

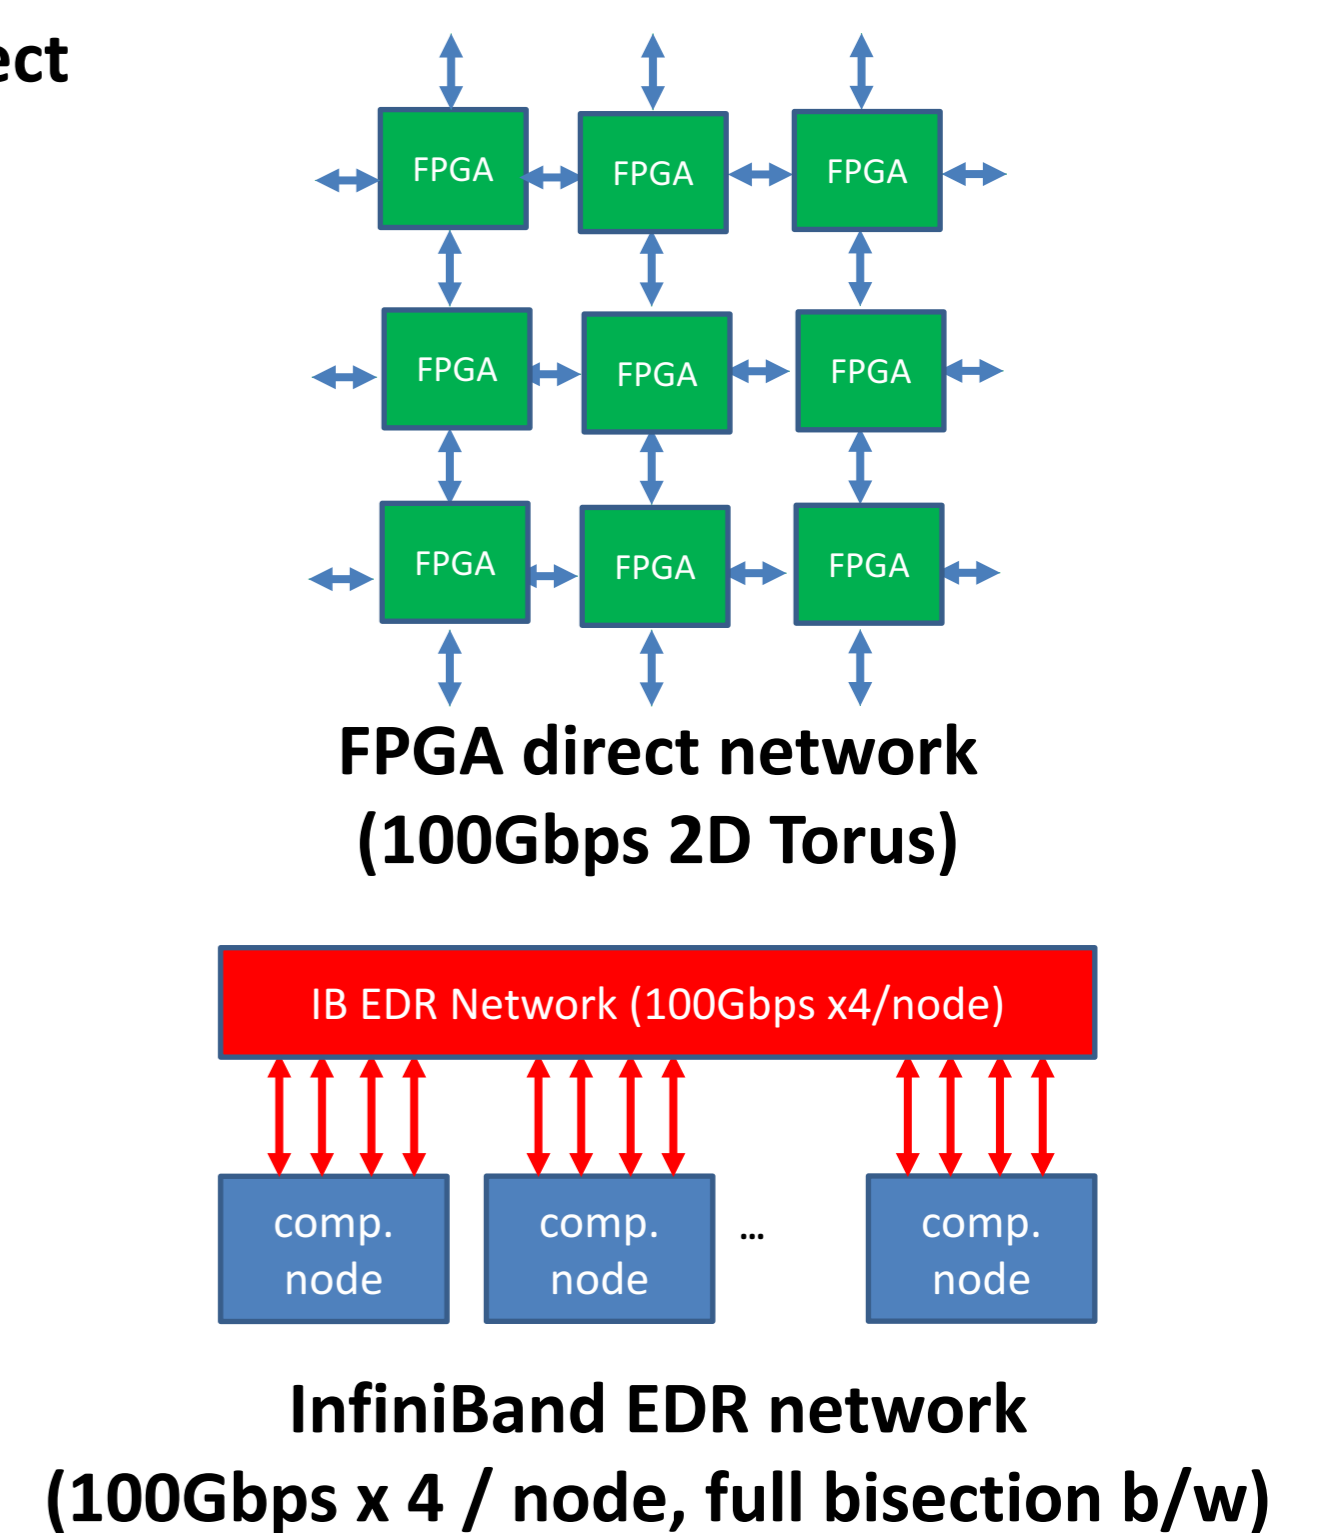
