

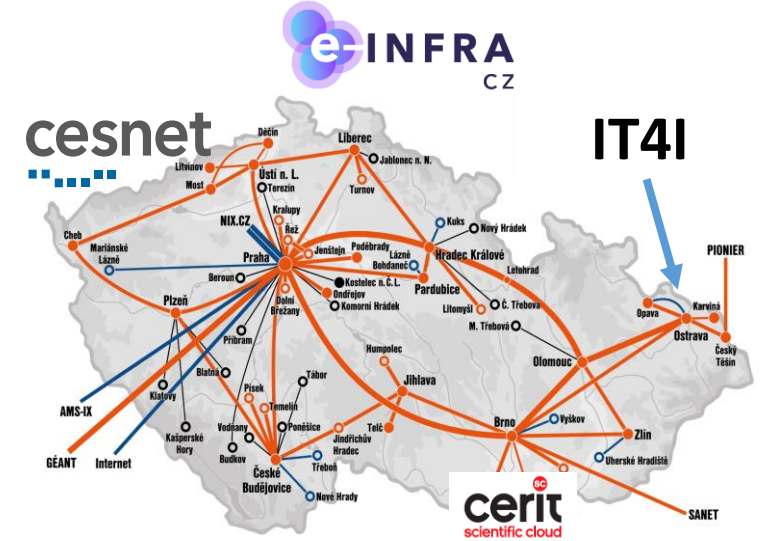
# IT4INNOVATIONS NATIONAL SUPERCOMPUTING CENTER

## CZECH REPUBLIC

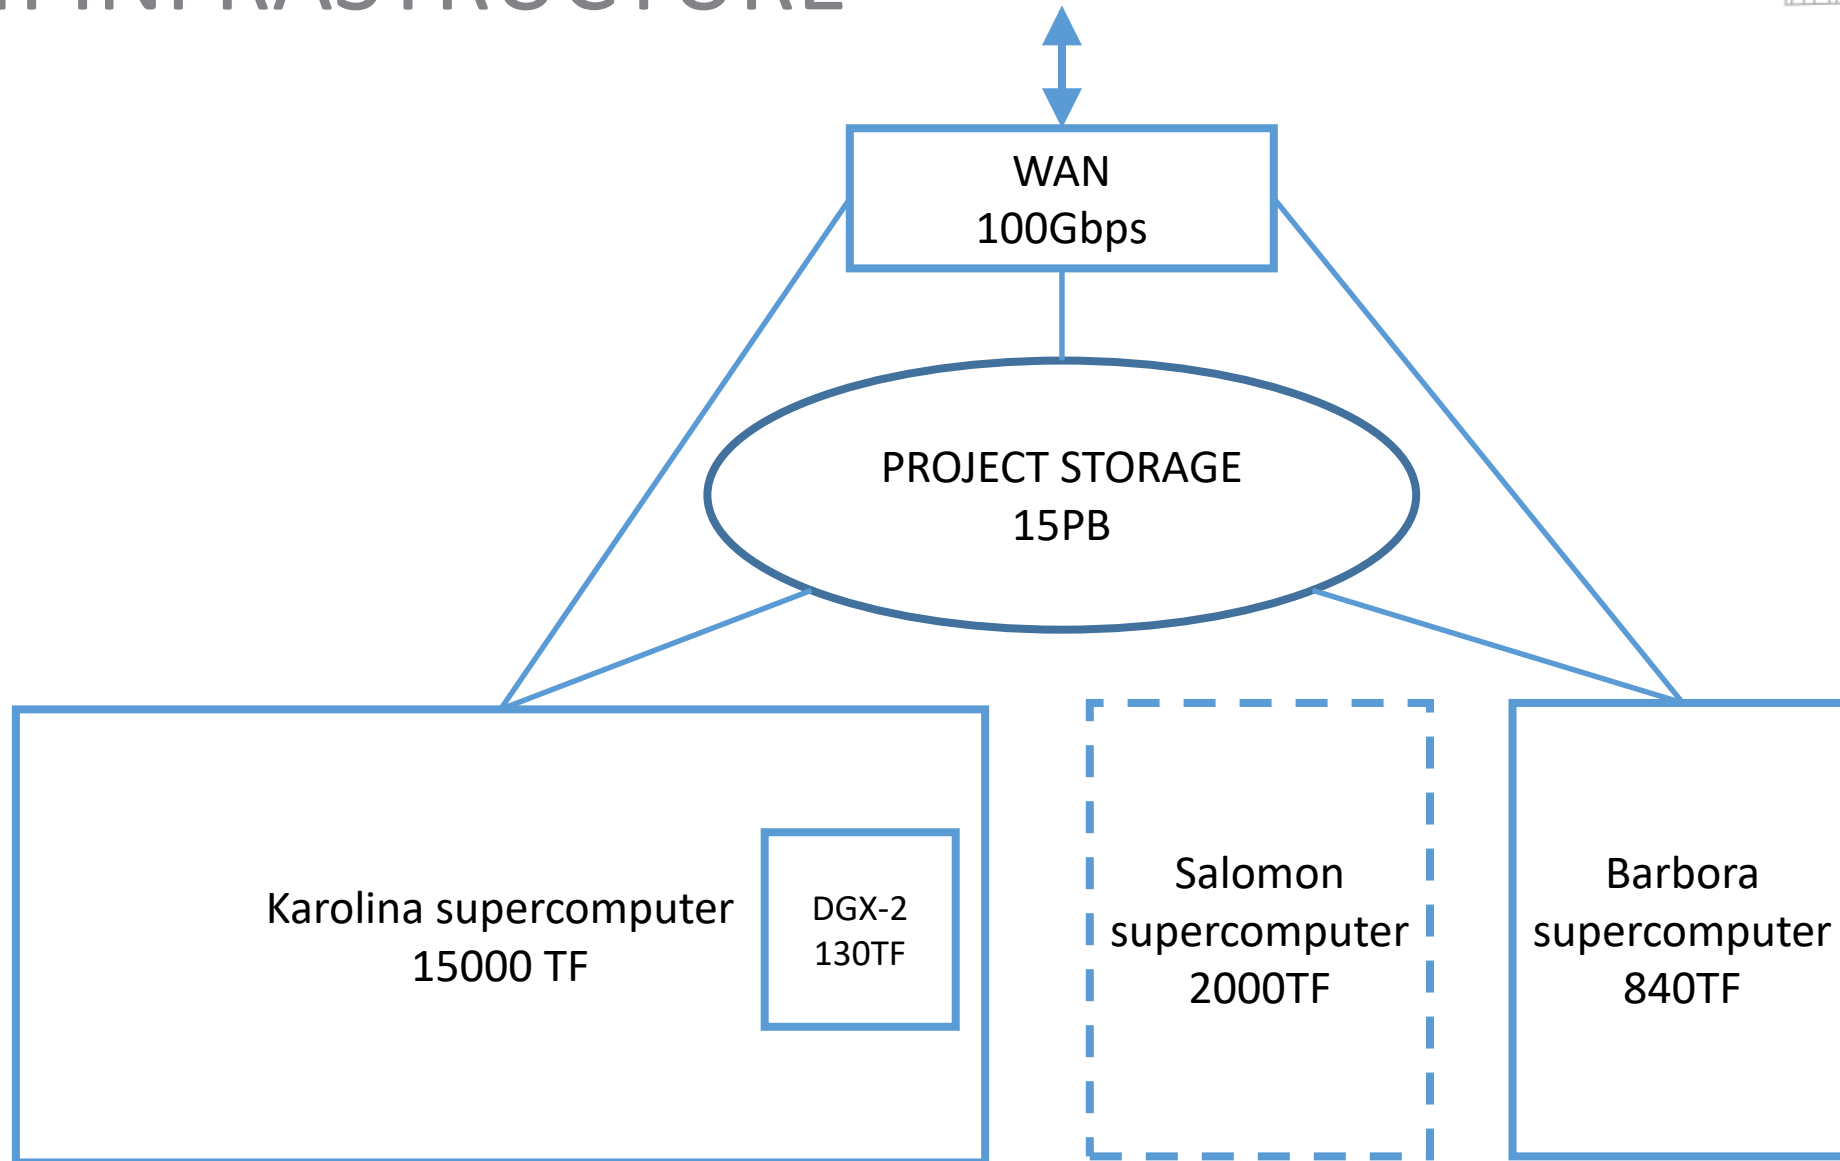


# IT4I INTRODUCTION

- Established in 2011 in Ostrava, Czech Republic
- Unit of the VSB – Technical University of Ostrava
- Member of e-INFRA CZ, a strategic research infrastructure
- Operating 3 supercomputers (Salomon, Barbora, Nvidia DGX-2)  
Provider of HPC resources for CR and EU
- 5 research laboratories, over 130 FTE
- Participating in EU HPC initiatives
  - EuroHPC, PRACE, EUDAT, ETP4HPC, BDVA
- Strong international collaboration, 14 H2020 projects, cooperation with industry
- Training and educational activities



