



# **PIK Computing Centre**

### Equipment, capabilities and current status

Andrey Kiryanov







## **General information**

- PIK Computing Centre is located at the ground floor of one of the buildings
- Computer Centre rooms are green, other technical rooms are blue
- Outside we have three chillers, coolant tanks and a relatively small diesel generator



NATIONAL RESEARCH CENTRE «KURCHATOV INSTITUTE»



### **Chillers under the roof** (we've got lots of snow in winter)







## Power and cooling infrastructure

- Both main power and cooling can sustain up to 300 kVA of load
  - With maximum load the computing equipment can run on batteries for about 10 minutes, which is enough for a graceful shutdown
  - Cooling system has its own UPS and runs on batteries for about 30 minutes
  - Diesel powers the cooling system pumps if everything else fails
  - No guaranteed power in case of a major failure, but power line is redundant
- Currently the Computing Centre equipment is worth 120 kVA of load
  - Roughly 40 minutes of runtime on batteries
  - Over 50% of rack space is unused (12 racks populated out of 28)
  - Network infrastructure is designed for full capacity (except InfiniBand)
  - We can move in new servers without any modifications to the infrastructure



«KURCHATOV INSTITUTE»



### CCV and monitoring (still in development)





# **Computing equipment**

- Peak theoretical performance is ~362 Tflops
- Real LINPACK results:
  - ~200 Tflops on Xeon CPUs (no AVX-512), effectiveness ~80%
  - ~68 Tflops on Xeon Phi (KNL) CPUs (AVX-512), effectiveness ~50%
- Computing equipment:
  - 160 nodes with Xeon CPUs: 2.4 GHz, 28 cores, 128 GB RAM per node (4.5 GB RAM per core) – 4 480 cores
  - 40 nodes with Xeon Phi (KNL) CPUs: 1.4 GHz, 68 cores (272 virtual), 112 GB RAM per node – 10 880 virtual cores
  - 16 nodes with Xeon CPUs: 2.4 GHz, 28 cores, 1 TB RAM + 1.6 TB NVMe SSD
  - 2 nodes with Xeon CPUs: 2.4 GHz, 28 cores, 1.5 TB RAM



NATIONAL RESEARCH CENTRE «KURCHATOV INSTITUTE»



## **CACS with HACS**









#### http://top50.supercomputers.ru/

6	Москва Центр обработки данных НИЦ "Курчатовский институт" 2015 г.	774/11082	узлов: 364 (2xXeon E5-2680v3 2.5 GHz 128 GB RAM) узлов: 23 (2xXeon E5-2680v3 [Acc: 3x NVidia K80] 2.5 GHz 128 GB RAM) оеть: Infiniband FDR/Gigabit Ethernet/Gigabit Ethernet	374.13	500.55	SuperMicro, Борлас
7	Нижний Новгород Нижегородский государственный университет им. Н.И. Лобачевского 2014 г.	360/30760	узлов: 100 (2xXeon E5-2660 [Acc: 3x Kepler K20X] 2.2 GHz 65.536 GB RAM) узлов: 50 (2xXeon E5-2660 [Acc: 2x Fermi 2090] 2.2 GHz 65.536 GB RAM) узлов: 10 (2xXeon E5-2660 [Acc: 2x Phi 5110P] 2.2 GHz 65.536 GB RAM) узлов: 10 (2xXeon E5-2660 [Acc: 3x Fermi 2090] 2.2 GHz 65.536 GB RAM) узлов: 10 (2xXeon E5-2660 2.2 GHz 65.536 GB RAM) сеть: Infiniband QDR/Infiniband QDR/Gigabit Ethernet	289.50	573.00	Ниагара Компьютерс
8	Челябинск Южно-Уральский государственный университет 2013 г.	768/28032	узлов: 384 (2xXeon X5680 [Acc: Xeon Phi SE10X] 3.33 GHz 24.576 GB RAM) оеть: Infiniband QDR/Infiniband QDR/Gigabit Ethernet	288.20	473.64	Группа компаний РСК
9	Гатчина ОИВРТ НИЦ "Курчатовский институт" - ПИЯФ 2017 г.	496/9024	узлов: 160 (2xXeon E5-2680v4 2.4 GHz 128 GB RAM) узлов: 40 (Xeon Phi 7250 1.4 GHz 112 GB RAM) узлов: 30 (2xXeon E5-2680v4 2.4 GHz 256 GB RAM) узлов: 20 (2xXeon E5-2650v4 2.2 GHz 256 GB RAM) узлов: 16 (2xXeon E5-2680v4 2.4 GHz 1024 GB RAM) узлов: 2 (2xXeon E5-2680v4 2.4 GHz 1540 GB RAM) ость: Infiniband EDR/Gigabit Ethernet/Gigabit Ethernet	200.44	362.38	NP-IT, Ниагара компьютерс





### **UPS room**







## Storage and network

- Storage systems
  - Lustre with 2.9 PB of raw disks (~2.3 PB of visible storage + 29 TB of metadata)
  - Ceph with 2.5 PB of raw disks (two full racks)
- Networks
  - 100 Gbps InfiniBand EDR interconnect
  - Ethernet network 1-10 Gbps
  - Management network for IPMI, etc.
  - Redundant external 10 Gbps connection





### Lustre disk boxes and servers









- Main OS: CentOS 7, 64-bit
  - Management: xCAT
  - Batch system: Slurm
  - Monitoring: Ganglia
  - MPI: Open MPI, Intel MPI, Platform MPI
  - Compilers: GCC, Intel Compiler
- Infrastructure
  - Schneider DCE & DCO
  - GENESIS64 SCADA





## InfiniBand switch and Ceph servers







## **Current status**

- All the equipment is up and running
- We're in the storage configuration phase
  - Application software will be deployed as soon as the storage is ready and reliable
- External connection is still 1 Gbps
  - Will be upgraded to 10 Gbps in the coming months
- PIK Computing Centre is expected to be open for users after Christmas



NATIONAL RESEARCH CENTRE «KURCHATOV INSTITUTE»



## **DMZ switches and routers**









# Thank you!



