

IHEP connectivity issues

Progress report

Fabio Hernandez

fabio@in2p3.fr

Fazhi Qi

qfz@ihep.ac.cn



January 10th, 2012

Context

- In the previous meeting (2011-12-12) we reported the low throughput we were observing when transferring data from IHEP to CC-IN2P3

<http://indico.in2p3.fr/conferenceDisplay.py?confId=6204>

- Today, we report on the progress made since then

Network throughput tests

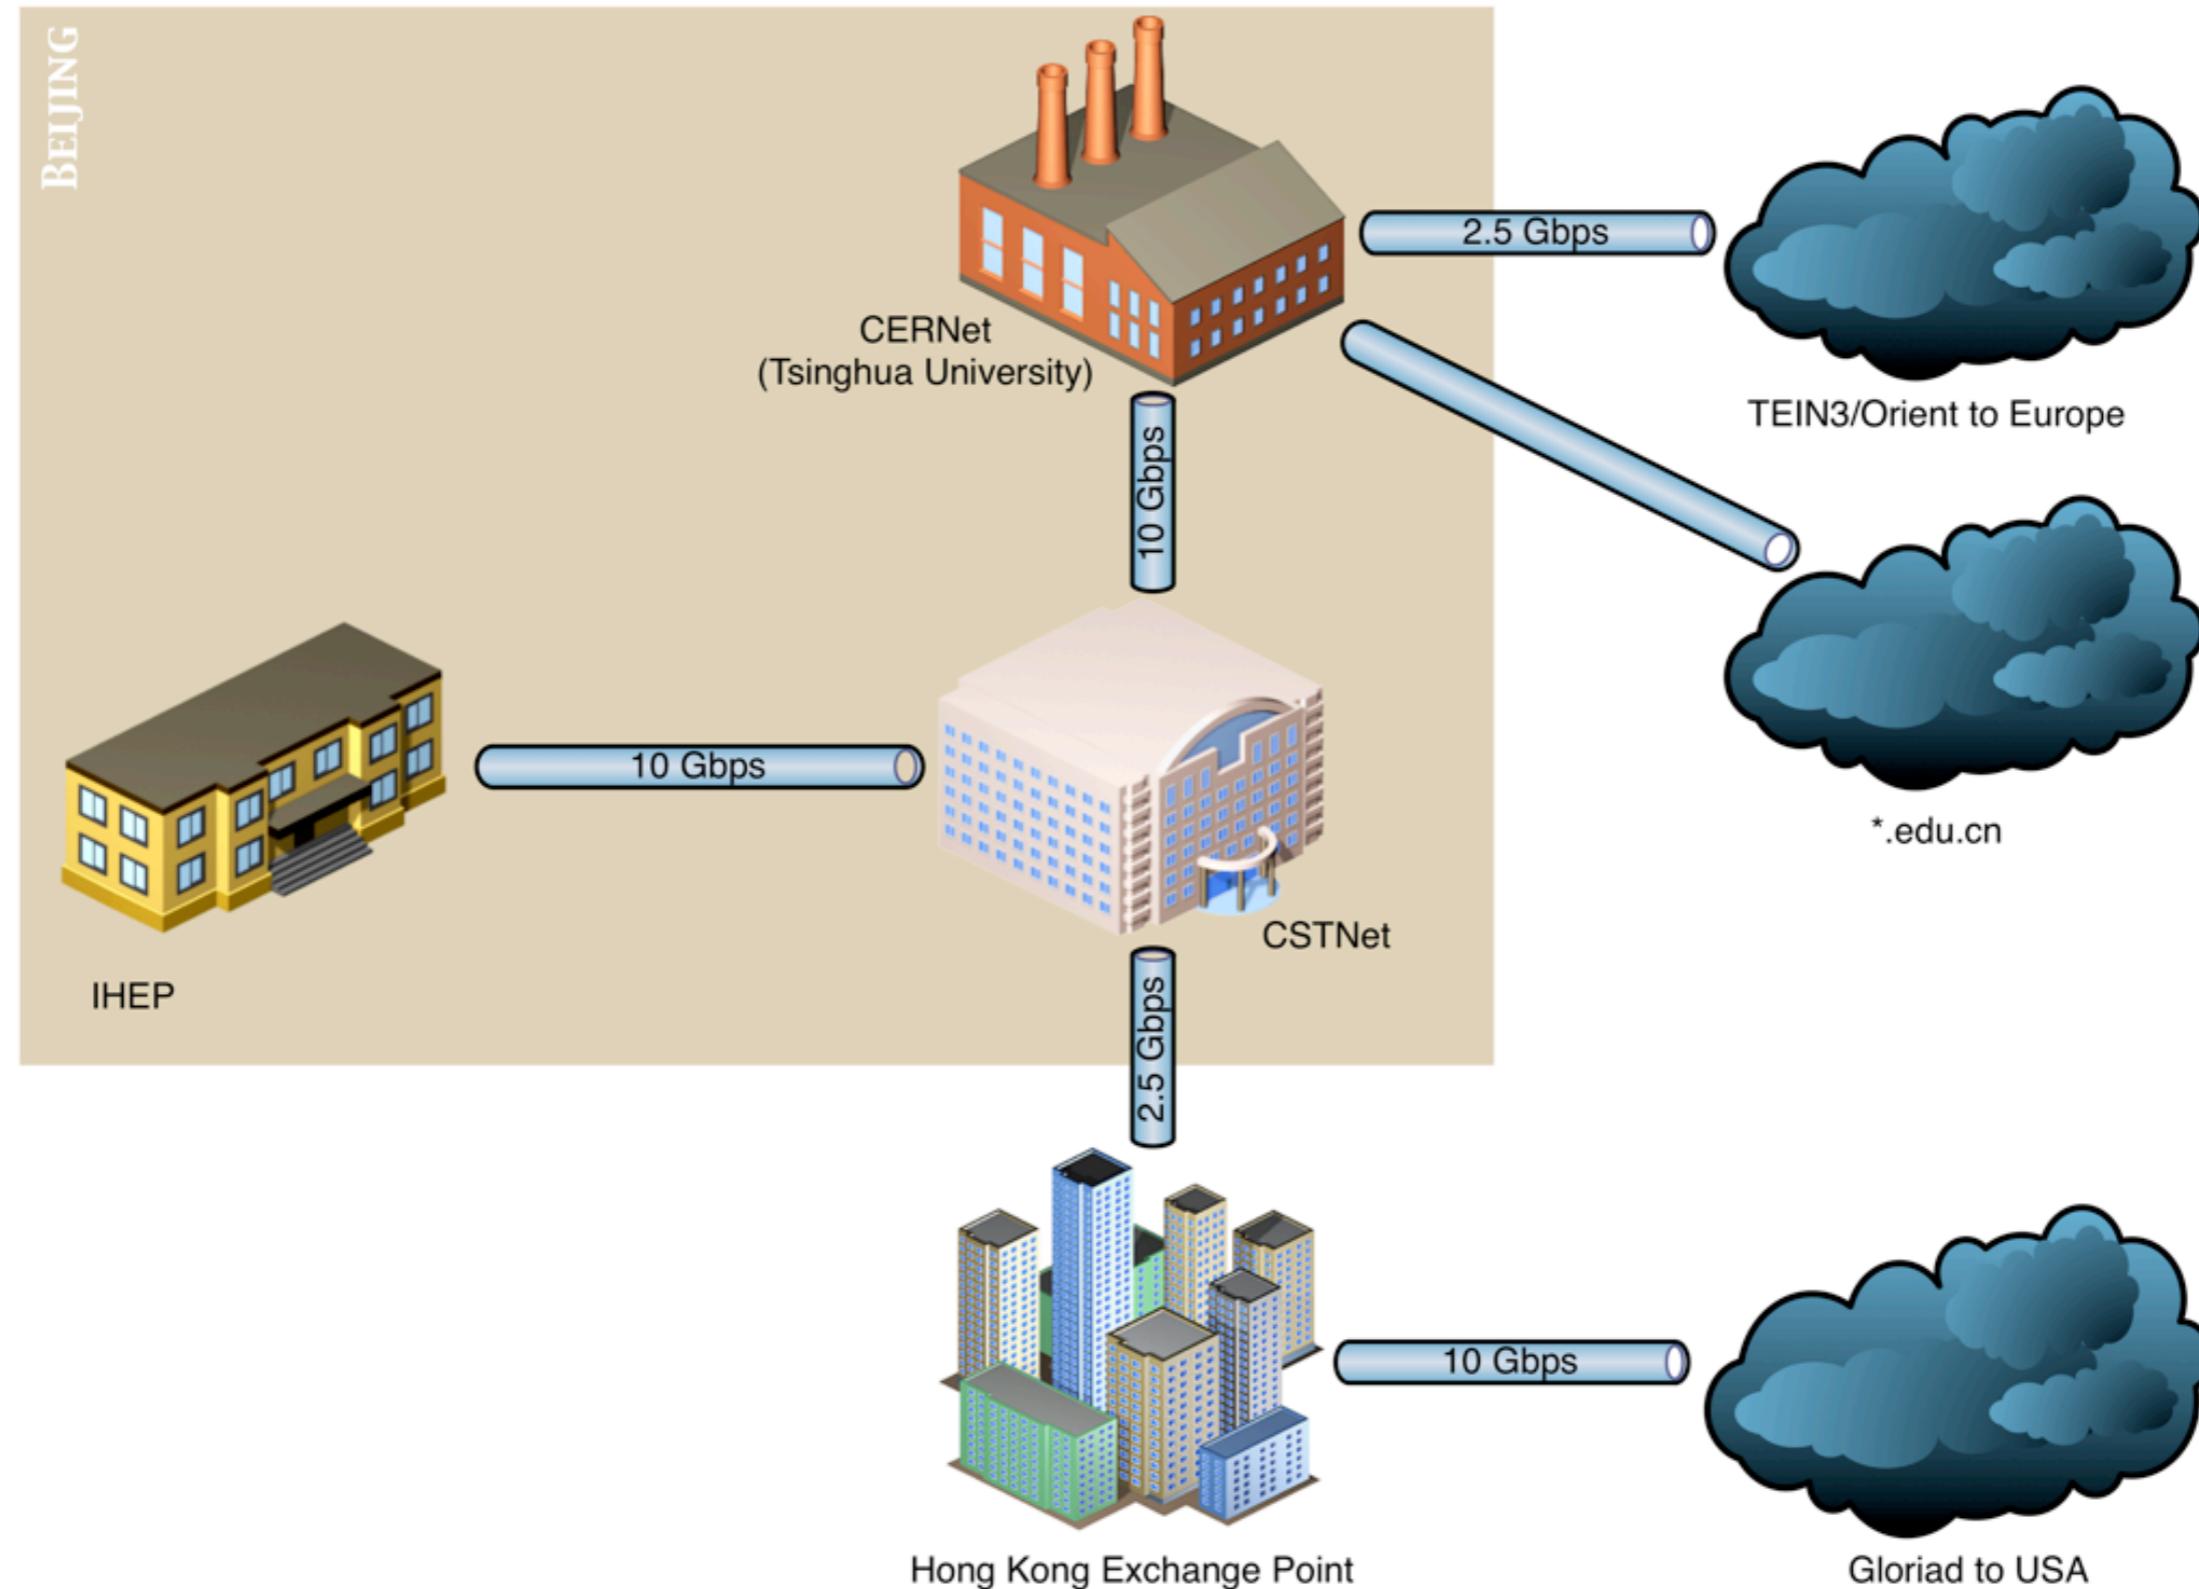
- Detailed report on the results of some network throughput tests from IHEP to several machines in the IN2P3 network