# IT4I INFRASTRUCTURE



# NVIDIA DGX-2




- Intel Xeon Platinum 8168 processor, 2x24, AVX-512
- 1.5 TB RAM, 512GB HBM
- 16x2560 Volta V100 GPGPU
- Unified Address space
- NVME SSD storage 30TB
- **130TF Peak!**



# THE BARBORA SUPECROMPUTER



- **192x Compute nodes**
- **1x SMP node**
- **8x GPU nodes, 4x Nvidia V100**
- **Infiniband HDR network**
- **SCRATCH storage**  
Burst buffer, 200TB, 28GB/s
- **14x NVMe, accesible remotely**
- **2x Remote vizualization, NVidia Quadro**
- **840TF Peak**

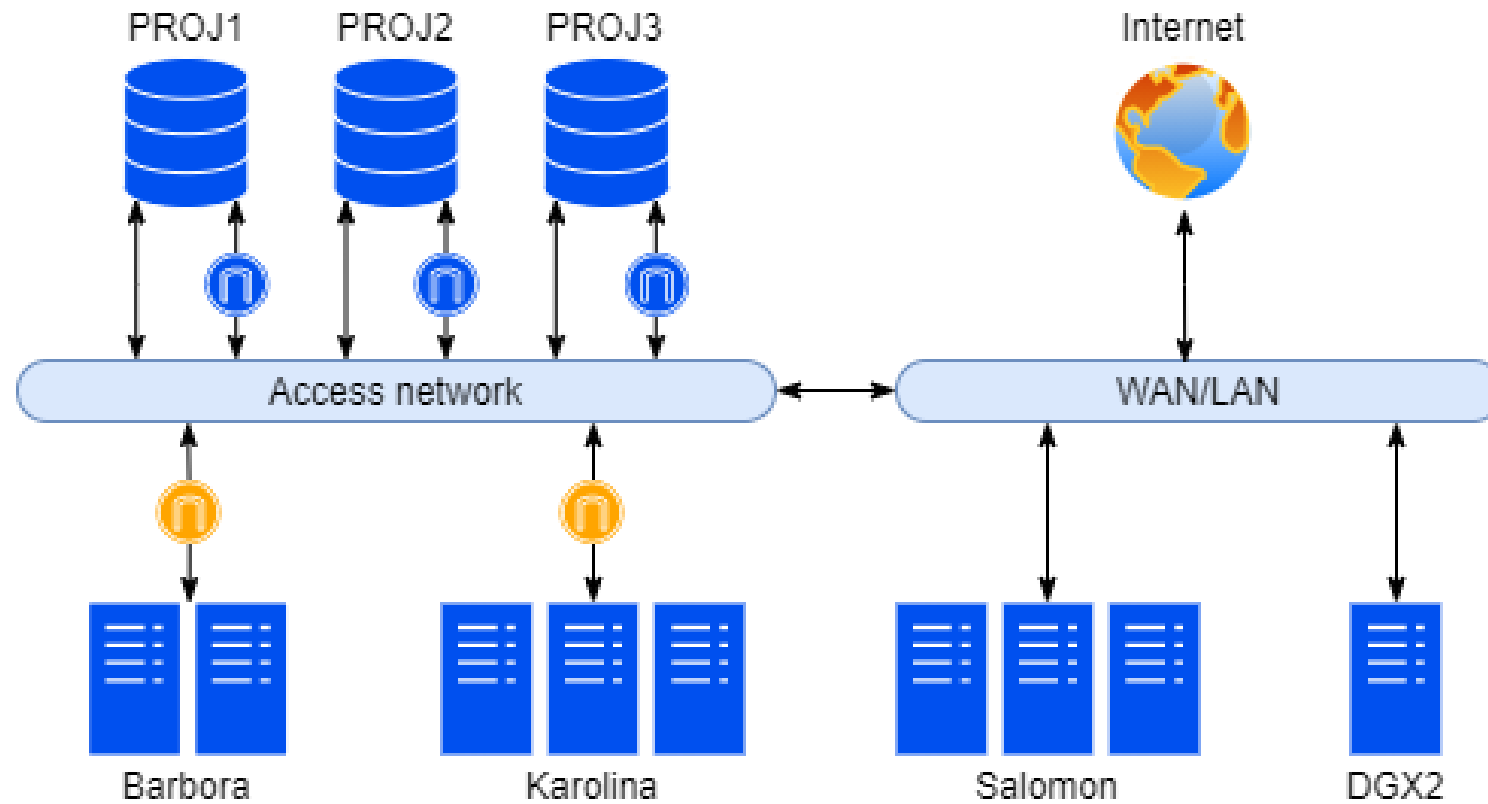


# THE PROJECT STORAGE





- Independent
- Extendable
- Scalable
- Redundant
  
- 3x4.9 PB
- 12GB/s aggregated
- NFS protocol
- Data gateways (GridFTP, RSYNC, etc)

