LHCONE for ATLAS Tokyo Tier-2

Tomoaki Nakamura
on behalf of ICEPP regional analysis center group
ICEPP, The University of Tokyo

東京大学
THE UNIVERSITY OF TOKYO

Evolution of disk storage capacity in Tokyo

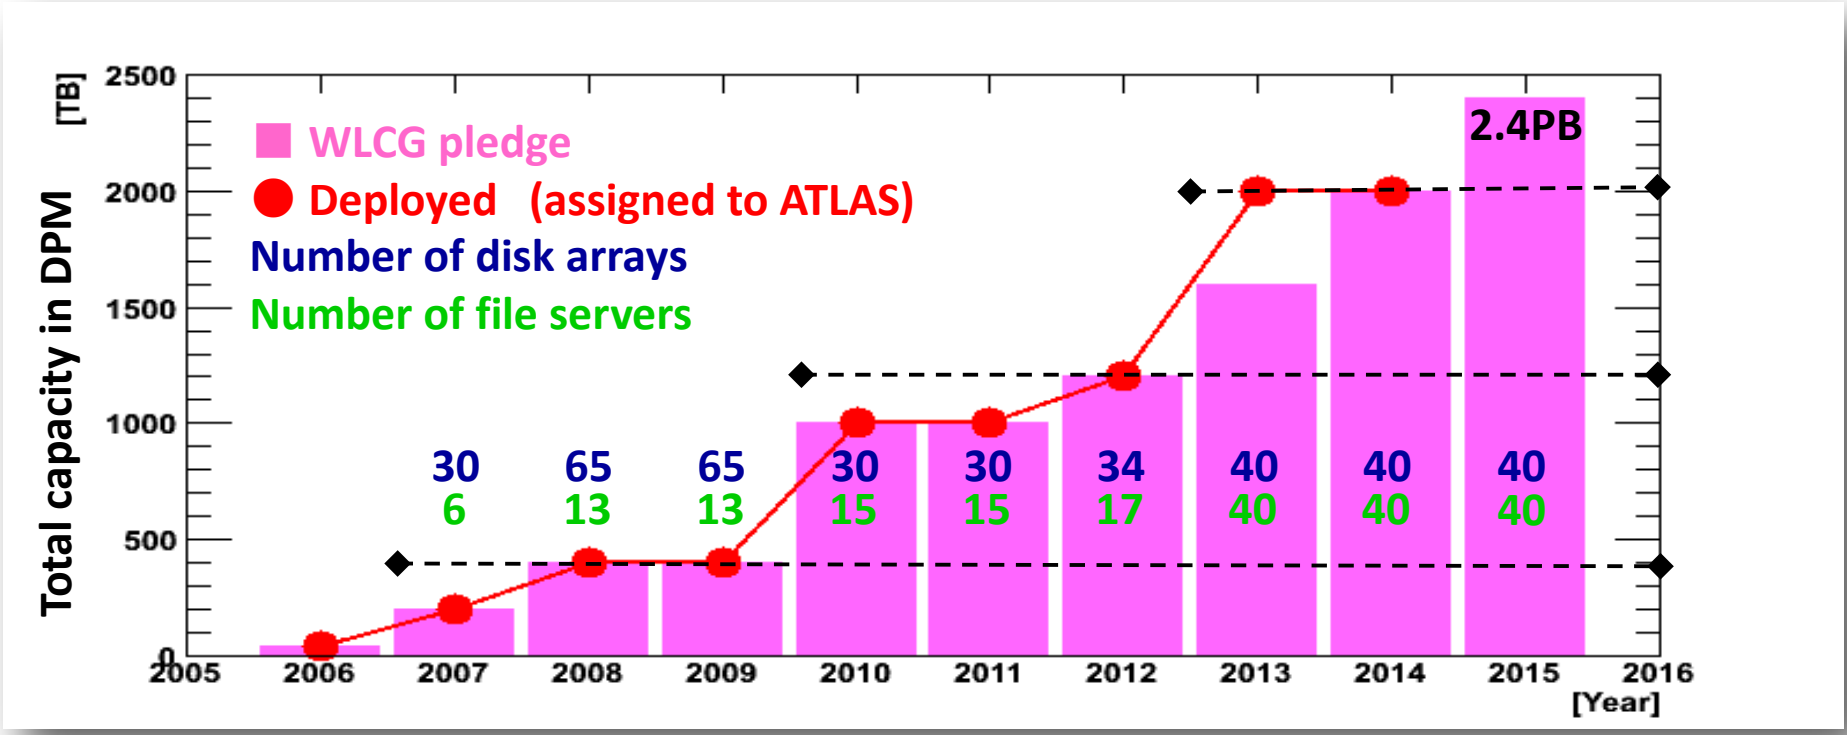

Pilot system for R&D

1st system 2007 - 2009

16x500GB HDD / array
5disk arrays / server
XFS on RAID6
4G-FC via FC switch
10GE NIC

2nd system 2010 - 2012

24x2TB HDD / array
2disk arrays / server
XFS on RAID6
8G-FC via FC switch
10GE NIC

3rd system 2013 - 2015

24x3TB HDD /array
1disk array / server
XFS on RAID6
8G-FC without FC switch
10GE NIC

