

CC-IN2P3 – IHEP connectivity issues

Progress report

Fabio Hernandez
fabio@in2p3.fr

October 11th, 2012

Context

- Summary of the previous meeting (2012-07-02)

Perfsonar host at CC-IN2P3 now using 1 Gbps instead of 10 Gbps network interface

significant performance increase between CC-IN2P3 and IHEP observed via the regular throughput tests results (more on this later)

details: <https://indico.in2p3.fr/conferenceDisplay.py?confId=7336>

- In the following slides we summarize the known progress made since last meeting

Information sharing

- Reference document at IHEP's Twiki
<http://twiki.ihep.ac.cn/twiki/bin/view/InternationalConnectivity/IHEP-CCIN2P3>
- All the historical information about this issue can be found there
- Timely updated

Perfsonar throughput tests



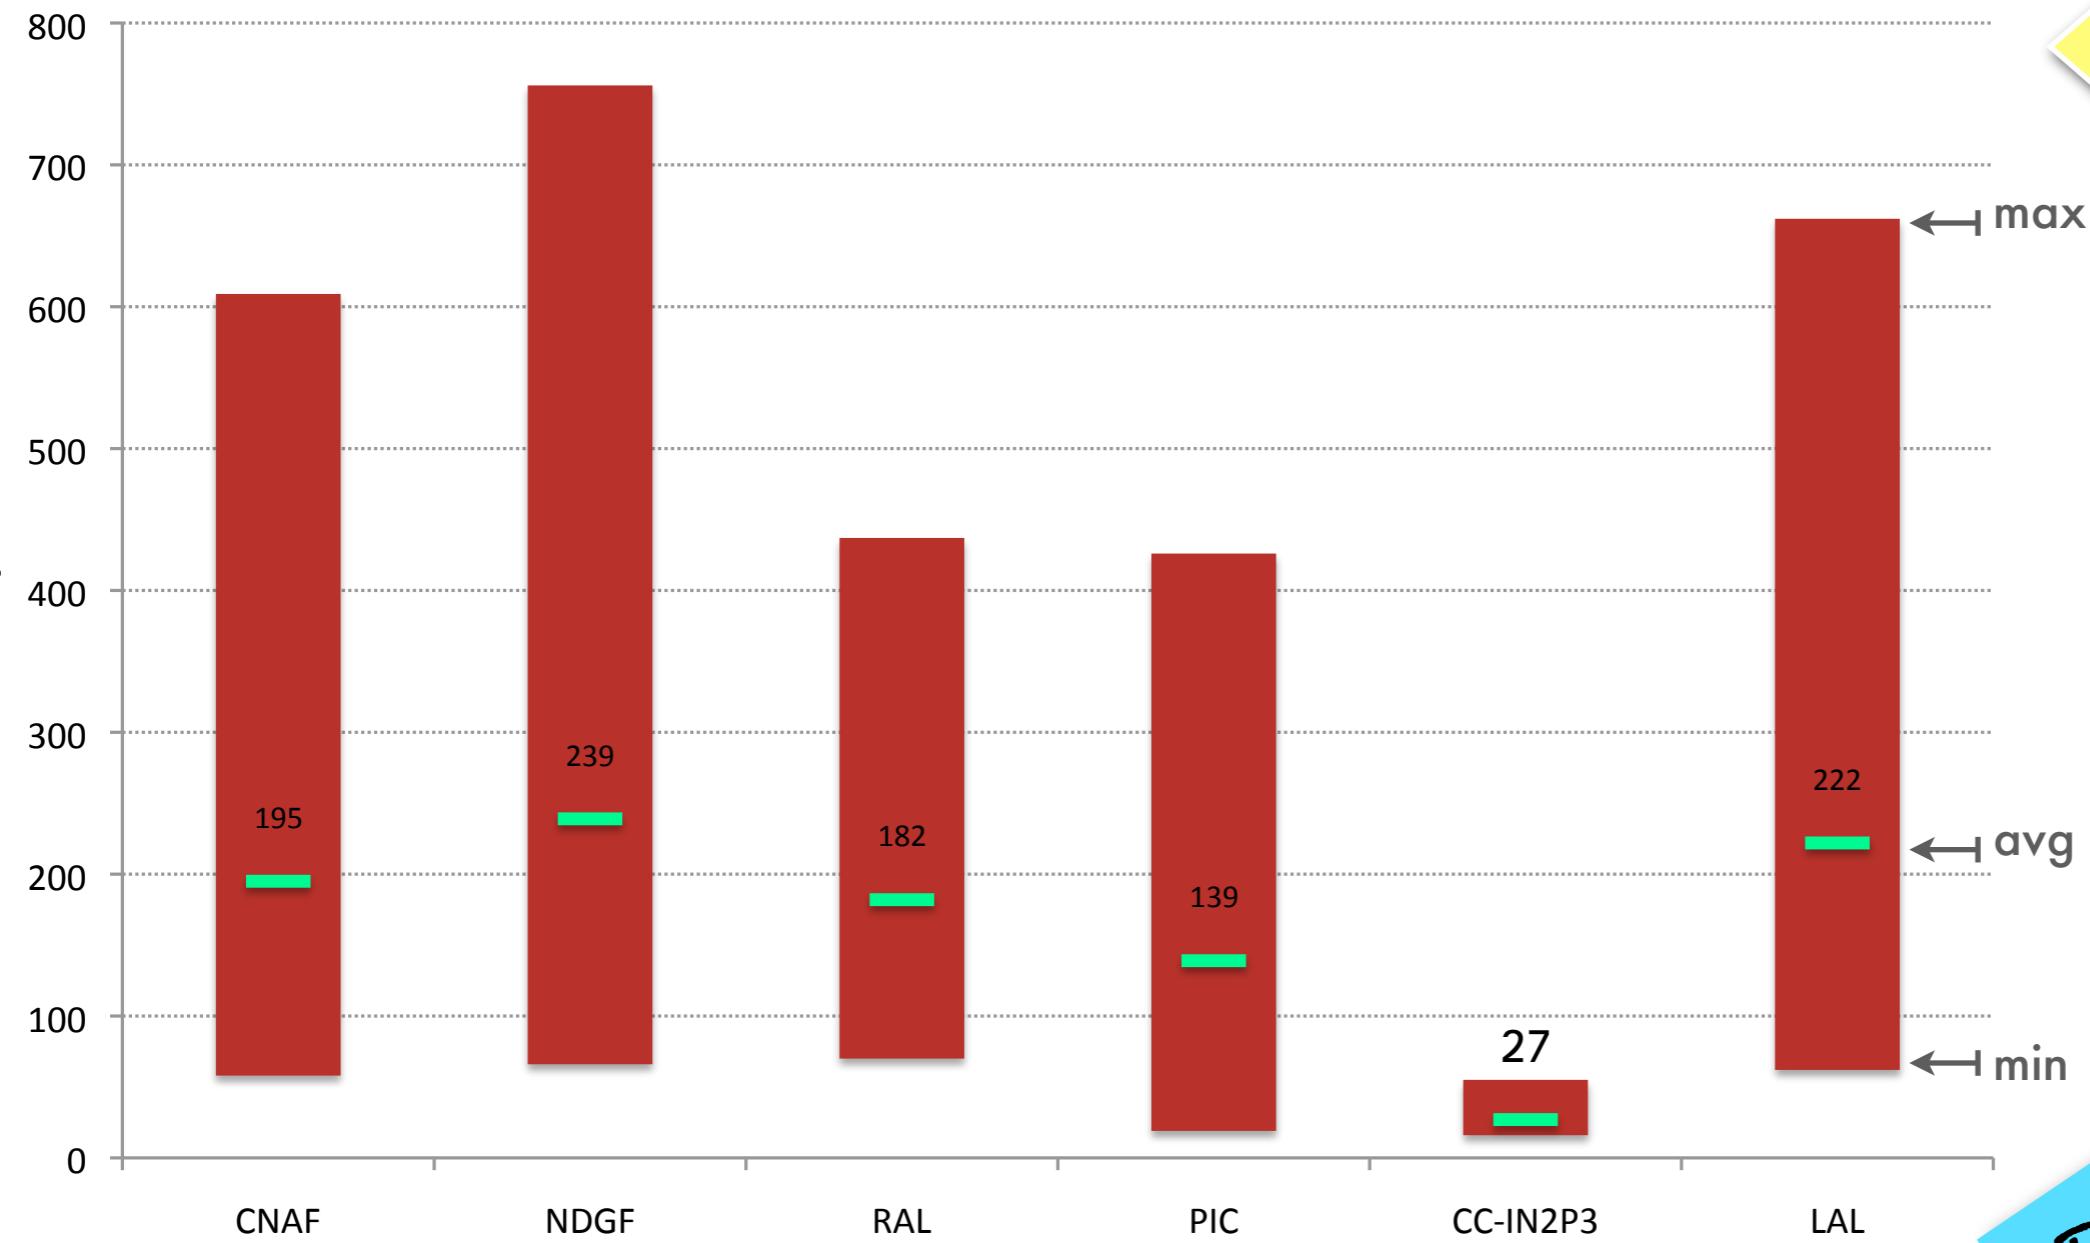
