

New JCAHPC Supercomputer System

Joint Center for Advanced HPC (JCAHPC)

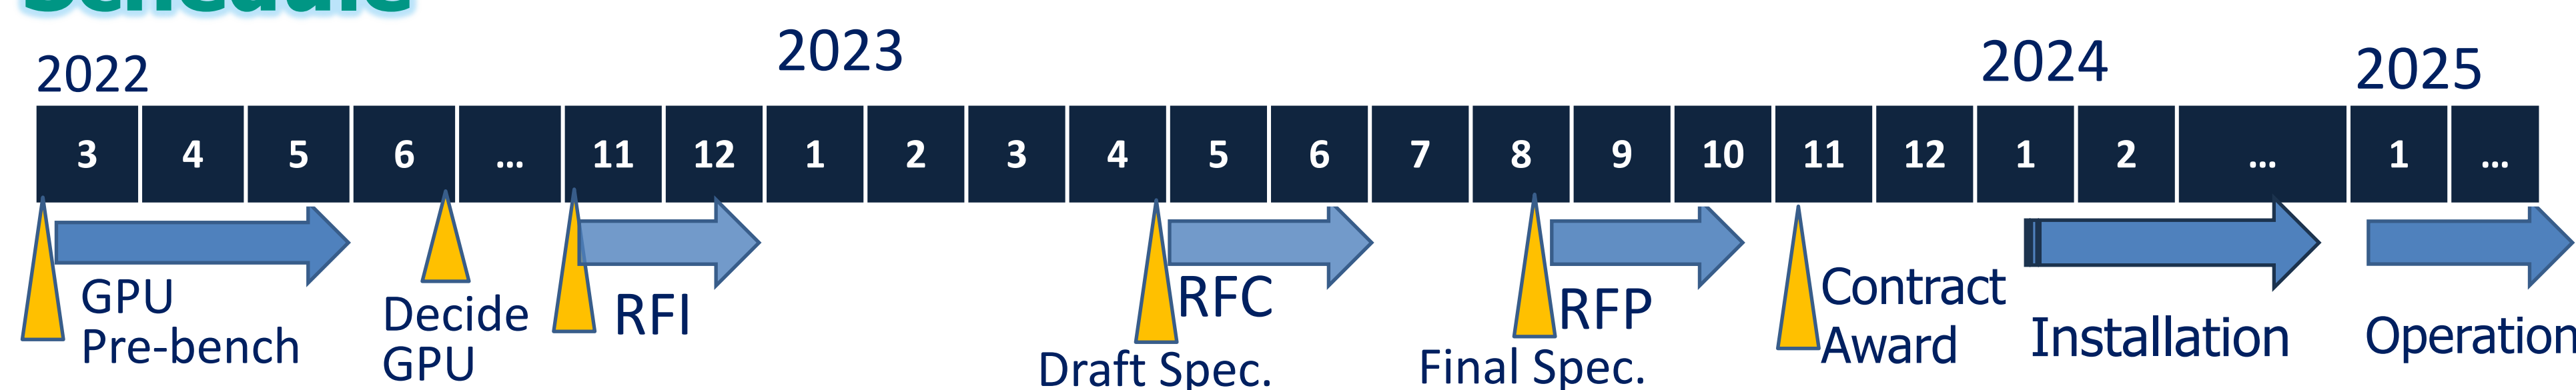
Center for Computational Sciences, University of Tsukuba
Information Technology Center, The University of Tokyo

Overview

We decided to introduce a new supercomputer system "OFP-II" (tentative, aka. Post-OFP) with a total peak performance of **79.5 PFLOPS** as a successor of Oakforest-PACS. We will start its operation in **January 2025**. This system will be installed in the Kashiwa campus, the University of Tokyo, at the exact same location as the Oakforest-PACS system.

The OFP-II will be offered to researchers in Japan and their international collaborators through various types of programs operated by HPCI, by MEXT's Joint Usage/Research Centers, and by each of us under original supercomputer resource sharing programs. Inheriting the philosophy of the introduction of Oakforest-PACS, OFP-II aims to promote novel computational science methods leveraging AI, such as AI for HPC/Science, to offer a platform supporting Society 5.0 by integration of simulation, data analysis, and machine learning, in addition to support users of large-scale applications.