Tests performed late December and early January

*Very low throughput observed: **2.5 Mbits/sec** on average*

See the report attached to the agenda

IHEP International connectivity



Actions

- Early January, got permission from several European WLCG tier-1s to regularly monitor the network throughput from and to IHEP

by using the Perfsonar infrastructure deployed at all those sites

- Currently monitoring the links between IHEP and the following sites

IN2P3 (FR), NDGF(DN), CNAF(IT), RAL(UK), PIC(ES), CERN

Pending: SARA (NL) and KIT(DE)

- Made arrangements to be able to autonomously perform iperf tests between IHEP and machines in CSTNet and CERNet

these organizations manage the routers between IHEP and then Orient/TEIN3 link (see diagram in the previous slide)

IHEP \longleftrightarrow CC-IN2P3 network route

- Symmetrical route for inbound and outbound traffic
- IHEP \longleftrightarrow CSTNet \longleftrightarrow CNGI6IX (CERNet) \longleftrightarrow TEIN3/Orient \longleftrightarrow GEANT2 \longleftrightarrow RENATER \longleftrightarrow CC-IN2P3

TEIN3/Orient is a direct link London - Beijing

Throughput tests performed

[See annex for details]

Segment	Throughput [Mbits/sec]	Status
IHEP → CSTNet	942	✓
CSTNet → CNGI6IX	965	✓
CNGI6IX → CC-IN2P3	588	✓
IHEP → CNGI6IX	45	✗
IHEP → GEANT2	6	✗
IHEP → CC-IN2P3	2	✗

Route: IHEP ↔ CSTNet ↔ CNGI6IX ↔ TEIN3/Orient ↔ GEANT2 ↔ RENATER ↔ CC-IN2P3

Diagnosis and solution

- Problem located at the CSTNet internal network
- Difference in these 2 routes:

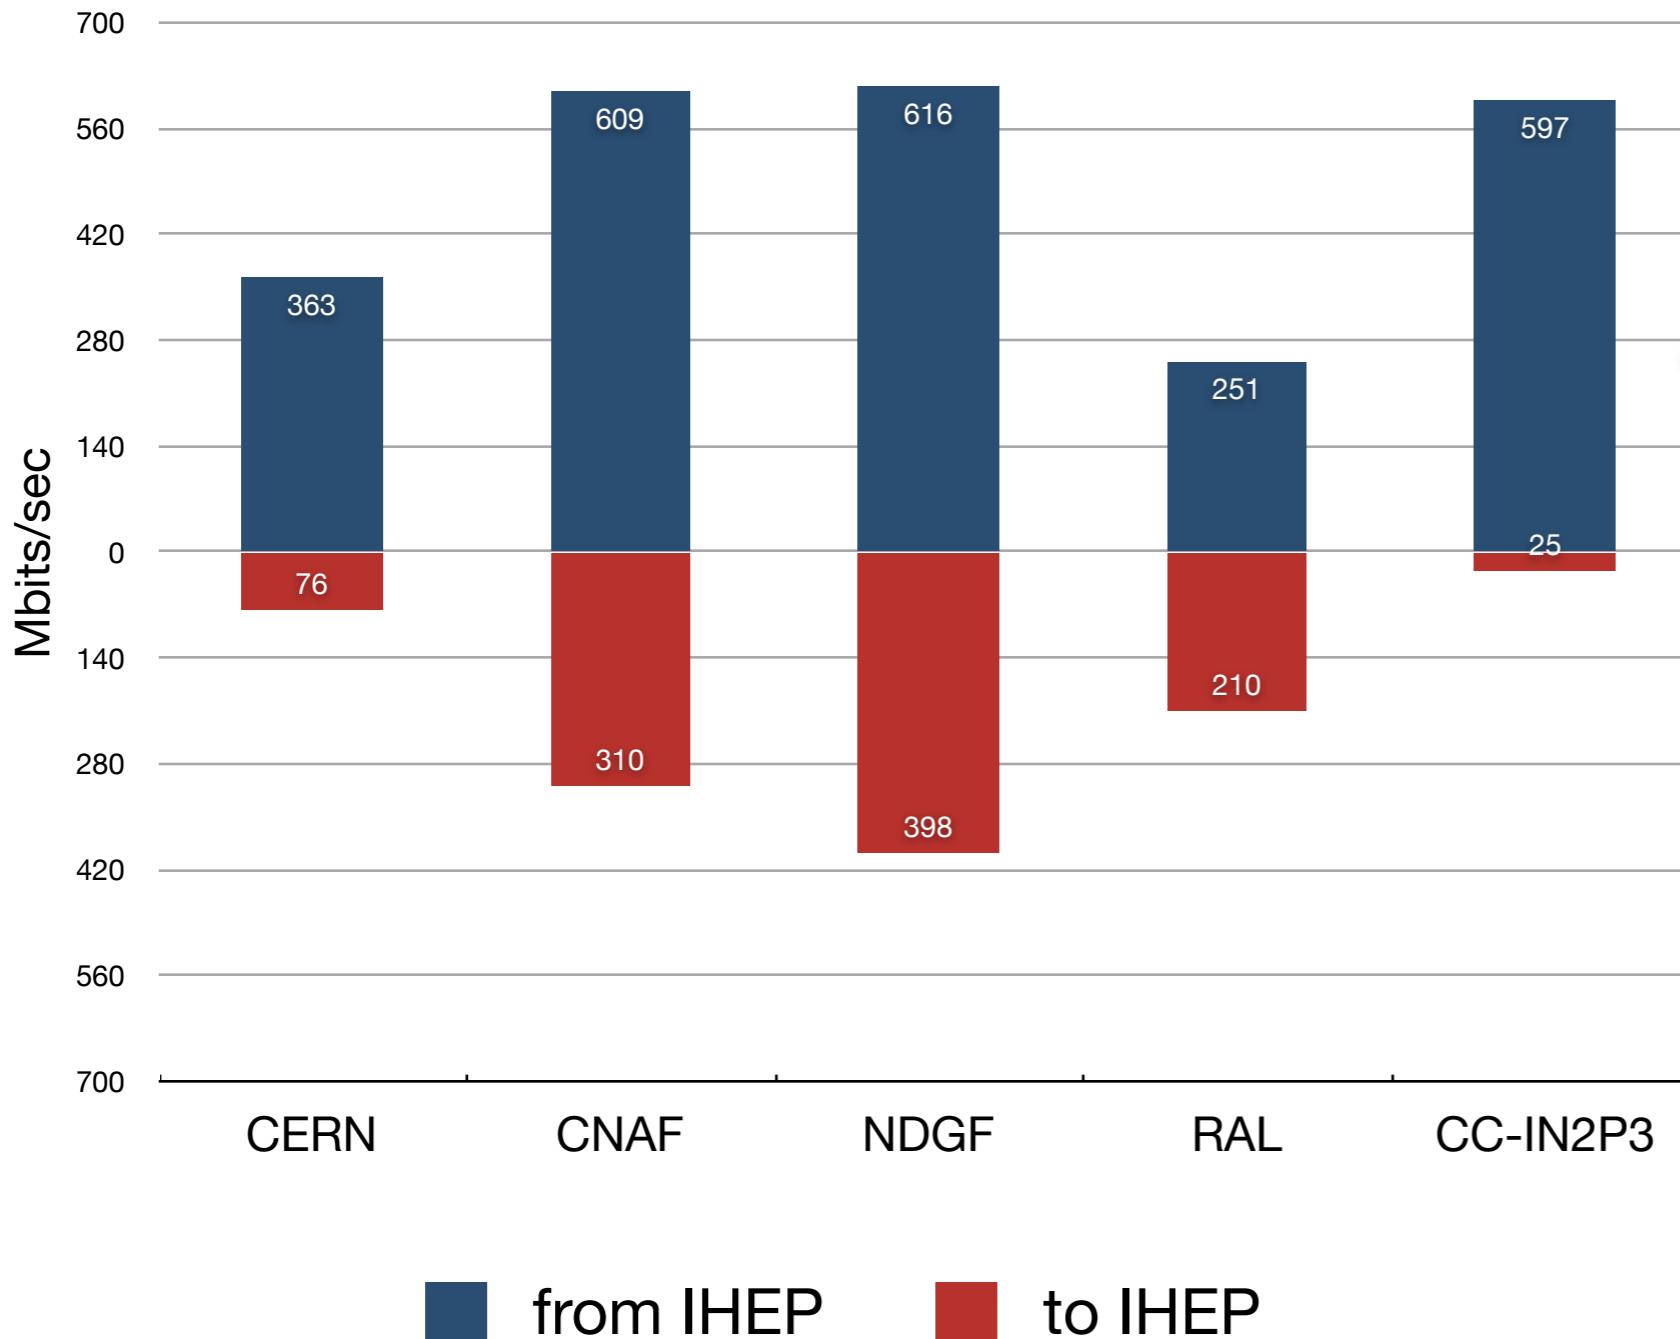
CSTNet → CNGI6IX (good performance)

IHEP → CNGI6IX (bad performance)

- CSTNet experts found that the network security policies were different between those 2 routes
policies modified by January 21st

Current throughput

Average Measured Network Throughput



Average network throughput over the last month up to 09/02/2012 as reported by IHEP's Perfsonar:

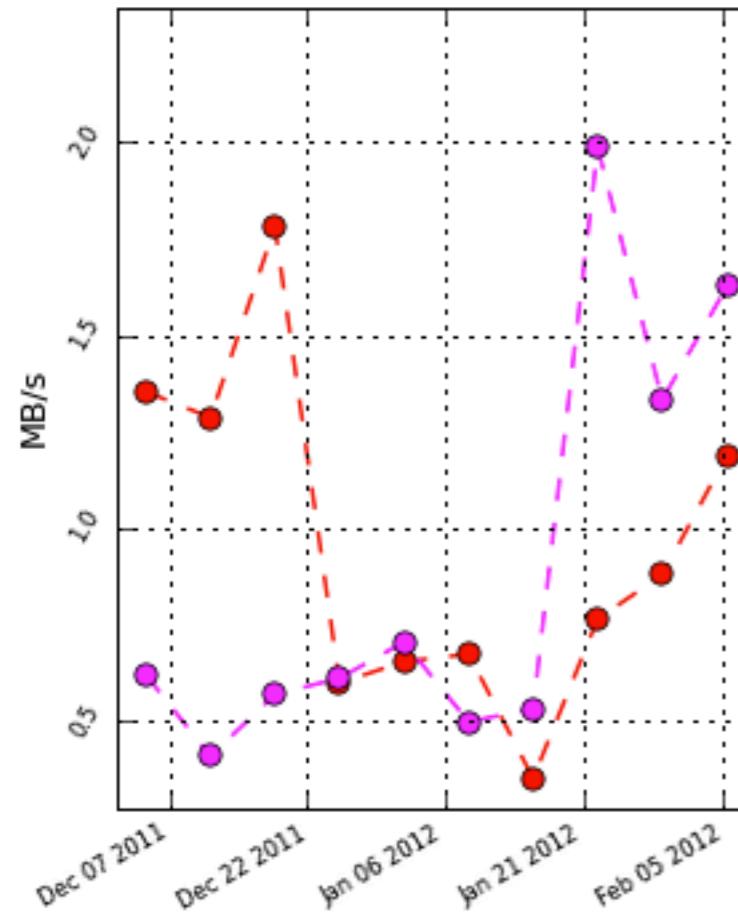
<http://perfsonar.ihep.ac.cn>

NOTE: the throughput from CC-IN2P3 to IHEP was obtained from CC-IN2P3's Perfsonar

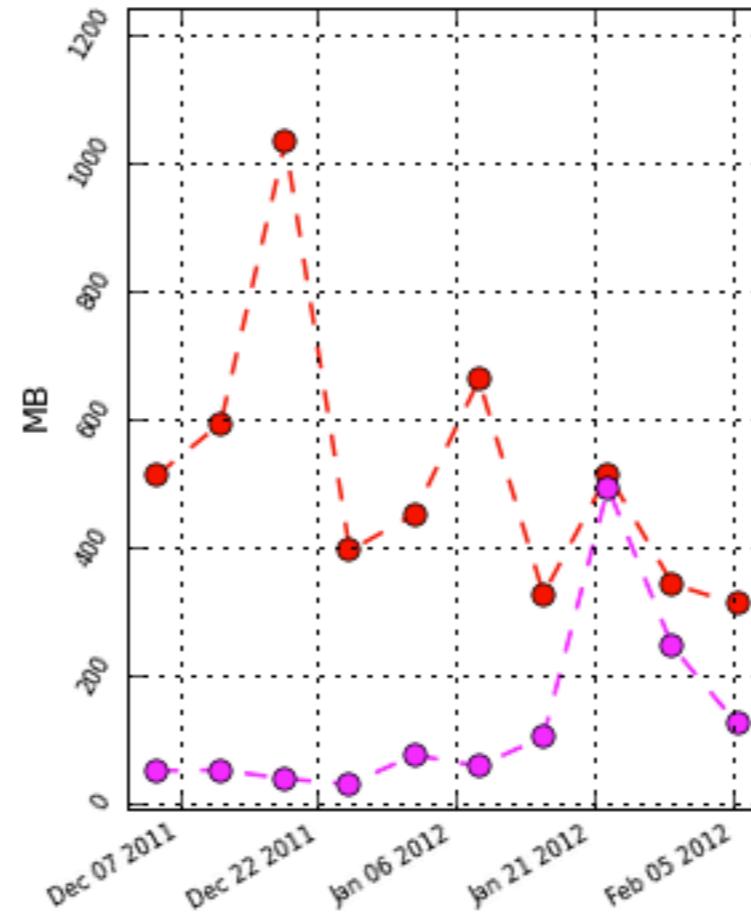
LHC Data Exchange with IHEP

ATLAS data transfer between CC-IN2P3 and IHEP

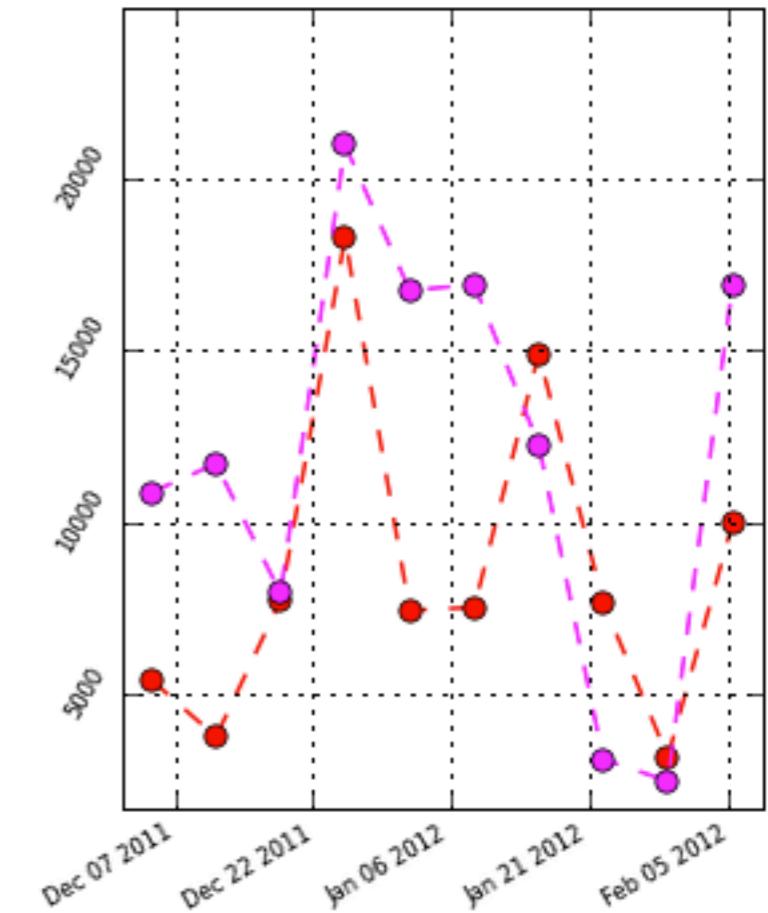
FTS transfer rates



File sizes



#files transferred



Total number of file transfers: 205424

We can observe that from January 21st on, the increased throughput from IHEP to CC-IN2P3 was beneficial for ATLAS transfers.