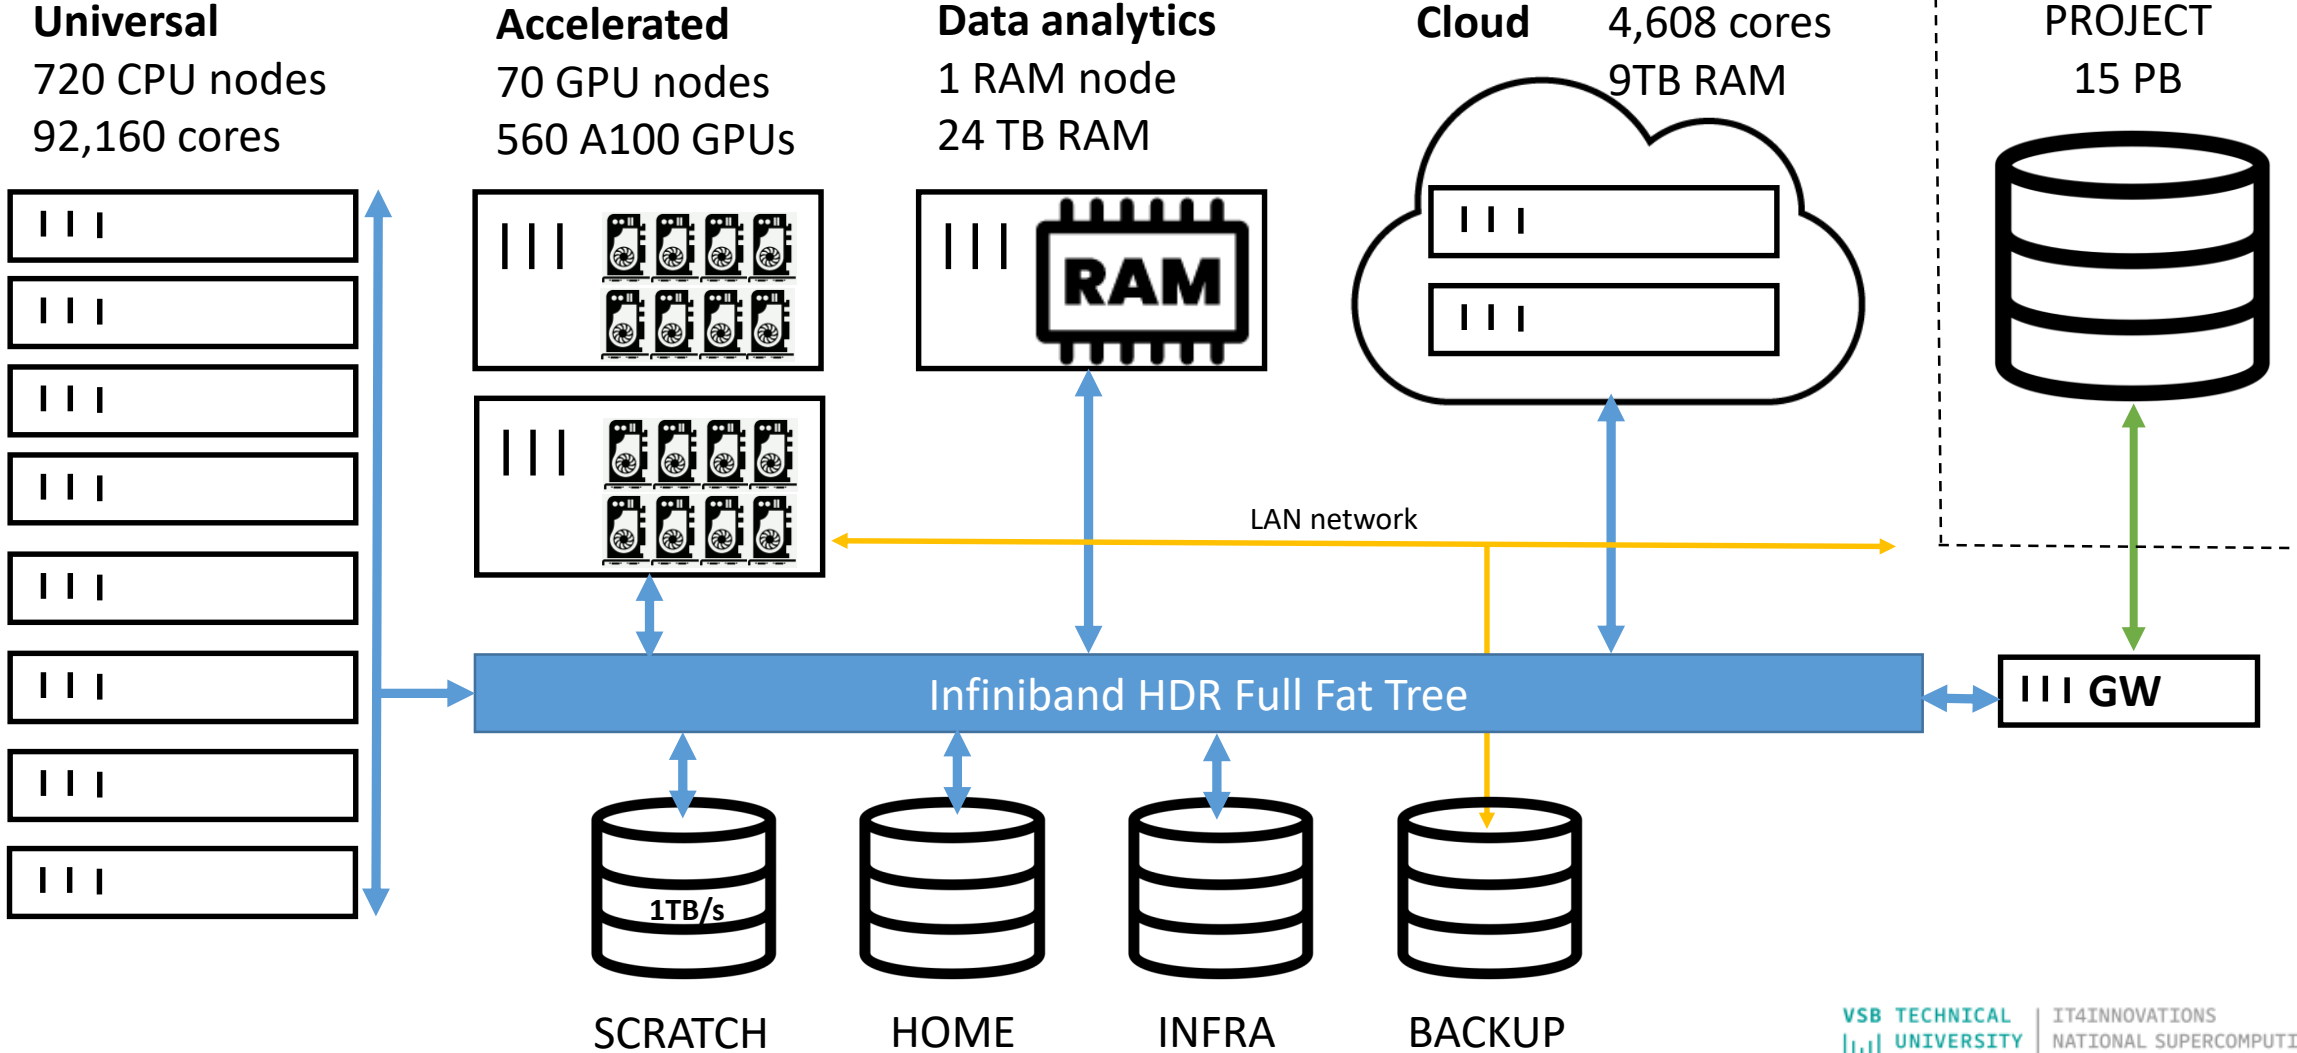
# THE PROJECT STORAGE



- 3x3 GPFS/NFS servers
- 3x2 Data gateways protocol
- 3x7 Disk arrays IBM Storwize V5030E
- 3x39TB SSD for small files

 : Data gateways  
 : Network gateways

# KAROLINA ARCHITECTURE





# UNIVERSAL PARTITION





- **700x HPE Proliant XL225n server 1x SMP node**
