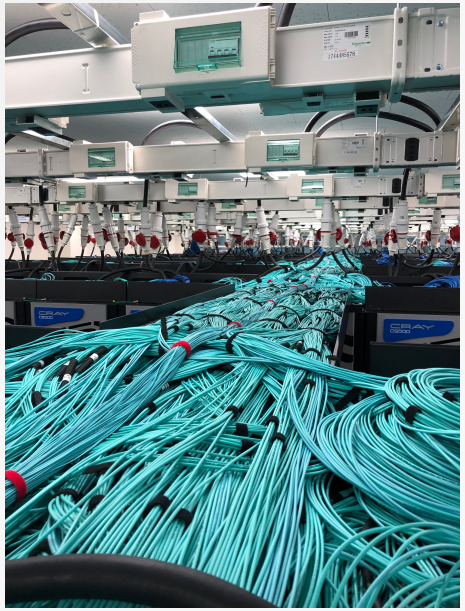

KISTI-5 supercomputer overview

Theoretical performance 25.7PF = 25.3PF CS400 w/KNL+0.4PF CS500 w/SKL

KISTI -5 project

Storage 20PB SFS@300GB/s, 10PB Archiving

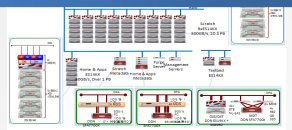

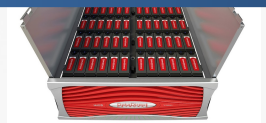

Computing nodes Cray 3112-AA000T(2U enclosure), 8,305 KNL Compute modules

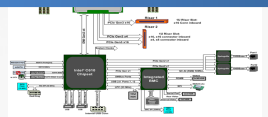

CPU-only nodes Cray 3111-BA000T(2U enclosure), 132 SKL Compute modules

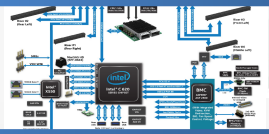

Interconnect OPA(Omni-Path Architecture), Fat-Tree, 50% Blocking

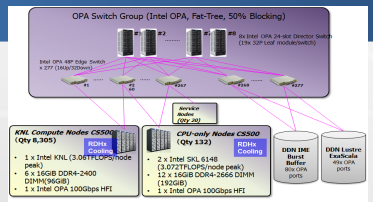

KISTI-5 Supercomputer (I)

Theoretical performance 25.7PF = 25.3PF CS400 (KNL) + 0.4PF CS500 (SKL)

Computing nodes

Cray 3112-AA000T(2U enclosure), 8,305 KNL Computing modules

- 1x Intel Xeon Phi KNL 7250 processor
- 96GB (6x 16GB) DDR4-2400 RAM
- 1x Single-port 100Gbps OPA HFI card
- 1x On-board GigE (RJ45) port

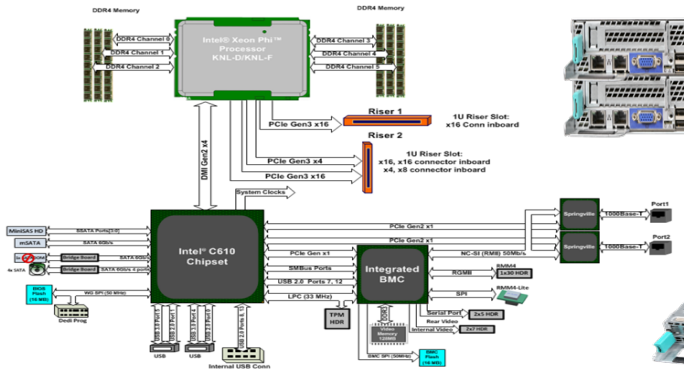

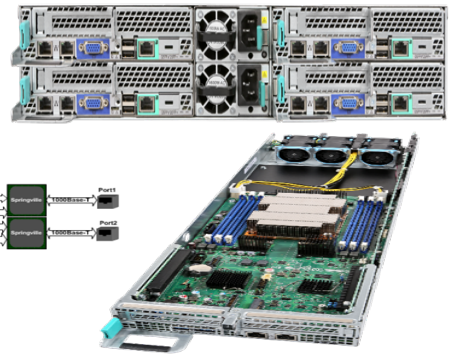

CPU-only nodes

Cray 3111-BA000T(2U enclosure), 132 Skylake Computing modules

- 2x Intel Xeon SKL 6148 processors
- 192GB (12x 16GB) DDR4-2666 RAM
- 1x Single-port 100Gbps OPA HFI card
- 1x On-board GigE (RJ45) port