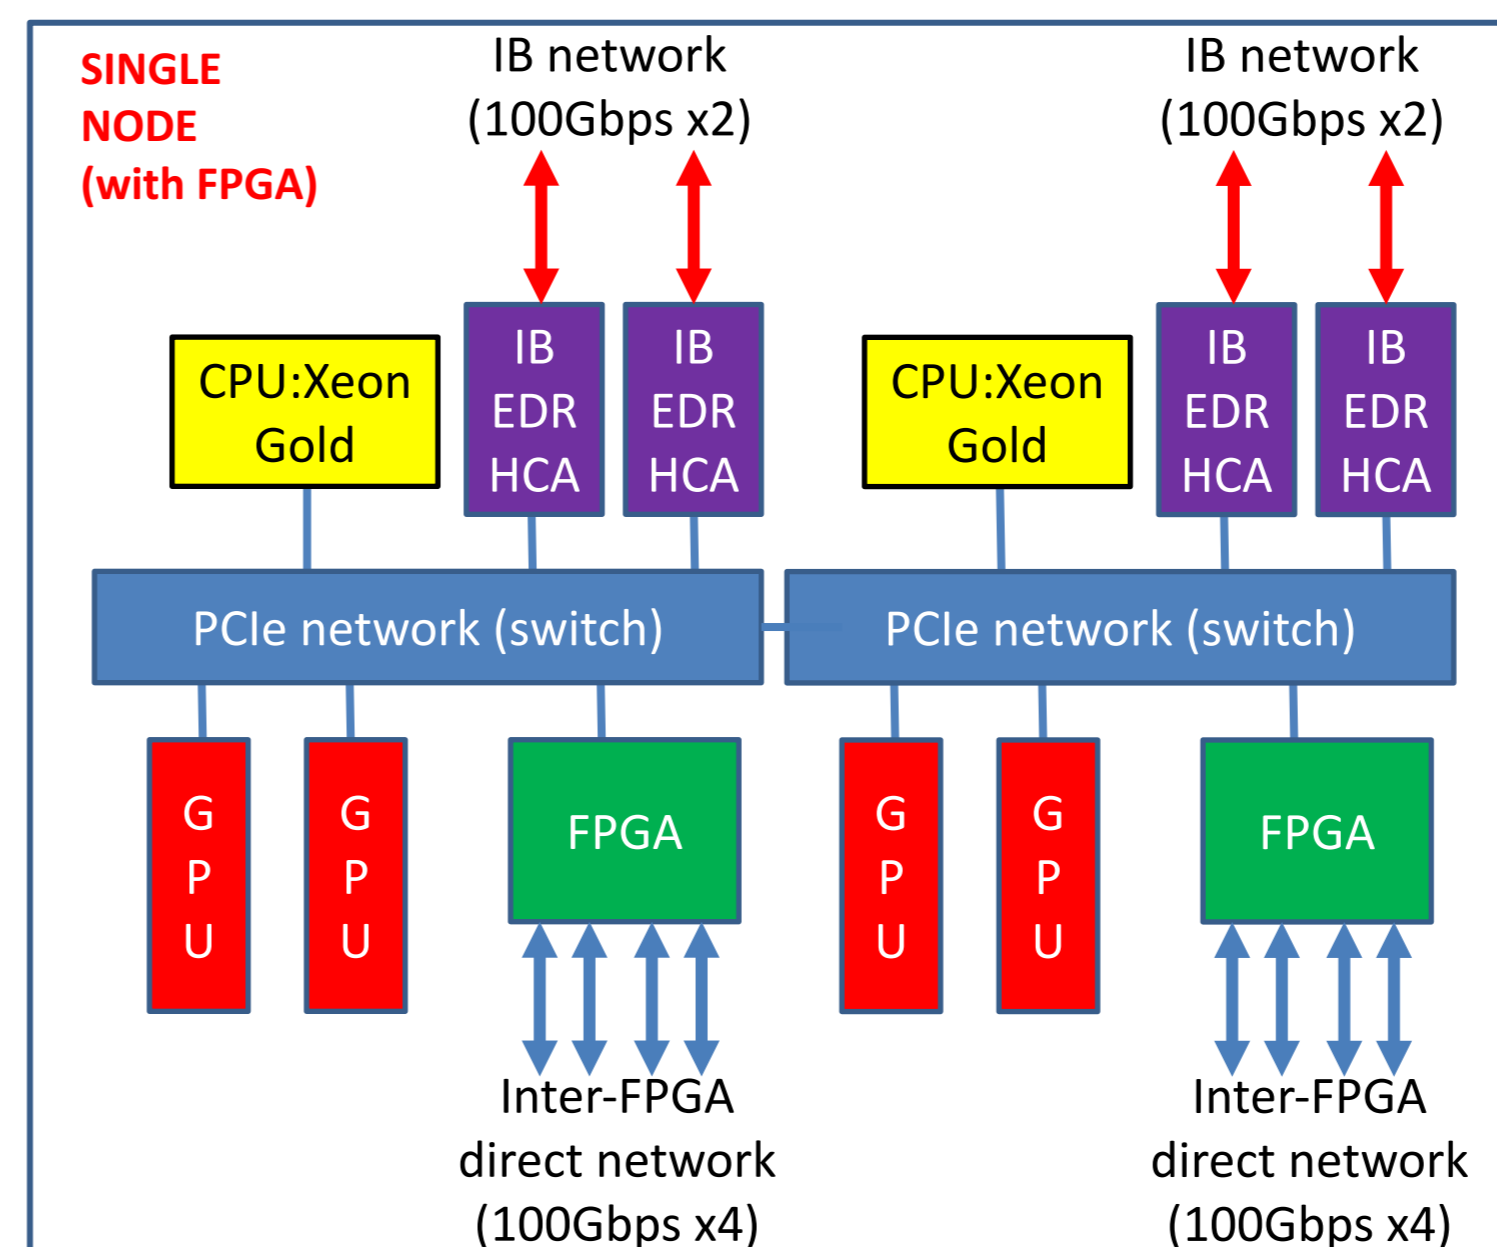
# Supercomputers at CCS: Cygnus and COMA