- **2x AMD EPYC 7H12, 2x64 cores**
- **256GB RAM DDR4**
- **100Gb/s (HDR100)**
- **CentOS 7**
- **5.3 TF Peak**



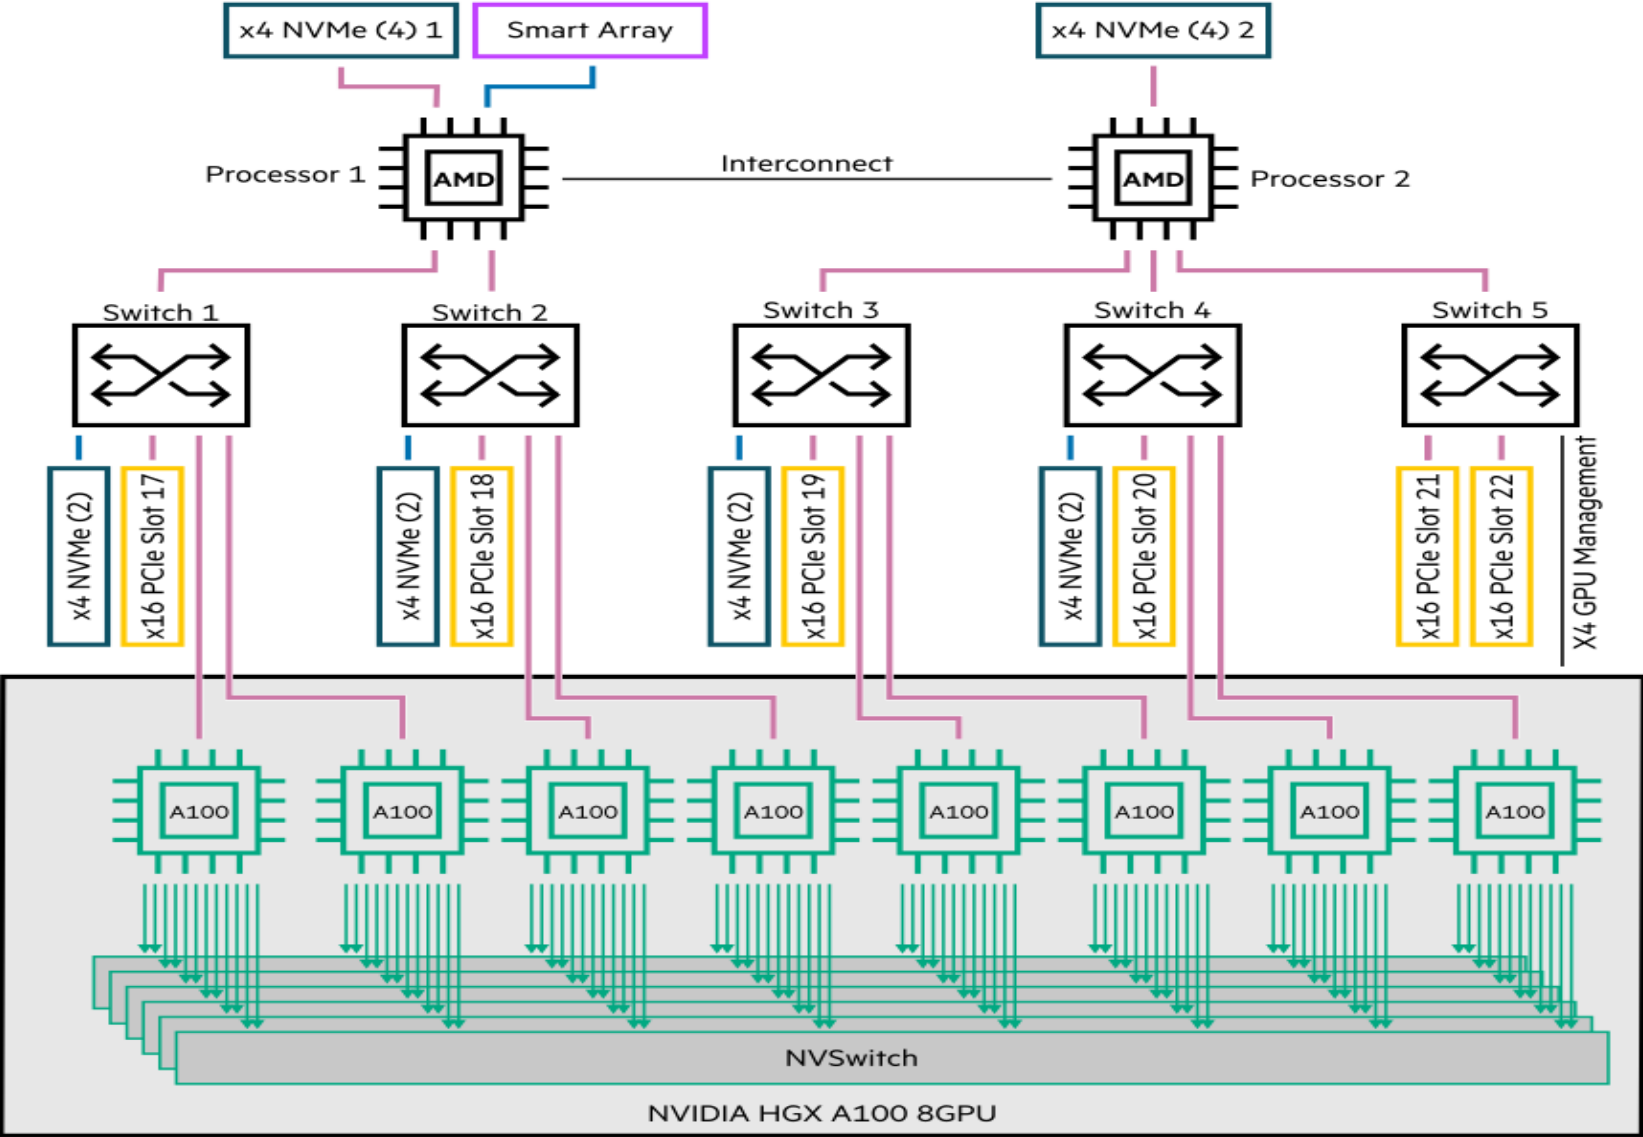
# ACCELERATED PARTITION



- **70x HPE Apollo 6500 G10+**
- **2x AMD EPYC 7452, 2x32 cores**
- **512GB RAM DDR4**
- **4x200Gb/s HDR**
- **CentOS 7**
- **158.4 TF Peak**



