# RU federated storage R&D: an overview

E. Ryabinkin et al

NRC “Kurchatov Institute”

RU fed. storage R&D

# Help → About

Motivation: small institutions have no enough people to support fully-fledged software stack. Distributed stuff like FAX, ALICE xrootd, EOS@CERN, etc (mostly) works. Let's try to analyze how we can set up distributed storage in our region and how it can be used from Grid sites, from HPC, etc.
Current layout of federation testbed:

- KI, Moscow;
- PNPI, Saint-Petersburg;
- SPbSU, Saint-Petersburg.

Network latency: $\approx$ 9 ms.
3 different instances, separated at the beginning, federated at the end. Hardware is like the mostly normal one (not a shiny new with latest CPUs, etc, but decent).
People: A. Kiryanov, A. Zarochentsev, D. Krasnopevtsev, R. Mashinistov, A. Klimentov, /me.

НИЦ «Курчатовский институт»

# Testsuite

Dmitry prepared an ATLAS job based on the TRT detector software. Properties:

- mostly stand-alone, input data is packaged and currently just copied locally;
- two-line patch to support Xrootd (stashed the patch somewhere, will write a new one, possibly 3-liner);
- runs on any machine where we have CVMFS for ATLAS (/cvmfs/atlas.cern.ch) and HTTP proxy that can do Frontier caching.

Current target clusters:

- Tier-1 in KI;
- Tier-2 in KI, PNPI, SPbSU;
- HPC in KI.

First metrics: simulation time at various clusters with local access to their own EOS instances. And simulation time for the federated EOS (with "global" redirector), when the data is taken from remote pools.
Why do we want federated EOS and not just say "copy from the remote X"?
We want full cycle of operations with redirector, The Real Thing$^{\mathrm{tm}}$.

НИЦ «Курчатовский институт»