Schedule



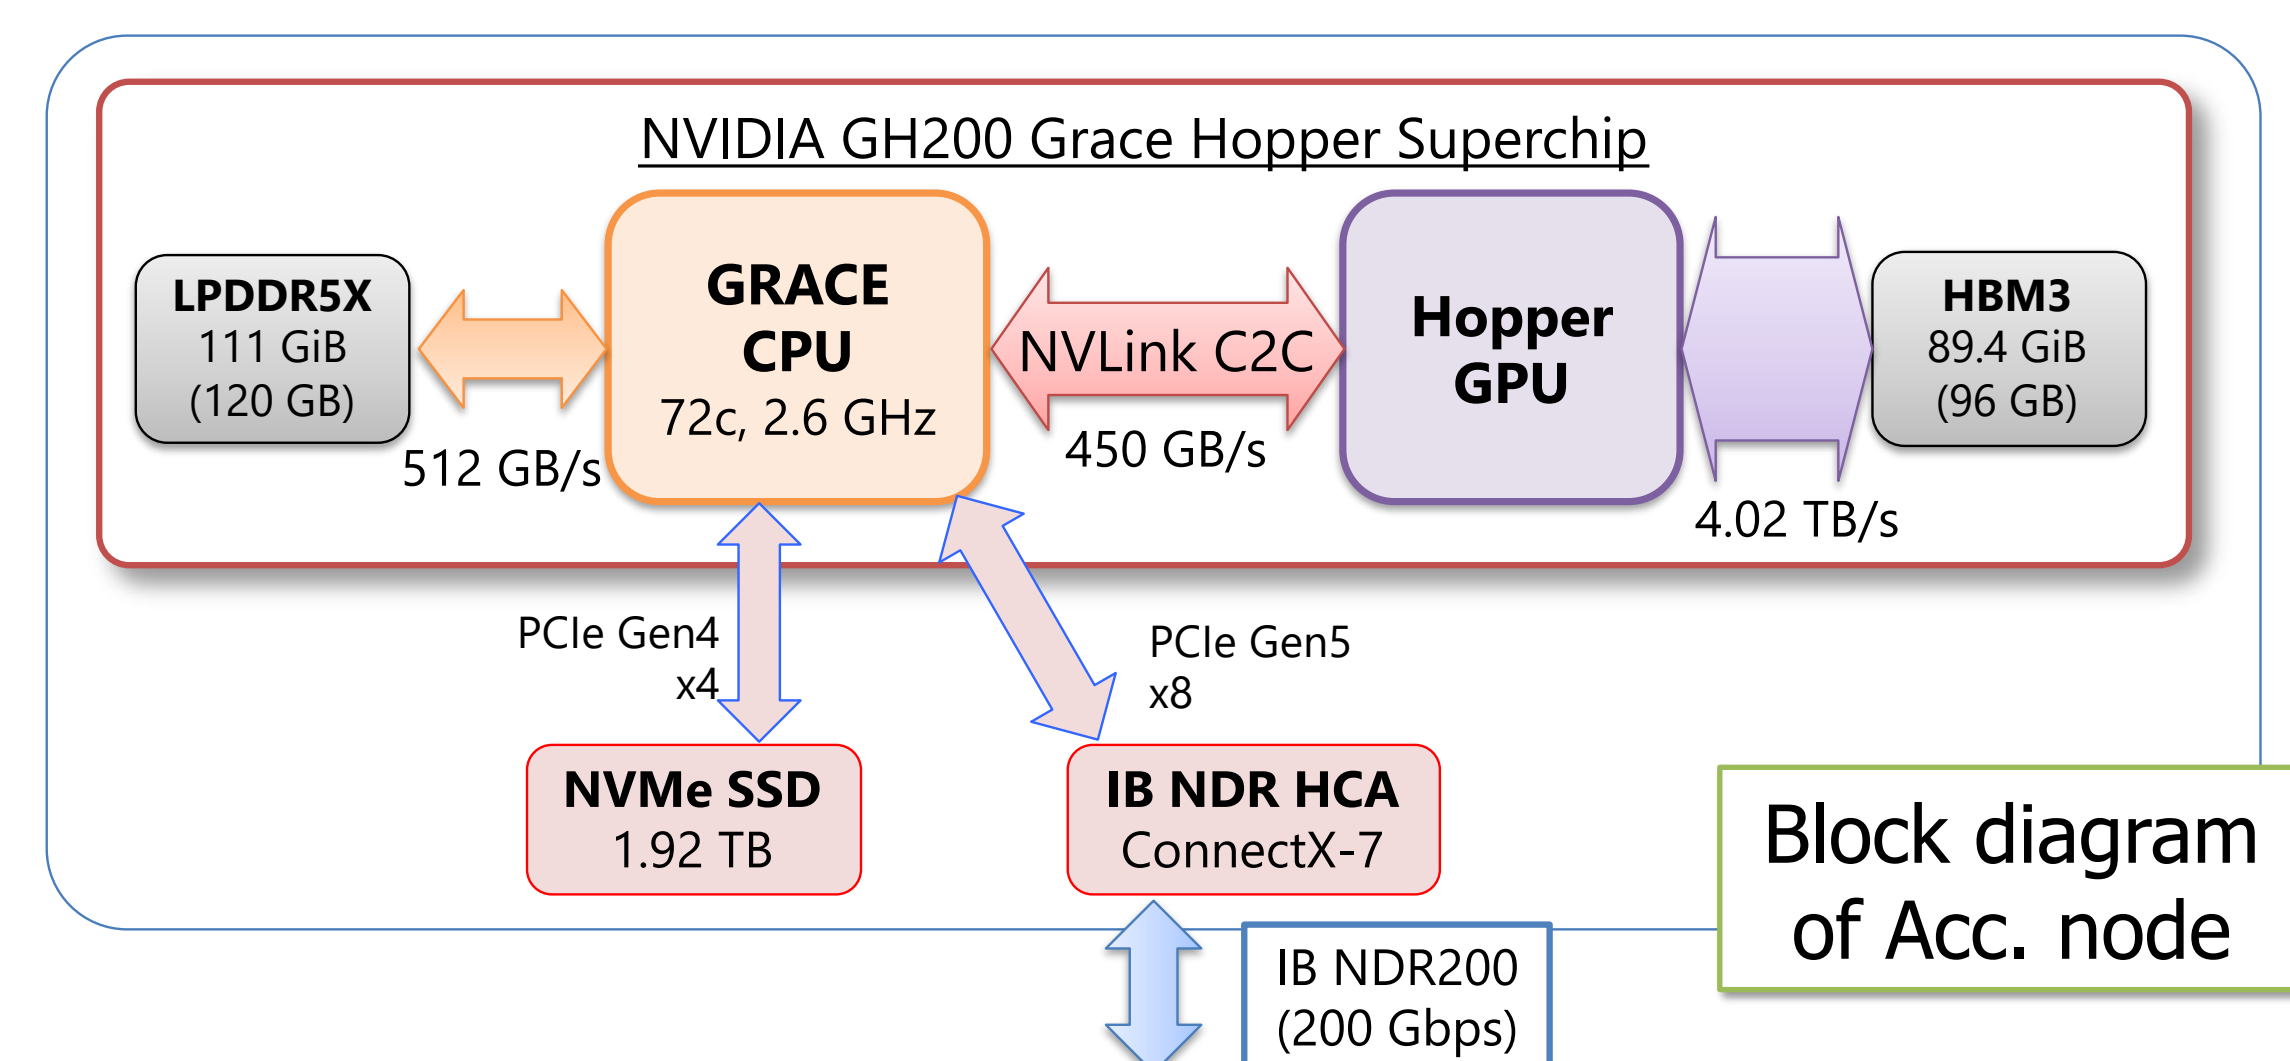
Configuration of OFP-II System

The OFP-II system has 1,120 nodes for Accelerator node, each of which consists of NVIDIA GH200 Grace Hopper Superchip (Arm Neoverse V2 core, cache-coherent with GPU), 190 nodes for General-purpose CPU node, and NVIDIA/Mellanox InfiniBand NDR as a high-performance interconnect. This is the first large-scale system with GH200 in Japan. The system will be installed by Fujitsu, and Supermicro and Fujitsu PRIMERGY servers are employed as compute nodes. The system employs the shared files system by All Flash with 10.3 PB provided by Data Direct Network (DDN). In addition, "Ipomoea-01" is also available as a large-scale external filesystem to be coupled to OFP-II.

The peak performance of OFP-II becomes 79.5 PFLOPS and the total memory BW is more than 5.6 PB/sec. All compute nodes and servers organizing the filesystem are connected by fat-tree topology based on InfiniBand-NDR, which provides full-bisection bandwidth. Therefore, flexible and efficient utilization and operation of compute nodes and file systems are available.