# NVIDIA A100 GPU



**Legend:**


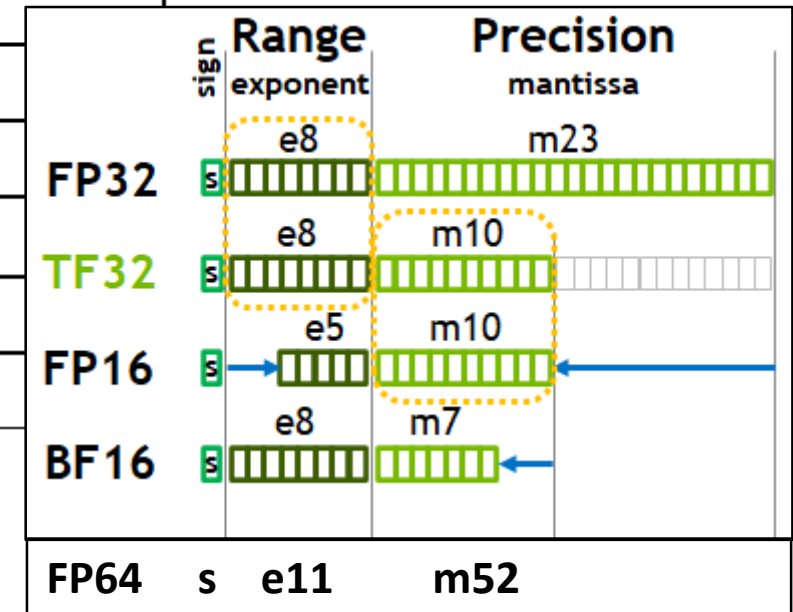
- x16 Gen 4
- x8 Gen 4
- 600 GB/s NVLink

**600GB/s GPU to GPU bandwidth**


# NVIDIA A100 GPU, 108SM





Peak FP64 <sup>1</sup>	9.7 TFLOPS
Peak FP64 Tensor Core <sup>1</sup>	19.5 TFLOPS
Peak FP32 <sup>1</sup>	19.5 TFLOPS
Peak FP16 <sup>1</sup>	78 TFLOPS
Peak BF16 <sup>1</sup>	39 TFLOPS
Peak TF32 Tensor Core <sup>1</sup>	156 TFLOPS   312 TFLOPS <sup>2</sup>
Peak FP16 Tensor Core <sup>1</sup>	312 TFLOPS   624 TFLOPS <sup>2</sup>
Peak BF16 Tensor Core <sup>1</sup>	312 TFLOPS   624 TFLOPS <sup>2</sup>
Peak INT8 Tensor Core <sup>1</sup>	624 TOPS   1,248 TOPS <sup>2</sup>
Peak INT4 Tensor Core <sup>1</sup>	1,248 TOPS   2,496 TOPS <sup>2</sup>



# DATA ANALYTICS PARTITION



- 1xHPE Superdome Flex
- 32x Intel Xeon 8268, 32x24 (768 cores)
- 24576GB RAM DDR4
- 2x200Gb/s HDR
- RedHat 7
- 71 TF Peak