CC-IN2P3 → IHEP —

IHEP → CC-IN2P3 —

Source: <http://ccperfsonar-lhcopn.in2p3.fr>

Network throughput variation: before

Network Throughput Variation — European sites to IHEP



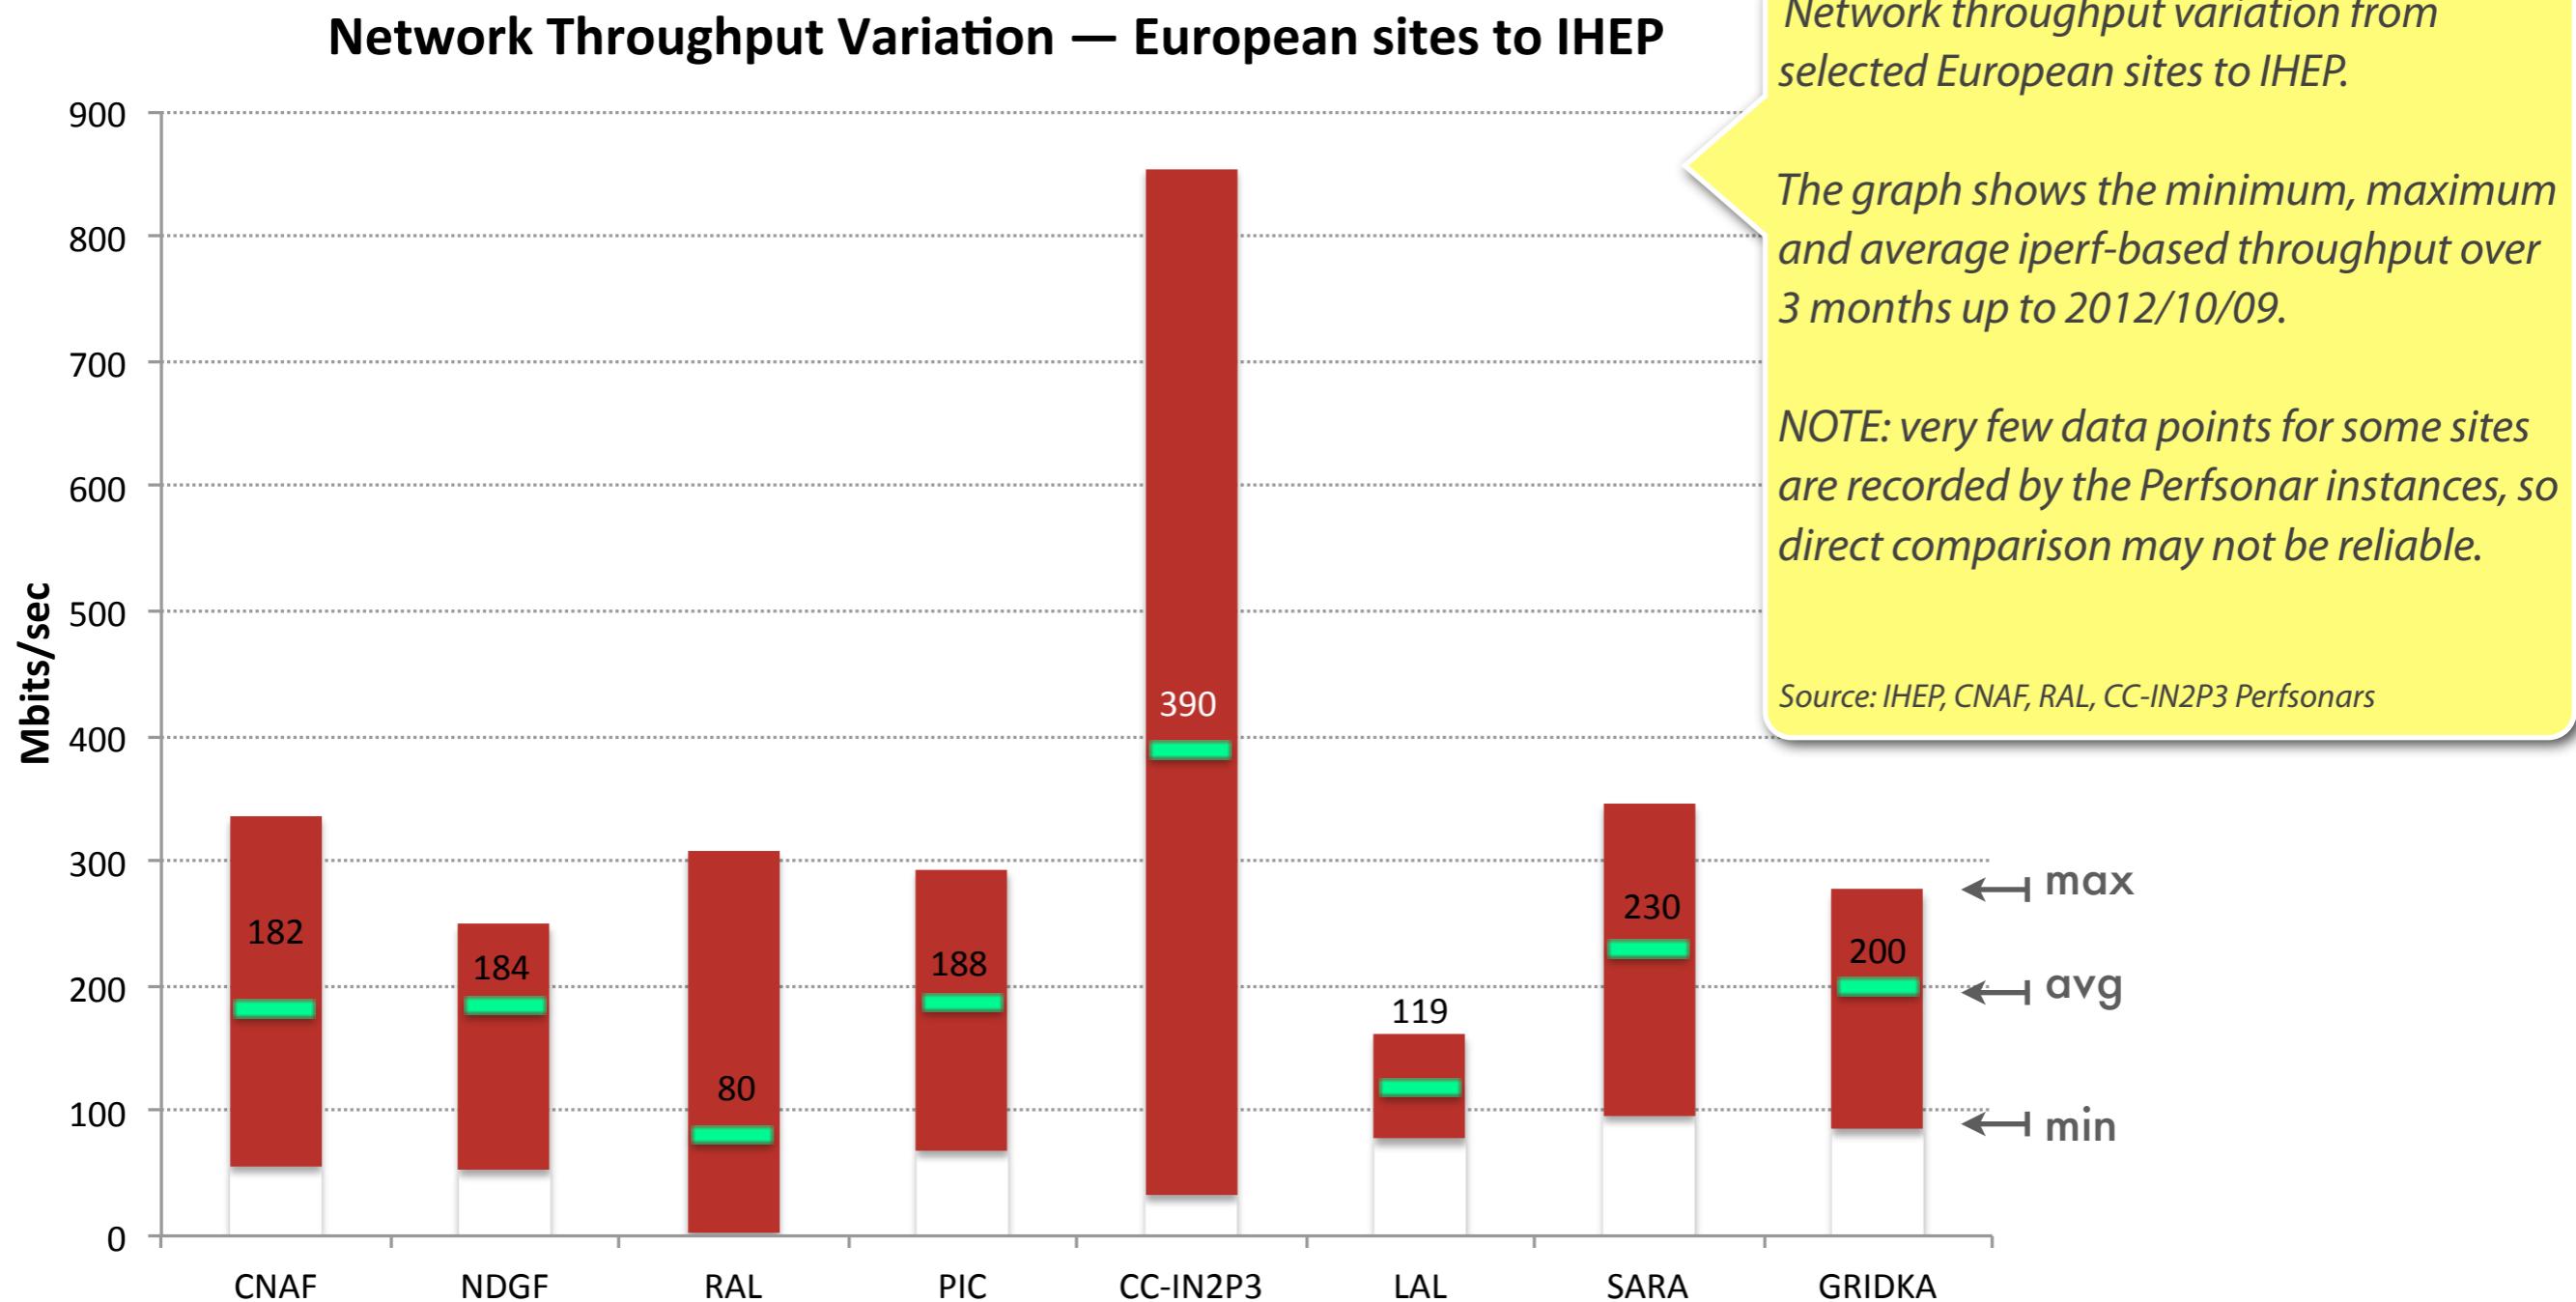
Network throughput variation from selected European sites to IHEP.