4th system 2016 - 2019

International network for Tokyo

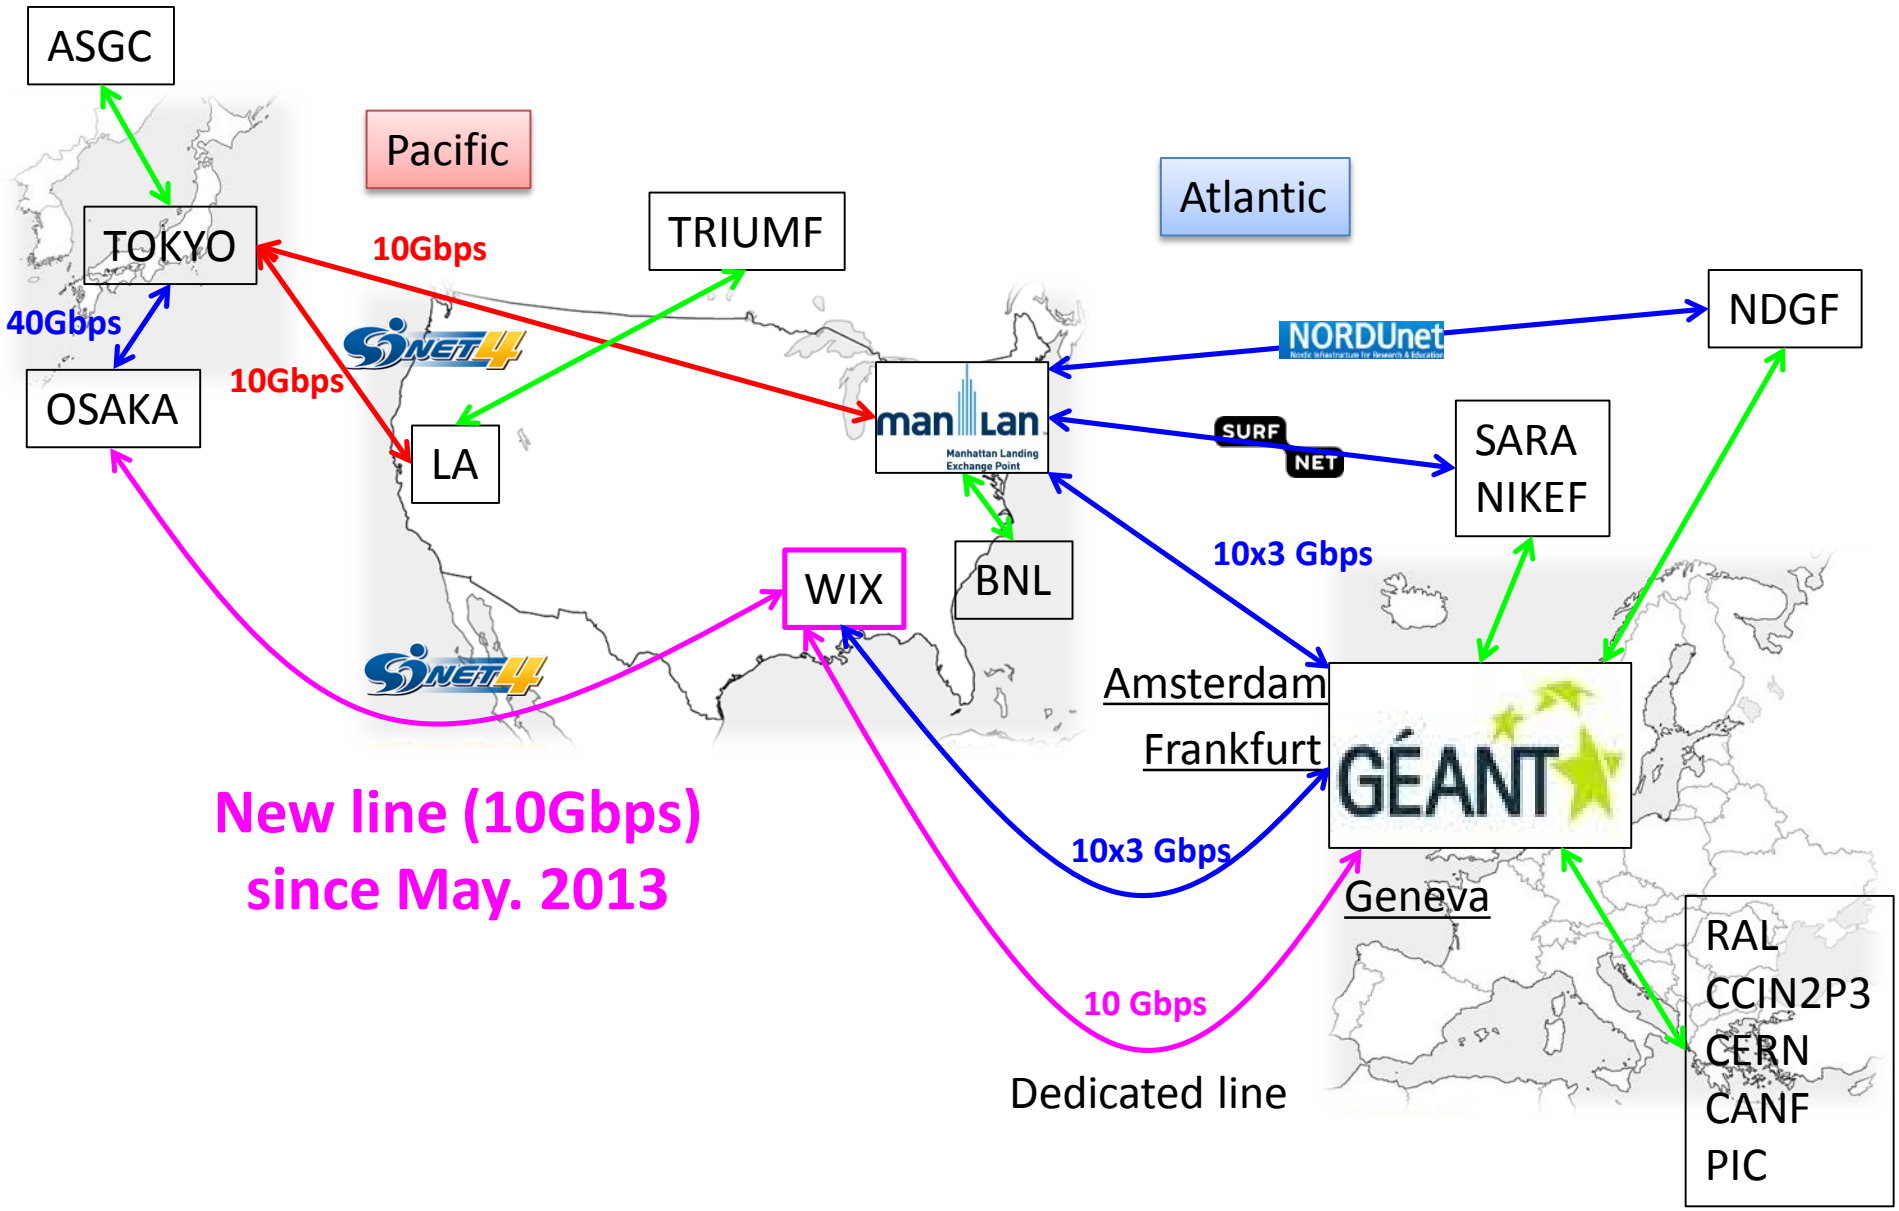

Transfer throughput in FY2013 (1 day average)

Monitored by network switch

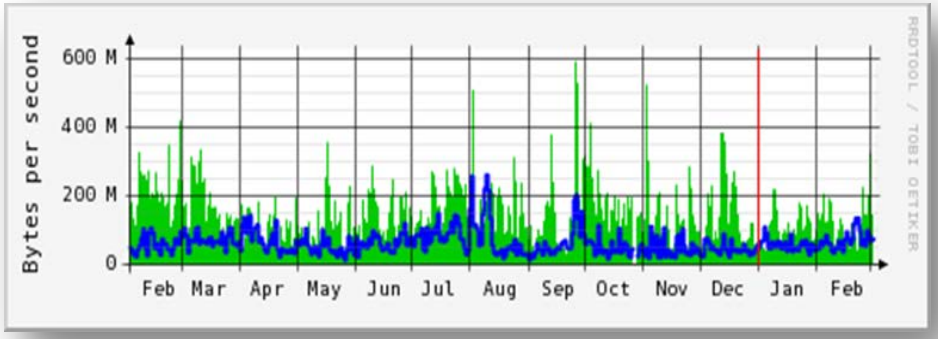

Incoming
Outgoing

Monitored by file servers (extracted from grid FTP logs)

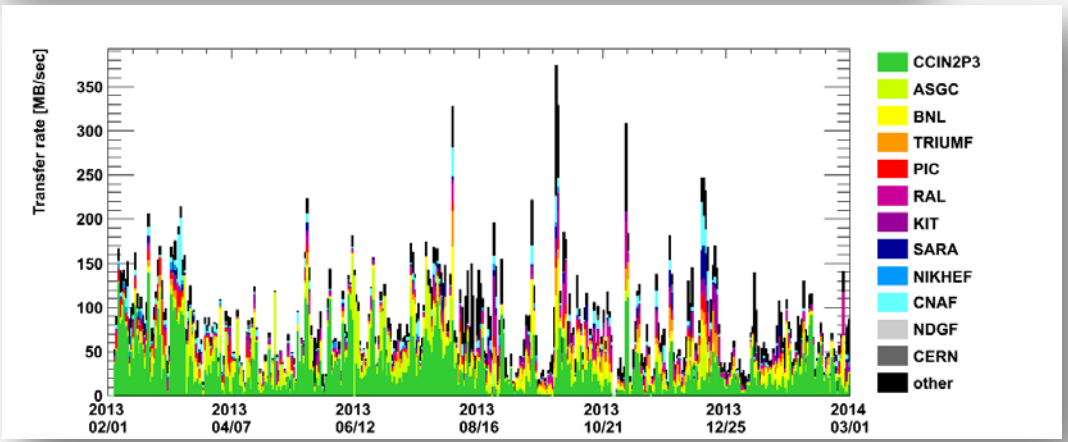

Incoming data

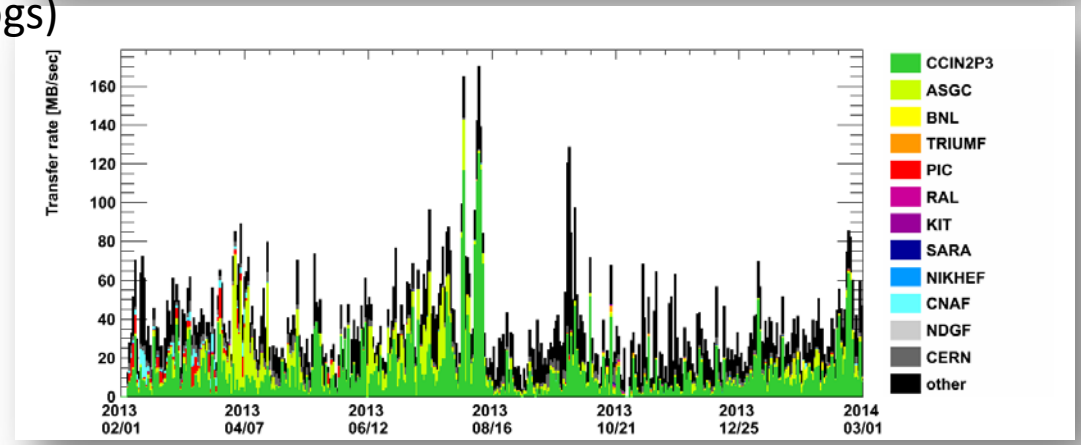

Outgoing data

Transfer throughput in FY2013 (10 min. average)

Sustained transfer rate

Incoming data: ~100MB/sec in one day average

Outgoing data: ~50MB/sec in one day average

300~400TB of data in Tokyo storage is replaced within one month!

Peak transfer rate

Almost reached to 10Gbps

Need to increase bandwidth and stability!

Monitored by file servers
(extracted from grid FTP logs)