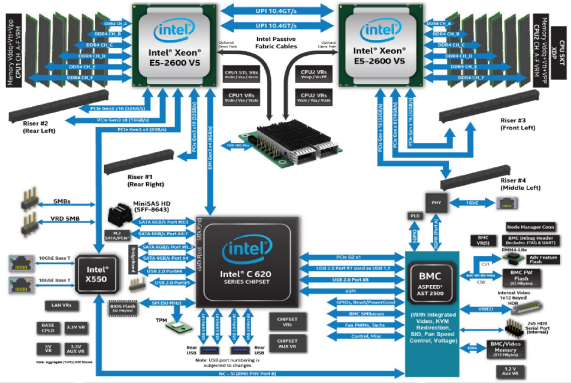

KISTI-5 Supercomputer (II)

Storage — 20PB SFS@300GB/s, 10PB Archiving

- Global scratch: 20PB, 0.3TB/s
  (DDN ES14KX 9ea, 360 x 8TB disk each)
- Home and application directory 1PB
- NVMe Burst Buffer: 0.8PB, 0.8TB/s
  (IME240 40ea 19 NVMe SSD each)
- Cray TSMSF and IBM TS4500

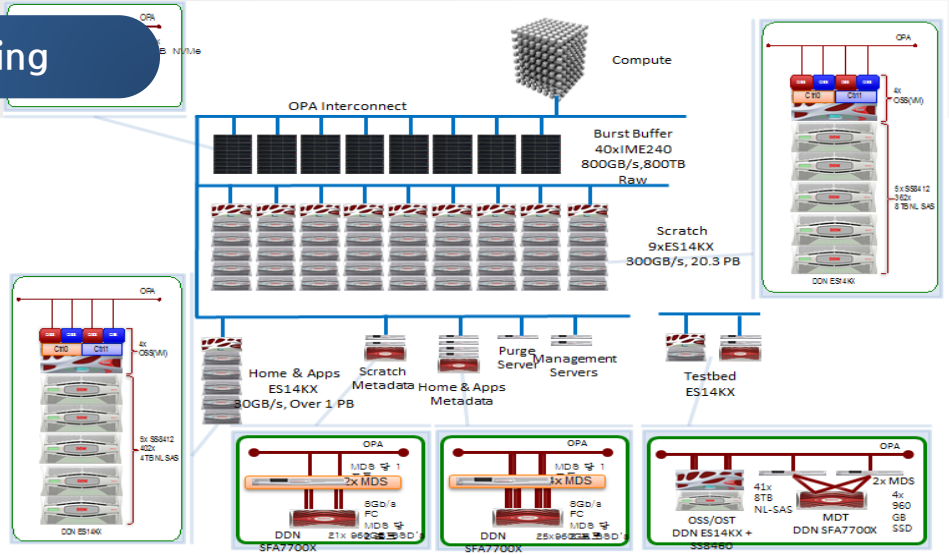

Interconnect — OPA(Omni-Path Architecture), Fat-Tree, 50% Blocking

- Intel OPA High-speed interconnect switch
  274x 48-port OPA edge switches
  8x 768-port OPA core switches
- Bandwidth: 12.3 GB/sec
- Bisectional Bandwidth : 27 TB/sec
- 10^{-16} BER(Bit Error Rate), Adaptive routing

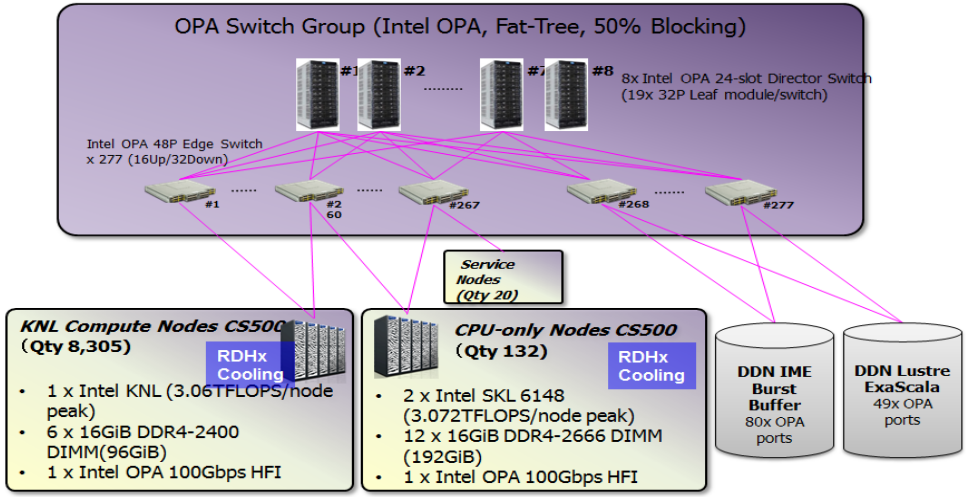

Thank you