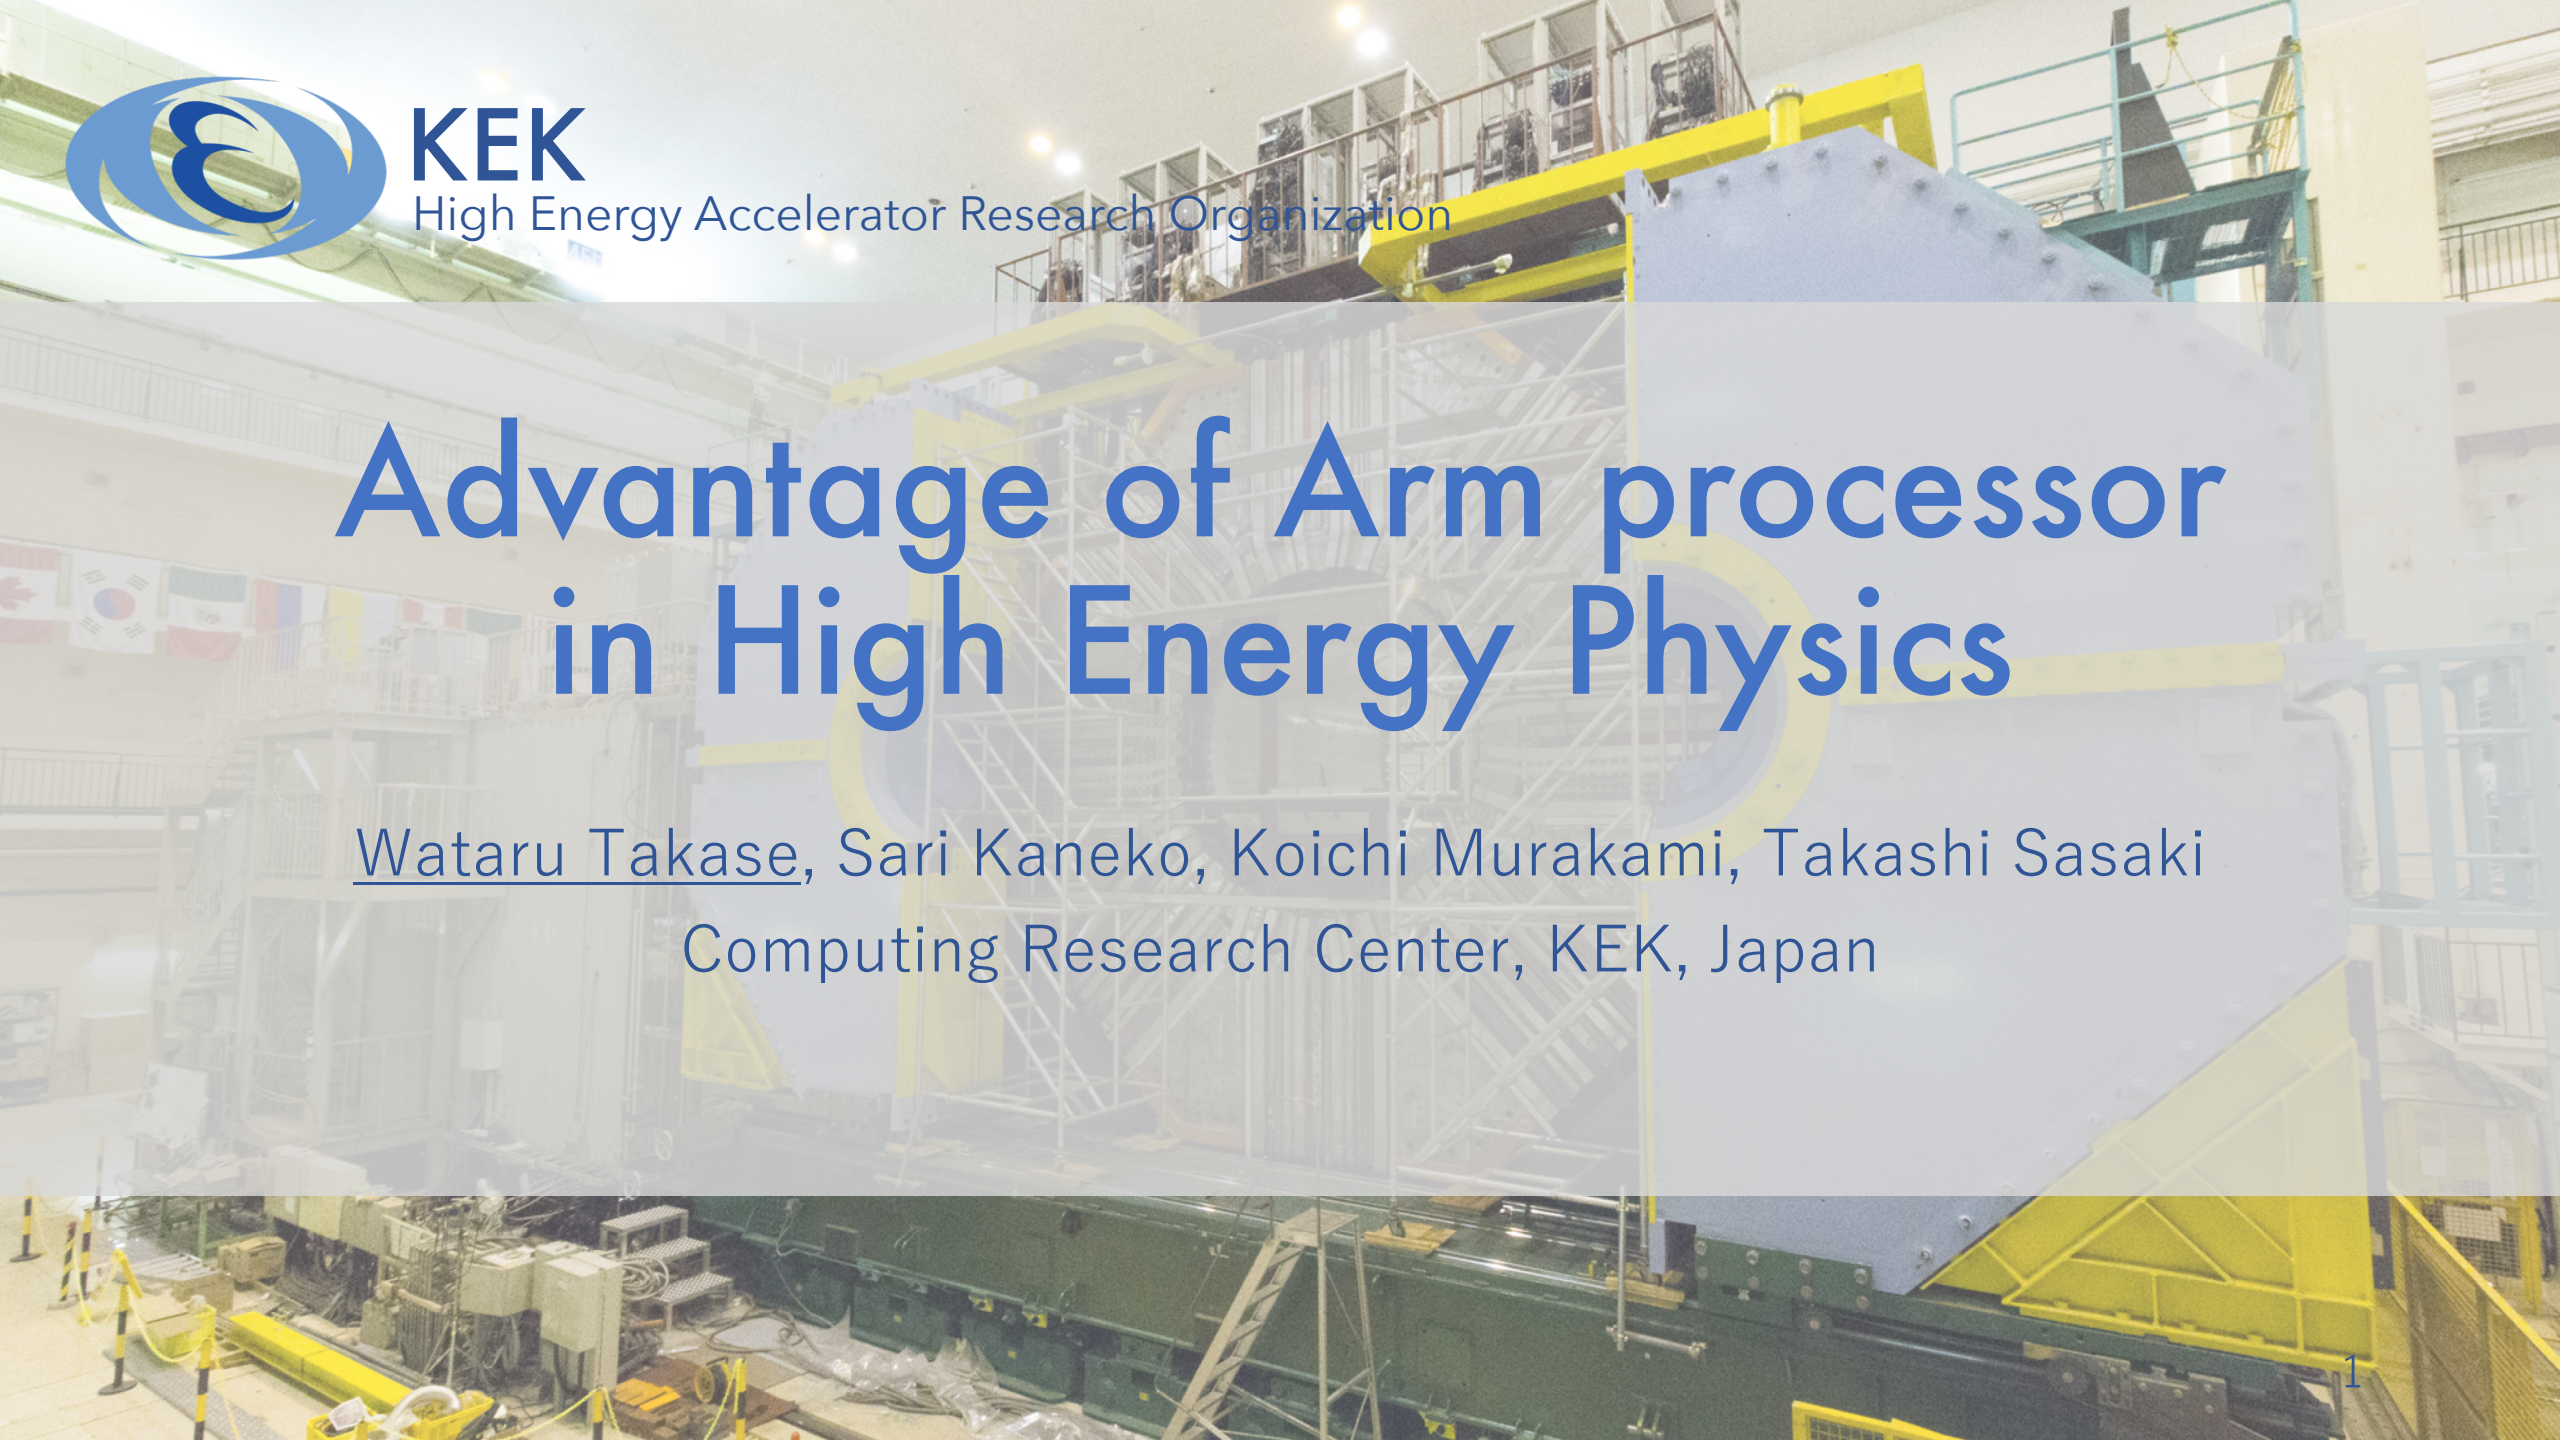

KEK
High Energy Accelerator Research Organization

Advantage of Arm processor in High Energy Physics

Wataru Takase, Sari Kaneko, Koichi Murakami, Takashi Sasaki
Computing Research Center, KEK, Japan

About KEK

- High Energy Physics research institute in Japan
- Founded in 1971
  - Will mark its 50th anniversary in 2021

- Mission:
  - Make discoveries that address the most compelling questions in a wide range of fields, including particle physics, nuclear physics, materials science, and life science.

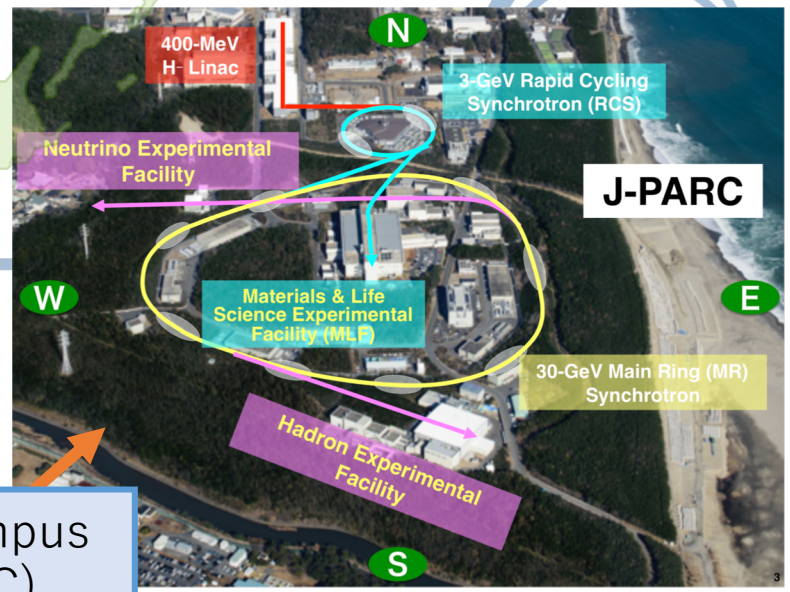

Source: Fujio NAITO

Tokai campus
(J-PARC)

Tokyo

Tsukuba campus

Particle Accelerators & Detectors

- Accelerates particles almost to the speed of light
- Smashes them at the collision point