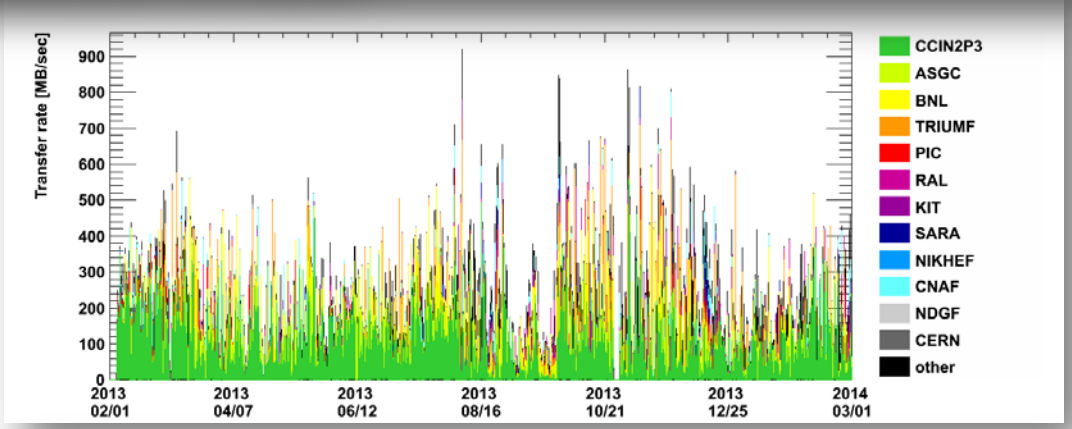

Incoming data

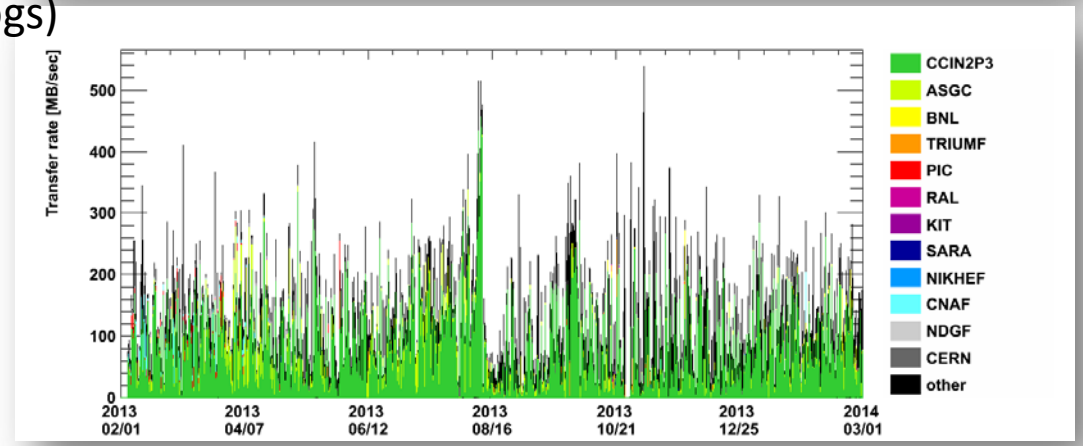

Outgoing data

Configuration for the LHCONE evaluation

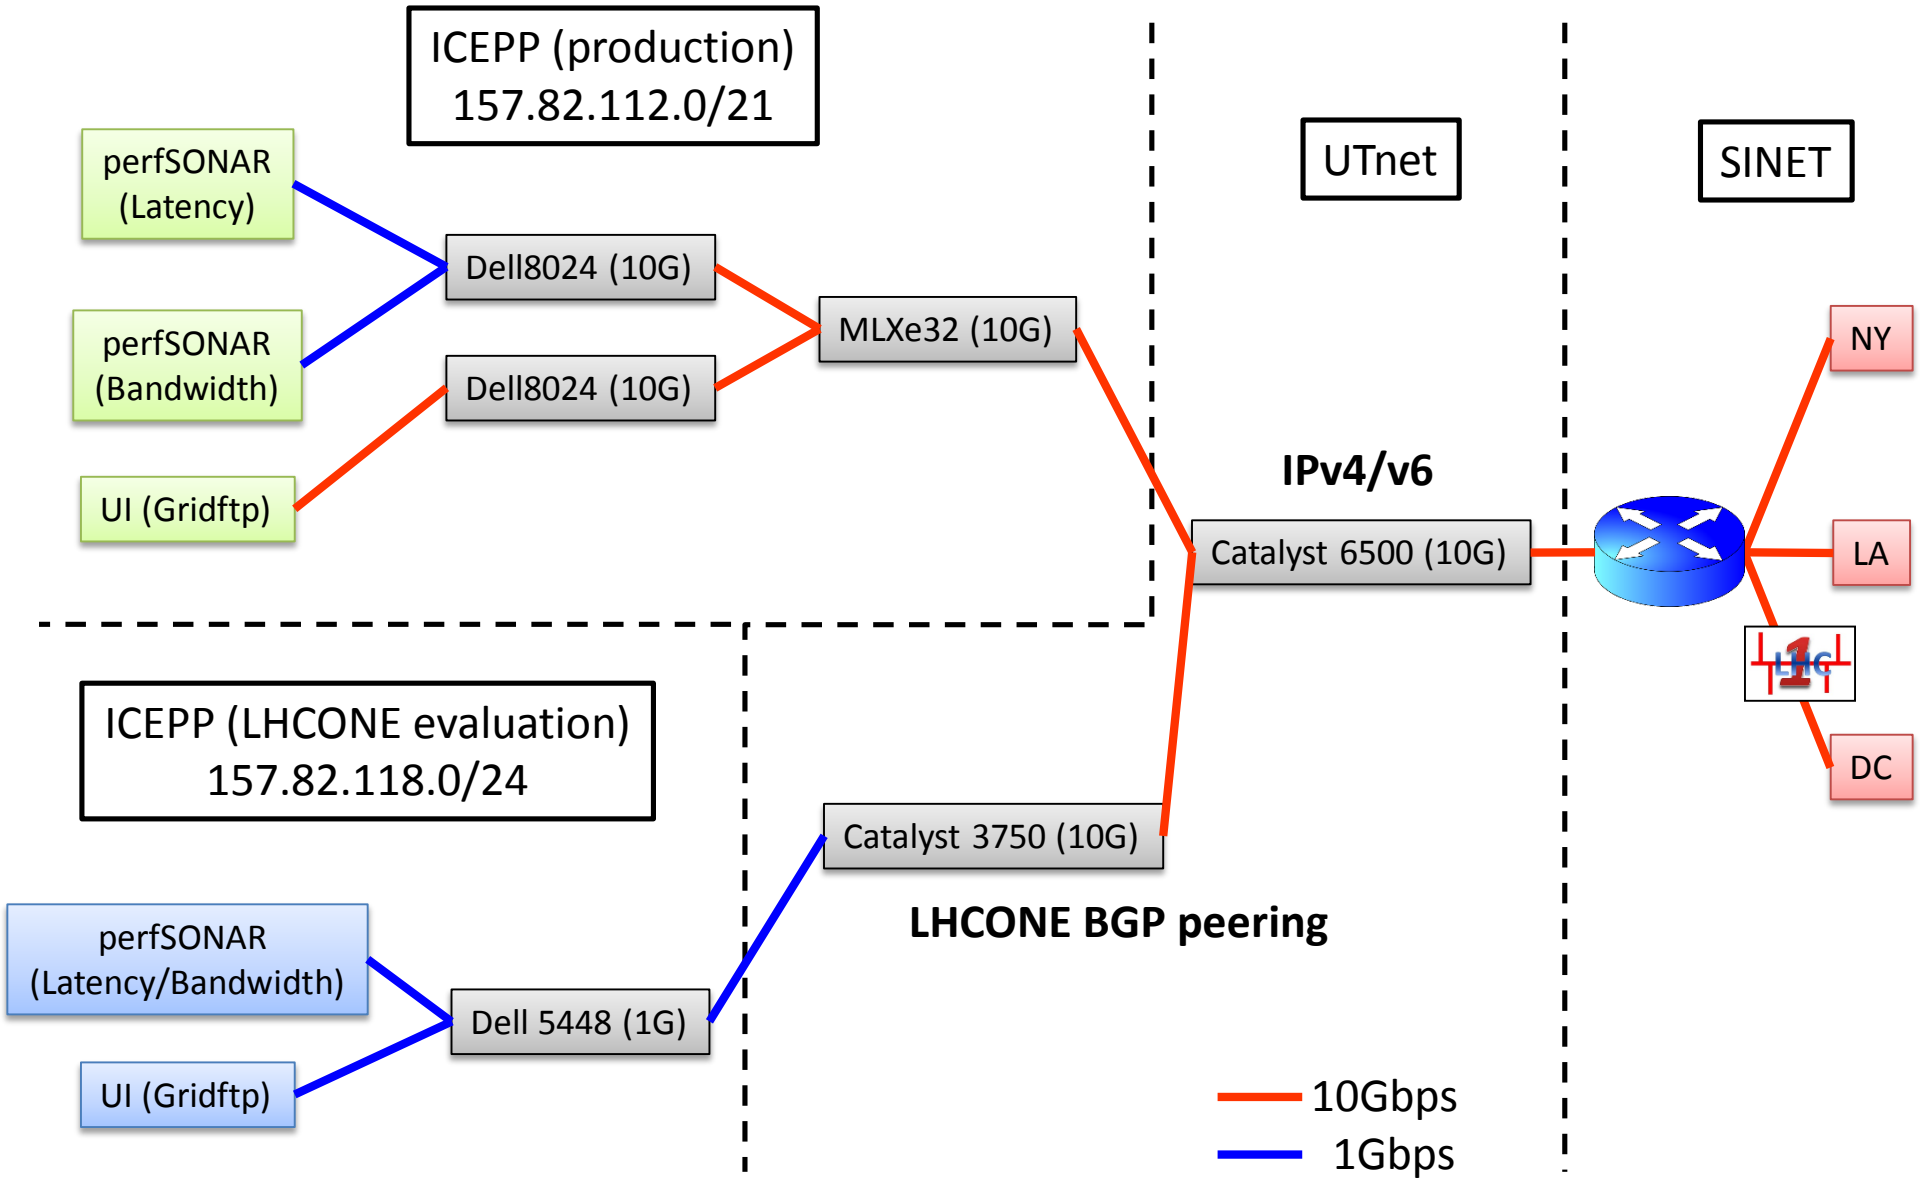

Stability on packet loss (CC-IN2P3)

Directly affect to transfer rate.