# COMPUTE NETWORK

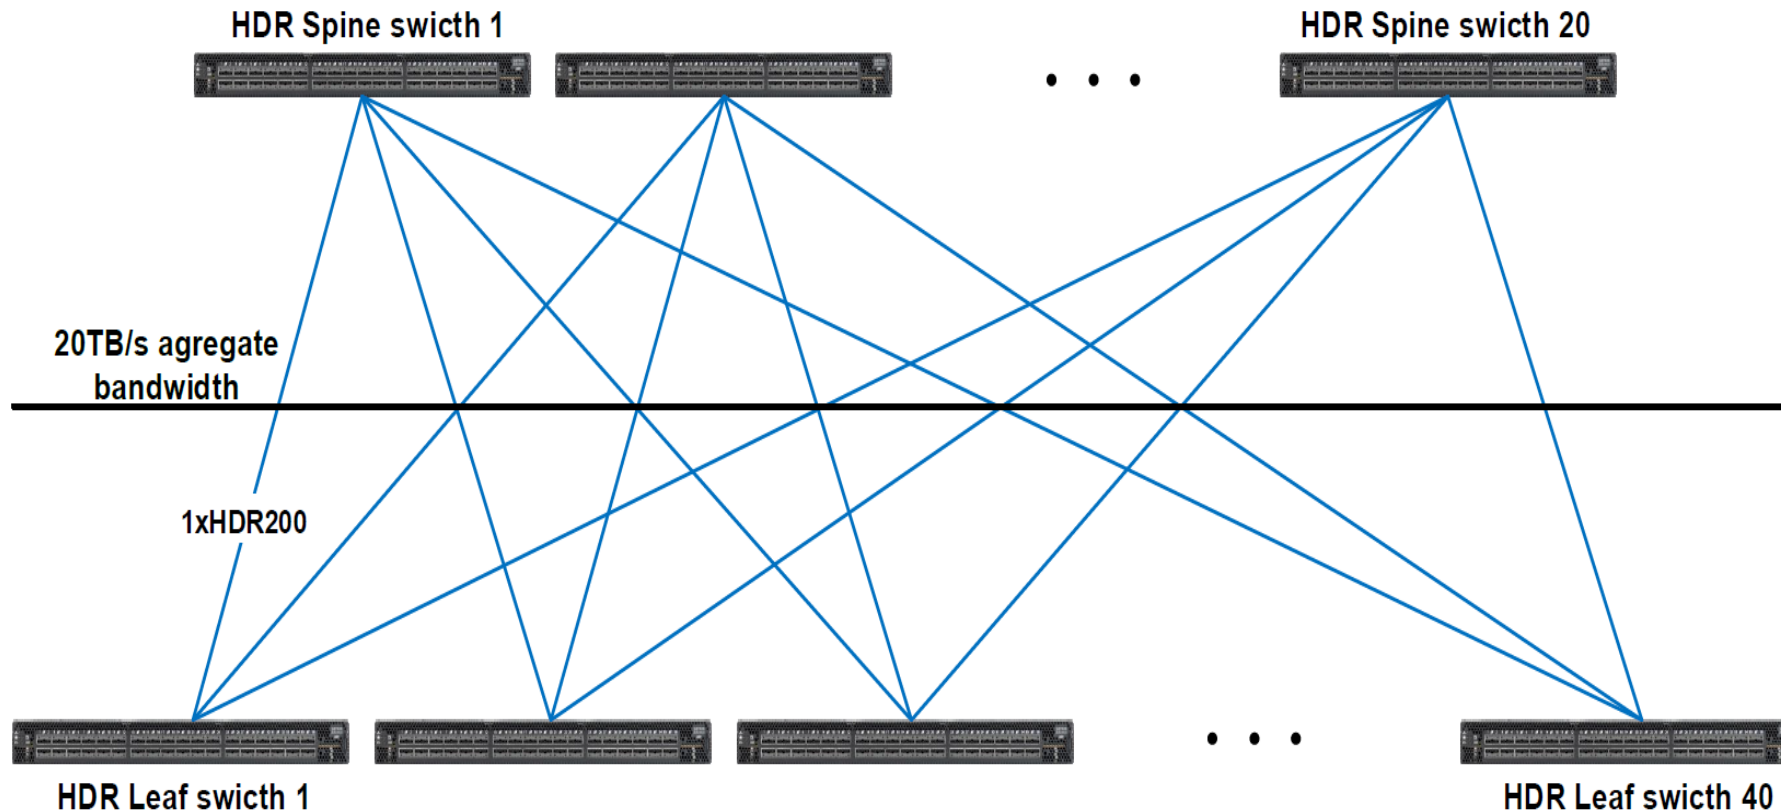


**Technology: HDR**

**Topology: Non-Blocking Fat Tree**

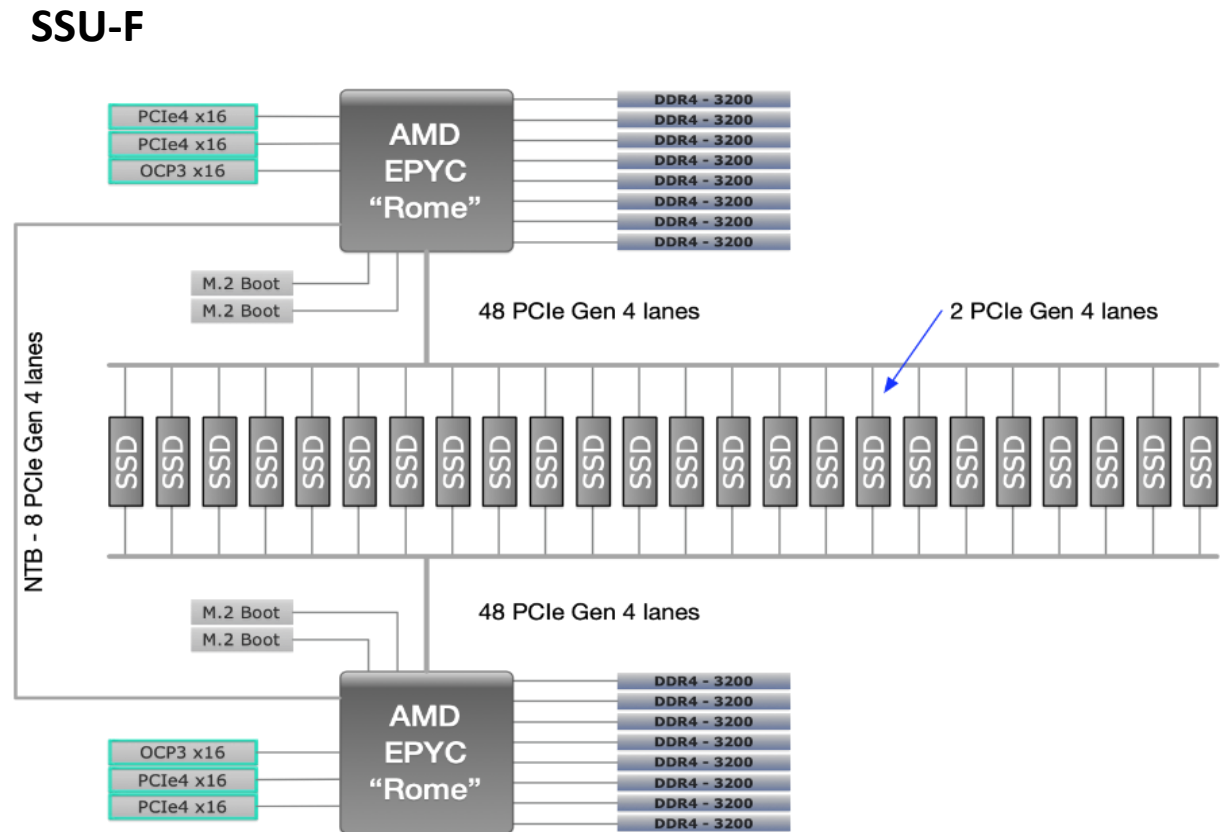
**Throughput: 200Gb/s for HDR200 connection,  
100Gb/s for HDR100 connection**

**Latency: Expected less than 3 microseconds**



# SCRATCH STORAGE

- ClusterStor E1000 All Flash
- 1xSMU (system mgmt)
- 1xMDU (metadata ctl)
- 24xSSU-F (storage unit)
- Size 1000TB
- Throughput 1000GB/s All flash
- LUSTRE Filesystem




# KAROLINA EXPECTED PERFORMANCE



## Performance to be installed:

- **R\_Peak: 15.2 PFlop/s**
- **R\_Max: 9.1 PFlop/s (LINPACK)**
- **R\_AI: 350 PFlop/s (DeepLearning)**
  
- **Universal partition: 2.3 PFlop/s (LINPACK) (720 nodes)**
- **Accelerated partition: 6.6 PFlop/s (LINPACK) (70 nodes)  
350 PFlop/s (DeepLearning)**
- **Data analytics partition: 40 TFlop/s (LINPACK)**
- **Cloud partition: 131 TFlop/s (LINPACK) (36 nodes)**