The graph shows the minimum, maximum and average iperf-based throughput over 1 month up to 2012/06/26.

*Source: IHEP's Perfonar
<https://perfsonar.ihep.ac.cn>*

*Presented in
July meeting*

Network throughput variation: after



Network throughput variation: after (cont.)

- Good news: there does not seem to be a inherent bandwidth limitation between CC-IN2P3 and IHEP
- Compared to the network throughput variation presented in July meeting, we observe that the maximum throughput from all EU sites but CC-IN2P3 to IHEP significantly decreased
- Could this be related to the replacement of IHEP's Perfsonar host operated on June 28th or to the duration of the throughput tests?

currently, the test duration for IHEP's Perfsonar is 20 seconds, while for CC-IN2P3's Perfsonar is 30 seconds

*unfortunately, most of IHEP's peer sites limit the throughput test duration to 20 seconds
20 seconds-long throughput tests are not long enough for trans-continental links*

Other actions performed

- At CC-IN2P3

Installation of a second Perfsonar host dedicated to latency tests

Throughput tests: ccperfsonar-lhcopn.in2p3.fr

Latency tests: ccperfsonar2-lhcopn.in2p3.fr

- At IHEP

Regular throughput tests with SARA

*Network reconfiguration at GridKa (initiated by Eric Lançon):
throughput with IHEP significantly improved, but still high variation*

Collecting data on the usage of Orient link (see next slides)

Orient link usage

- Fazhi managed to get permission to collect data on the usage of the Orient network link between Beijing and London