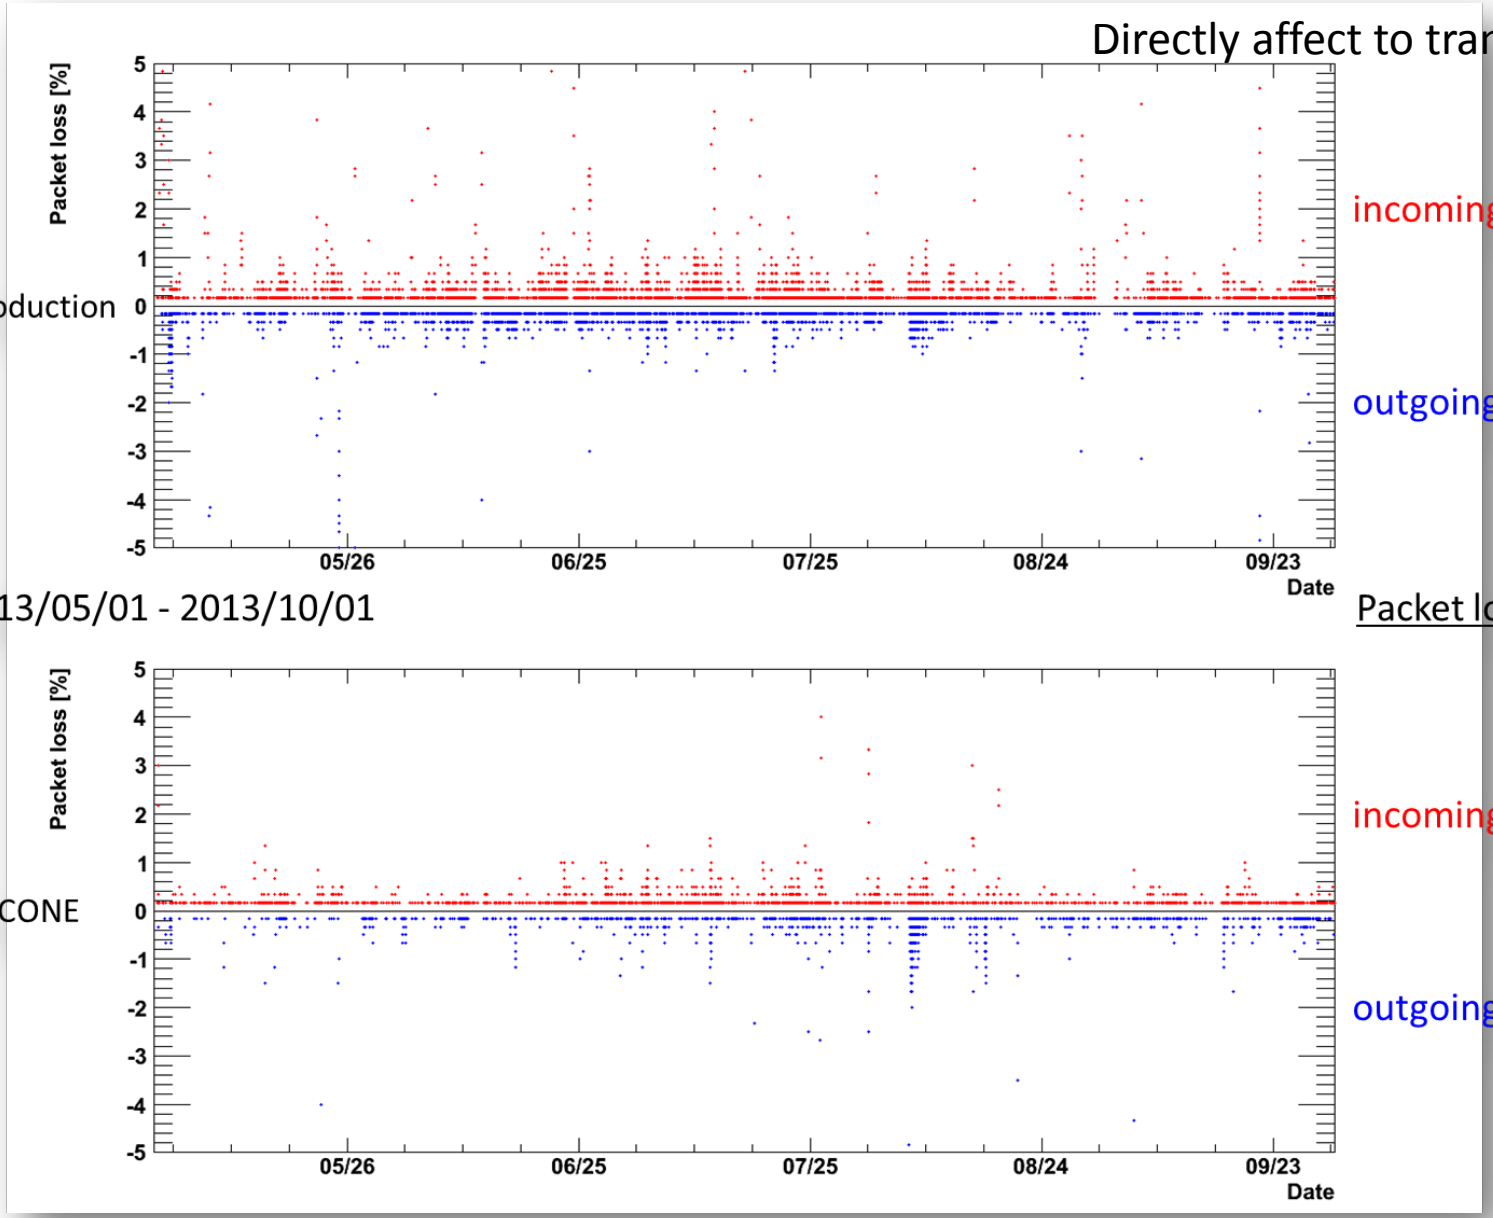

Fraction of packet loss (NY vs. DC)

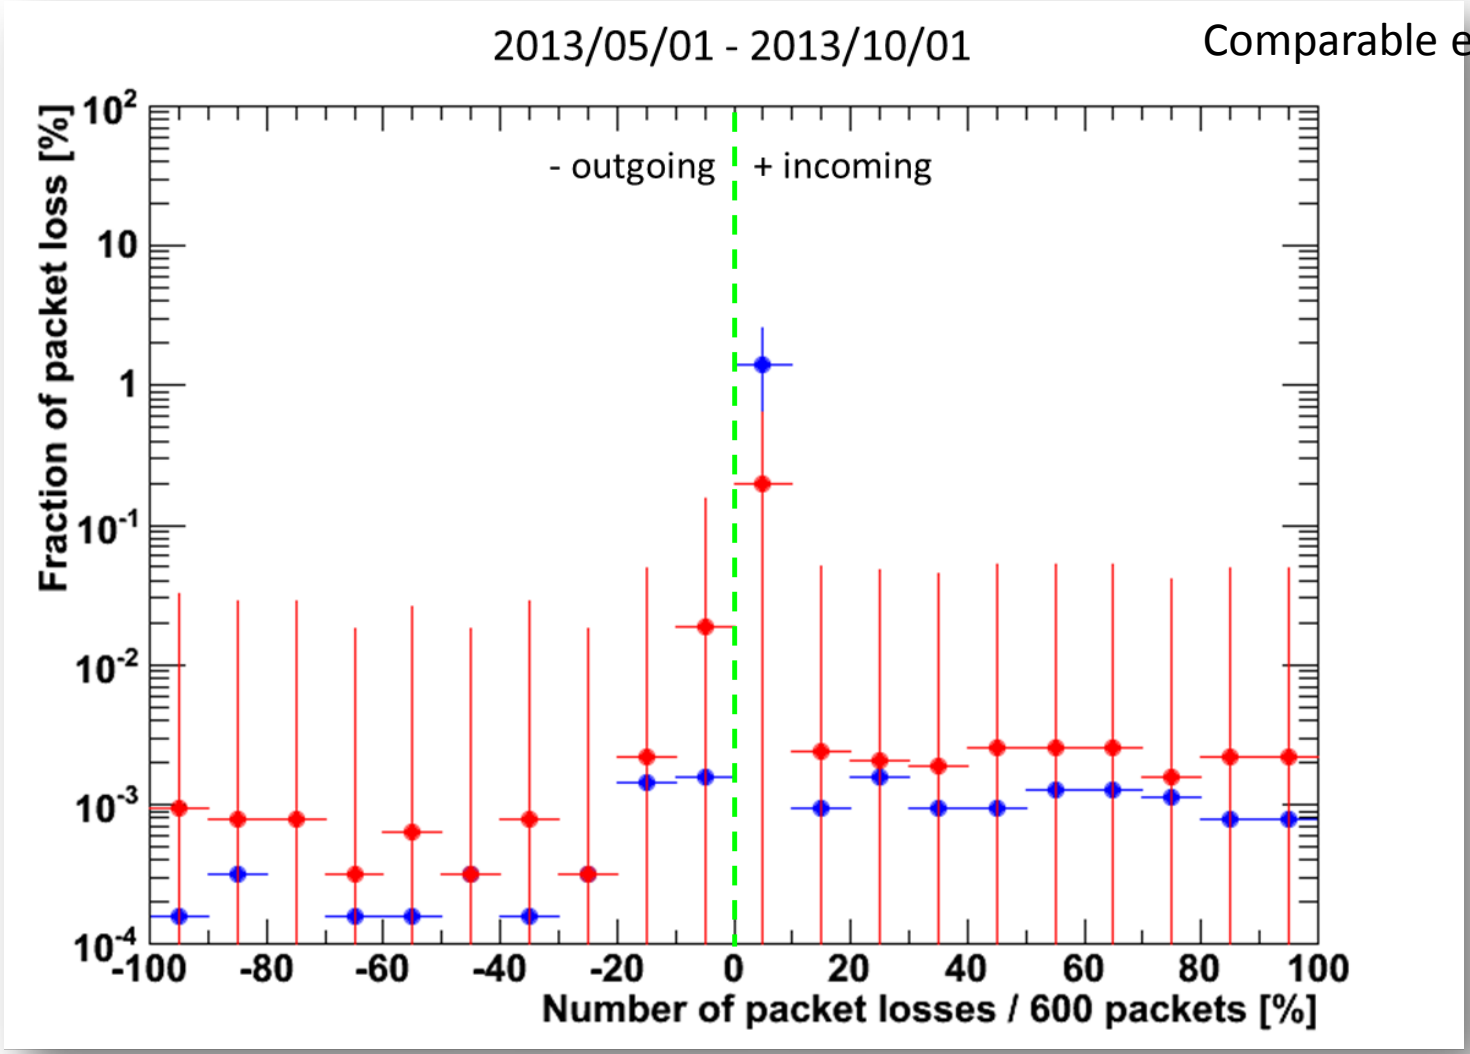

Comparable each other.

- NII perfsonar in NY – TOKYO (production)
- NII perfsonar in DC – TOKYO (LHCONE)

Minimum latency (CC-IN2P3)

Useful to know the typical latency and stability.