Specification of OFP-II System

Node Type	Accelerator	General-purpose CPU
Theoretical Peak Performance	78.2 PFLOPS	1.29 PFLOPS
Number of Nodes	1,120	190
Total Memory Capacity	220.0 TiB	23.75 TiB
Total Memory BW	5.07 PB/sec	608 TB/sec
Interconnect Topology	InfiniBand NDR200 (200 Gbps) Full-bisection Fat Tree	
Shared Filesystem	Lustre FS	
MDS	Server	DDN ES400NVX2
	# of Servers (VM)	1 (4)
	# of inodes	appx. 23.5 B
OSS	Server	DDN ES400NVX2
	# of Servers	10 set
	Capacity	10.3 PB (All Flash)
	Theoretical BW	1.0 TB/sec



Node Specification

Node Type	Accelerator node	General-Purpose CPU node
Server	Supermicro ARS-111GL-DNHR-LCC	FUJITSU Server PRIMERGY CX2550 M7
CPU	Processor	NVIDIA GH200 Grace Hopper Superchip (CPU: NVIDIA Grace)
	# of CPUs (Core)	1 (72)
	Frequency	2.6 GHz
	Theoretical Peak Perf.	2.99 TFLOPS
	Memory Capacity /node	111 GiB (120 GB)
GPU	Processor	NVIDIA Hopper
	# of GPUs	1
	# of SMs	132
	Theoretical Peak Perf.	66.9 TFLOPS
	Memory Capacity	HBM3 89.4 GiB (96 GB)
SSD	Memory BW /node	512 GB/s
	CPU-GPU conn.	NVLink C2C 450 GB/sec Cache-coherent
	SSD	NVMe SSD 1.92 TB, PCIe G4 x4