## Estimated **TOP 500** ranking:

- **Estimated ranking (1H2021):  
around #40 (worldwide) #10 (Europe)**



# KAROLINA TIMELINE



Partial Acceptance

Final Acceptance

December 2020

January 2021

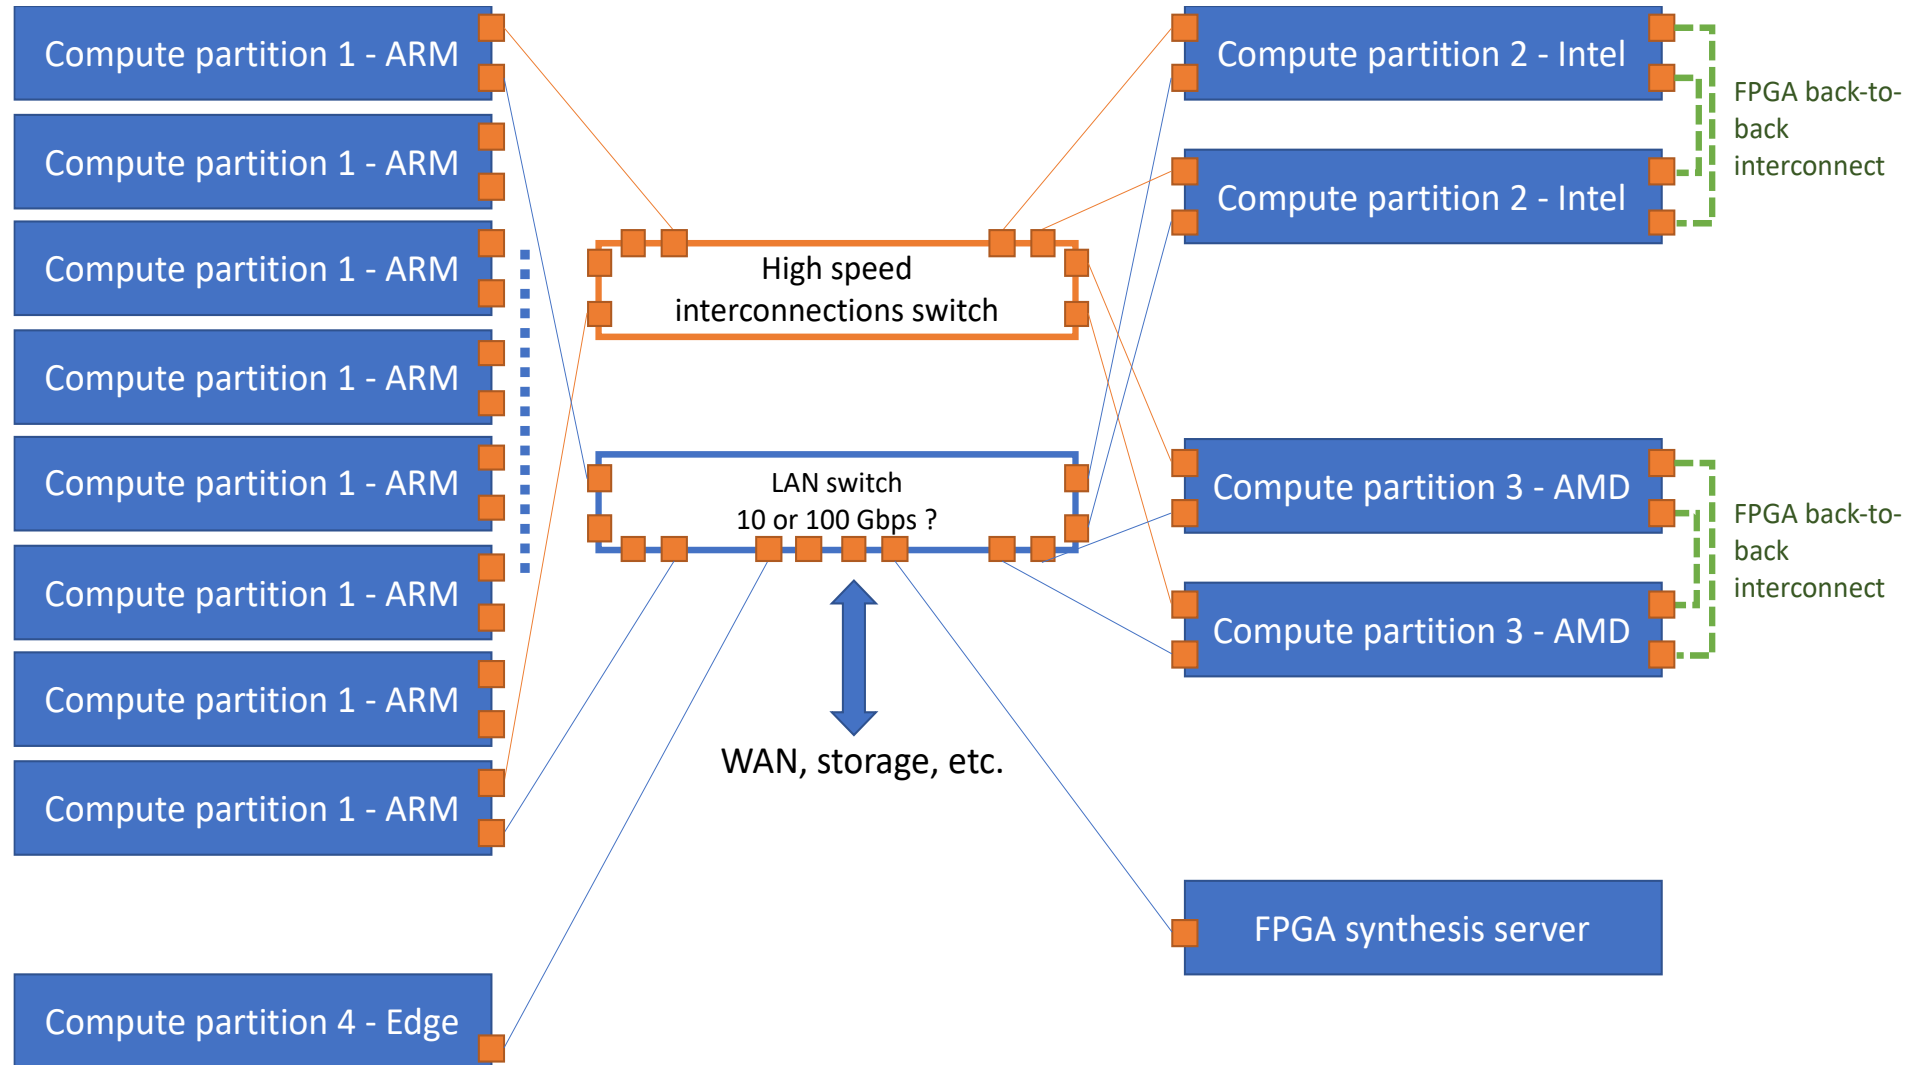
February 2021

April 2021

June 2021



# COMPLEMENTARY SYSTEM I




# LUMI CONSORTIUM

- | Unique consortium of 10 countries with strong national HPC centers
- | The resources of LUMI will be allocated per the investments
- | The share of the EuroHPC JU (50%) will be allocated by a peer-review process (cf. PRACE Tier-0 access) and available for all European researchers
- | The shares of the LUMI partner countries will be allocated by local considerations and policies – seen and handled as extensions to national resources



# DATACENTER IN KAJAANI



100% hydroelectric energy up to 200 MW

Very reliable power grid: Only one 2 min outage in 38 years

100% free cooling available, PUE 1.03



Waste heat reuse: effective energy price 35 €/MWh,  
negative CO<sub>2</sub> footprint: 13500 tons reduced every year