as seen by the network router at the Chinese side

data points every 2 minutes

time series plots and raw data available at

<http://nms.ihep.ac.cn/mrtg/CNGI-6IX-Bj-London.html>

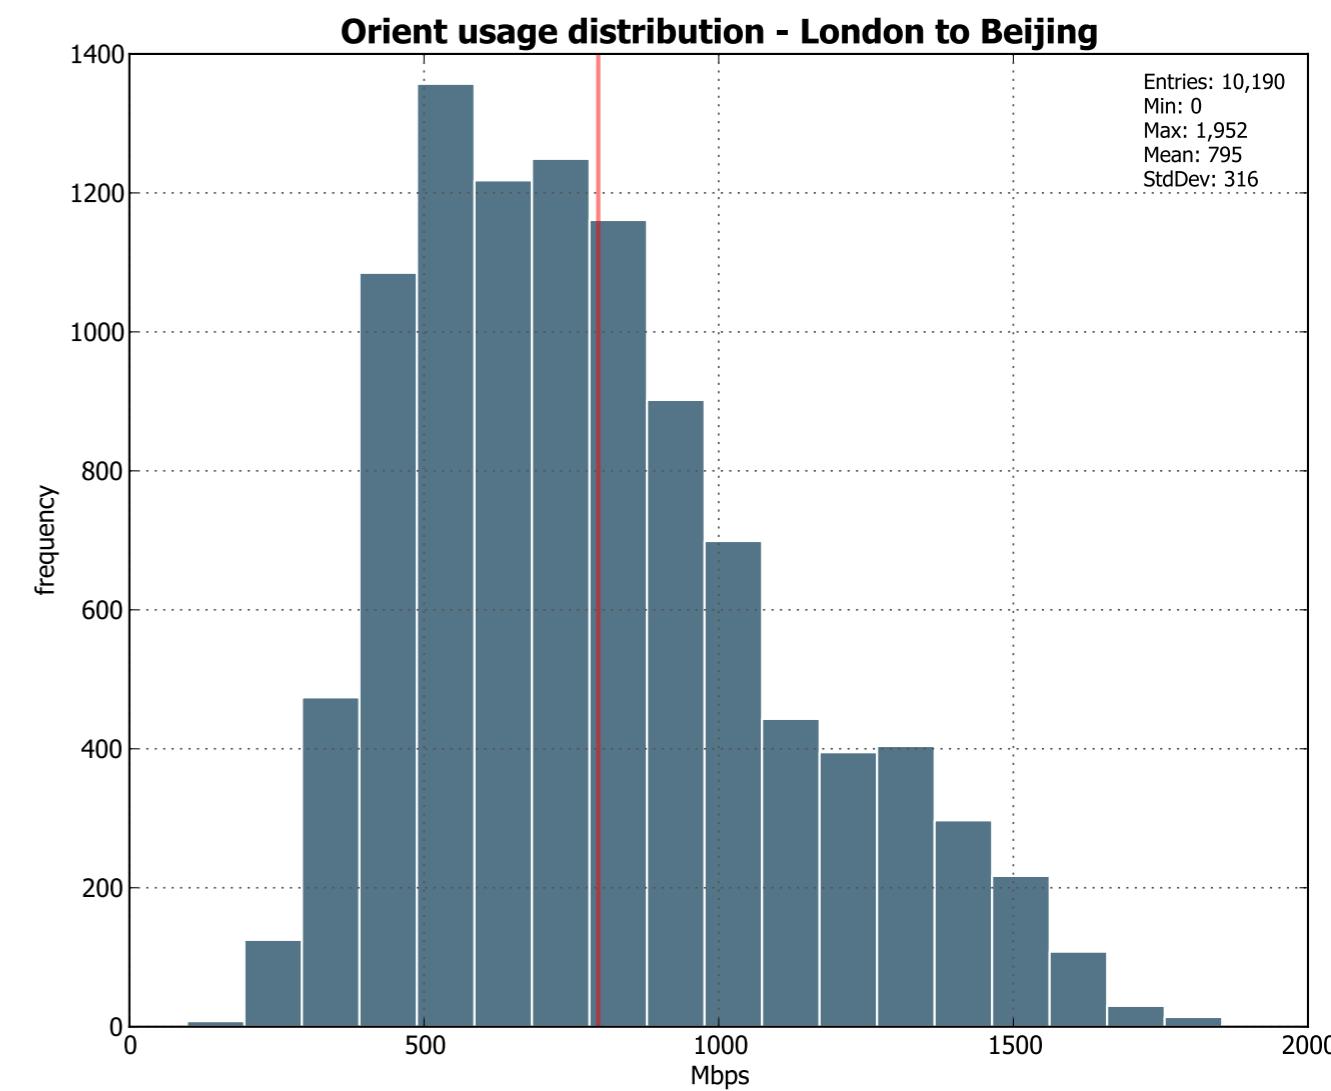
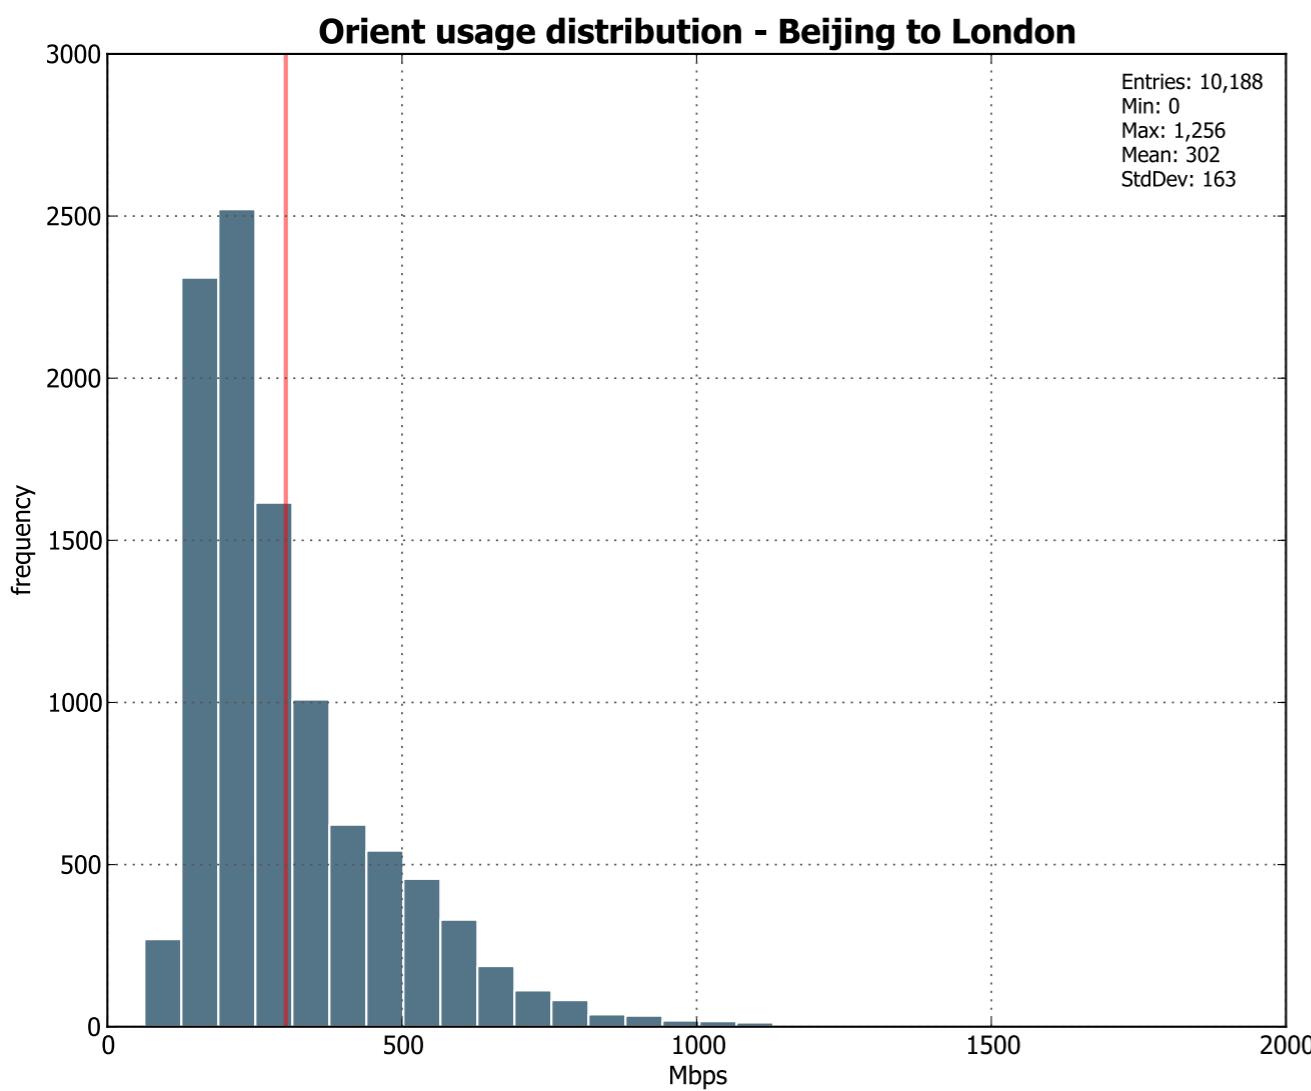
- No detailed information is available to us on the network traffic through this link

but we can see that the link is really used

Orient link usage (cont.)