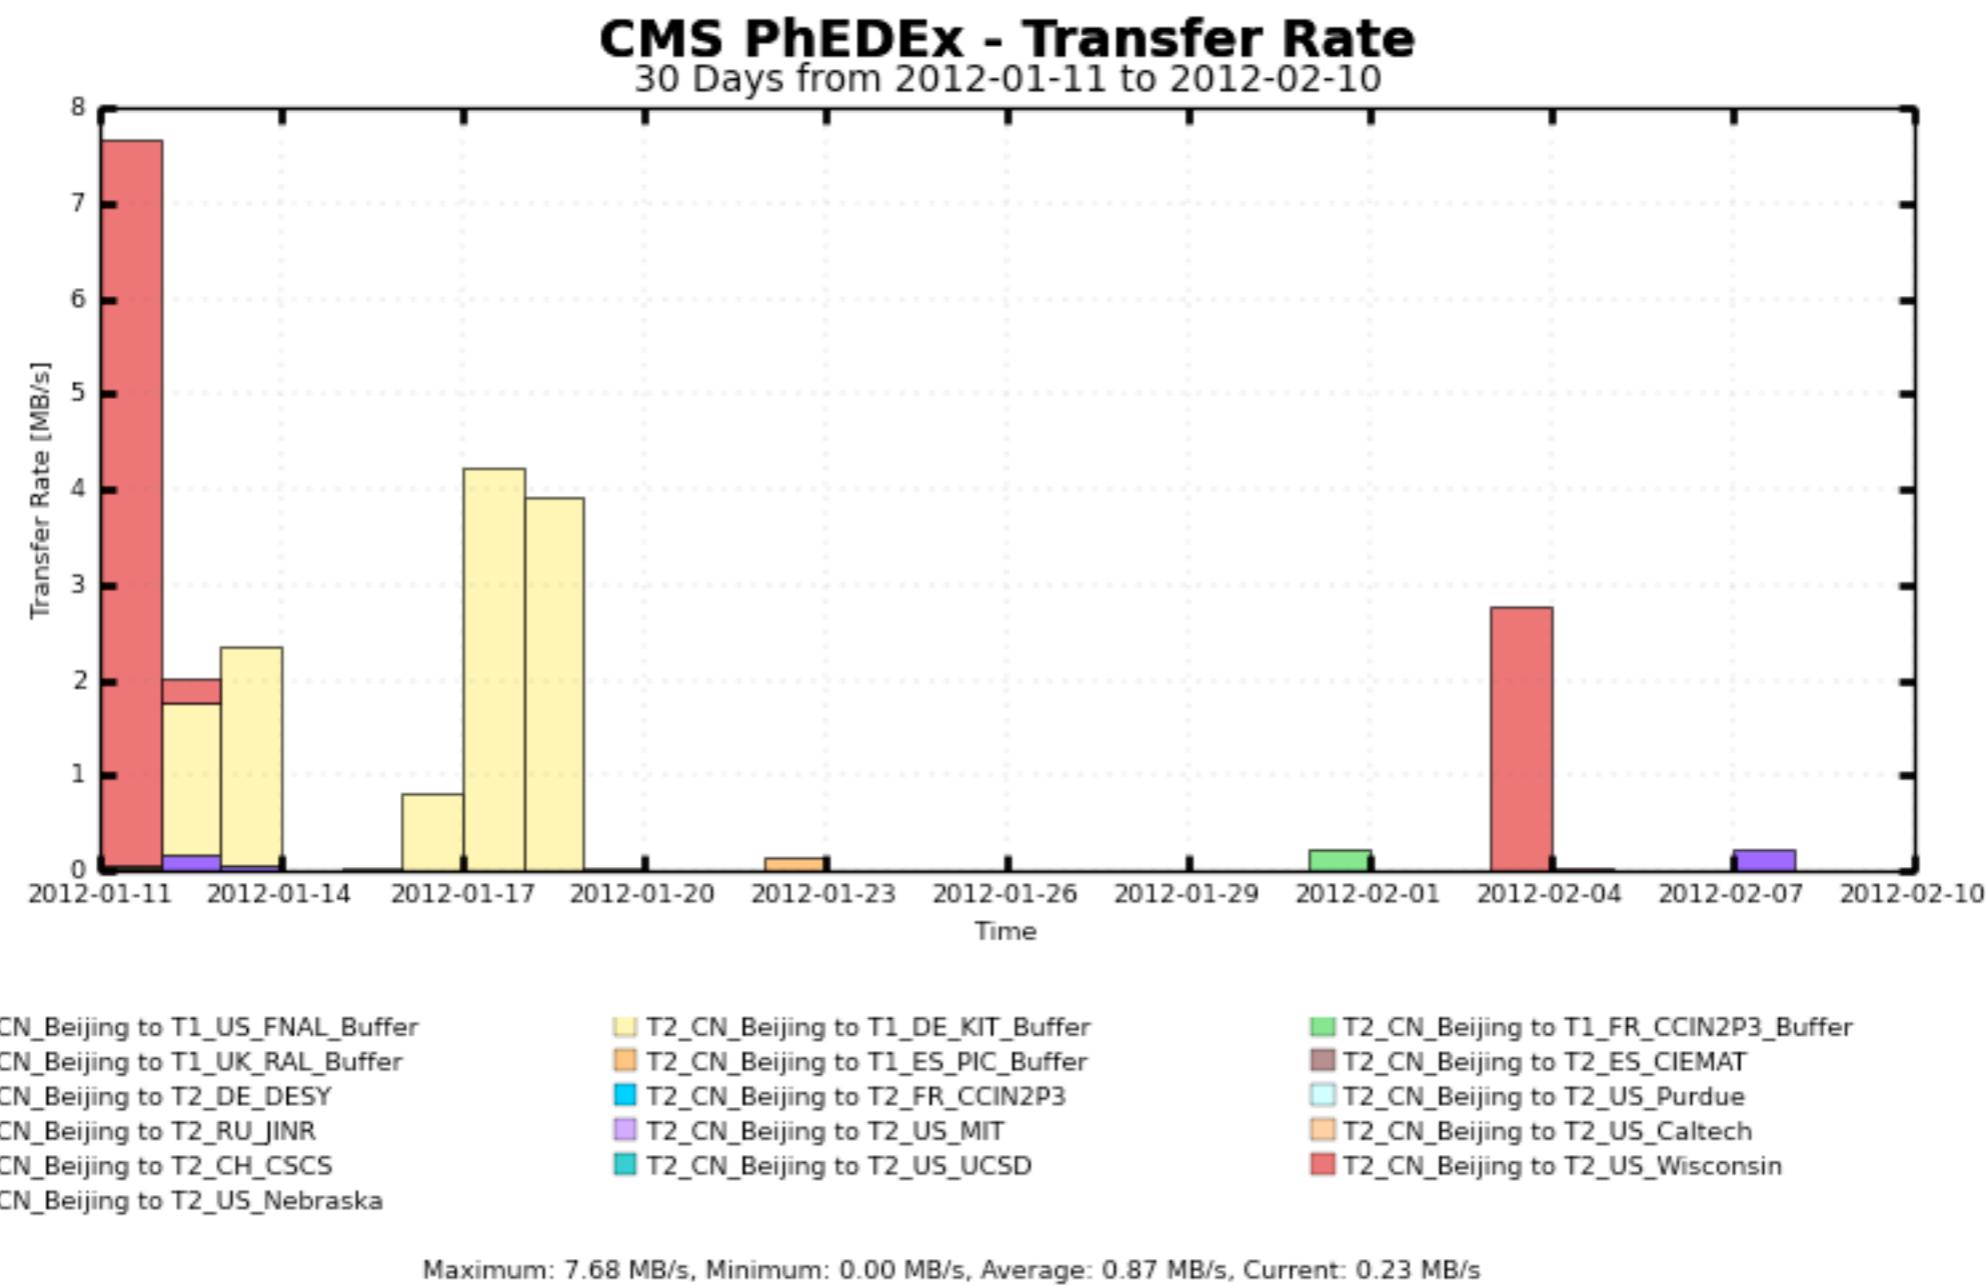
However, the observed absolute transfer rates are very low compared to the available bandwidth. Is some configuration on the FTS channel needed to take advantage of the network capacity?

● IN2P3-CC - BEIJING-LCG2 (85653 files)
● BEIJING-LCG2 - IN2P3-CC (119771 files)

Source: [ATLAS FTS Monitor](#)

LHC Data Exchange with IHEP (cont.)

