

# Supercomputer at CCS: Cygnus

## Multi-Hybrid Accelerated Computing Platform

### Combining goodness of different type of accelerators: GPU + FPGA

- GPU is still an essential accelerator for simple and large degree of parallelism to provide ~10 TFLOPS peak performance
- FPGA is a new type of accelerator for application-specific hardware with programmability and speeded up based on pipelining of calculation
- FPGA is good for external communication between them with advanced high speed interconnection up to 100Gbps x4 chan.

### Construction of “Cygnus”

- Operation started in May 2019
- 2x Intel Xeon CPUs, 4x NVIDIA V100 GPUs, 2x Intel Stratix10 FPGAs
- Deneb: 49 CPU+GPU nodes
- Albireo: 32 CPU+GPU+FPGA nodes with 2D-torus dedicated network for FPGAs (100Gbpsx4)



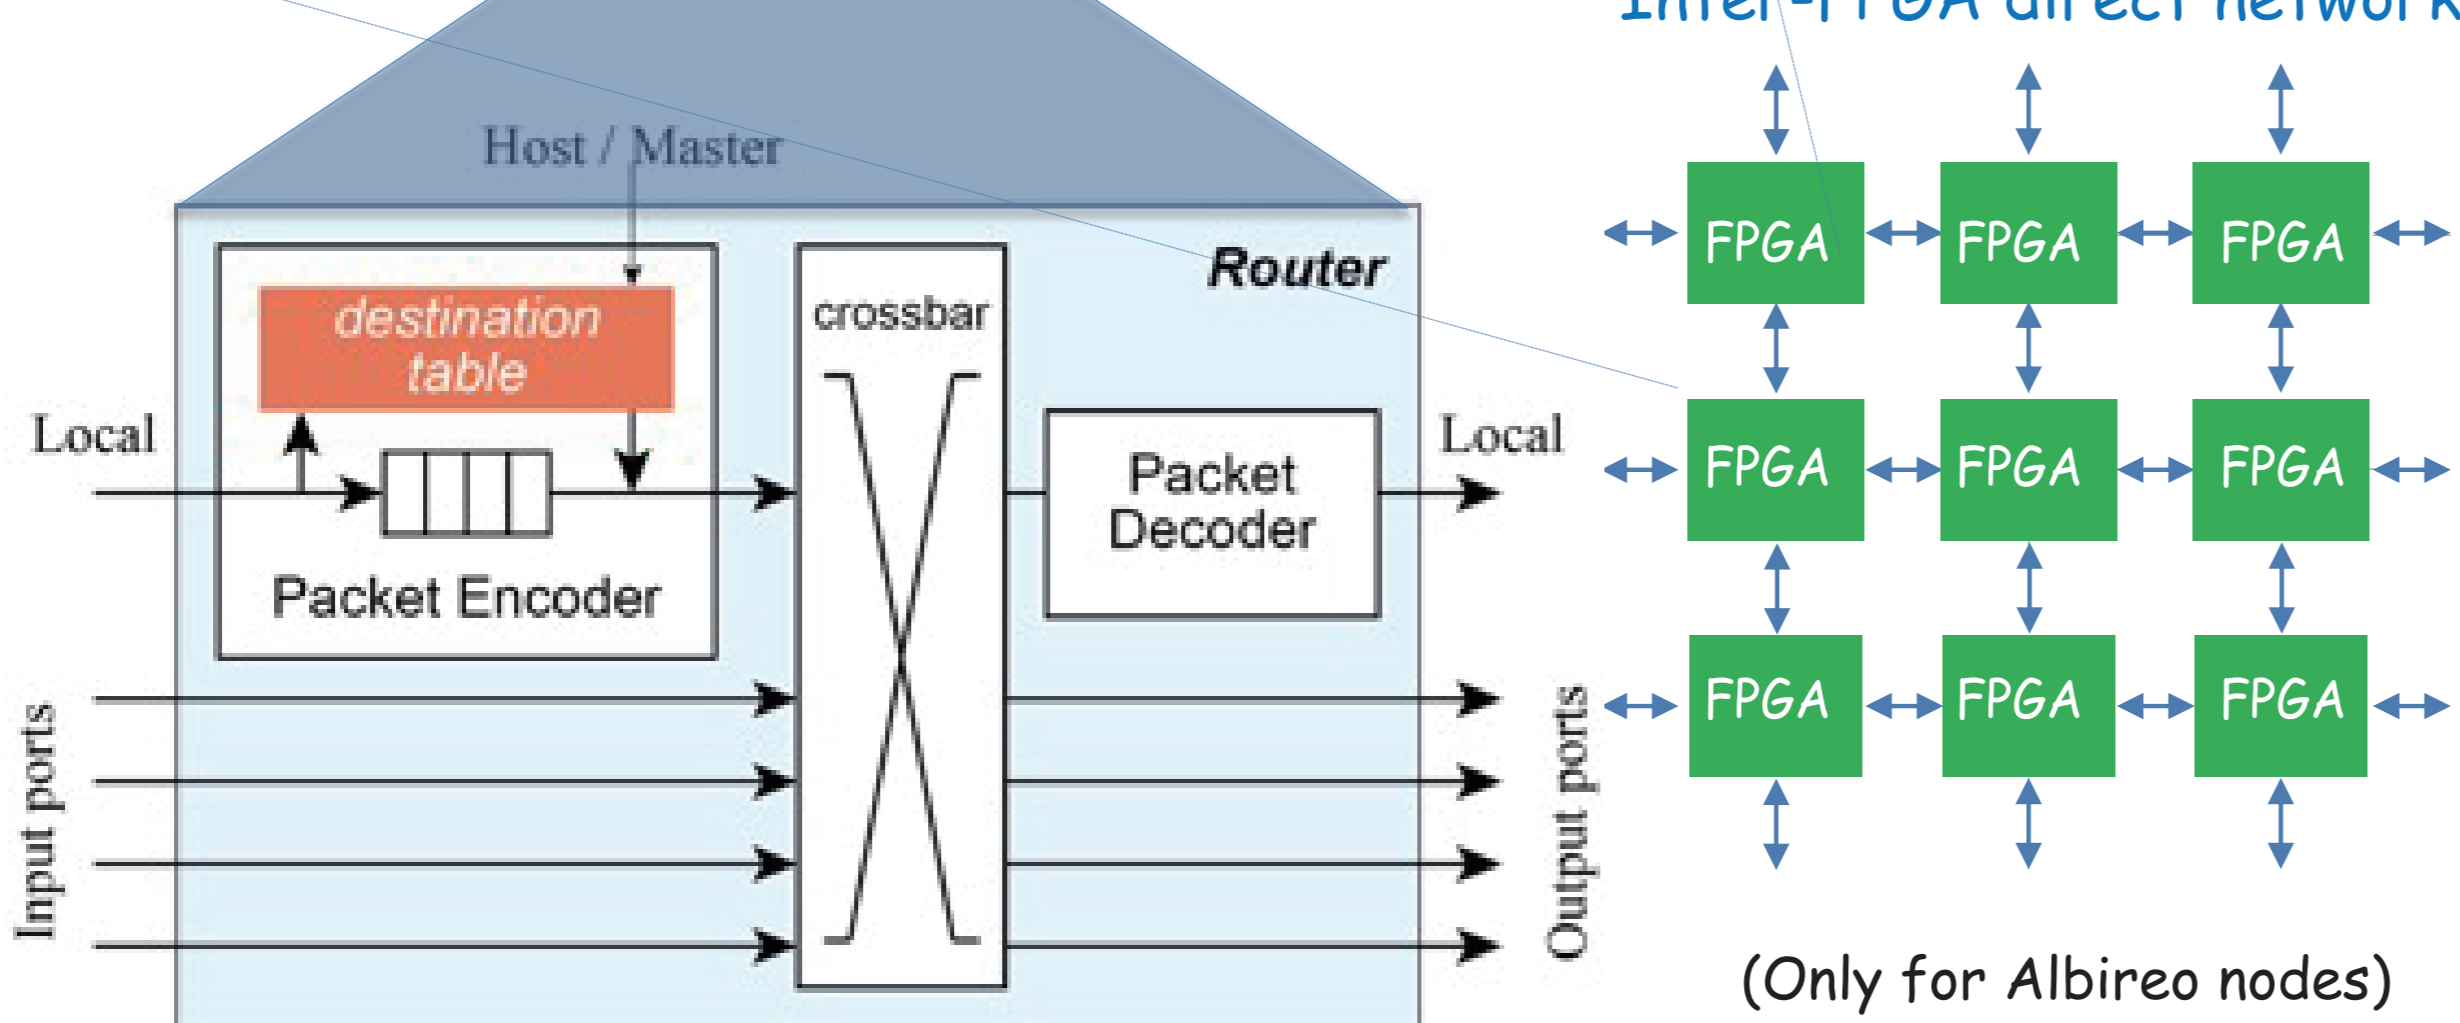
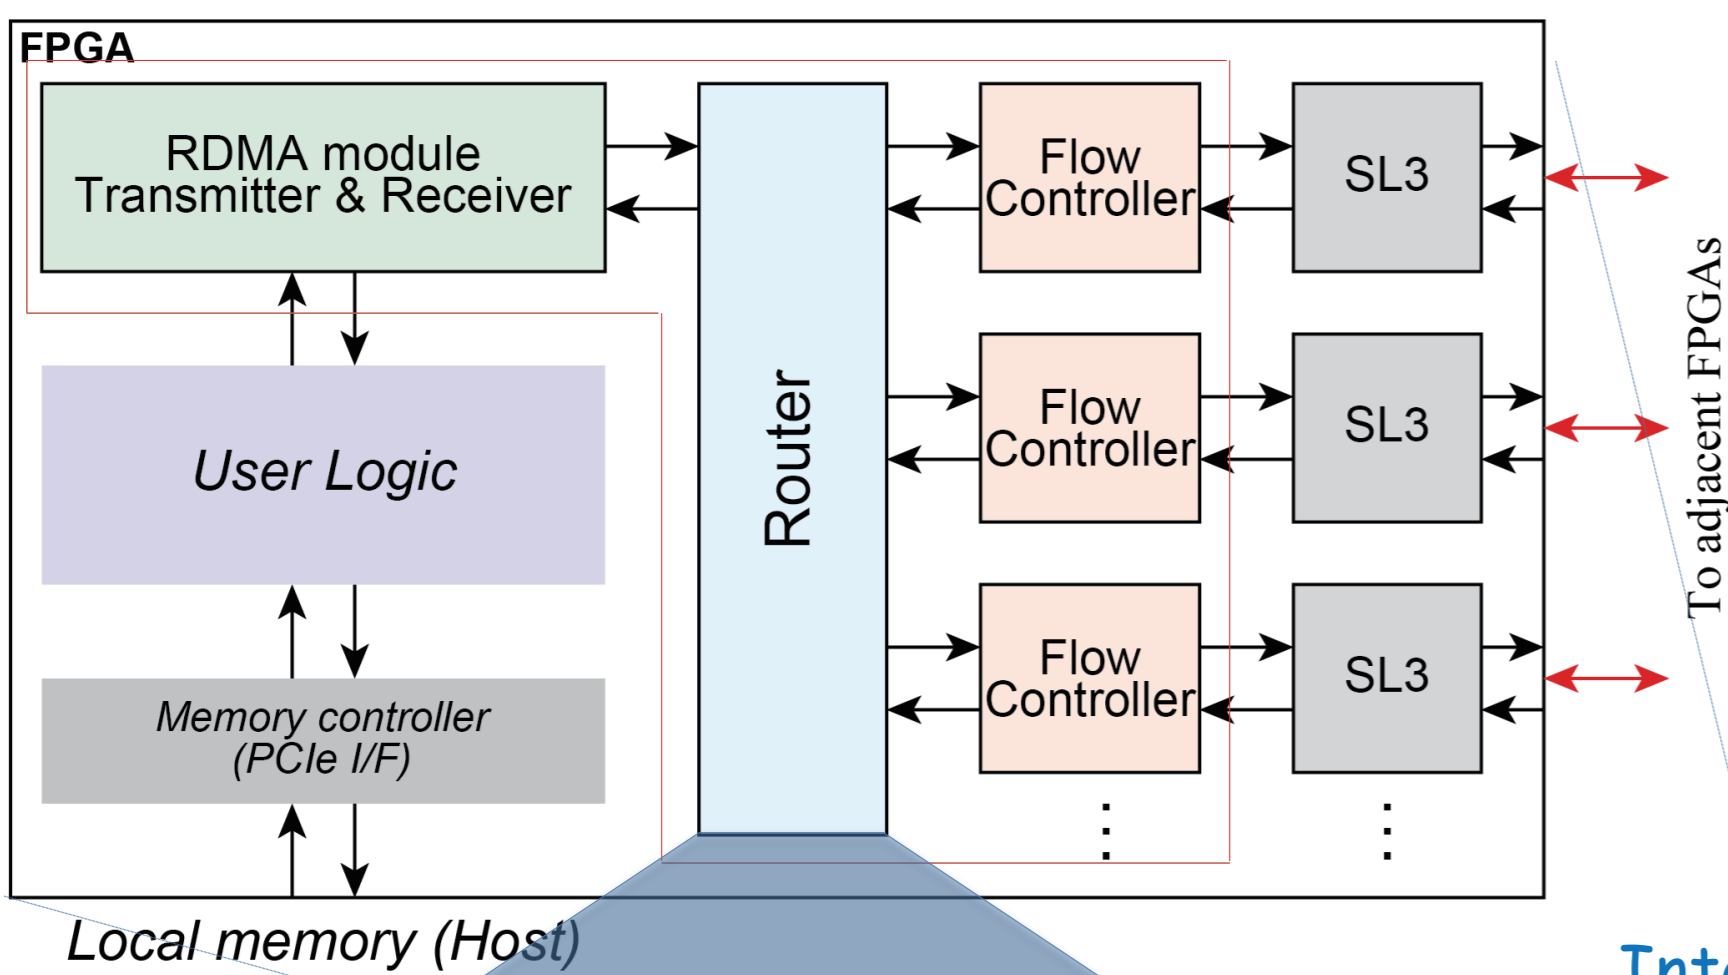
Target GPU: NVIDIA Tesla V100