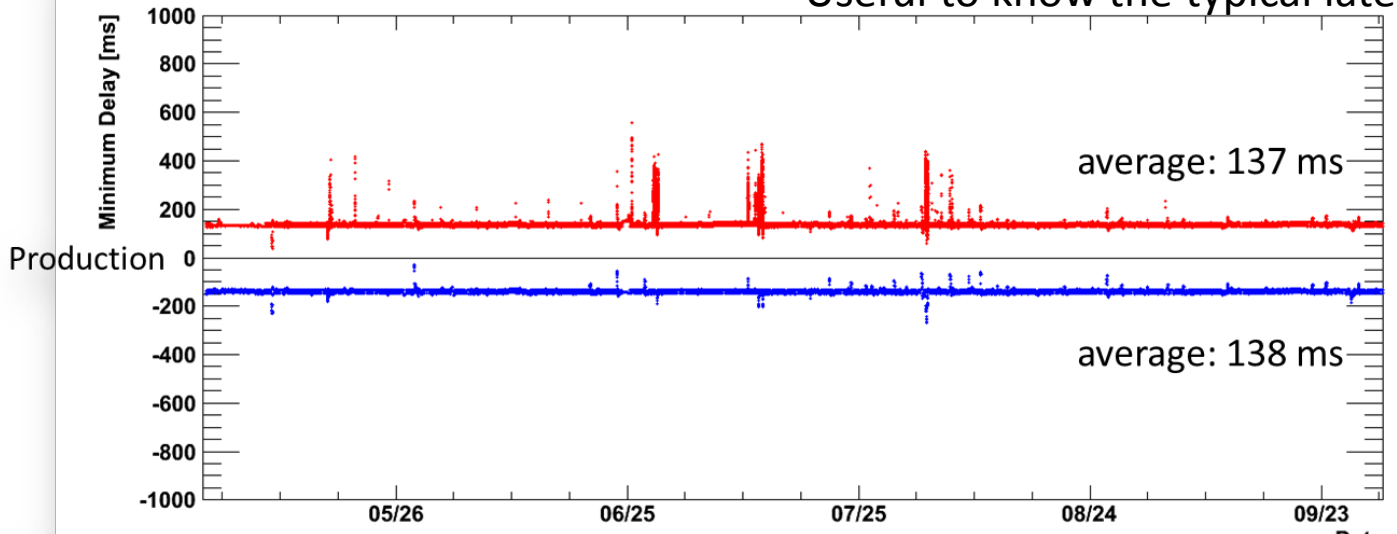

Production

2013/05/01 - 2013/10/01

Minimum latency

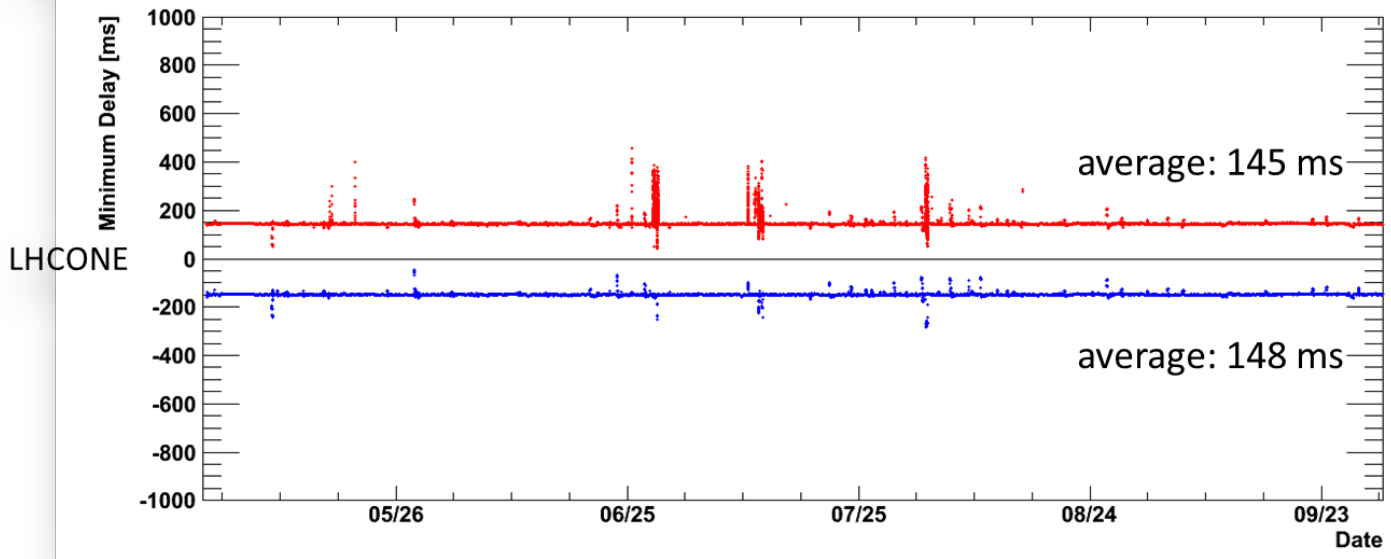

LHCONE

Minimum latency (CC-IN2P3)

Originating from other group in Univ. of Tokyo.

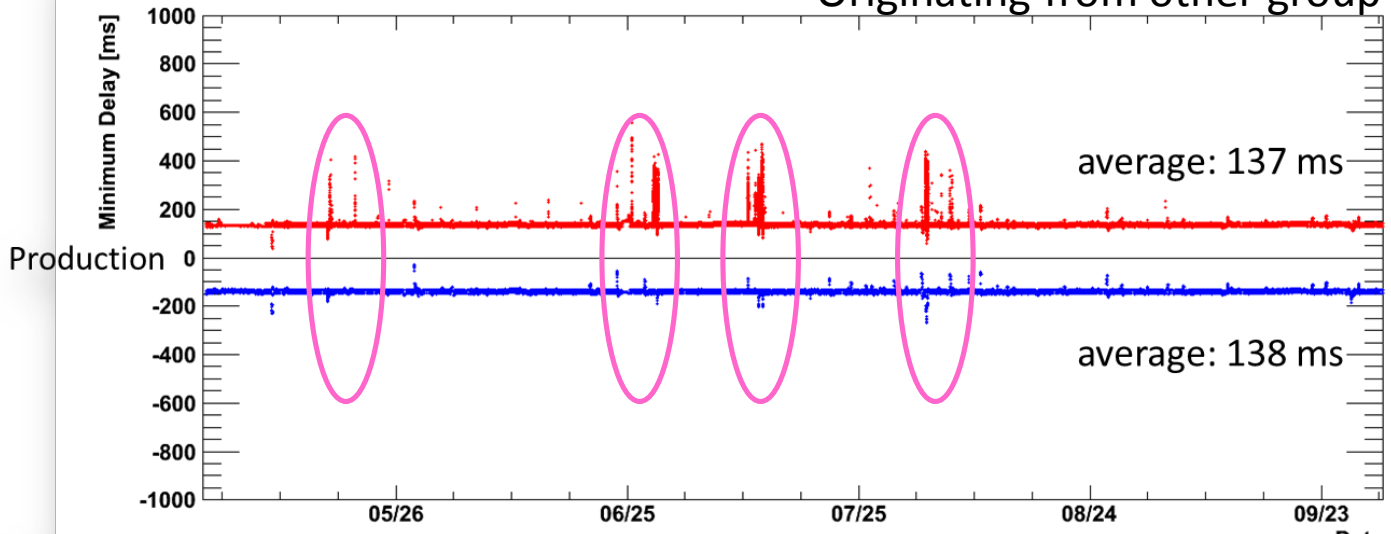

Production

2013/05/01 - 2013/10/01

Minimum latency

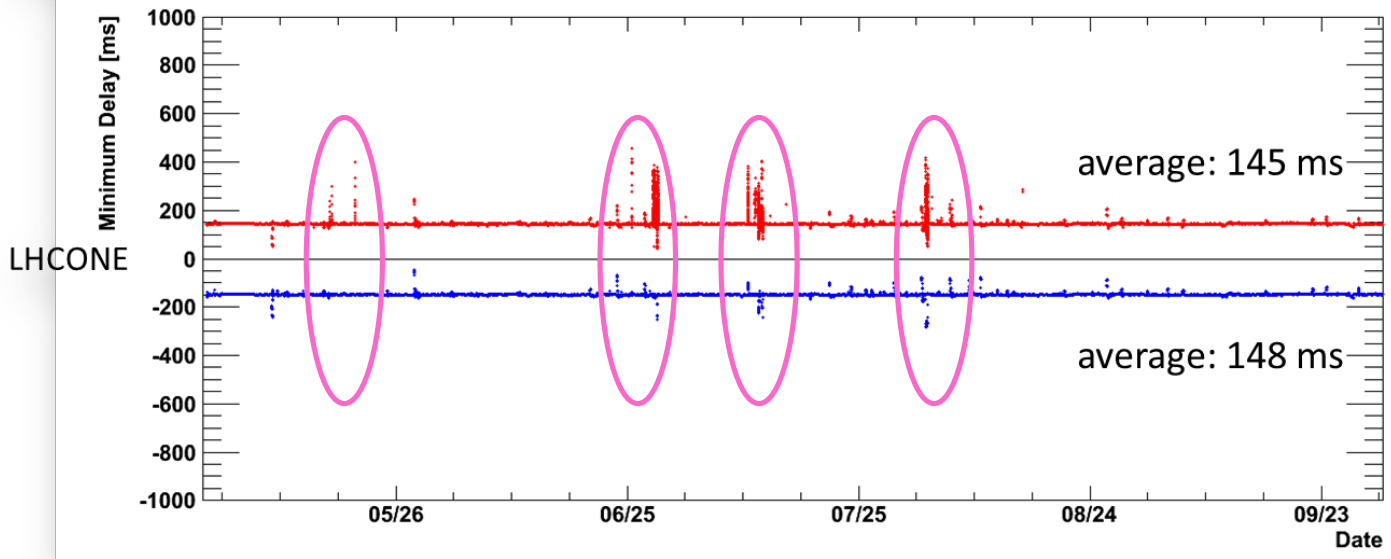

LHCONE

Distribution of Minimum latency (CC-IN2P3)

Production

RMS: 11.9

RMS: 9.0

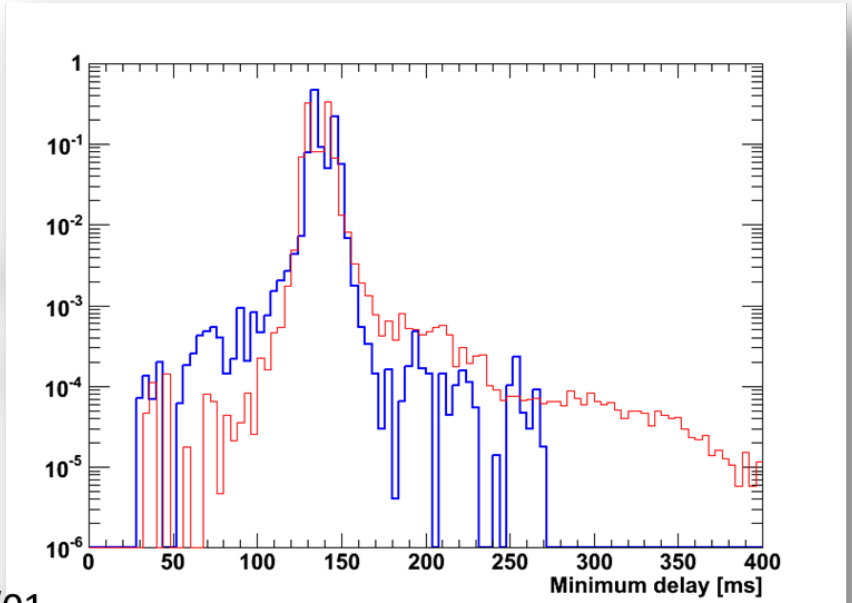

incoming

outgoing

2013/05/01 - 2013/10/01

Minimum latency

LHCONE

RMS: 9.0

RMS: 7.9

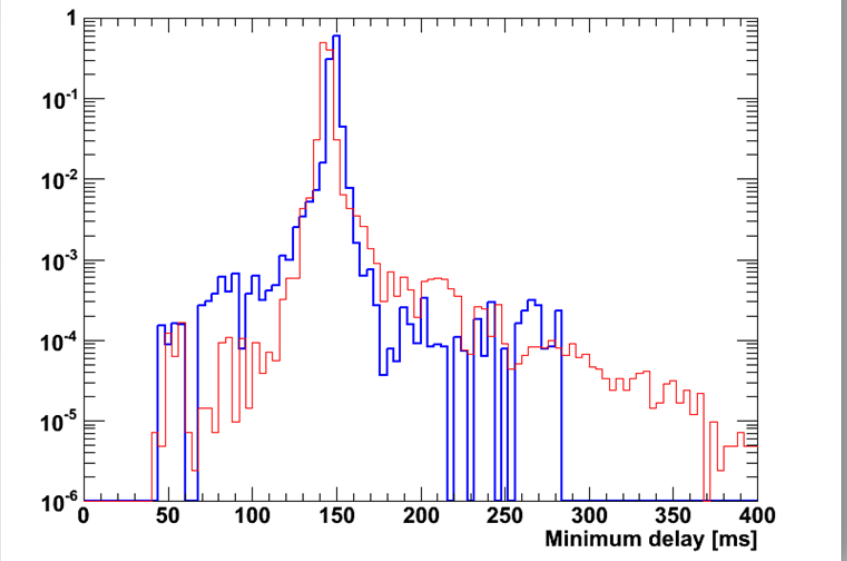

incoming

outgoing

Distribution of Minimum latency (CC-IN2P3)

miss measurement. originating from other group.

