

# Complex for mega-science data modelling and processing

V. Velikhov, E. Ryabinkin, V. Ilyin

NRC "Kurchatov Institute"

Mega-science complex

## Quick overview of current facilities

Our centre at a glance:

- More than 50 people.
- 4 hosting rooms with  $\approx 750$  KW of power/cooling each.
- Tier-2 and Tier-1 centers for WLCG: ALICE, ATLAS and LHCb are supported at the full scale, CMS support is present just for T2.
- 3 production HPC clusters ( $\approx 300$  TFLOPS of total computing power), various smaller facilities.
- 2 new HPC clusters this year (adding  $\approx 1.1$  PFLOPS).



Our data center currently almost meets Tier II criteria. We have

- two independent power feeds of 7.5 KW each,
- uninterruptible power supplies which will run the centre at its full capacity for 1/2 hour,
- diesel generator with fuel for 24 hours of operation (can be refilled on the fly),
- both precision and in-rack cooling,
- 24x7 infrastructure-on-duty and monitoring systems for all engineering equipment.

We're doing various routine exercises, such as going to the diesel generator bi-monthly, turning off one of the power feeds, switching to UPS to measure for how long it will run our current capacities.

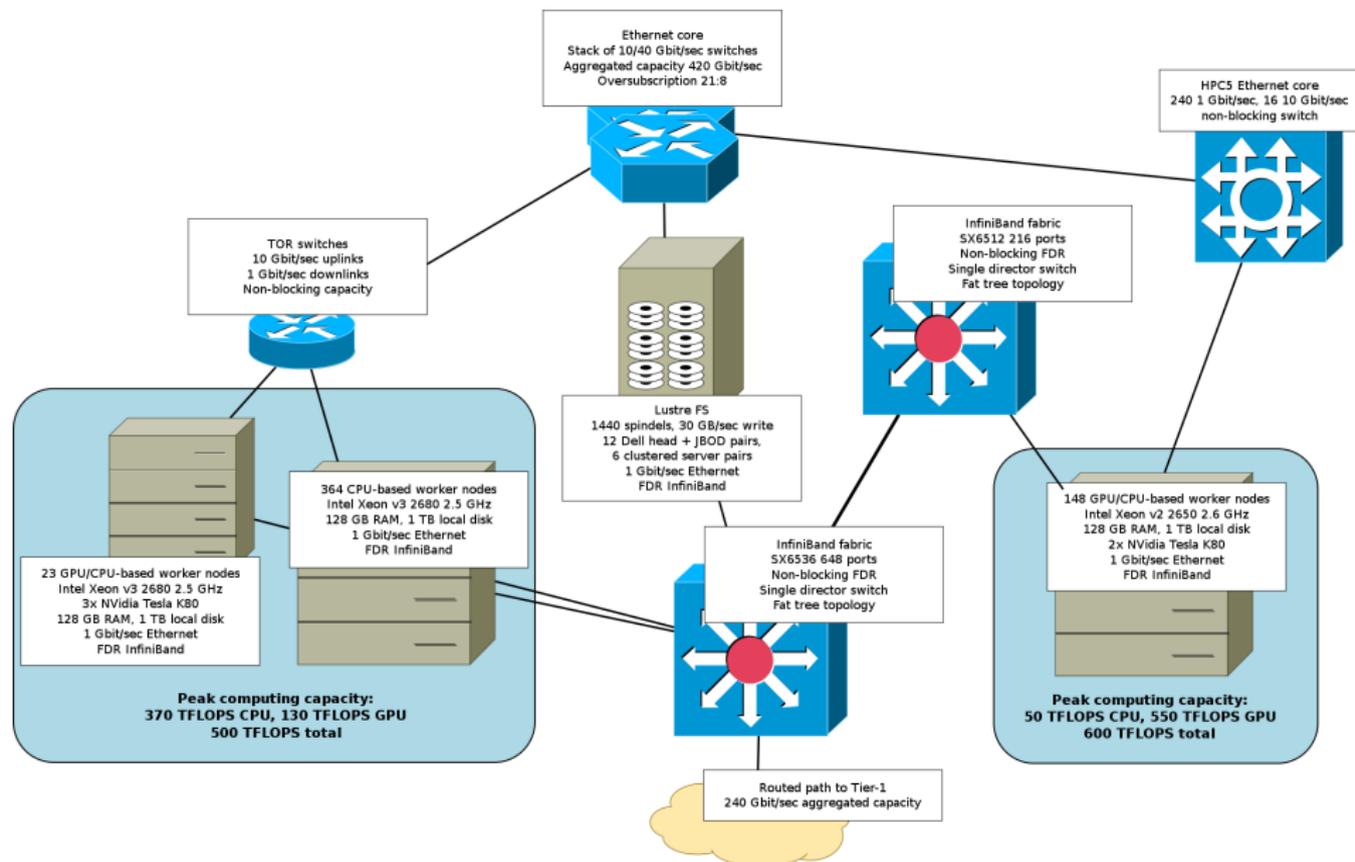
One of the recent findings was that not all incoming power feeds are born equal and feeds that supply 10 KA from 100 KA generators are better (in Russia) than those who supply 10KA directly.

So, we're trying to change our secondary feed to 100/10 KA to make situation to be symmetric.

Later today there will be more complete talk about our Tier-1, but we're glad that since 2003 it progressed a lot.



# New HPC machines: 4<sup>th</sup> and 5<sup>th</sup> generation



We run many types of workload. Some of them in a no particular order:

- modelling of ship body dynamics;
- computational astrophysics;
- modelling of biological objects;
- DNS resequencing and analysis;
- studies of condensed matter;
- ATLAS and (forthcoming) ALICE workloads;
- nuclear engine modelling;
- engineering tasks for ITER project;
- data mining for XFEL.

Therefore we're very happy that our new resources (HPC4/HPC5) are on-site and being currently deployed: users are waiting and can utilize them fully.



Two years ago Alexei Klimentov started his BigData laboratory in Kurchatov Institute and we were taken aboard to the software development (not only) for ATLAS.

Currently we participate in the following projects:

- adaptation of PanDA monitor to the NoSQL backends;
- data product catalog for ATLAS;
- development of federated storage (and not just storage).

It was (and is) a big supply of a fresh air for us and I want to say big “Thanks!” to Alexei and other people who make it possible.

We’re also involved in the XFEL project; a data mining group within our complex works on the classification of images obtained at this facility.

And since the most of software we running open-source software and believing in local expertise rather than in support contracts, we do a considerable amount of work to adopt these systems to our needs, to fix bugs and to try the new functionality.

Next two years will be devoted to get the most from our new HPC resources, to withstand Run-2 and to begin our preparations to Run-3 in all areas.

We got bigger than we were and the growth was large from the resource side, but not so many new people were added to our team. So we will need to improve in this area: it is the good and healthy team that makes any task to progress in a needed direction and that can explore new areas in any field.

Thanks for your attention!

Questions?

Suggestions?

Rotten tomatoes?