

# Russian Tier-1 at NRC KI

Eygene Ryabinkin

National Research Centre "Kurchatov Institute"

ATLAS T1/T2/T3 Jamboree, CERN

I will talk about the present status of Tier-1 prototype in Kurchatov Institute, but I will also give an overview of the Tier-1 project, because it is the first time I am presenting this to such a wide audience.

## Topics:

- Political stuff
- Tier-1 project
- Tier-1 facilities and infrastructure
- Tier-1 people
- Tier-1 networking
- Tier-1 storage and tape system
- Current status of services
- Plans for ATLAS

The whole Tier-1 in Russia is split between NRC KI and JINR:

- NRC KI: ALICE, ATLAS and LHCb
- JINR: CMS

No shared infrastructure excluding WAN networking.

NRC KI will stop its Tier-2 once Tier-1 will enter the production phase (somewhere in 2014) and ATLAS pledges from it will likely be transferred to IHEP (Protvino).

WLCG overview board accepted Russian Tier-1 proposal both for NRC KI and JINR in September 2012. We're currently preparing the MoU, have the plan for Tier-1 rollout and following its milestones.

Tier-1 project consists of two phases:

- project to build Tier-1 prototype (with 10% of the whole capacity); it is funded under contract with Russian ministry of education, started in 2011 and will end in the middle of 2013. Leading organisation is Kurchatov Institute with JINR as a sub-contractor (for CMS part).
- the Tier-1 itself; it will be funded by the Russian ministry of education and JINR; preliminary agreements are already there, though no signed contracts yet (it is a bit too early for that).

It is officially seen that the prototype project will yield the technical project on the whole Tier-1<sup>1</sup>.

Non-officially, it will also give already installed resources and initial Tier-1 staff team with some experience for Tier-1 business.

---

<sup>1</sup>We will just need to buy the rest of resources and everything will work ;))

A quick overview:

- 4 hosting rooms, 750 KW each<sup>2</sup>
- Class A power feeds (two independent inputs)
- UPS facilities that will run for 27 minutes at full load of all modules and no external power
- Diesel generators for 36 hours of uninterruptable run (can be more, just a question of fuel availability)
- Precision and in-row cooling systems with liquid cooling agent on the exchange contour and air in the primary contour
- Infrastructure monitoring in 24x7 mode with rotating duty staff
- Controlled access and video surveillance
- Action plans for handling of disasters (like no cooling and systems running at full throttle or big leaks of cooling liquid underfloor)

---

<sup>2</sup>Not only for Tier-1, for all computing facilities

Our current biggest problem.

We had started with two seasoned administrators and 4 new ones. The idea was to train new people at the small prototype and then quickly install all services mainly by seasoned people with new ones watching and following the process...

Our current biggest problem.

We had started with two seasoned administrators and 4 new ones. The idea was to train new people at the small prototype and then quickly install all services mainly by seasoned people with new ones watching and following the process. . .

This was failed: new administrators weren't really able to get the ideas about services at the acceptable speed and level of understanding.

So, we had started from scratch:

- rolled out a new installation/configuration system and made new administrators to take active part in its development
- switched to mode of supervision of new people by seasoned ones and almost everything being installed by new folks

This works better, though still not perfect.

WAN for 2012 and 2013:

- 2012: 2 Gbit/sec GÉANT IP shared circuits KI-CERN and JINR-CERN; trying to include it to LHC OPN
- 2013: 10 Gbit/sec ring KI  $\Rightarrow$  AMS  $\Rightarrow$  FRA  $\Rightarrow$  BUD  $\Rightarrow$  JINR  $\Rightarrow$  KI; also exploring the possibilities to have additional GEANT capacity

We had a meeting with GÉANT and Terena people in August 2012 and they are aware of our requirements for 10 Gbit/sec in 2013 and up to 40 Gbit/sec in 2014 (dedicated links for LHC OPN): they are currently evaluating various proposals from regional carriers.

We're working on bumping the existing connectivity with Russian Tier-2 sites.

Yet no plans for making the WAN channels to other Tier-1 sites (can be done via GÉANT or by arranging dedicated links).



## Storage and tape system

dCache: ready for testing, polished the installation and performed local transfer tests. Planning for distinct VO instances, evaluating various strategies for placing protocol doors.

EOS: installed 0.1, played with it a bit and then scratched it out during transition of training strategies for new administrators; just installed 0.2, still experiencing some problems with user authentication.

Trying fat (10 Gbit/sec) and thin (2 Gbit/sec) configurations for internal networking of pool nodes.

IBM TS3500 was delivered in September, next 1.5 months passed in basic tests and replacements of broken parts. Currently we're testing CASTOR components with it.

FTS: installed and getting familiar with it since August 2012.

## Current status of services

- dCache: ready
- FTS: ready
- EOS: struggling with user authentication
- Scheduler for CREAM: ready, uses Torque
- BDII: should be ready at this Friday
- CREAM CE: installing 3 instances with different shared area configuration (NFSv4 from different vendors)
- WNs: new administrators are in the process of their installation (CentOS 5.x and UMD-2)
- APEL: ready
- LB: pending
- perfSonar: ready
- CASTOR: started to evaluate
- Enstore with dCache: pending

## KI Tier-1 resource growth roadmap

We have the following roadmap (presented at WLCG overview board and being agreed to):

Year:	2012	2013	2014
CPU (HS06):	10720	54000	71000
Disk (TB):	1300	4800	6300
Tape (TB):	2000	5700	7400

Resources for 2013 are calculated as 10% of aggregate Tier-1 capacity at 2012.

Resources for 2014 are 30% increase of 2013 ones to catch up with the increased pledges for Tier-1.

## Plans for ATLAS

During Alexei's visit to NRC KI we agreed that we will provide minimal environment for data processing at the end of December.


Once it will be sufficiently stable, ATLAS will run processing of the dark data.

We're not increasing our pledge for Tier-2, but will give resources from Tier-1 in 2013 for ATLAS processing.

At the middle of 2013 we should come up with the model of Tier-2 support for our Tier-1.

Resource split between VOs: something like 45/45/10 for ALICE/ATLAS/LHCb, but it is still to be discussed and likely will involve 3 interested VOs.

## Questions, comments?



IF YOU KEEP SAYING "BEAR WITH ME FOR A MOMENT", PEOPLE TAKE A WHILE TO FIGURE OUT THAT YOU'RE JUST SHOWING THEM RANDOM SLIDES.