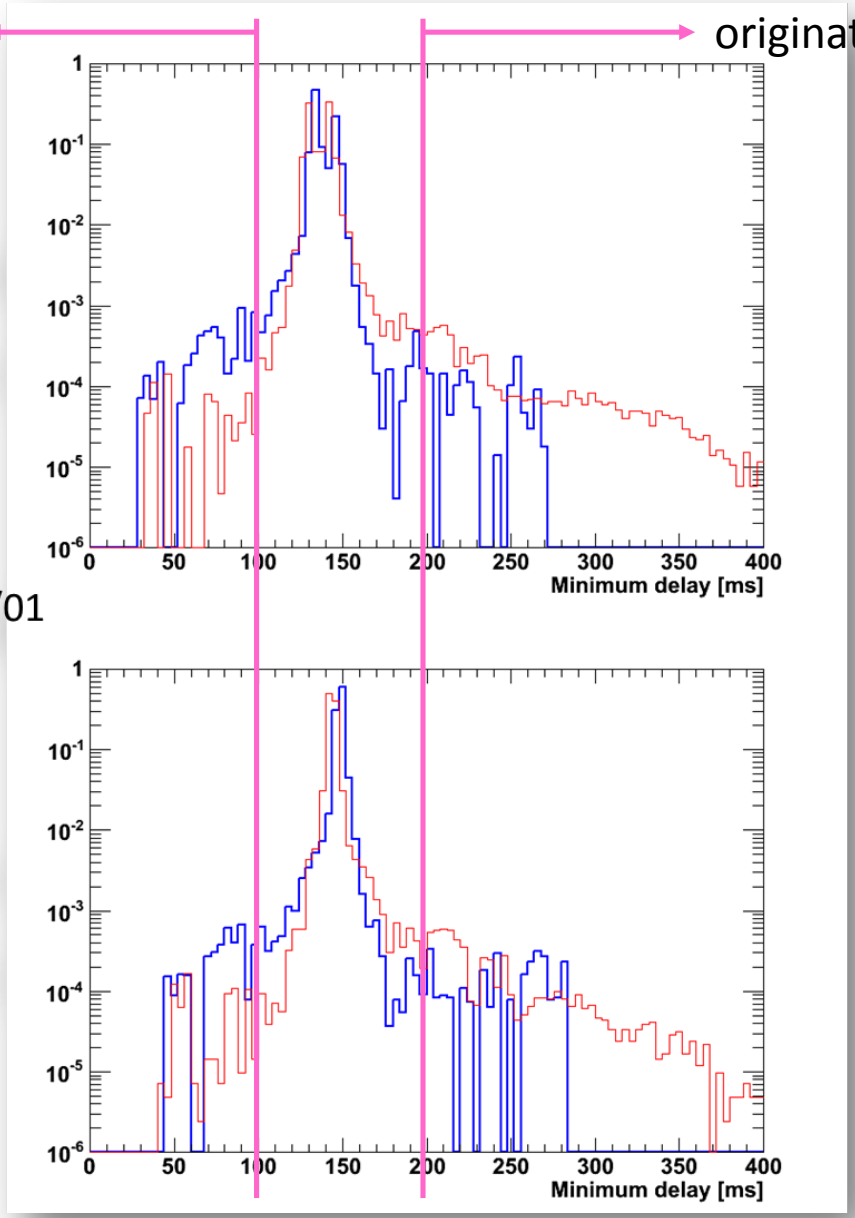

Production

RMS: 11.9

RMS: 9.0

incoming

outgoing

2013/05/01 - 2013/10/01

Minimum latency

LHCONE

RMS: 9.0

RMS: 7.9

incoming

outgoing

Maximum latency (CC-IN2P3)

Useful to find problems.

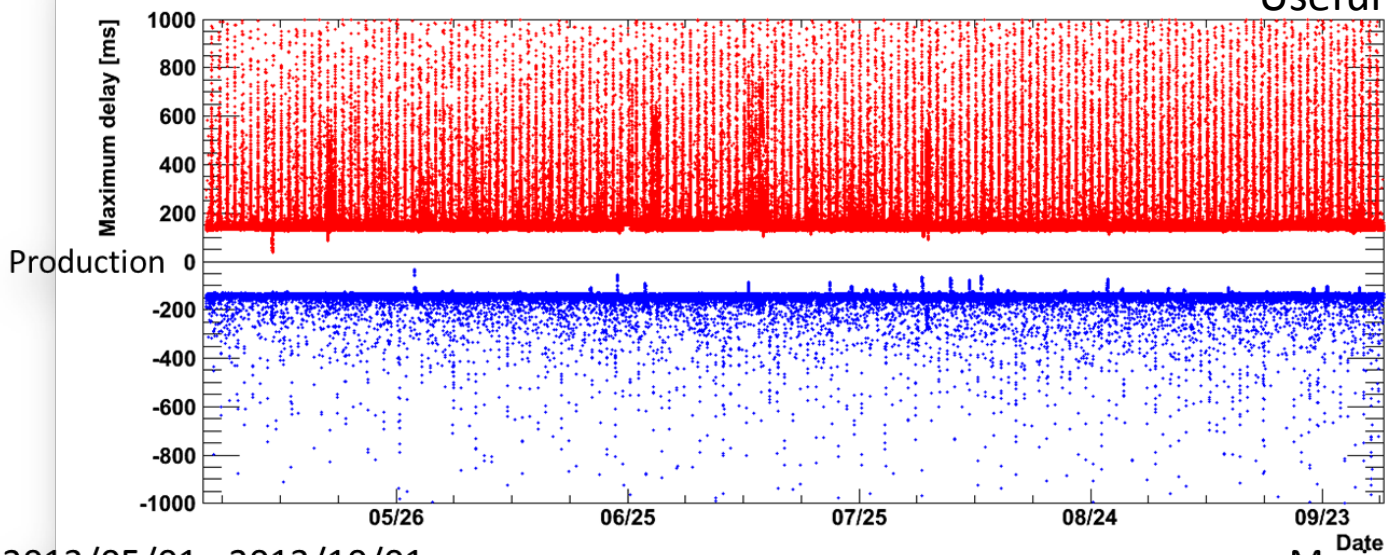

incoming

outgoing

Production

2013/05/01 - 2013/10/01

Maximum latency

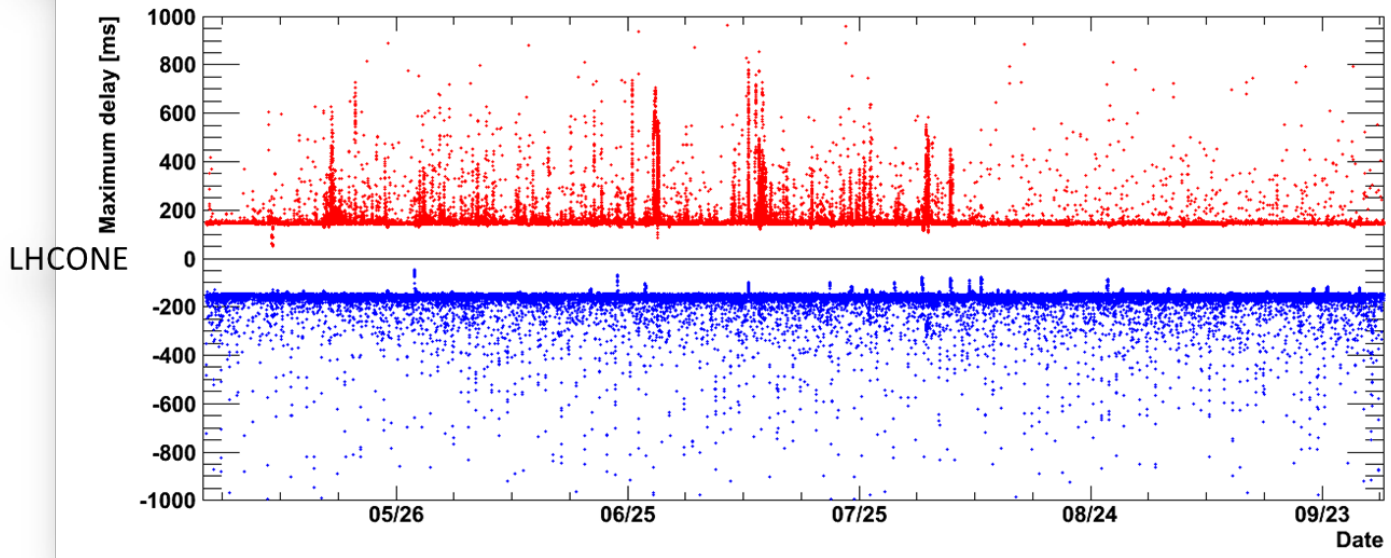

incoming

outgoing

LHCONE

Maximum latency (CC-IN2P3)

Also have spikes.
Additional periodic noise.