Target FPGA: Nallatech 520N

### FPGA design plan

- Router
  - For the dedicated network, this impl. is mandatory.
  - Forwarding packets to destinations
- User Logic
  - OpenCL kernel runs here.
  - Inter-FPGA comm. can be controlled from OpenCL kernel.
- SL3
  - SerialLite III : Intel FPGA IP
  - Including transceiver modules for Inter-FPGA data transfer.
- Users don't need to care



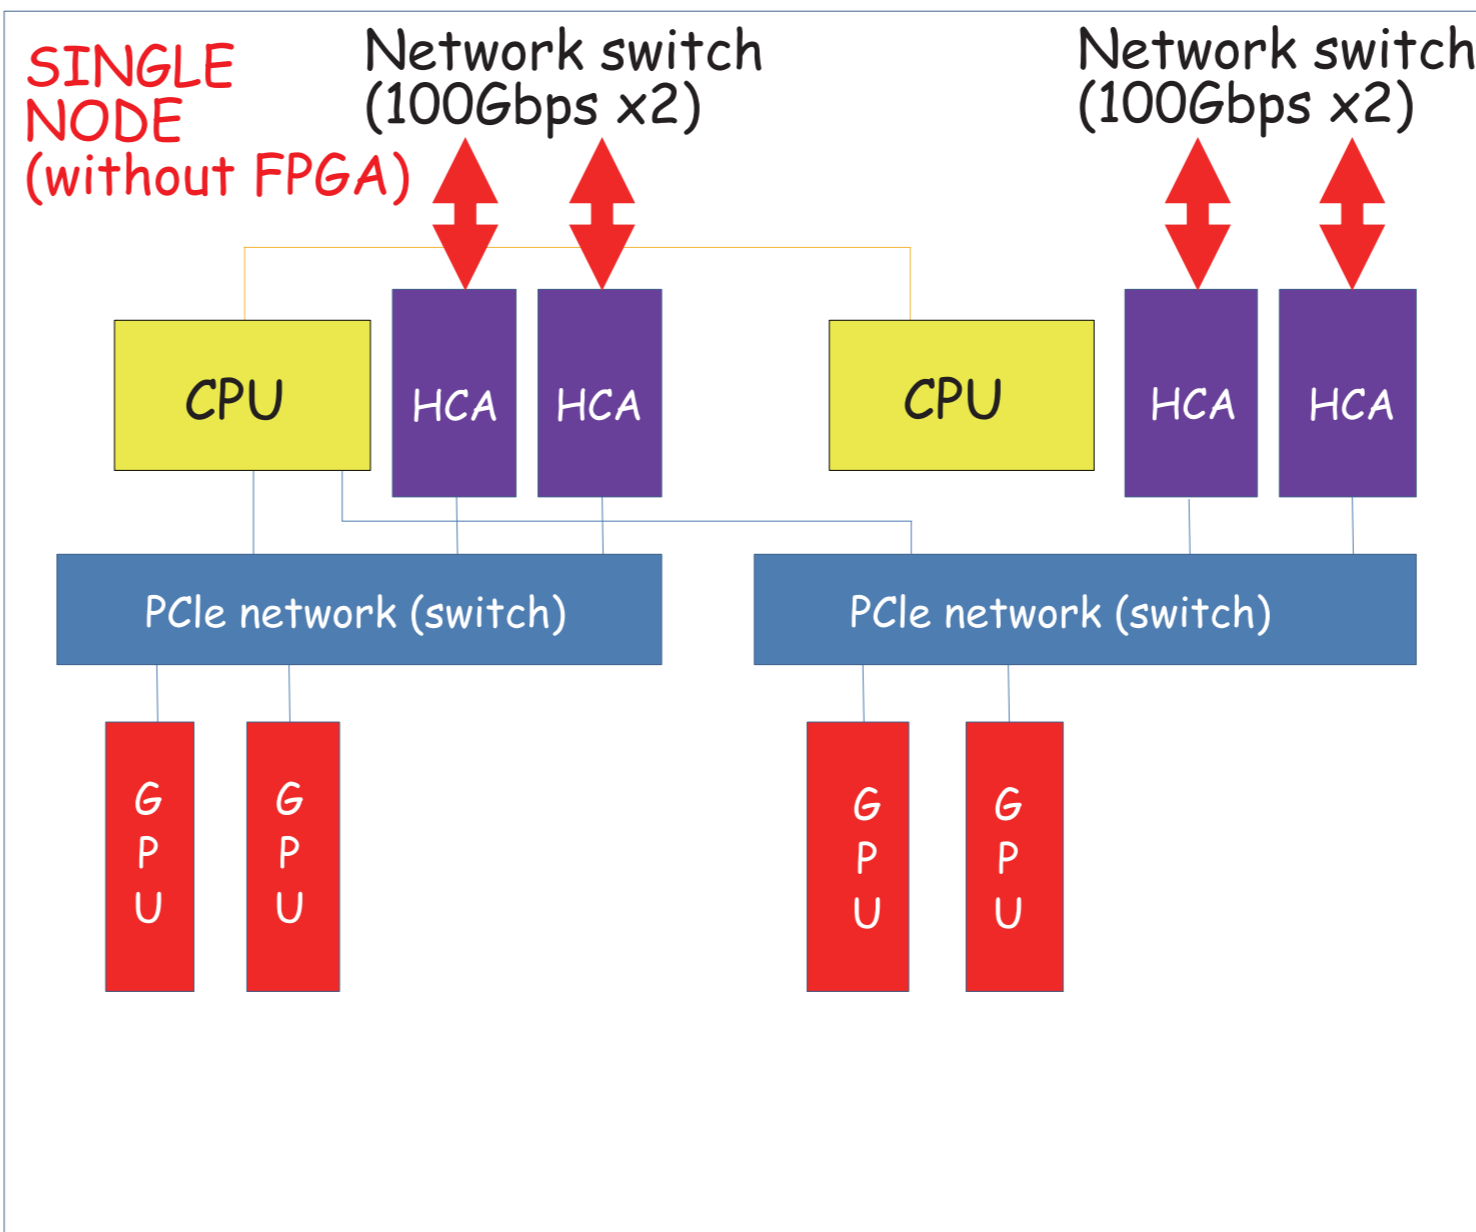
64FPGAs on Albireo nodes are connected directly as 2D-Torus configuration without Ethernet sw.