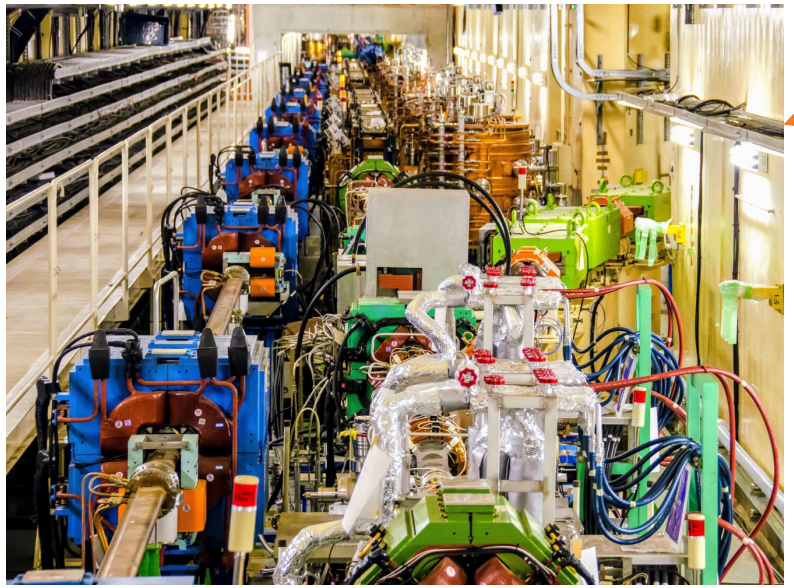

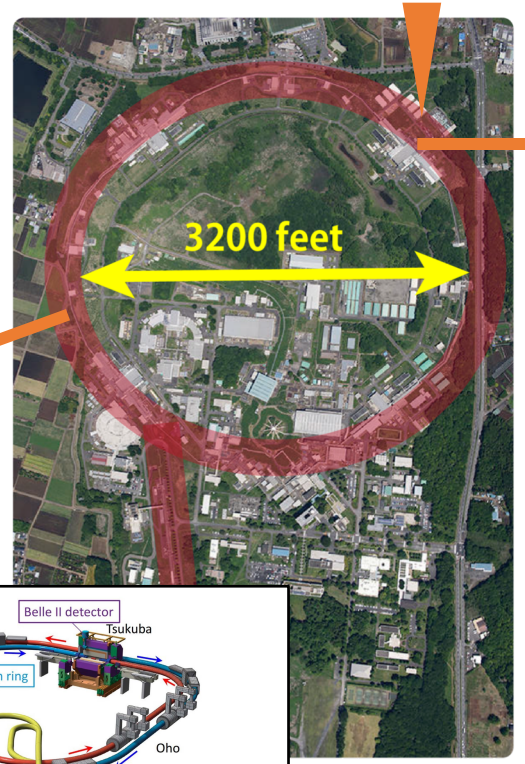

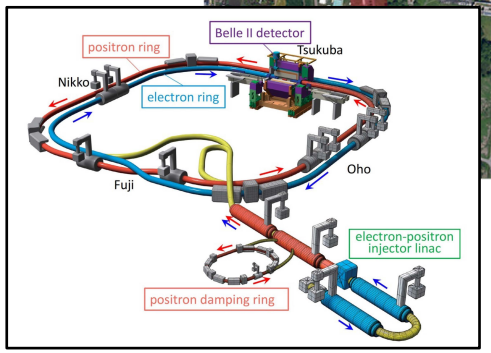

Source: SuperKEKB Design Report

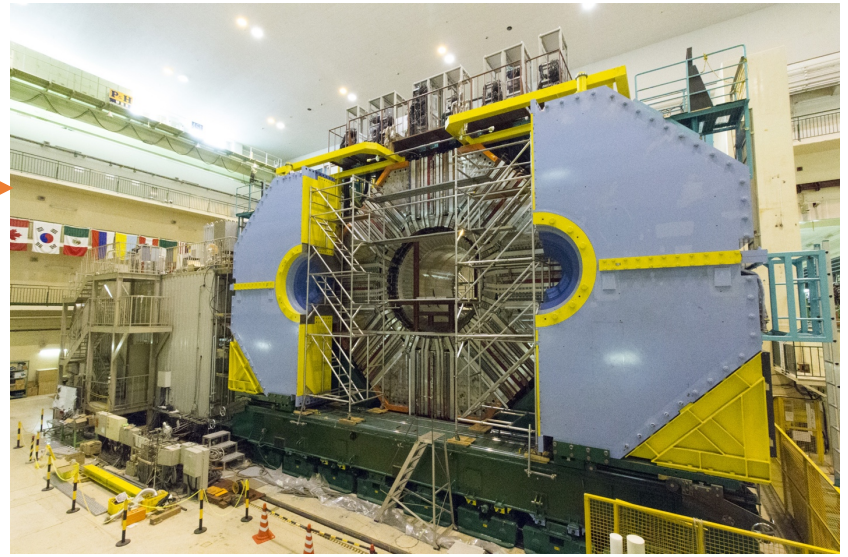

- Installed at the collision point
- Detects particle interactions

KEK
High Energy Accelerator Research Organization

KEK Central Computer System

- Linux cluster + Data storage (GPFS/HPSS)
  - **CPU: 10,024 cores (Intel Xeon E5-2693 v3)**
- In production since 2016
- Used for data analysis and numerical simulations in experiments
- More than 80% CPU utilization in average

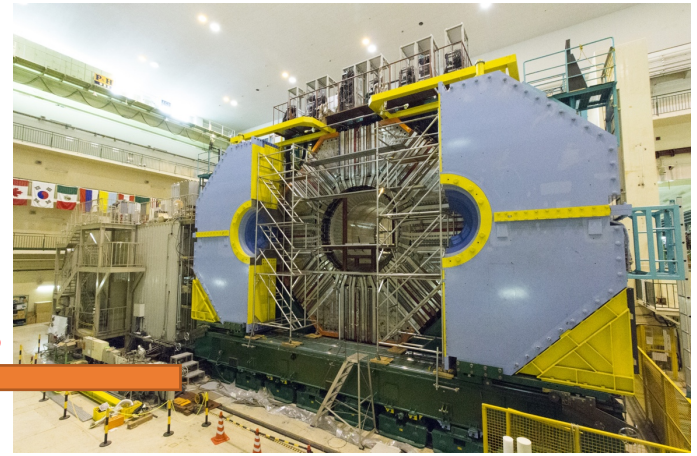

Up to 4GB/s

Detected particle interactions

Motivation

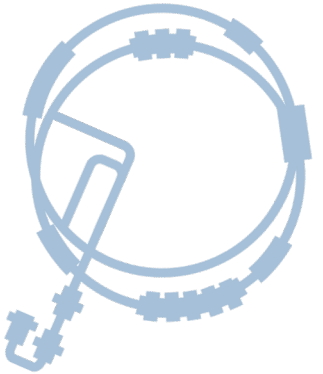

- TCO (Total Cost of Ownership) is our concern
  - Costs of electrical power grow rapidly after the 3.11 earthquake in 2011
  - Cooling facility costs also impact a lot
- We seek the most efficient platform for detector simulation
  - More than 50% of computing resources are used
- Arm processors can be a solution?

First Test Environment

- Marvell ThunderX2 Sabre platform
- CPU: Marvell ThunderX2 CN9980 x2
  - 32 cores/CPU
- Memory: 256 GB
  - 16GB DDR4 2666 DIMM x16
- Chipset configuration in BIOS
  - Turbo/CPPC Mode: Autonomous - CPPC on
  - Symmetric Multi-Thread: 1, 2, 4
    - Measured with changing this parameter
- Kernel: 4.14.0-115.el7a.0.1.aarch64
- OS: CentOS 7.6.1810 (AltArch)

Geant4: Used for Detector Simulation

- A software toolkit to simulate the interaction of particles with matter by Mote Carlo method.
  - Integer operations are more important than floating operations.
  - Loops an event which shoots primary particle and simulates particle interactions.