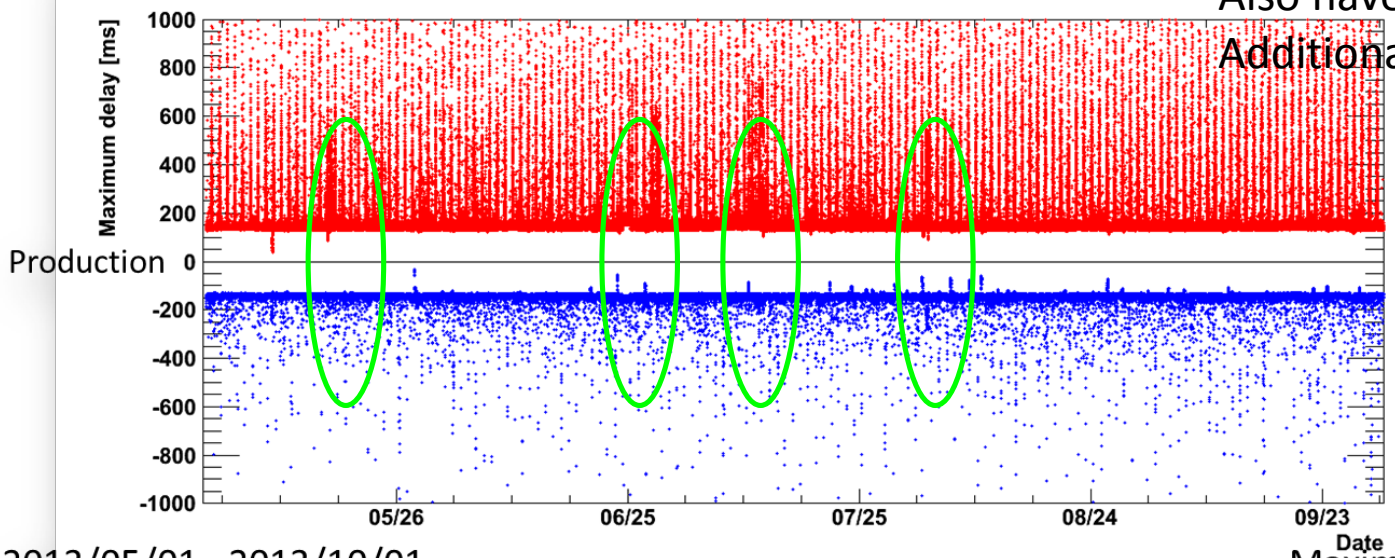

Production

2013/05/01 - 2013/10/01

Maximum latency

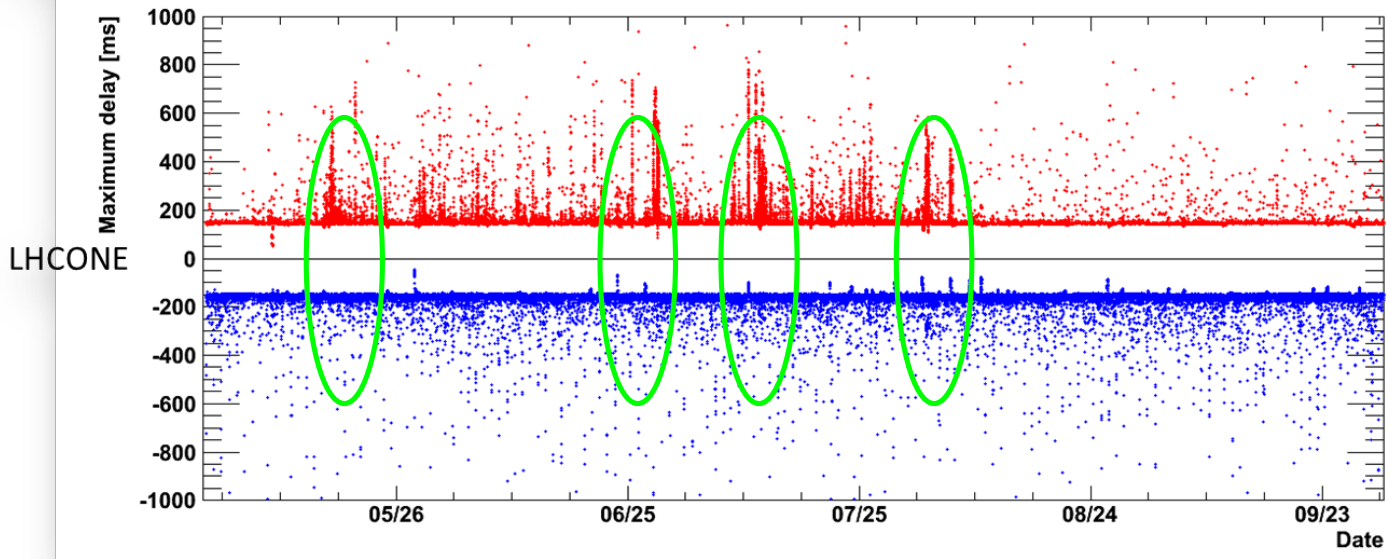

LHCONE

Distribution of Maximum latency (CC-IN2P3)

Production

2013/05/01 - 2013/10/01

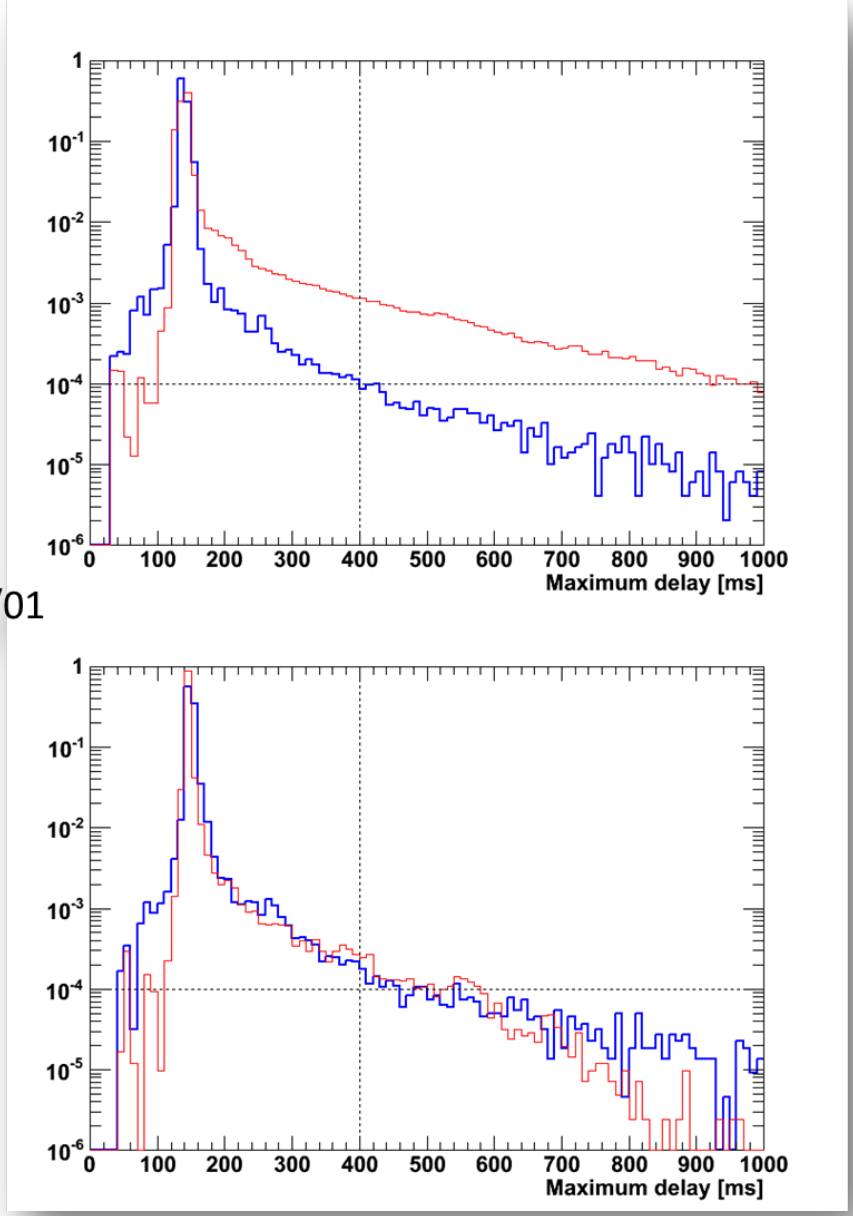

incoming

outgoing

Maximum latency

LHCONE

incoming

outgoing

Distribution of Maximum latency (CC-IN2P3)

Production

2013/05/01 - 2013/10/01

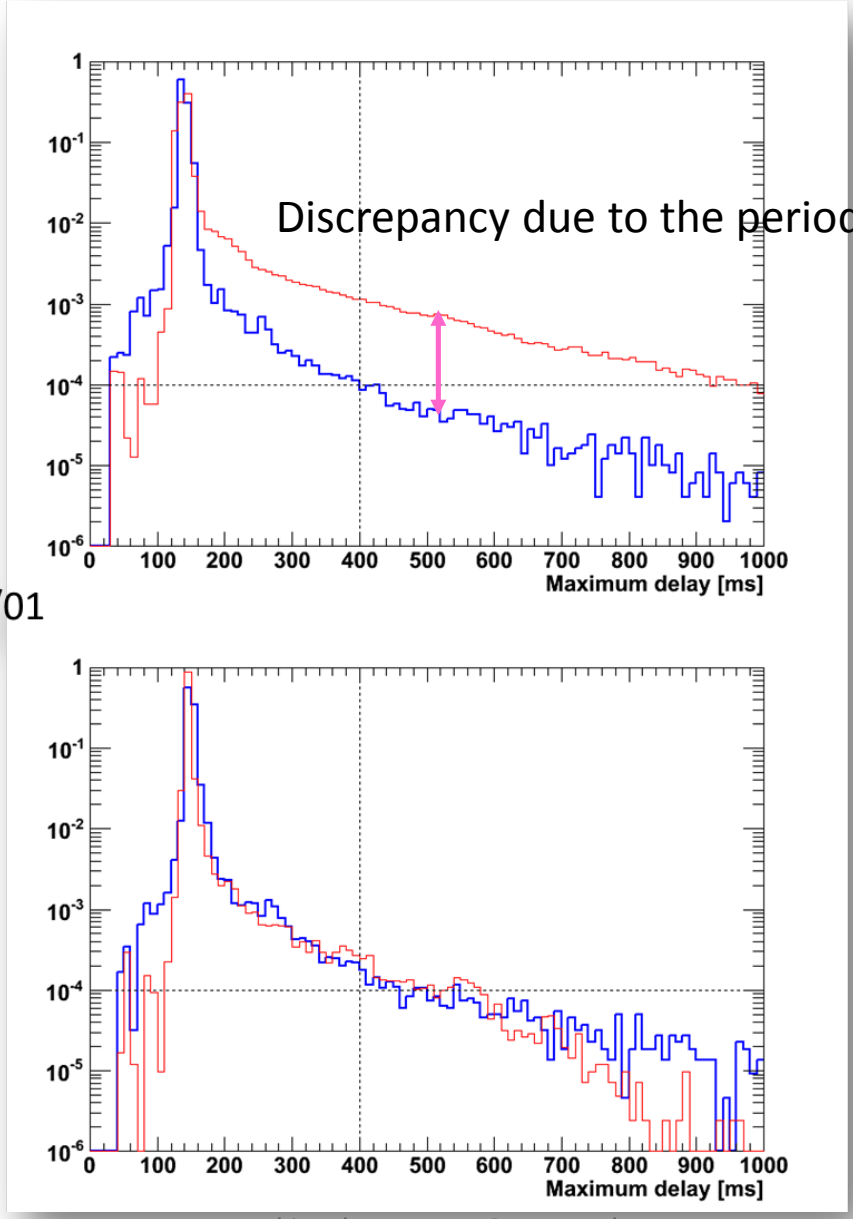

Maximum latency

incoming

outgoing

LHCONE

incoming

outgoing

Also for the other sites

LA (NII)