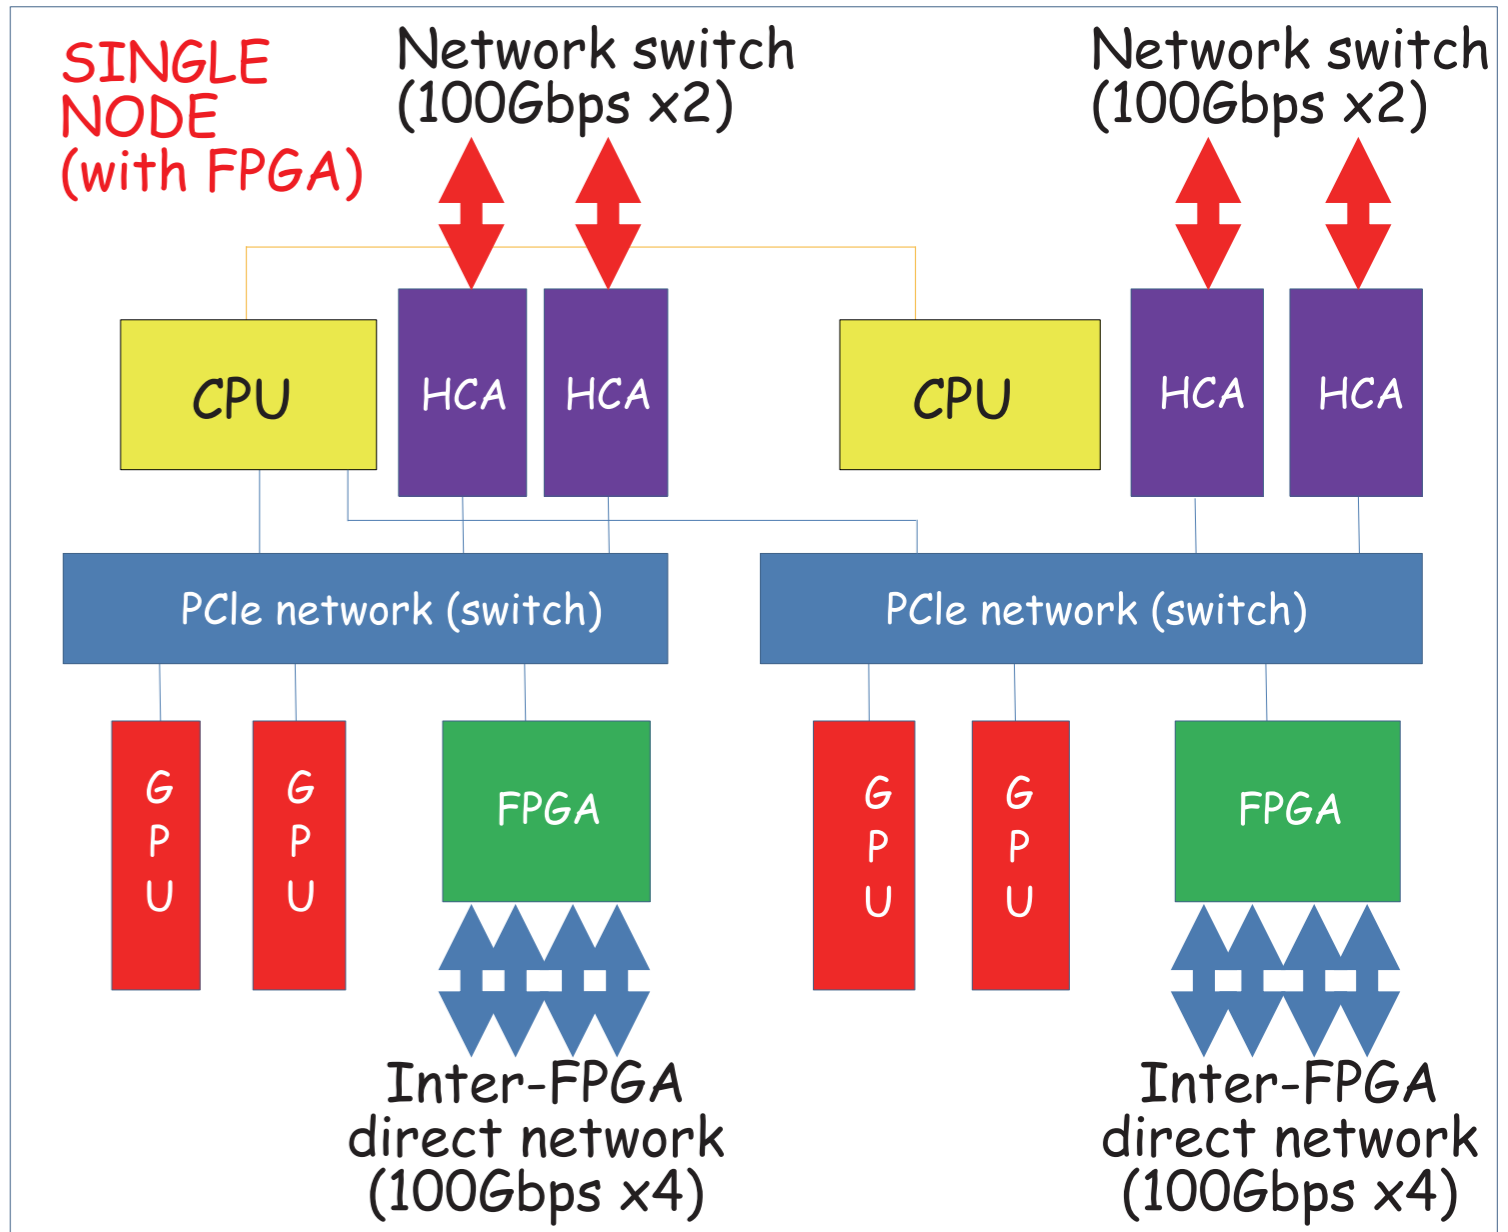
Cygnus

### Specification of Cygnus

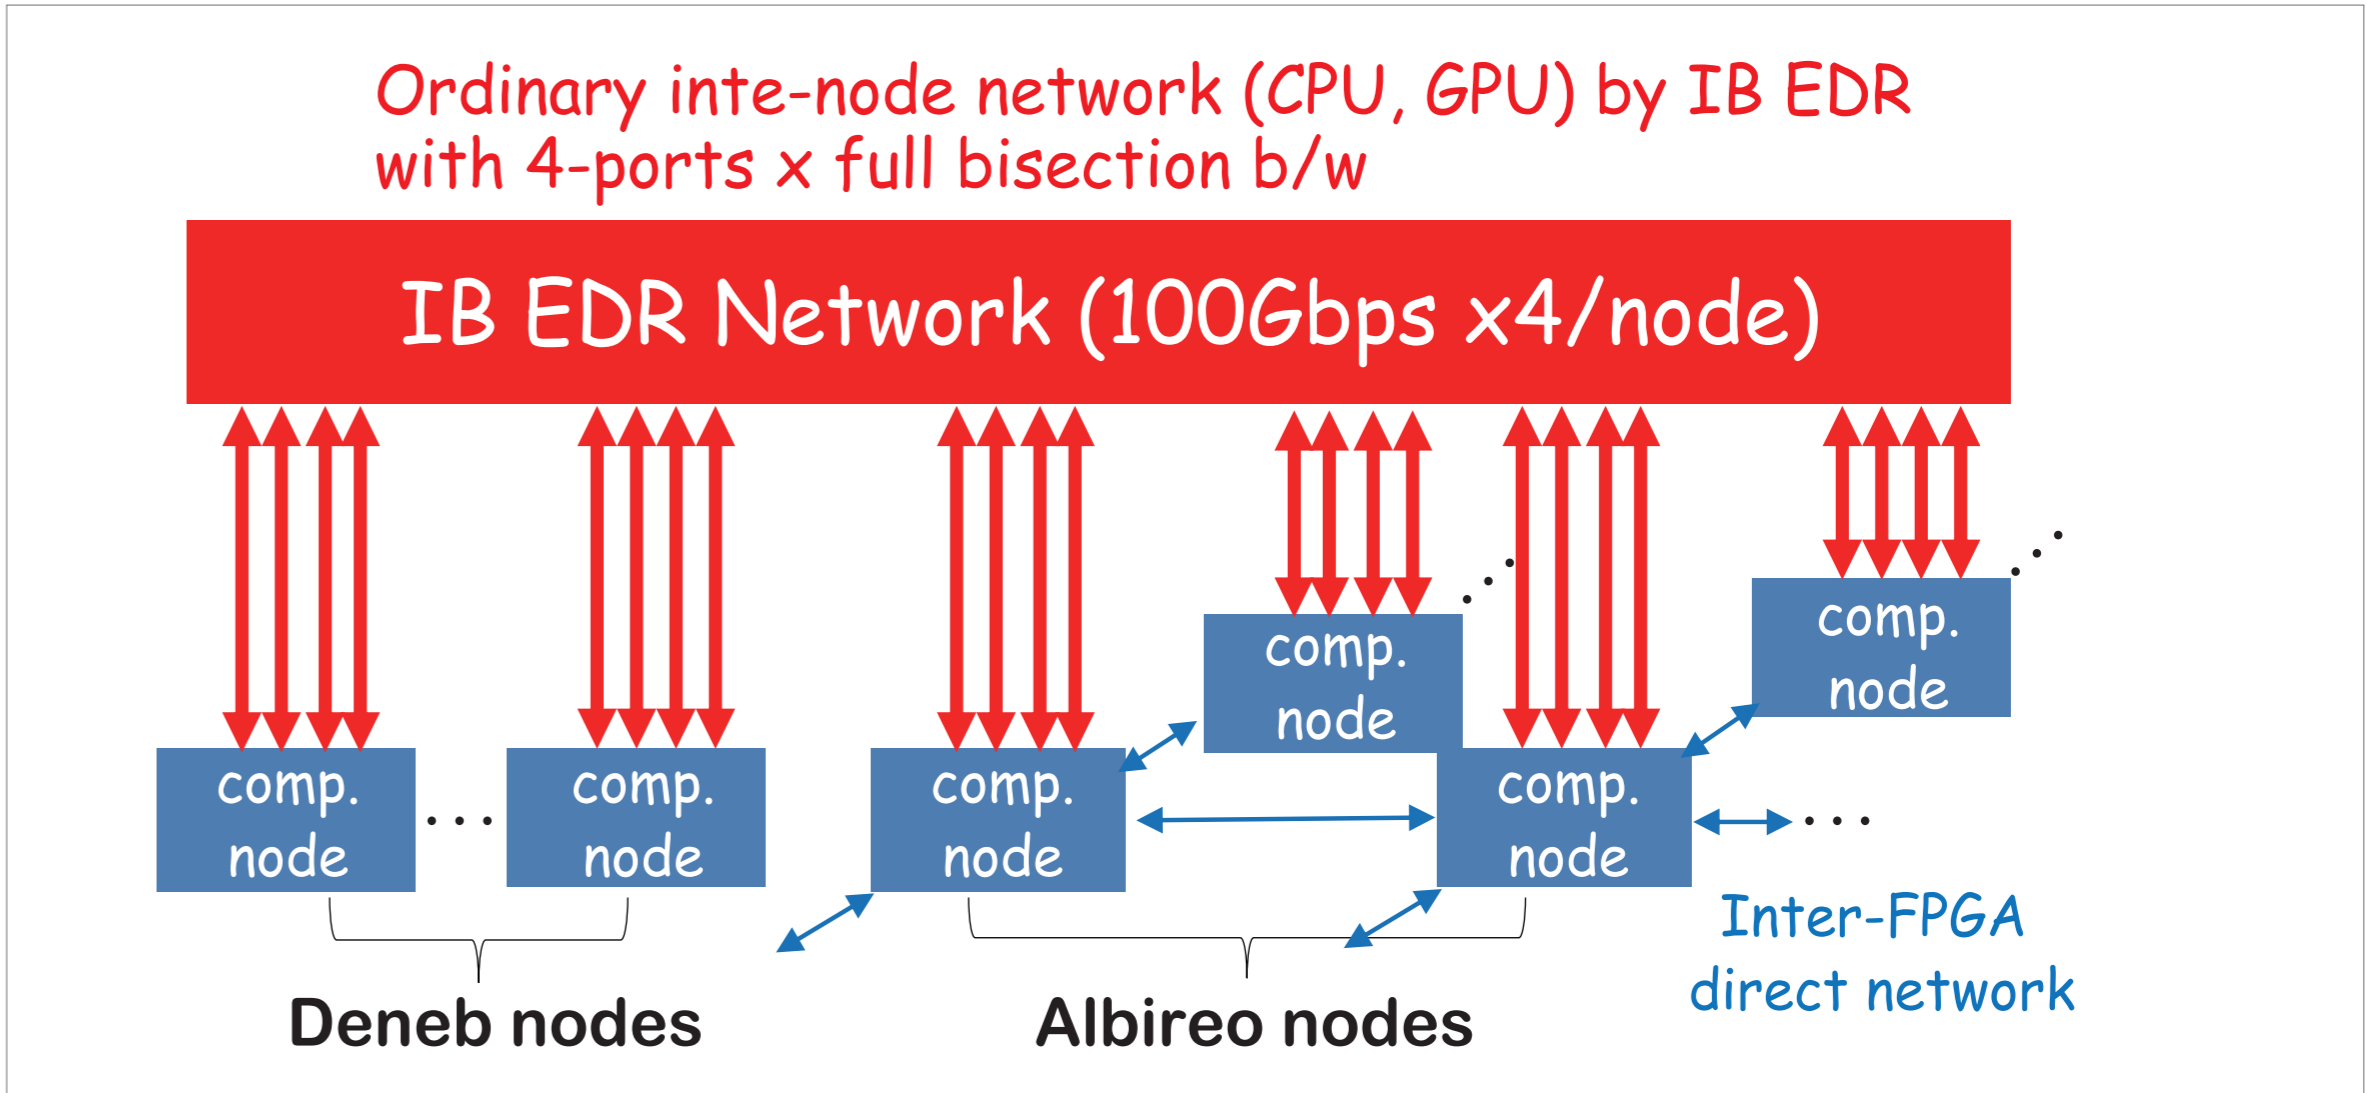
Item	Specification
Peak performance	<b>2.4 PFLOPS DP</b> (GPU: 2.2 PFLOPS, CPU: 0.2 PFLOPS, FPGA: 0.6 PFLOPS SP) ⇒ enhanced by mixed precision and variable precision on FPGA
#of nodes	<b>81</b> (32 Albireo (GPU+FPGA) nodes, 49 Deneb (GPU-onlu) nodes)
Memory	<b>192 GiB</b> DDR4-2666/node = <b>256GB/s</b> , <b>32GiB</b> x 4 for GPU/node = <b>3.6TB/s</b>
CPU / node	<b>Intel Xeon Gold</b> (SKL) x2 sockets
GPU / node	<b>NVIDIA V100</b> x4 (PCIe)
FPGA / node	<b>Intel Stratix10</b> x2 (each with <b>100Gbps x4 links/FPGA</b> and <b>x8 links/node</b> )
Global File System	Luster, RAID6, <b>2.5 PB</b>
Interconnection Network	Mellanox InfiniBand <b>HDR100</b> x4 (two cables of <b>HDR200 / node</b> ) <b>4 TB/s aggregated bandwidth</b>
Programming Language	CPU: C, C++, Fortran, OpneMP, GPU: OpenACC, CUDA, FPGA: OpenCL, Verilog HDL
System Vendor	NEC



Deneb node (x48)



Albireo node (x32)



Ordinary inte-node communication channel for CPU and GPU, but they can also request it to FPGA



Check it now!

“Cygnus Movie”