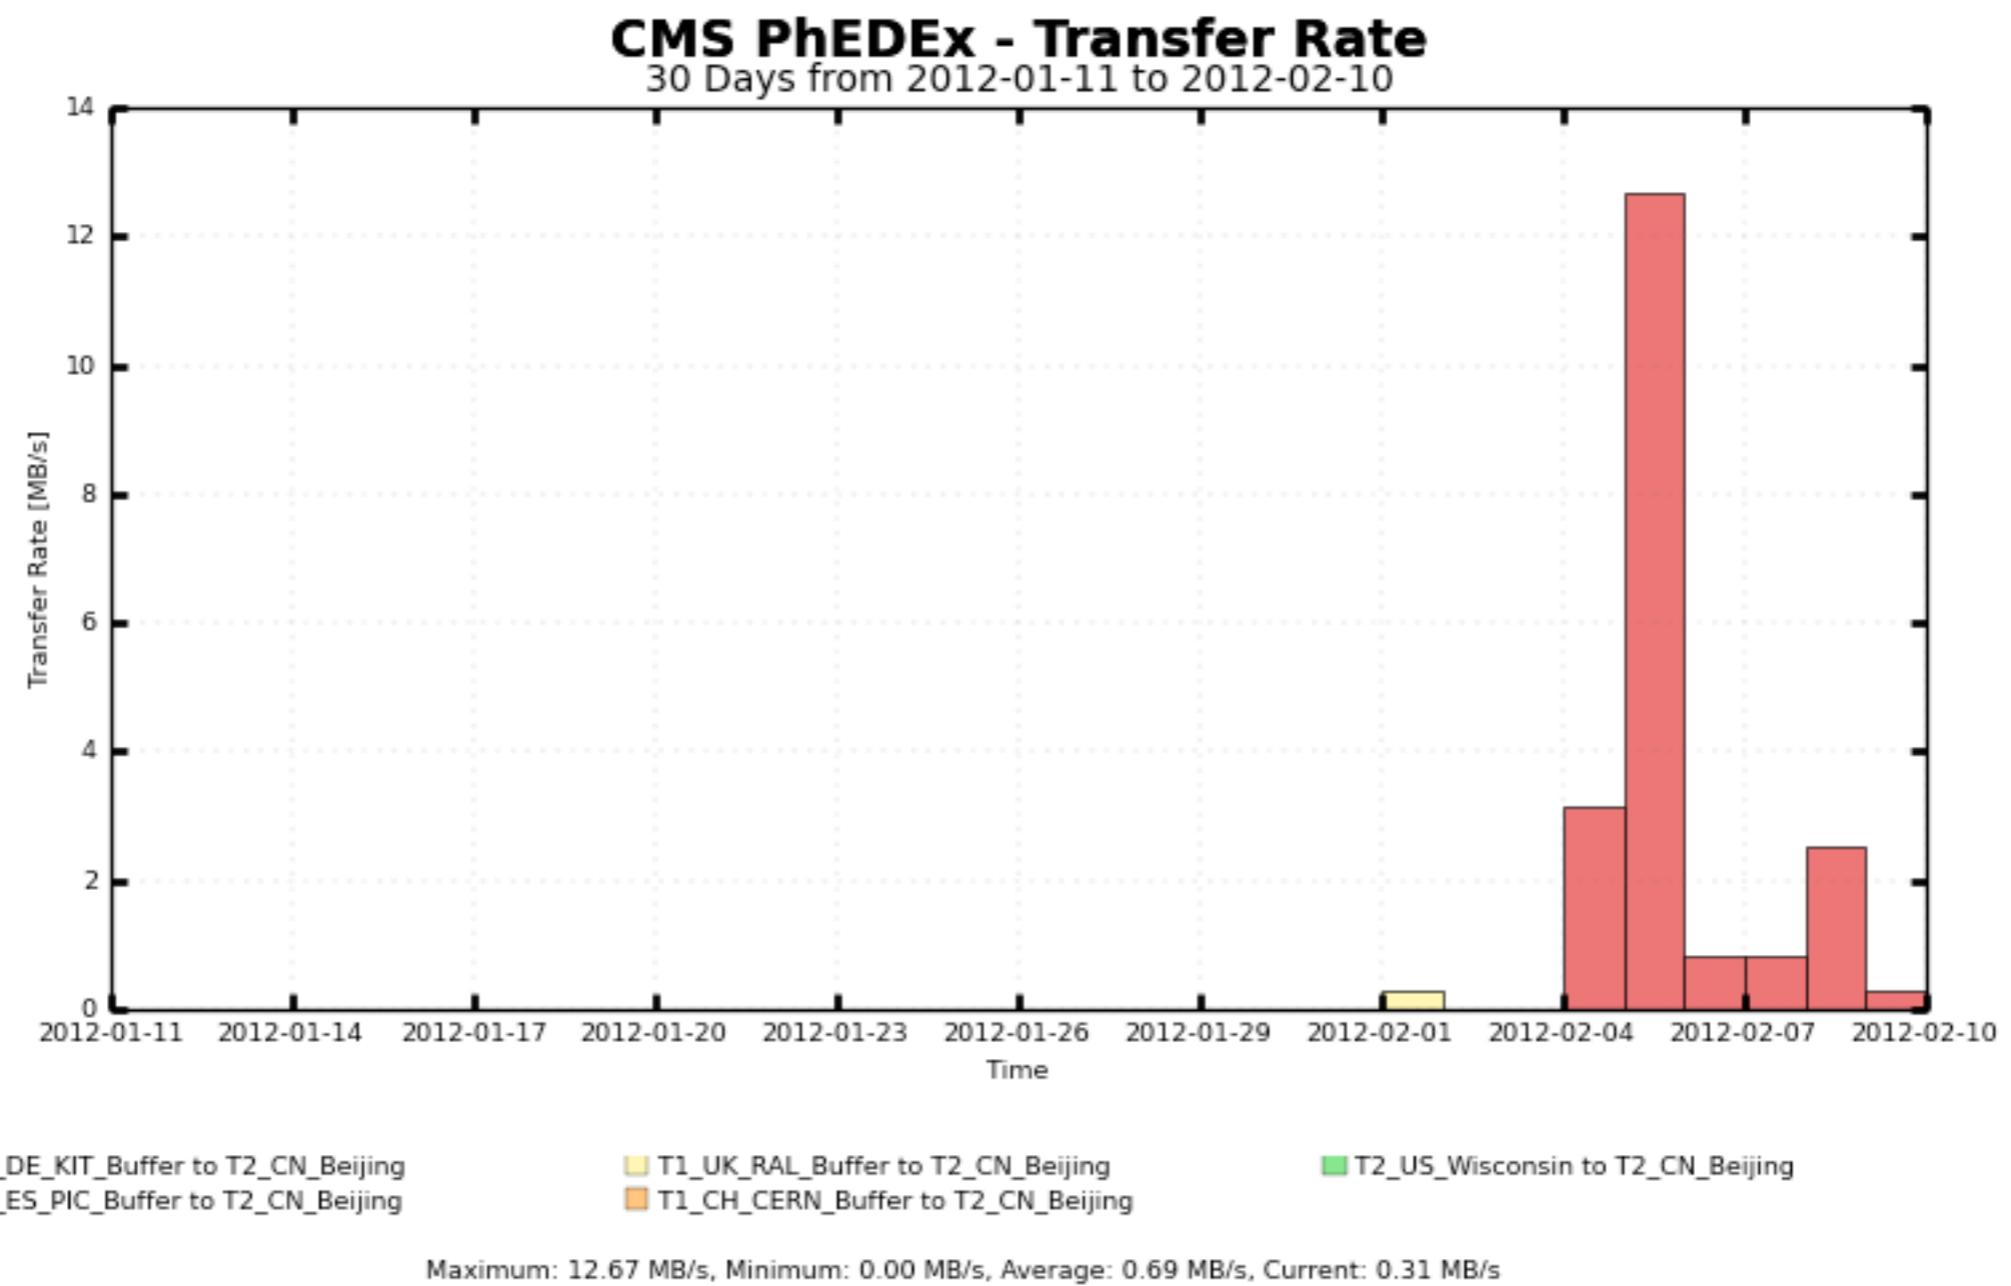
CMS: transfer rate from IHEP



Source: [CMS Phedex](#)