Extreme connectivity: Kajaani DC is a direct part of the Nordic backbone.  
4x100 Gbit/s to GÉANT in place, can be easily scaled up to multi-terabit level

Elevated security standards guaranteed by ISO27001 compliancy




# DATACENTER IN KAJAANI



- | LUMI will be an **HPE Cray EX** supercomputer manufactured by **Hewlett Packard Enterprise**
- | Peak performance over **550 petaflop/s** makes the system one of the world's fastest
  - | Fastest today is Fugaku supercomputer in Japan with 513 petaflop/s, second fastest Summit in USA with 200 petaflop/s)

1 system  
**550**  
**Pflop/s**  
Peak Performance


Computing power  
equivalent to  
**1 500 000**  
Modern laptop computers



Size of a tennis court

Modern platform for  
High-performance  
computing,  
Artificial intelligence,  
Data analytics  
Based on GPU technology

# LUMI SUPERCOMPUTER

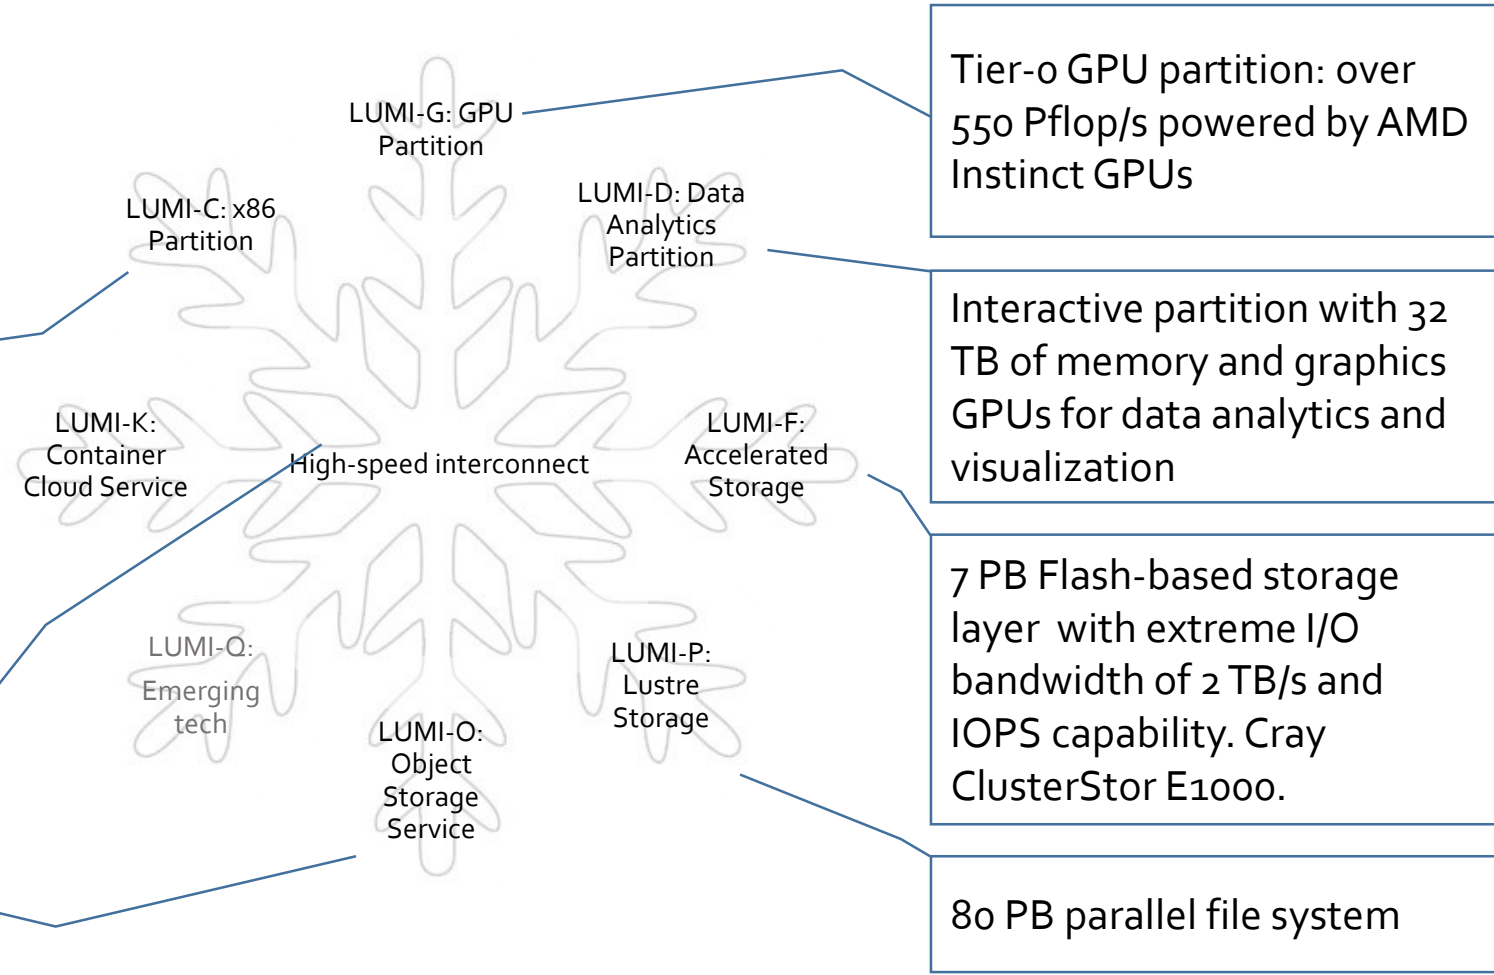


*LUMI is a Tier-0 GPU-accelerated supercomputer that enables the convergence of high-performance computing, artificial intelligence, and high-performance data analytics.*



- Supplementary CPU partition
- ~200,000 AMD EPYC CPU cores

Possibility for combining different resources within a single run. HPE Slingshot technology.


30 PB encrypted object storage (Ceph) for storing, sharing and staging data



# LUMI NODE



2 nodes per blade



COPYRIGHT 2020 HPE

# SUMMARY



- IT4INNOVATIONS – Czech national supercomputing center
- **Karolina EUROHPC supercomputer** - 9.1 PFlop/s Linpack
- Massively accelerated - 8x Nvidia Ampere A100 per node (6.6 PFlop/s)
- Partial acceptance April 2021, full acceptance June 2021
- LUMI – Most powerful supercomputer in Europe, Dedicated share for CR

Our supercomputers support science, industry, and society





Branislav Jansík  
branislav.jansik@vsb.cz

IT4Innovations National Supercomputing Center  
VSB – Technical University of Ostrava  
17. listopadu 2172/15  
708 00 Ostrava-Poruba, Czech Republic  
www.it4i.cz

