

# Hardware Platform for the SPD experiment

Andrey Kiryanov, NRC KI – PNPI

IX SPD collaboration meeting, 12-16 May 2025

# Offline computing baselines

- No need for MPI, RDMA and other HPC stuff
- No need for GPUs and AI-optimized hardware
- No need for hardware RAID (EOS and Ceph can handle redundancy)
- We are mostly interested in performance and memory throughput **per core**
- AMD platforms are the most promising because of the memory subsystem performance
  - <https://www.servethehome.com/memory-bandwidth-per-core-and-per-socket-for-intel-xeon-and-amd-epyc/>
- Fat tree network topology,

# Worker node

- Compute server 2U
  - 2Twin2 (4 nodes), 4 x 96 cores (UP) @3.1GHz, 4 x 384GB RAM, 4 x (2x10GbE, 2x25GbE)
  - 12 DIMM slots, 12x32=384 GB, 96 cores per node (4 GB per core)
  - 2x2000W PSU, expected load >1500W per chassis
- Switches
  - Spine: 32x100GE QSFP28, 4x200GE QSFP56
  - Leaf: 48x25GE SFP28, 6x100GE QSFP28
- Rack 42U
  - 12 server with 1U interval
  - TOR switch 1U
  - 4 608 cores per rack
- For 20000 cores we need 4-5 racks (52 servers)
  - 450 MRub

# Storage server

- Storage server 4U
  - 2 x 28 cores @2.0GHz, 512GB RAM, 2x10GbE, 2x100GbE
  - 16 DIMM slots, 3 PCI-e x16 slots, 2 NVMe slots
  - 60 3.5" SAS hot-swap bays, with 28TB per disk total capacity is 1680TB
  - 2x2000W PSU
- Switches
  - Spine/Leaf: 32x100GE QSFP28, 4x200GE QSFP56
- Rack 42U
  - 8 server with 1U interval
  - TOR switch 1U
  - 13.4PB per rack
- For 100PB we need 7-8 racks (60 servers)
  - 480 MRub