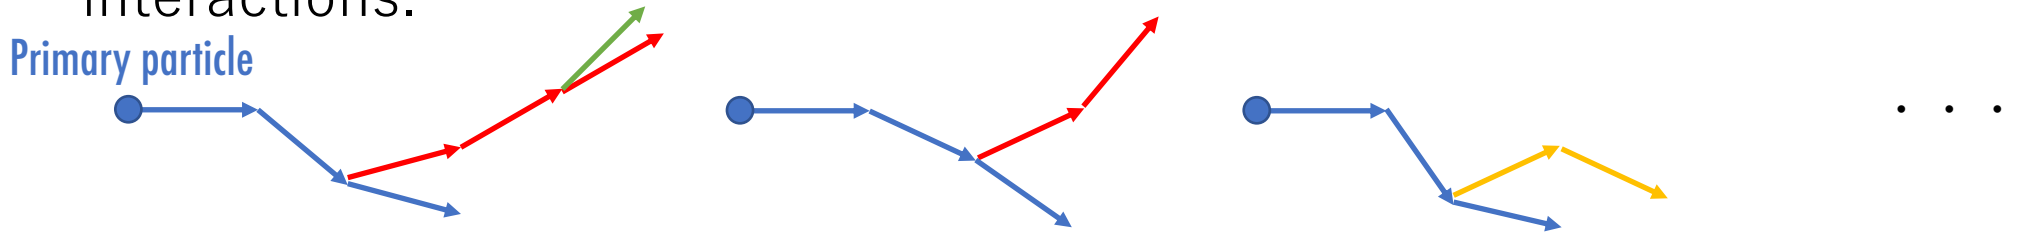

- Widely used in particle physics, space science, medicine, etc…

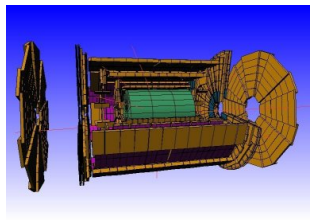
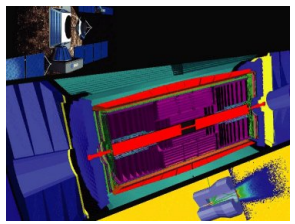
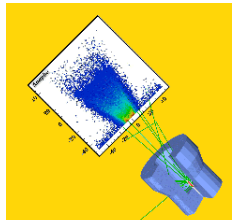
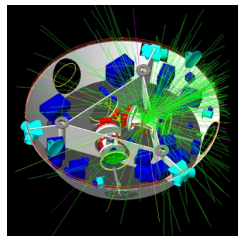

G4Bench: Geant4-based Benchmark Tool

- Provides three types of Geant4-based Monte Carlo simulation workloads:
  - Vgeo: Simulation of voxel geometry water phantom for medical application.
  - Ecal: Simulation of Electromagnetic calorimeter array.
  - Hcal: Simulation of hadron sandwich calorimeter of Lead and Scintillator.
- We executed the applications and measured performance by changing the number of threads.

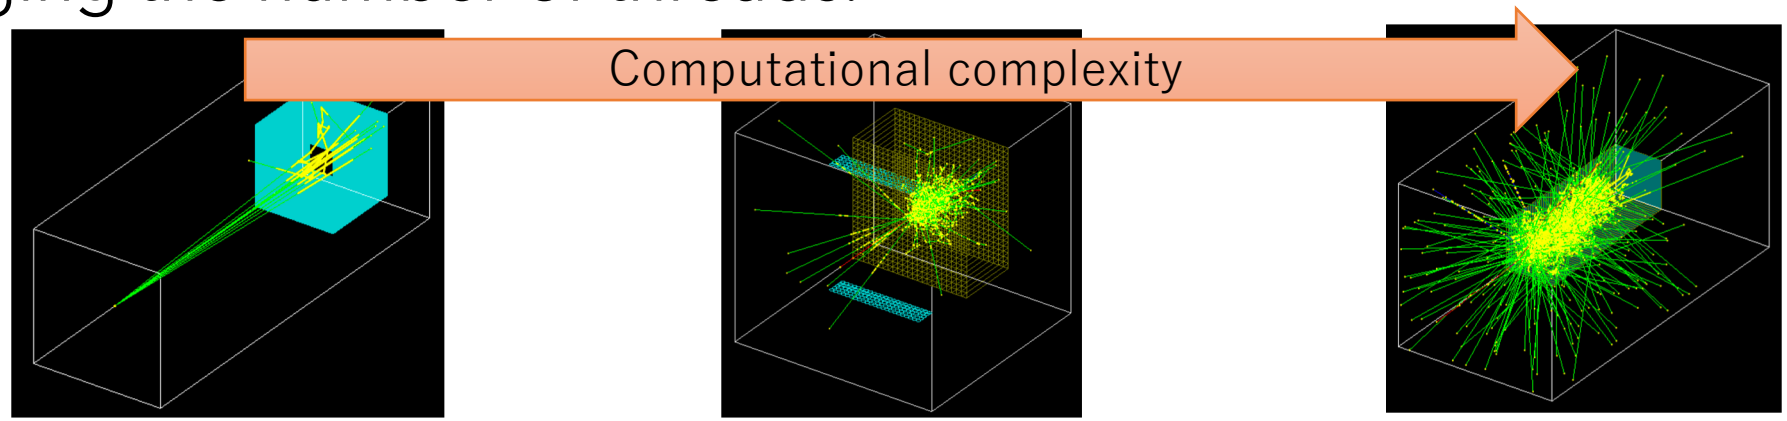

Vgeo: light-weight workload

Ecal: Middle-weight workload

Hcal: Heavy-weight workload

Results: Comparison of SMT-Mode Performances

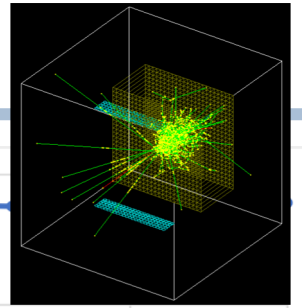

Middle-weight workload

- 1SMT mode is scaled ideally.
- Total throughputs of 2SMT and 4SMT modes are better than 1SMT's.
  - 2SMT: 1.43 times better
  - 4SMT: 1.86 times better
- We haven't found out the reason of the performance degradations yet.

Evaluation of Power Consumption

- We performed G4Bench on the ThunderX2 server and the Intel E5 family server and measured the power consumptions