## Cygnus: Multi-Hybrid Cluster for Next Generation's Accelerated Computing (March 2019)

Specification of Cygnus system (these numbers and performance are currently planned, but subject to be changed.)

Item	Specification
Peak performance	~3.0 PFLOPS (GPU: 2.3 FPGA: 0.5 CPU: 0.2)
# of nodes	~80
CPU/node	Intel Xeon Gold x 2
GPU/node	NVIDIA V100 x4
FPGA/node	Intel Stratix10 x2 (for 32 nodes)
File system	Lustre, ~3 PB user area (DDN)
Interconnection Network	Mellanox InfiniBand EDR x4 chan.
Total network bandwidth	~4 TB/s
Language	CPU: Fortran90, C, C++; GPU: OpenACC; FPGA: OpenCL, Verilog HDL

Each node is equipped with both IB EDR and FPGA-direct network (32 nodes are equipped with both FPGAs and GPUs, and other nodes are with GPUs only)

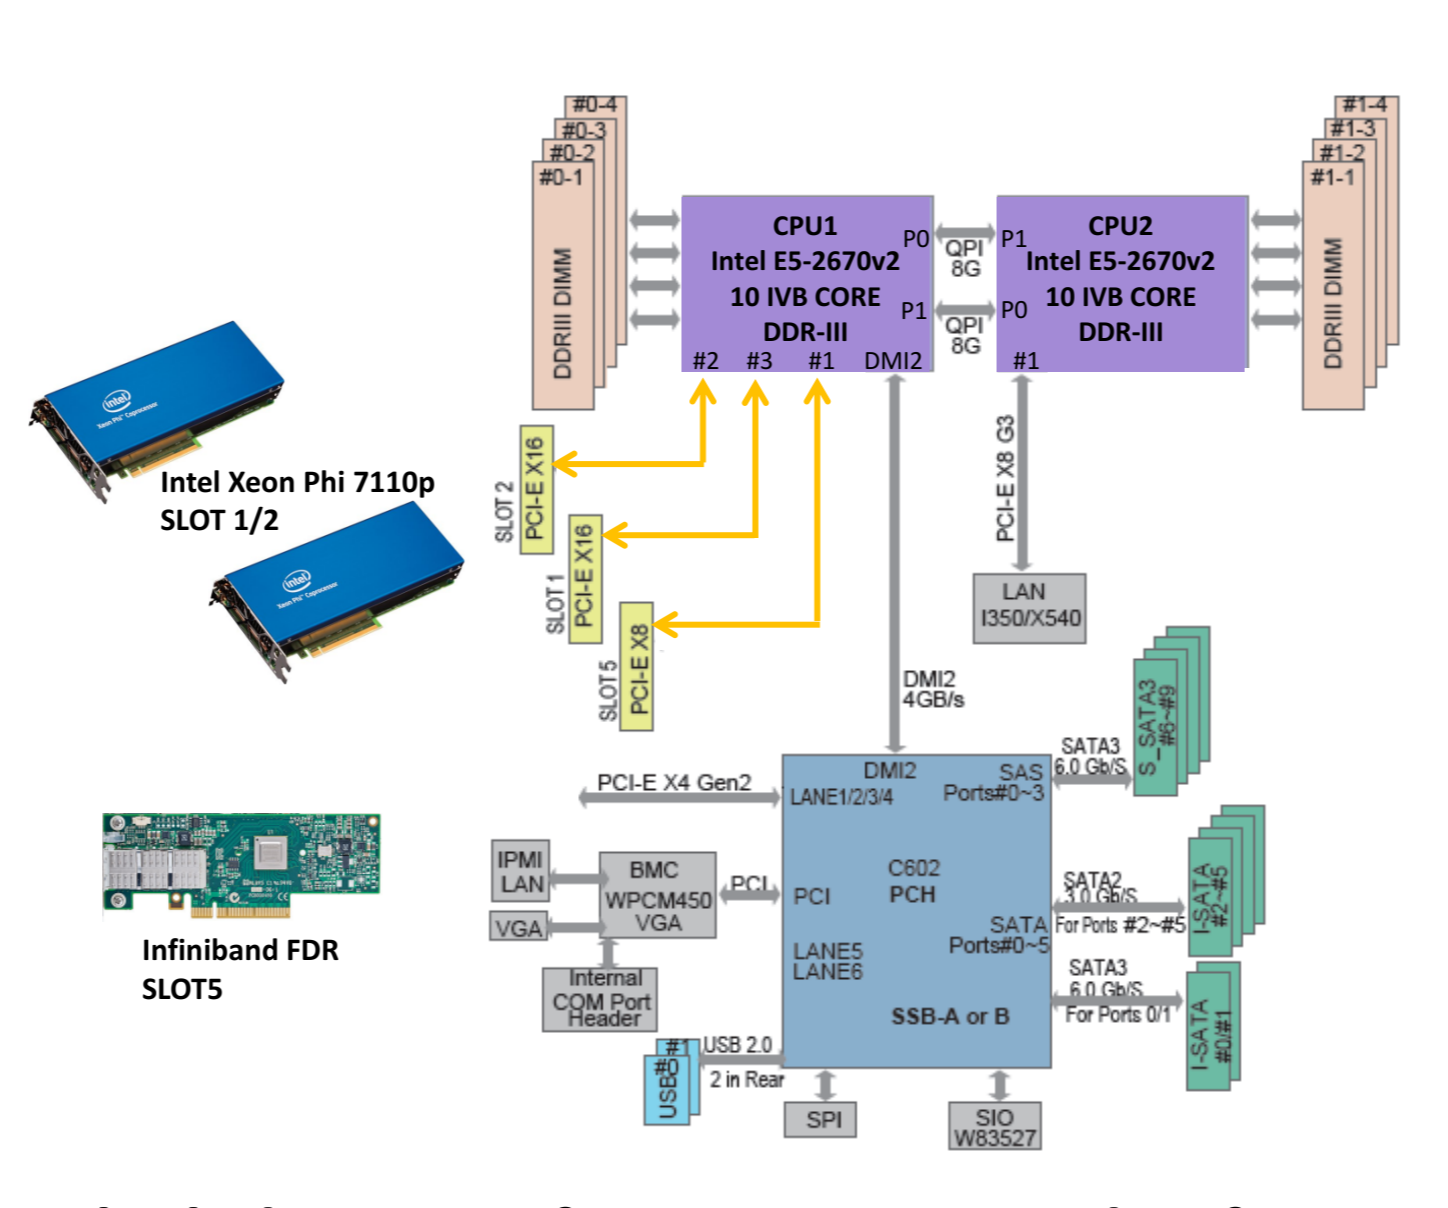


Cygnus is the next generation's supercomputer in Center for Computational Sciences introducing multiple accelerated computing devices with GPU and FPGA together in each computation node. It is conceptually called "Multi-Hybrid Accelerating Cluster" since we apply different type of accelerators according to the computation characteristics of applications or even in a single application. The absolute performance of GPU is still very high where it depends on a large degree of SIMD-type of data parallelism. However, GPU sometime degrades its performance when the parallelism shrinks or by frequent communication over multiple devices/nodes due to its communication performance weakness.

FPGA is expected as a new accelerating device for HPC since it provides high flexibility on hardware designing and fine-grained pipeline manner of computation which provides different way of accelerating feature from GPU although its absolute floating point performance is lower than GPU. Moreover, the latest FPGA can support not only the computation and reconfigurable architecture but also multiple