

Russian Tier-1: current status and plans

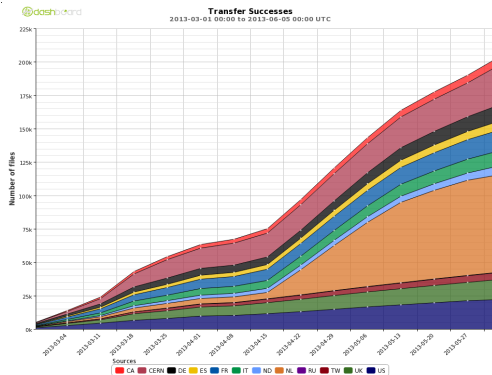
Eygene Ryabinkin, Alexandra Berezhnaya, Vasily Velikhov,
Viatcheslav Ilyin, Roman Kolchin, Yury Lazin, Ilya Lyalin,
Igor Tkachenko, Fedor Checherov

National Research Centre “Kurchatov Institute”

ALICE T1/T2/T3 meeting, Lyon

This presentation will outline the current status and plans for the Tier-1 (prototype) in NRC "Kurchatov Institute" that is being built to serve ALICE, ATLAS and LHCb.

First of all, let's talk about current status.




Currently we have the following hardware and service layout:

- 740 working CPUs with 4 GB RAM and 50 GB of local disk per core and 1440 more such CPUs are currently being installed;
- Disk-based storage: dCache (production) and EOS (almost ready) instances; 2.2 PB of disks are divided between these two;
- Tape: IBM TS3500 library;
- WLCG services: CREAM CE (2), APEL, BDIIs (site and top-level, both are clustered), VO-BOX, PerfSonar bandwidth and latency nodes;
- standard IT services: DBs (PostgreSQL, MySQL), firewalls, NATs, HTTP proxy farm (CVMFS, Frontier and friends), installation nodes, etc;
- batch system: currently based on Torque/Maui with our own patches; evaluating Slurm-based one.

Current status: networking

- Two network types, WAN and LAN split across different physical hardware sets;
- WAN is organised to carry $N \times 10$ Gbit/sec of external bandwidth with $N \approx 4 - 10$;
- LAN is designed to scale for 10^{th} of Tbit/sec: core is based on Extreme Networks Black Diamond X8 and active-active inter-switch links; storage is connected directly to the core devices, WNs and services are going through distribution switches;
- LAN is routed, but multicast-friendly (ALICE Torrents!);
- WAN is connected to GEANT (see next slide), to the Russian academic networks (10 Gbit/sec), to various research and educational networks and to LHC OPN;
- new NOC, additional networking people, own routers and autonomous system were brought in to support all of this.

- As you probably know, current Russian GEANT link was cut to 100 Mbit/sec due to NREN debt;
- nevertheless, we're transferring the data¹ from/to other sites all over the world:



- LHC OPN peering was established with 2 Gbit/sec channel; currently we're (still) working on the 10 Gbit/sec dedicated link to CERN;
- Our networking could be better (and we're trying hard to bump it), but we already built a good connectivity.

¹that cavity on tuesday was the time just after the cut

Current status: storage

- Disk instances carry directly-attached, iSCSI and SAS storage; after some experiments we had converged to plain directly-attached storage boxes with 10 Gbit/sec network and phasing out all other storage types;
- Standard 2-4 TB SATA disks, RAID 6 everywhere, storage nodes are loaded with 16 disks each;
- Tape library has 2500+ tape slots (up to 10 PB with IBM enterprise 3592 cartridges); current capacity – 2 PB, 3 TS 1140 drives;
- Current tape library software is Enstore; running it for some months and getting experience with plain Enstore and dCache/Enstore SE. Works mostly fine up to now.

Current status: people

- Infrastructure support team: 24x7 monitoring (2 people with rota; 7 people in total), support people for power and cooling infrastructure: 11 people.
- Hardware support team: currently we have 6 people who handle all replacements, diagnostics, etc for all our hardware (not only Grid-related one).
- Monitoring: one dedicated person to do monitoring in 12x7; and our fellow sysadmins do the monitoring (with the help from automated systems).
- 1 FTE for ALICE and ATLAS support persons who will mediate the sysadmin teams and VO-specific stuff; something like 0.5 FTE for LHCb support:
 - allocated Ilya Lyalin for the ATLAS business;
 - have a good understanding of what ALICE support person should do; still searching for one;
 - LHCb support person isn't yet determined, though his tasks are mostly understood.

Current status: what we actually already did

- ATLAS part of the current Tier-1 was built, tested and now it is working as the production resource (but without tapes);
- with ATLAS we did reprocessing of 2011's 2.76 TeV data, so our infrastructure was sufficiently stressed with real-world workload;
- reprocessing stats: 57 TB of input data received and processed, 37000 jobs were run, 900 were failed \Rightarrow 98% of jobs were successful.
- Site is almost always passing standard SAM tests, despite of our experiments with various services.
- Support people are oriented at the production-grade support and we're pushing them hard to achieve this.
- Many services were clustered or made highly-available: this allows for a bit longer allowable response time to the individual failures and provides better resilience.

Let's talk about future



Plans: new hardware and software

- Turns out that we won't buy any new hardware this year (despite our rollout plan that included twice as much storage and CPUs at the end of 2013).
- This is connected to the budget cuts; nevertheless we are working on getting all planned hardware during first half of 2014.
- Life is hard and everything happens :(We take it as the opportunity to focus on the software and operational issues ;)
- We're planning to deploy Slurm-based CEs at the end of this year: they should be more stable and easier to support than Torque-based ones.
- More solutions to manage our infrastructure are needed: mature hardware inventory system, more fine-grained monitoring systems that will be interconnected and can be used to cross-system analysis of found problems.

Plans: resources for ALICE

- We are currently installing new hardware and redoing our LAN core to be routed; expected ETA is 2-3 weeks.
- After this we will be ready to start to closely work to provide ALICE resources and test our Tier-1 without tapes.
- Integration of tape with Xrootd-based storage is planned, but not before we will get some real-world operational experience with Enstore and ATLAS.
- Primary disk-based storage element type for ALICE will be EOS; if it will suit everyone, we will stick with it, if not – dCache is the next candidate (we already have a great deal of dCache experience with storage for ATLAS).
- We expect a bit of breakage at the beginning of operations, so don't throw valuable data at us on the first stages, please ;)

Plans: people and process

- We are looking for more people that will administer systems and services; it is hard to find good persons for this role, but we're always working on that.
- We will delegate more operational duties to our 24x7 infrastructure team; it is thought to be enlarged to allow for such additional load.
- Monitoring of Grid resources and services will be one of the 24x7 duties, for this we need simpler dashboards with easily-identifiable problems; working on that.
- After getting experience with ATLAS Tier-1 business we're upgrading our internal procedures and views on how things should be done to allow our resources to work smoothly and efficiently.

That's all, folks!

Questions, comments?