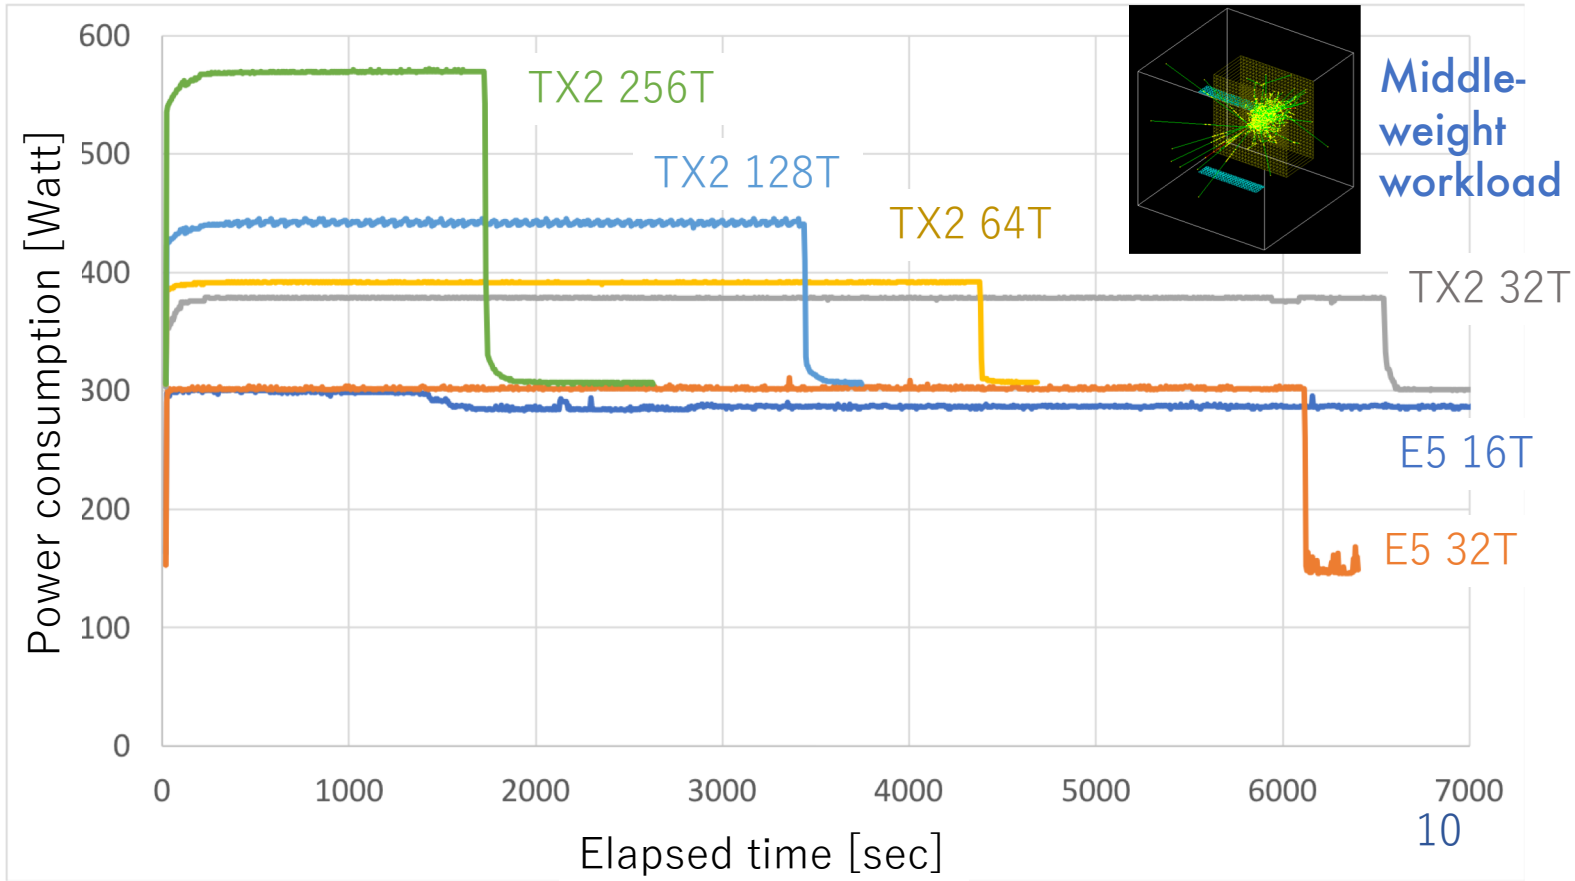

Evaluation of Power Consumption

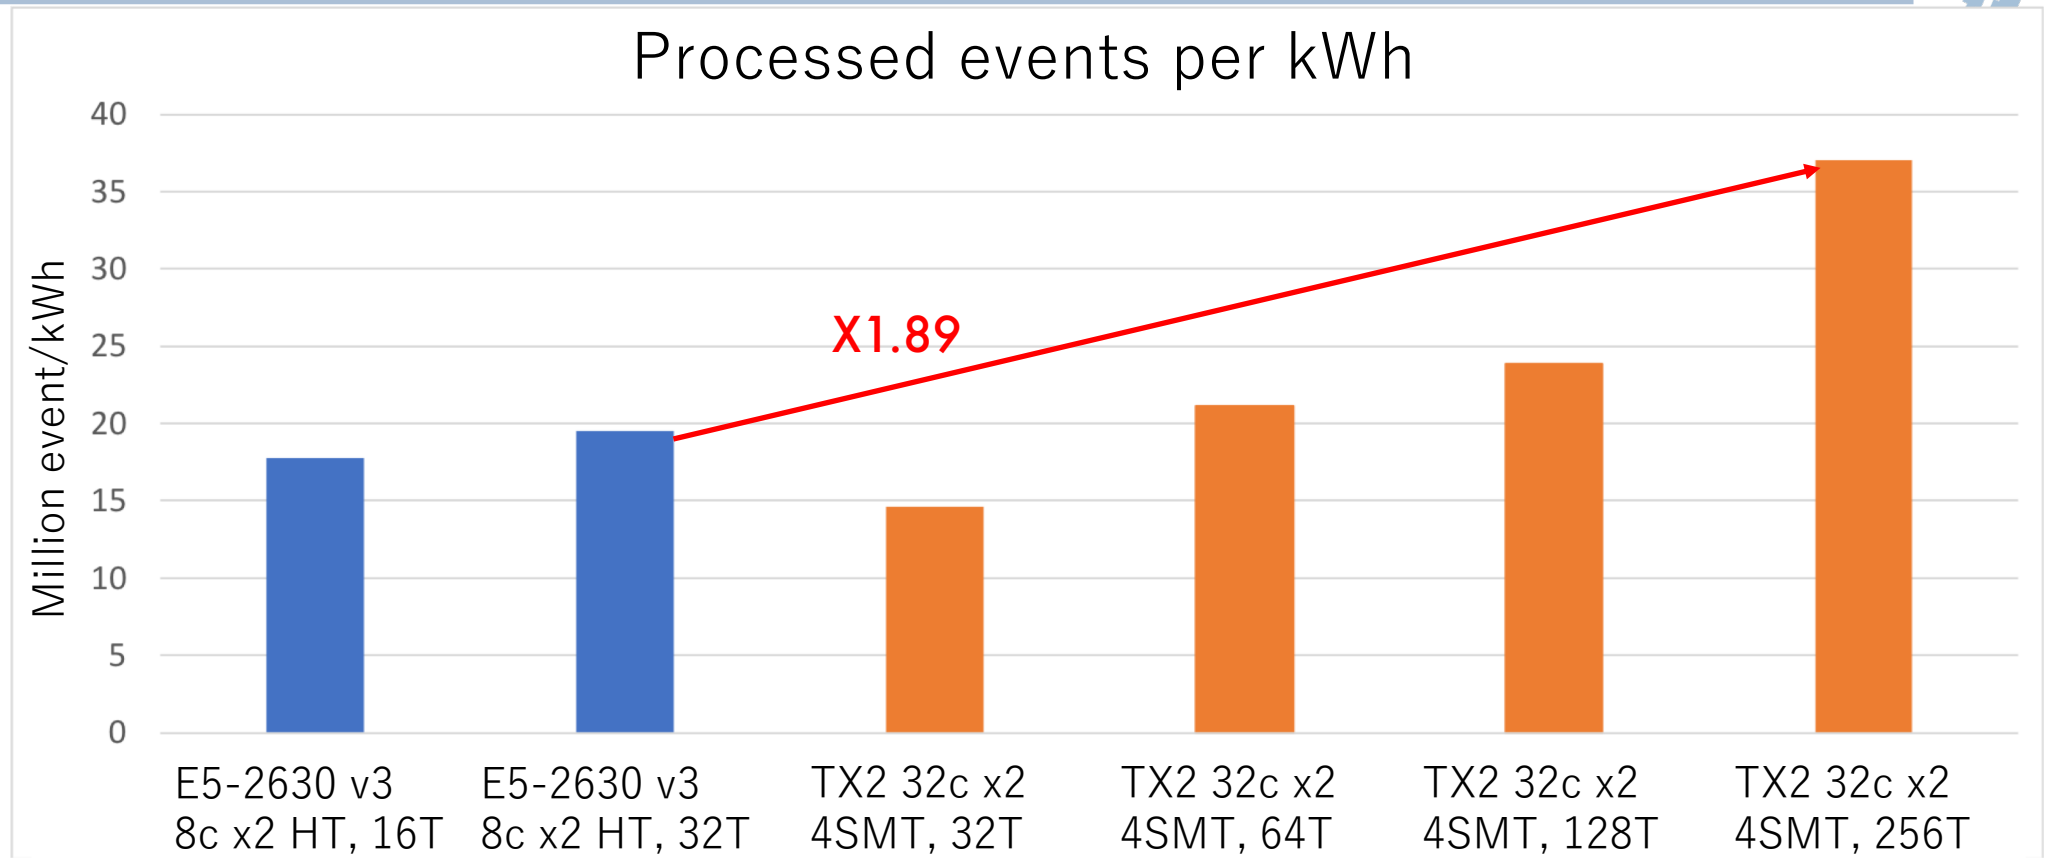

- ThunderX2 got better results than Xeon E5 in terms of performance and power efficiency.

Bought HPE Apollo 70 for Further Tests

- The G4Bench performance of Apollo 70 is equivalent to the Marvell test server
- Installed an HPE AR64z 2U Arm node on an HPE Apollo z70 chassis
- CPU: Marvell ThunderX2 CN9980 x2
  - 32 cores/CPU
- Memory: 256 GB
  - 16GB DDR4 2666 DIMM x16
- Chipset configuration in BIOS
  - Turbo/CPPC Mode: Autonomous – Turbo
  - Symmetric Multi-Thread: 1, 2, 4
- Kernel: 4.18.0-80.7.2.el7.aarch64
- OS: CentOS 7.7.1908 (AltArch)