ports of external communication links where each of them reaches up to 100Gbps to make FPGA-oriented interconnection network. We consider that GPU and FPGA can compensate with each other to cover each side of weakness and become a new solution to attack the tough problems for GPU clusters and develop a prototype of next generation's hybrid supercomputer with these accelerating technologies. Cygnus will be delivered on March 2019.

## COMA: Cluster Of Many-core Architecture Processor

Item	Specification
Peak performance	1.001 PFLOPS (MIC: 843.9TF, CPU: 157.2TF)
# of nodes	393
File system	Lustre, 1.5 PB user area (DDN SFA12000)
Infiniband network switch	393 port FDR (Mellanox SX6536)
Total network bandwidth	2.75 TB/s
Language	Fortran90, C, C++
MPI	MVAPICH2, Intel MPI
System Management	Cray Advanced Cluster Engine, SLURM



Block diagram of computation node of COMA



COMA (Cluster Of Manycore Architecture) is the 9th generation of PACS series supercomputer at CCS, of which computation node is equipped with two set of Intel Xeon Phi co-processor (Knights Corner) as well as two sockets of Intel Xeon CPU (Ivy Bridge). The interconnection network topology is a Fat-Tree with full bisection bandwidth supported by Mellanox InfiniBand/FDR. To avoid communication bottleneck between multiple PCIe devices over Intel QPI, all these devices are connected to single socket of Xeon CPU.

Either of "offloading", "native" and "symmetric" modes of Xeon Phi co-processor is available for users. Especially, the code development in the native mode is recommended to users to shift them to our new platform, Oakforest-PACS under management of JCAHPC as the largest Xeon Phi (Knights Landing) cluster in Japan with 25 PFLOPS of peak performance. Several computational science codes developed and tuned on COMA are ported to Oakforest-PACS with several times faster performance on each node.