LHC Data Exchange with IHEP (cont.)

CMS: transfer rate to IHEP



Source: [CMS Phedex](#)

What we know so far

- **Routes**

from traceroute results, we observe that NDGF, CNAF, RAL and CC-IN2P3 all use the same link from London to Beijing.

CERN seems to use another route

See annex for details

- The observed throughput from both NDGF and CNAF to IHEP are reasonable good
 - even if the inbound and outbound throughputs are not symmetrical*
- Then, the extremely low throughput from CC-IN2P3 to IHEP seems very suspicious and could be the symptom of an abnormal situation

Acknowledgements

- Thanks to all the network experts who helped us diagnose and who fixed the issue

CSTNet

CNGI6IX and Orient NOC (CERNet, Beijing)

GEANT2

CC-IN2P3

European tier-1s: CNAF, NDGF, PIC, RAL, CERN

Questions & Comments

Annex

NOTE

- The following slides show traceroute results from the peer site's Perfsonar to IHEP's Perfsonar host
- They were collected on Friday 10/02/2012 by 12h15 Beijing time

Traceroute: NDGF to IHEP

```
traceroute to 202.122.32.166 (202.122.32.166), 30 hops max, 140 byte packets
1 dk-ndgf.nordu.net (109.105.124.65)  0.196 ms  0.174 ms  0.162 ms
2 dk-ore2.nordu.net (109.105.102.49)  0.188 ms  0.180 ms  0.169 ms
3 dk-ore.nordu.net (109.105.97.21)  0.243 ms  0.270 ms  0.309 ms
4 dk-uni.nordu.net (109.105.97.14)  0.810 ms  0.800 ms  0.805 ms
5 nordunet-bckp2.rt1.ams.nl.geant.net (62.40.125.205)  13.848 ms  13.886 ms  13.921 ms
6 as1.rt1.lon.uk.geant2.net (62.40.112.138)  21.890 ms  21.916 ms  21.884 ms
7 210.25.189.53 (210.25.189.53)  198.911 ms  198.970 ms  198.996 ms
8 210.25.189.17 (210.25.189.17)  200.534 ms  201.508 ms  202.289 ms
9 210.25.189.75 (210.25.189.75)  198.966 ms  198.994 ms  199.016 ms
10 8.130 (159.226.253.41)  199.321 ms  199.305 ms  199.390 ms
11 8.207 (159.226.253.74)  199.376 ms  199.328 ms  199.316 ms
12 * * *
13 perfsonar.ihep.ac.cn (202.122.32.166)  199.395 ms  199.384 ms  199.377 ms
```

Traceroute: CNAF to IHEP

traceroute to 202.122.32.166 (202.122.32.166), 30 hops max, 140 byte packets

1	gw-mdm.cnaf.infn.it (131.154.254.30)	0.435 ms	0.571 ms	0.706 ms
2	cisco7600.cnaf.infn.it (192.168.150.156)	0.392 ms	0.444 ms	0.533 ms
3	ru-cnaf-rt1-bo1.bo1.garr.net (193.206.128.17)	0.607 ms	0.617 ms	0.680 ms
4	rt1-bo1-rt1-mi1-l2.mi1.garr.net (193.206.134.21)	3.829 ms	3.989 ms	3.976 ms
5	garr.rt1.mil.it.geant.net (62.40.124.129)	3.724 ms	3.784 ms	3.802 ms
6	as0.rt1.gen.ch.geant2.net (62.40.112.33)	11.290 ms	11.302 ms	11.281 ms
7	as1.rt1.par.fr.geant2.net (62.40.112.30)	19.953 ms	19.968 ms	20.015 ms
8	as0.rt1.lon.uk.geant2.net (62.40.112.106)	27.314 ms	27.352 ms	27.388 ms
9	210.25.189.53 (210.25.189.53)	204.511 ms	204.470 ms	204.410 ms
10	210.25.189.17 (210.25.189.17)	205.924 ms	206.735 ms	207.452 ms
11	210.25.189.75 (210.25.189.75)	204.613 ms	204.566 ms	204.610 ms
12	8.130 (159.226.253.41)	204.648 ms	232.038 ms	231.587 ms
13	8.207 (159.226.253.74)	204.457 ms	204.475 ms	204.550 ms
14	perfsonar.ihep.ac.cn (202.122.32.166)	204.807 ms	204.803 ms	204.784 ms

Traceroute: CC-IN2P3 to IHEP

```
traceroute to 202.122.32.166 (202.122.32.166), 30 hops max, 140 byte packets
1 Lyon-OPN.in2p3.fr (193.48.99.100)  0.313 ms  0.312 ms  0.399 ms
2 Lyon-INTER.in2p3.fr (192.70.69.129)  0.875 ms  0.860 ms  0.885 ms
3 * * *
4 te0-0-0-3-paris1-rtr-001.noc.renater.fr (193.51.189.37)  6.664 ms  6.647 ms  9.385 ms
5 renater-LB1.rt1.par.fr.geant.net (62.40.124.69)  5.647 ms  5.679 ms  5.652 ms
6 as0.rt1.lon.uk.geant2.net (62.40.112.106)  13.105 ms  13.081 ms  13.055 ms
7 210.25.189.53 (210.25.189.53)  190.145 ms  190.139 ms  190.123 ms
8 210.25.189.17 (210.25.189.17)  191.558 ms  191.030 ms  191.819 ms
9 210.25.189.75 (210.25.189.75)  190.060 ms  189.989 ms  189.968 ms
10 8.130 (159.226.253.41)  190.469 ms  190.670 ms  190.644 ms
11 8.207 (159.226.253.74)  190.632 ms  190.214 ms  190.191 ms
12 perfsonar.ihep.ac.cn (202.122.32.166)  190.487 ms  190.458 ms  190.571 ms
```