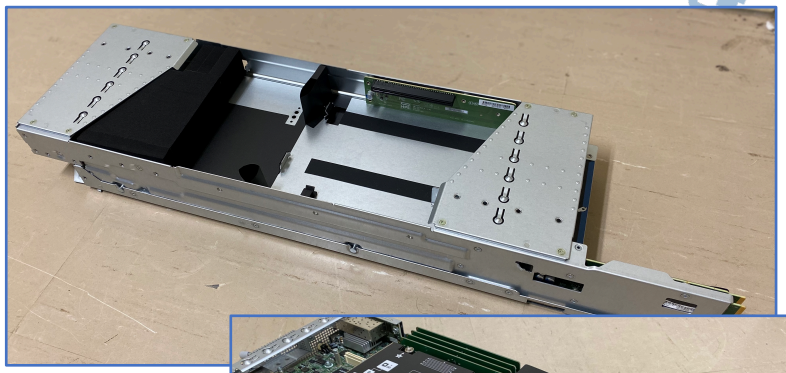

ThunderX2 vs. Intel Xeon Gold 6148: Performance comparison

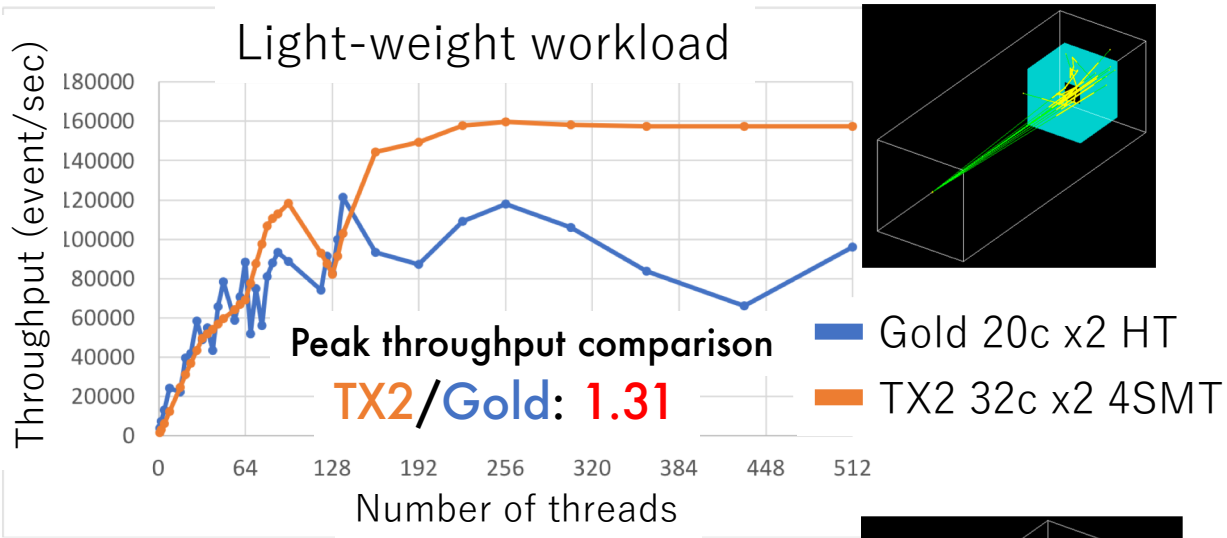

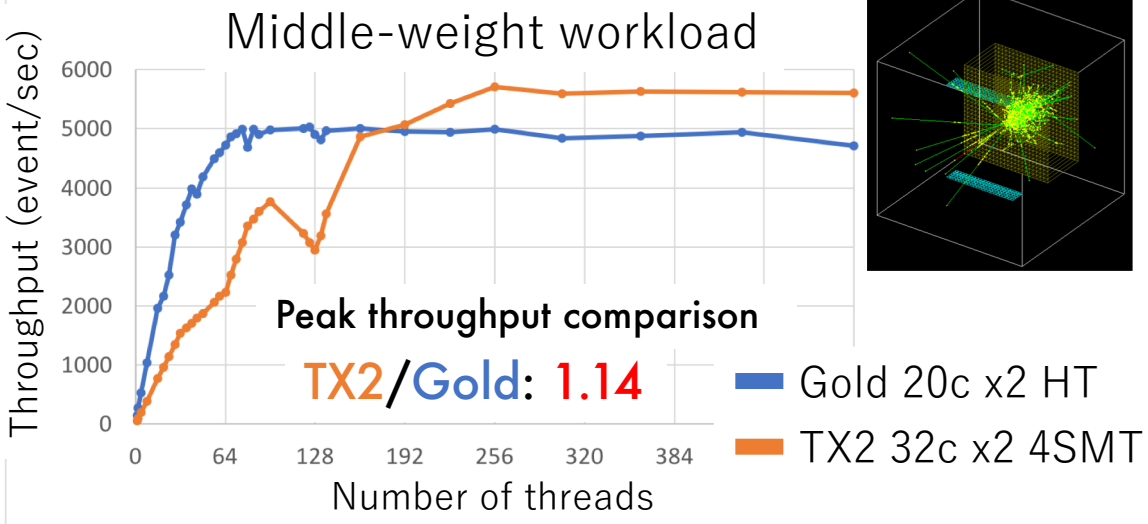

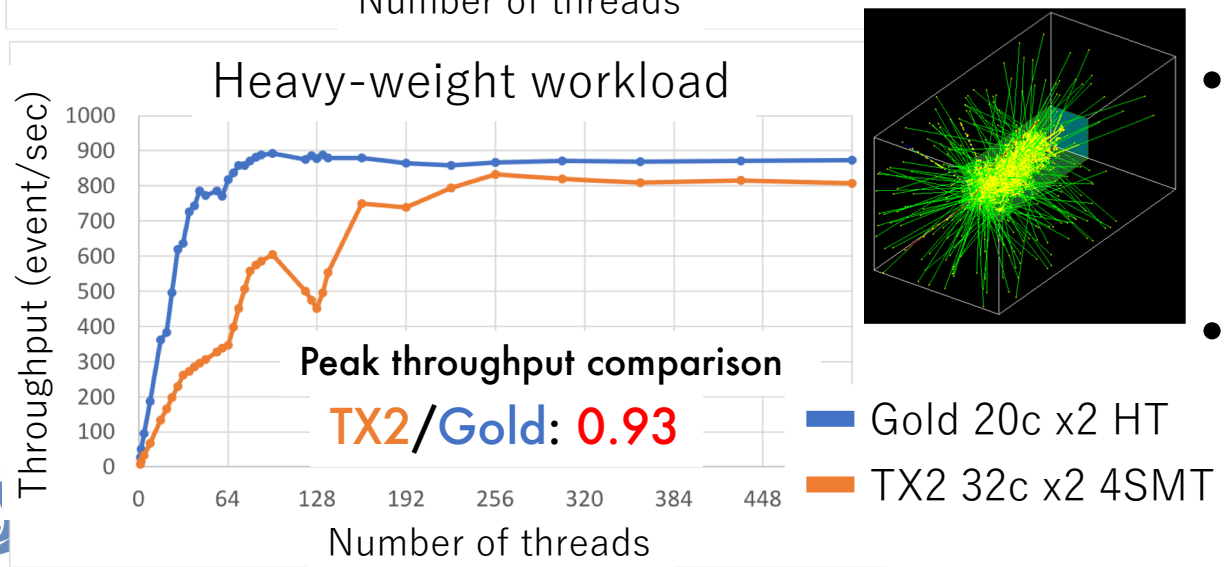

- Xeon Gold and ThunderX2 performances are scaled up to 80 and 256 threads respectively.
- On light and middle -weight workload tests, total throughput of ThunderX2 exceeds the Xeon Gold.