Beijing → London

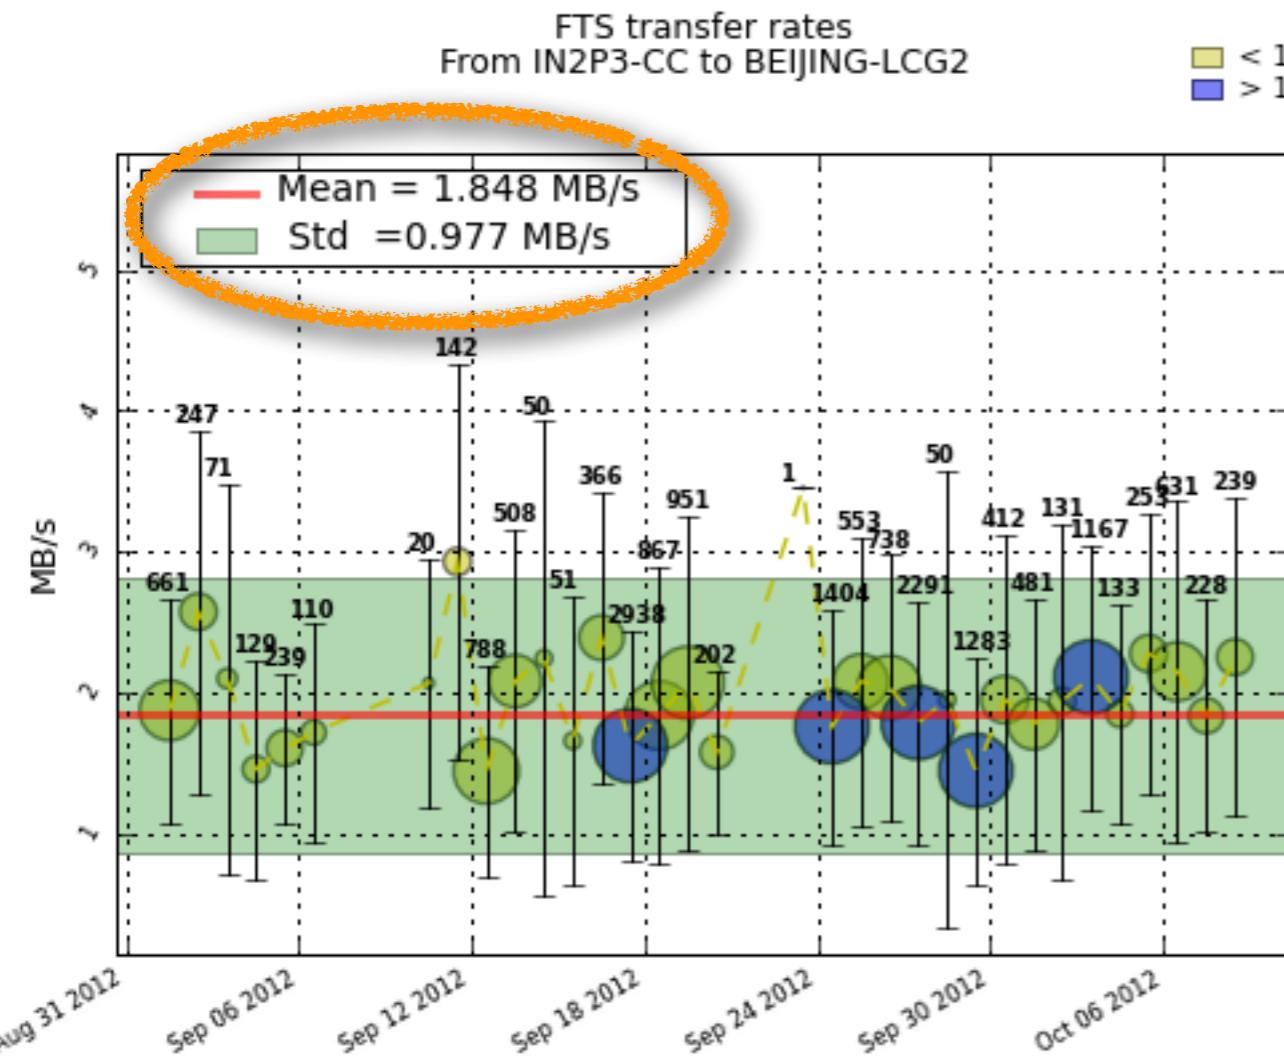
London → Beijing



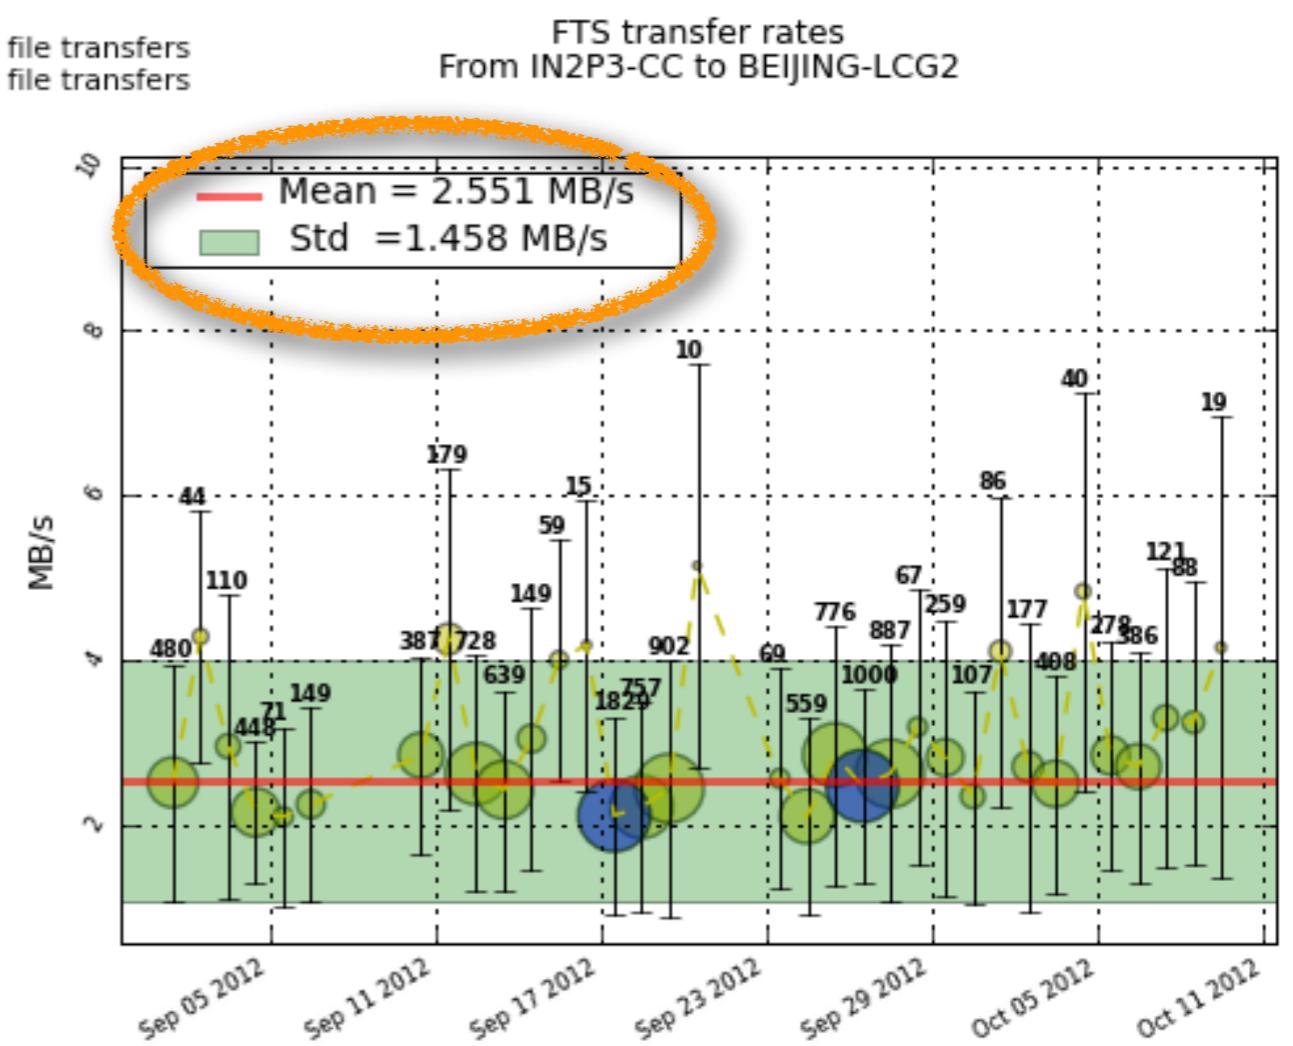
Period: 2012-09-22 to 2012-10-09

Source: <http://nms.ihep.ac.cn/mrtg/CNGI-6IX-Bj-London.html>

Atlas file transfer throughput: CC-IN2P3 → IHEP



file sizes: 0.1 - 1GB



file sizes: >1GB

Source: <http://bourricot.cern.ch/dq2/ftsmon>

What is next?

- Why the file transfer application can not reach the throughput the link can deliver?
from Atlas measurements we observe that the file transfers get an average of 10% of the throughput measured by Perfsonar between the border routers of the two sites
- Are there bandwidth constraints at the LAN levels of CC-IN2P3 or IHEP sites?
Fazhi is preparing a modification to the LAN configuration to increase the bandwidth up to the firewall: this is relevant to the storage machines involved in receiving and sending files
- Is access to disk of the storage machines the limiting factor?
- Is the file transfer application able to (almost fully) exploit the available bandwidth when using transcontinental links?
- We should aim at regularly measuring the raw throughput **between transfer machines** at both sites
- Should we perform again another test campaign as the one performed by Yvan Calas early this year?
to measure the real throughput that can be reached via gridFTP file transfers in the current network conditions

Questions & Comments