Traceroute: RAL to IHEP

```
traceroute to 202.122.32.166 (202.122.32.166), 30 hops max, 140 byte packets
1 130.246.179.254 (130.246.179.254)  0.743 ms  1.054 ms  1.374 ms
2  ymp8.rl.ac.uk (192.100.78.18)  0.726 ms  1.041 ms  1.411 ms
3  ge-2-0-0.read-sbr1.ja.net (146.97.41.165)  0.693 ms  0.693 ms  0.703 ms
4  ae13.lond-sbr3.ja.net (146.97.33.146)  2.049 ms  2.059 ms  2.075 ms
5  janet.rt1.lon.uk.geant.net (62.40.124.197)  7.792 ms  7.795 ms  7.789 ms
6  210.25.189.53 (210.25.189.53)  179.181 ms  179.132 ms  179.122 ms
7  210.25.189.17 (210.25.189.17)  180.382 ms  181.286 ms  181.959 ms
8  210.25.189.75 (210.25.189.75)  179.266 ms  179.269 ms  179.262 ms
9  8.130 (159.226.253.41)  179.674 ms  179.677 ms  179.671 ms
10 8.207 (159.226.253.74)  179.489 ms  179.469 ms  179.532 ms
11 perfsonar.ihep.ac.cn (202.122.32.166)  179.764 ms  179.736 ms  179.731 ms
```

Traceroute: CERN to IHEP

traceroute to 202.122.32.166 (202.122.32.166), 30 hops max, 140 byte packets

- 1 l513-c-rftec-6-ip8.cern.ch (128.142.223.225) 0.415 ms 0.413 ms 0.426 ms
- 2 l513-b-rbrmx-3-gl35.cern.ch (194.12.137.57) 0.289 ms l513-b-rftex-4-rl15.cern.ch (194.12.138.57) 0.556 ms l513-b-rbrmx-2-hl6.cern.ch (194.12.138.25) 0.280 ms
- 3 b513-b-rbrml-1-ml1.cern.ch (194.12.148.13) 0.275 ms b513-b-rbrml-1-ml3.cern.ch (194.12.148.17) 0.313 ms b513-b-rbrml-1-ml1.cern.ch (194.12.148.13) 0.304 ms
- 4 g513-e-rci76-1-pg1.cern.ch (194.12.144.9) 0.667 ms 0.719 ms 0.545 ms
- 5 g513-e-rci76-1-fe1.cern.ch (192.65.184.165) 0.755 ms 0.796 ms 0.850 ms
- 6 e513-e-rbrxl-1-pe1.cern.ch (192.65.184.161) 0.674 ms 0.650 ms 0.642 ms
- 7 e513-e-rbrxl-2-ne0.cern.ch (192.65.184.38) 0.603 ms 0.708 ms 0.700 ms
- 8 kreonet2.cern.kreonet.net (134.75.108.57) 176.317 ms 177.879 ms 176.296 ms
- 9 134.75.108.209 (134.75.108.209) 299.397 ms 299.724 ms 299.373 ms
- 10 134.75.105.2 (134.75.105.2) 349.853 ms 351.249 ms 351.236 ms
- 11 134.75.108.102 (134.75.108.102) 328.882 ms 328.867 ms 329.830 ms
- 12 8.192 (159.226.254.165) 328.465 ms 329.352 ms 329.802 ms
- 13 8.198 (159.226.254.253) 365.473 ms 365.402 ms 365.900 ms
- 14 8.131 (159.226.253.53) 284.625 ms 286.040 ms 285.618 ms
- 15 8.207 (159.226.253.78) 298.332 ms 298.322 ms 298.305 ms
- 16 perfsonar.ihep.ac.cn (202.122.32.166) 285.915 ms 285.582 ms 287.297 ms

Segmented throughput tests

iperf:IHEP->62.40.123.114

```
[root@performance ~]# iperf -c 62.40.123.114 -w 16m -t 180
```

Client connecting to 62.40.123.114, TCP port 5001

TCP window size: 32.0 MByte (WARNING: requested 16.0 MByte)

[3] local 202.122.32.162 port 33636 connected with 62.40.123.114

port

5001

[ID] Interval Transfer Bandwidth

[3] 0.0-196.6 sec 150 MBytes 6.38 Mbits/sec

IHEP → GEANT2

: -c 103.2.208.66 -w 16m -t 180

ting to 103.2.208.66, TCP port 5001

ze: 32.0 MByte (WARNING: requested 16.0 MB)

!5.189.112 port 42720 connected with 103.2.20

Transfer Bandwidth

sec 19.7 GBytes 942 Mbits/sec

IHEP → CSTNet

```
[root@192_168_0_3 ~]# iperf -c 210.25.189.112
```

Client connecting to 210.25.189.112, TCP port 5001

TCP window size: 42.9 MByte (default)

[3] local 103.2.208.66 port 59640 connected with 210.25.189.112 port 5001

[ID] Interval Transfer Bandwidth

[3] 0.0-10.0 sec 1.13 GBytes 965 Mbits/sec

CSTNet → CNGI6IX

ihep@debian:\$ iperf -c ccxfert01.in2p3.fr -i 2 -t 60 -w 15M

Client connecting to ccxfert01.in2p3.fr, TCP port 5001

TCP window size: 30.0 MByte (WARNING: requested 15.0 MByte)

[3] local 210.25.189.112 port 53821 connected with 193.48.99.201 port

5001

[ID] Interval Transfer Bandwidth

[3] 0.0- 2.0 sec 20.0 MBytes 83.9 Mbits/sec

[3] 2.0- 4.0 sec 0.00 Bytes 0.00 bits/sec

[3] 4.0- 6.0 sec 14.2 MBytes 59.8 Mbits/sec

[3] 6.0- 8.0 sec 23.9 MBytes 100 Mbits/sec

[3] 8.0-10.0 sec 31.9 MBytes 134 Mbits/sec

[3] 10.0-12.0 sec 51.9 MBytes 218 Mbits/sec

[3] 12.0-14.0 sec 91.4 MBytes 383 Mbits/sec

[3] 14.0-16.0 sec 129 MBytes 541 Mbits/sec

[3] 16.0-18.0 sec 128 MBytes 537 Mbits/sec

[3] 18.0-20.0 sec 133 MBytes 558 Mbits/sec

[3] 20.0-22.0 sec 135 MBytes 565 Mbits/sec

[3] 22.0-24.0 sec 140 MBytes 588 Mbits/sec

CNGI6IX → CC-IN2P3

[root@performance ~]# iperf -c 210.25.189.112 -i 2 -t 60 -w 15M

Client connecting to 210.25.189.112, TCP port 5001

TCP window size: 30.0 MByte (WARNING: requested 15.0 MByte)

[3] local 202.122.32.162 port 51849 connected with 210.25.189.112

port 5001

[ID] Interval Transfer Bandwidth

[3] 0.0- 2.0 sec 42.1 MBytes 177 Mbits/sec

[3] 2.0- 4.0 sec 0.00 Bytes 0.00 bits/sec

[3] 4.0- 6.0 sec 10.5 MBytes 44.0 Mbits/sec

[3] 6.0- 8.0 sec 20.9 MBytes 87.6 Mbits/sec

[3] 8.0-10.0 sec 10.2 MBytes 43.0 Mbits/sec

[3] 10.0-12.0 sec 10.8 MBytes 45.1 Mbits/sec

[3] 12.0-14.0 sec 10.6 MBytes 44.6 Mbits/sec

[3] 14.0-16.0 sec 10.6 MBytes 44.6 Mbits/sec

[3] 16.0-18.0 sec 20.8 MBytes 87.0 Mbits/sec

[3] 18.0-20.0 sec 10.8 MBytes 45.1 Mbits/sec

[3] 20.0-22.0 sec 10.8 MBytes 45.1 Mbits/sec

IHEP → CNGI6IX