ThunderX2 vs. Intel Xeon Gold 6148: Events/kWh comparison

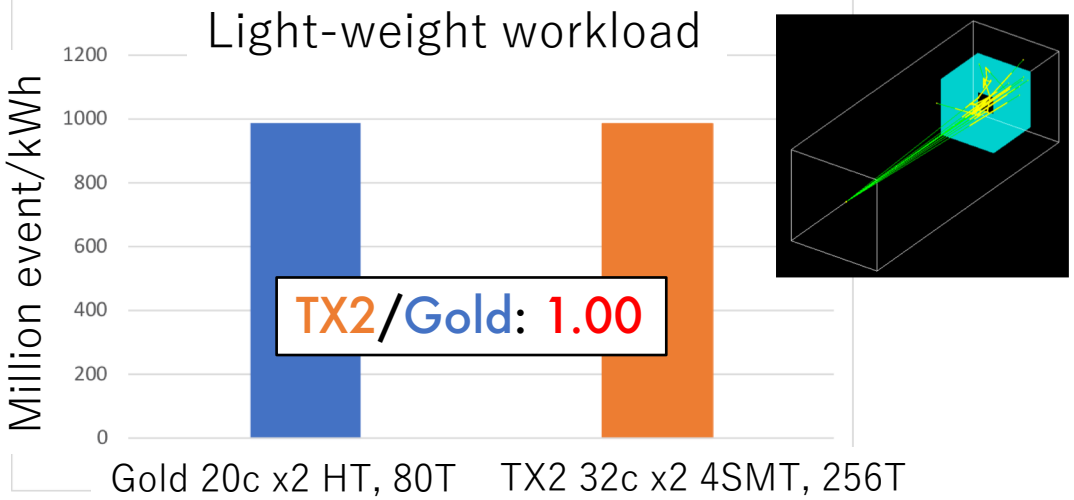

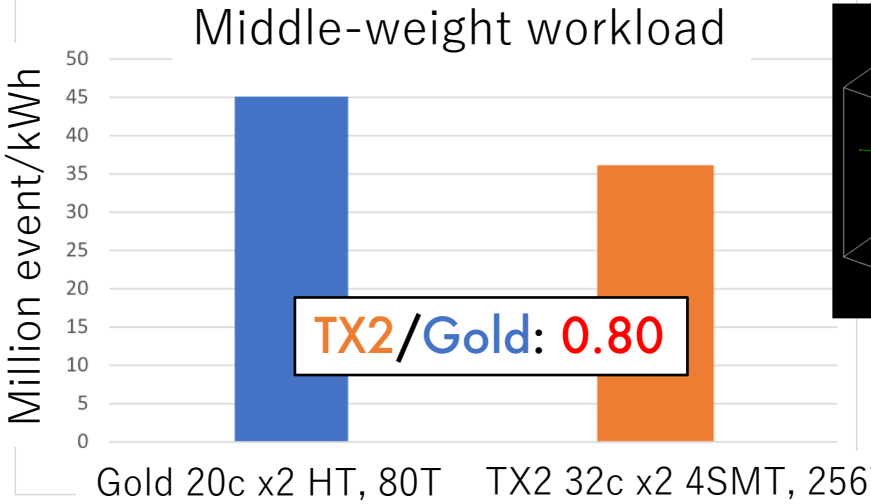

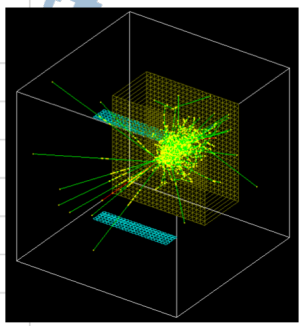

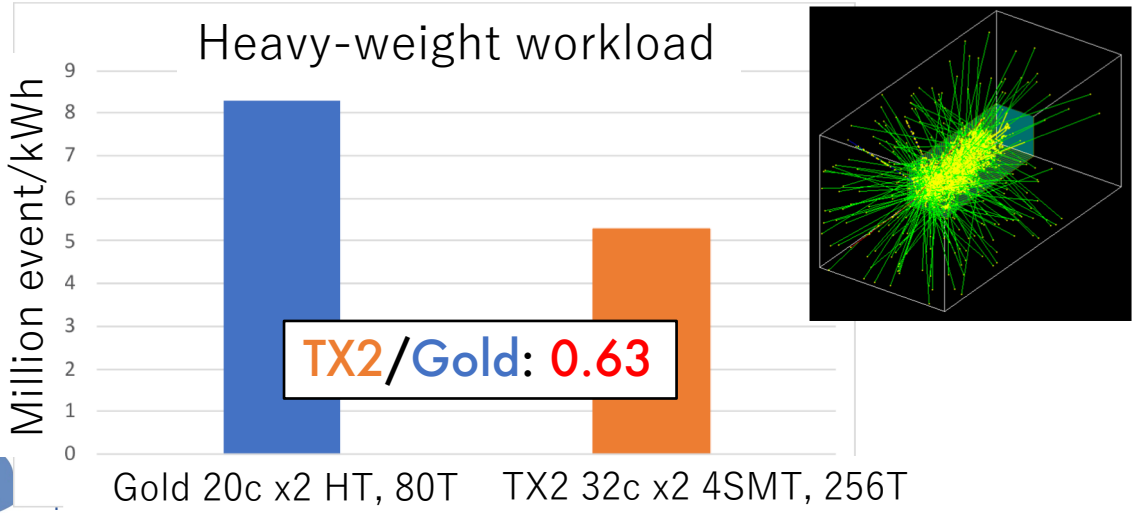

- Xeon Gold got better results.
- ThunderX2 server consumes much power than Xeon Gold server.

Who Consumes Much Power?

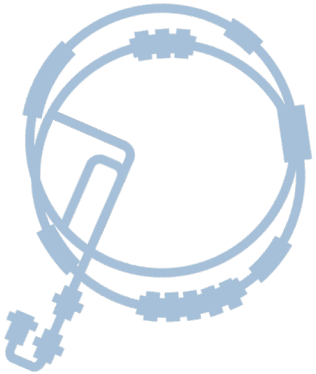

	CPU	Total TDP	Cooling-fan Wattage on the tested server	Fan Wattage/ Total TDP	Server-level power consumption*1
Intel Xeon E5 server	Intel Xeon E5 2630 v3 x2	85W x2 = 170W	14.5W x6 + 23W x2 = 133W	0.78	299W
ThunderX2 server (HPE Apollo 70)	Marvell ThunderX2 CN9980 x2	180W x2 = 360W	54.5W x4 = 218W	0.61	588W
Intel Xeon Gold server	Intel Xeon Gold 6148 x2	150W x2 = 300W	18W x5 + 12.6W x2 = 115W	0.38	392W

- Cooling fans of ThunderX2 server consume much power.
- Although total TDP of the ThunderX2 is only 1.2 times higher than the Xeon Gold, the server power consumption difference is much more.
- The Xeon Gold can be cooled by less fan power than the Xeon E5.

*1 Average power consumption during G4Bench Ecal simulation.

Comparison of Processors

		Xeon E5-2630 v3	Xeon Gold 6148	ThunderX2 CN9980
Number of cores		8	20	32
SMT		1, 2	1, 2	1, 2, 4
Frequency		2.4 GHz	2.4 GHz	2.2 GHz
Memory channel		4	6	8
TDP		85 W	150 W	180 W
USD/core		USD667[*1]/8=USD83	USD3072[*1]/20=USD154	USD1795[*2]/32=USD56
Cost performance Events/second/2CPU price	Light	1.53	1	2.25
	Middle	1.44	1	1.94
	Heavy	1.39	1	1.60
Server-level performance Events/second/node	Light	0.33	1	1.32
	Middle	0.31	1	1.14
	Heavy	0.30	1	0.93
Energy efficiency Events/kWh	Light	0.46	1	1.00
	Middle	0.39	1	0.80
	Heavy	0.37	1	0.63

*1 Recommended Customer Price on https://ark.intel.com

*2 https://www.cavium.com/news/cavium-announces-thunderx2-general-availability

Summary

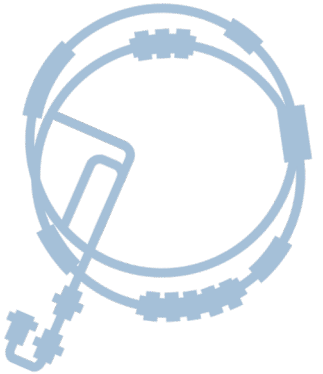

- We have evaluated Marvell ThunderX2 performance using Geant4-based Monte Carlo simulation benchmark.
  - ThunderX2 cost performance is very attractive.
  - ThunderX2 server-level performance is better than Xeon Gold server on light and middle -weight workload tests.
  - Intel Xeon Gold server is more power efficient than ThunderX2 server.
- We are looking forward to the next Arm processor!
- Future work:
  - Execute memory intensive benchmark.
  - Execute other HEP software.
- Thanks to Marvell for the preparation of the ThunderX2 test server and for giving us useful advice.

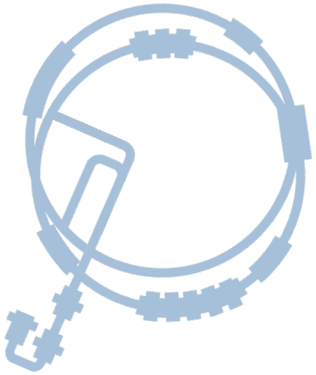

IBM System x3550 M5

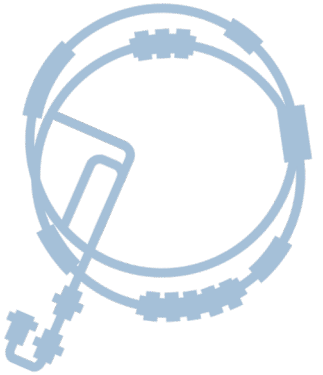

- CPU: Intel Xeon E5-2630 v3 x2
  - 8 cores/CPU
- Memory: 64 GB
  - 16GB DDR4 1866 DIMM x4
- Kernel: 3.10.0-693.21.1.el7.x86_64
- OS: CentOS 7.6.1810 (Core)

HPE ProLiant DL360 Gen10

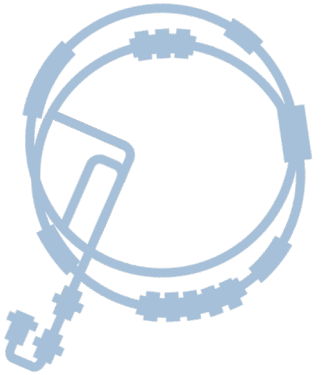

- CPU: Intel Xeon Gold 6148 x2
  - 20 cores/CPU
- Memory: 256 GB
  - 16GB DDR4 2666 DIMM x16
- Kernel: 4.15.0-66-generic
- OS: Ubuntu 18.04.3 LTS

Cooling Fans

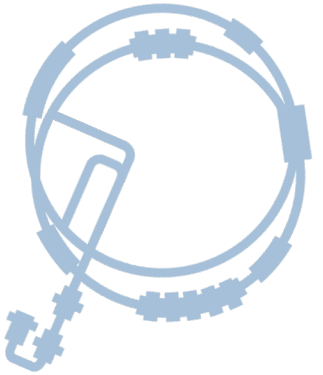

- IBM System x3550 M5
  - Nidec UltraFlo 12V, 1.21A x6
  - DELTA ELECTRONICS 12V, 1.9A x2
- HPE Apollo 70:
  - Nidec UltraFlo: 12V, 4.54A x4
- HPE ProLiant DL360 Gen10:
  - DELTA ELECTRONICS: 12V, 1.5A x5
  - SUNON: 12V, 1.05A x2