(US)

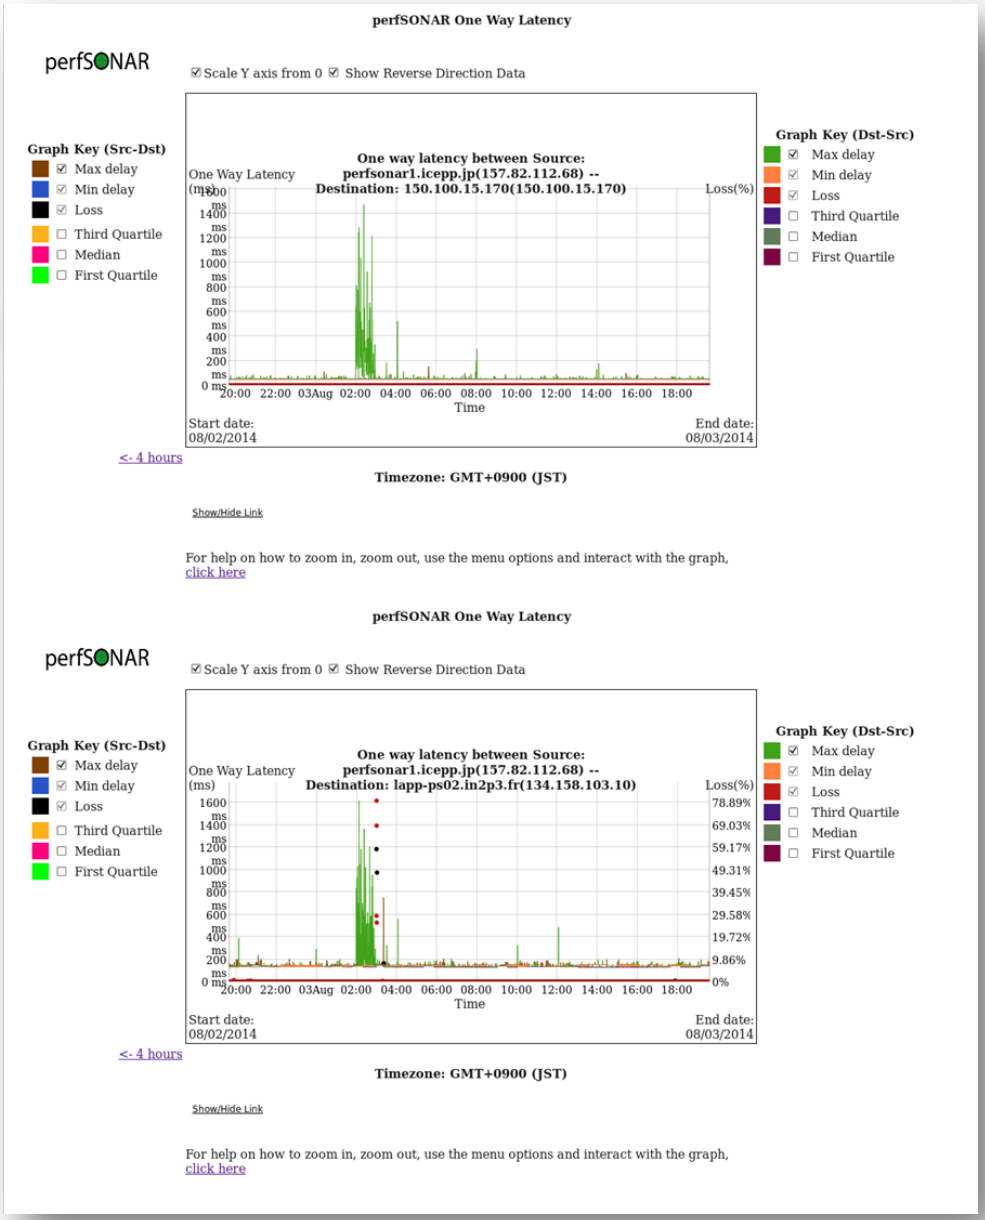

LAPP

(FR)

- One of the perfsonar instance in Tokyo seems to fall into the busy state once in a day.
- It is independent of source sites.
- But, no significant errors in system and service logs.

Maximum latency (masked by time)

Periodic nose can be cleaned up.

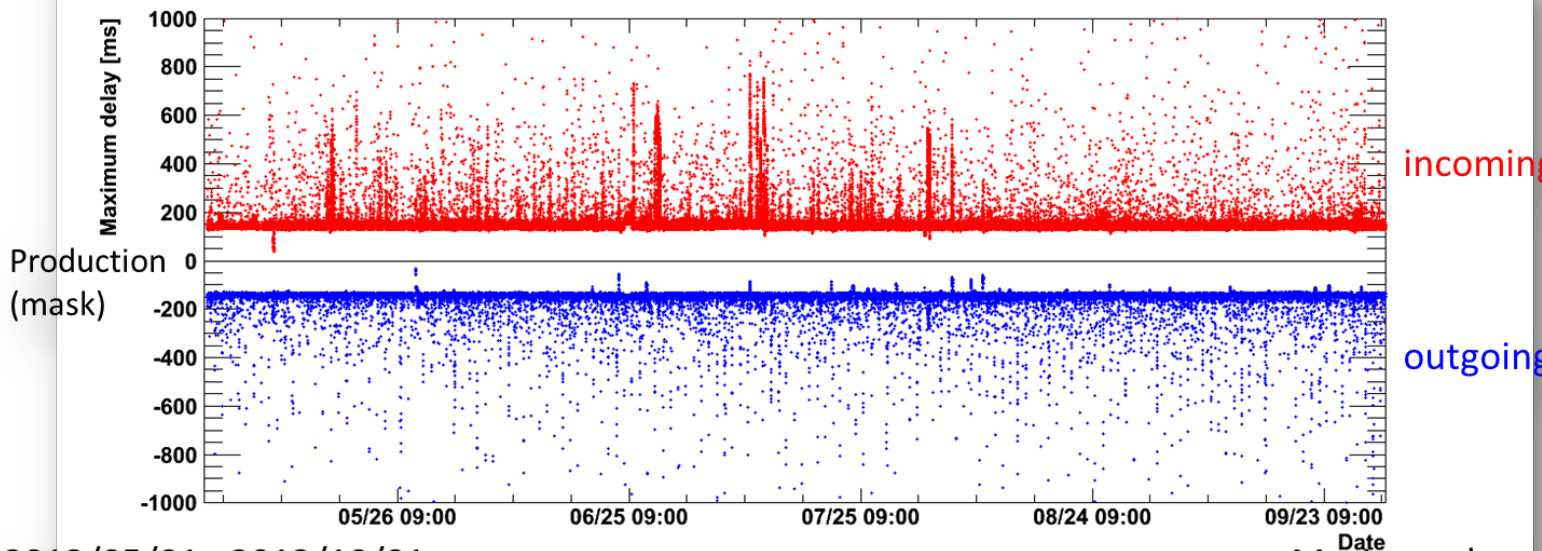

Production (mask)

2013/05/01 - 2013/10/01

Maximum latency

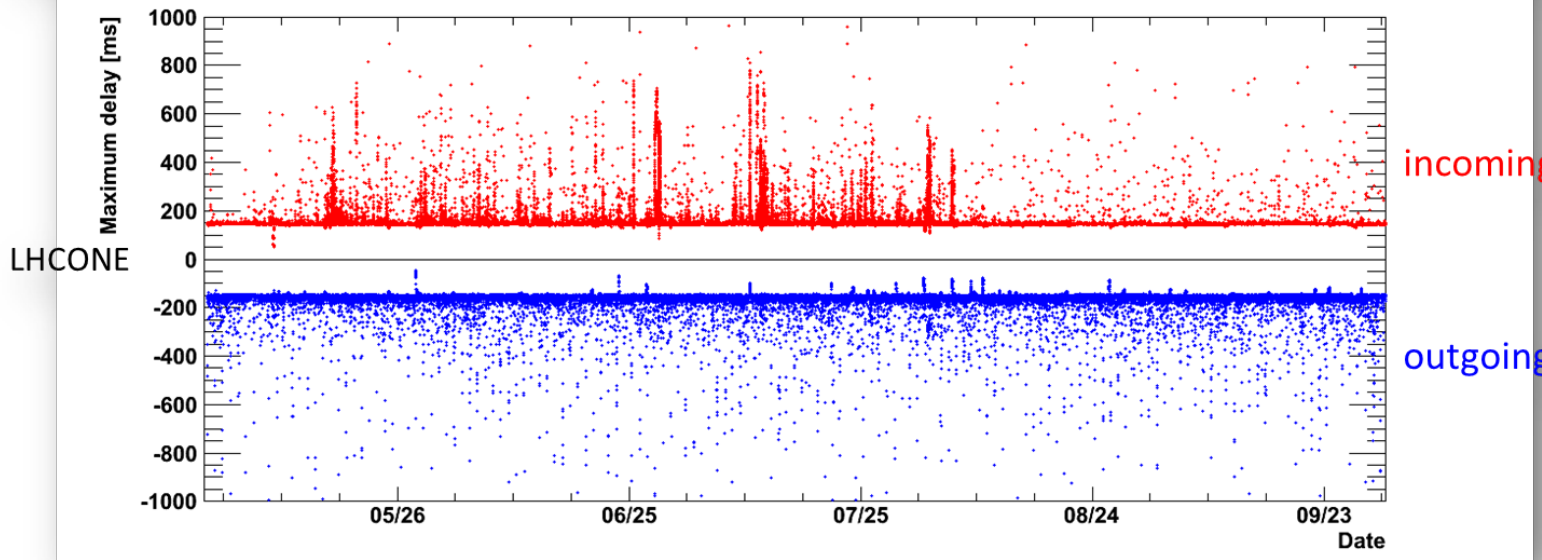

LHCONE

Maximum latency by mask (CC-IN2P3)

Production
(mask)

2013/05/01 - 2013/10/01

LHCONE

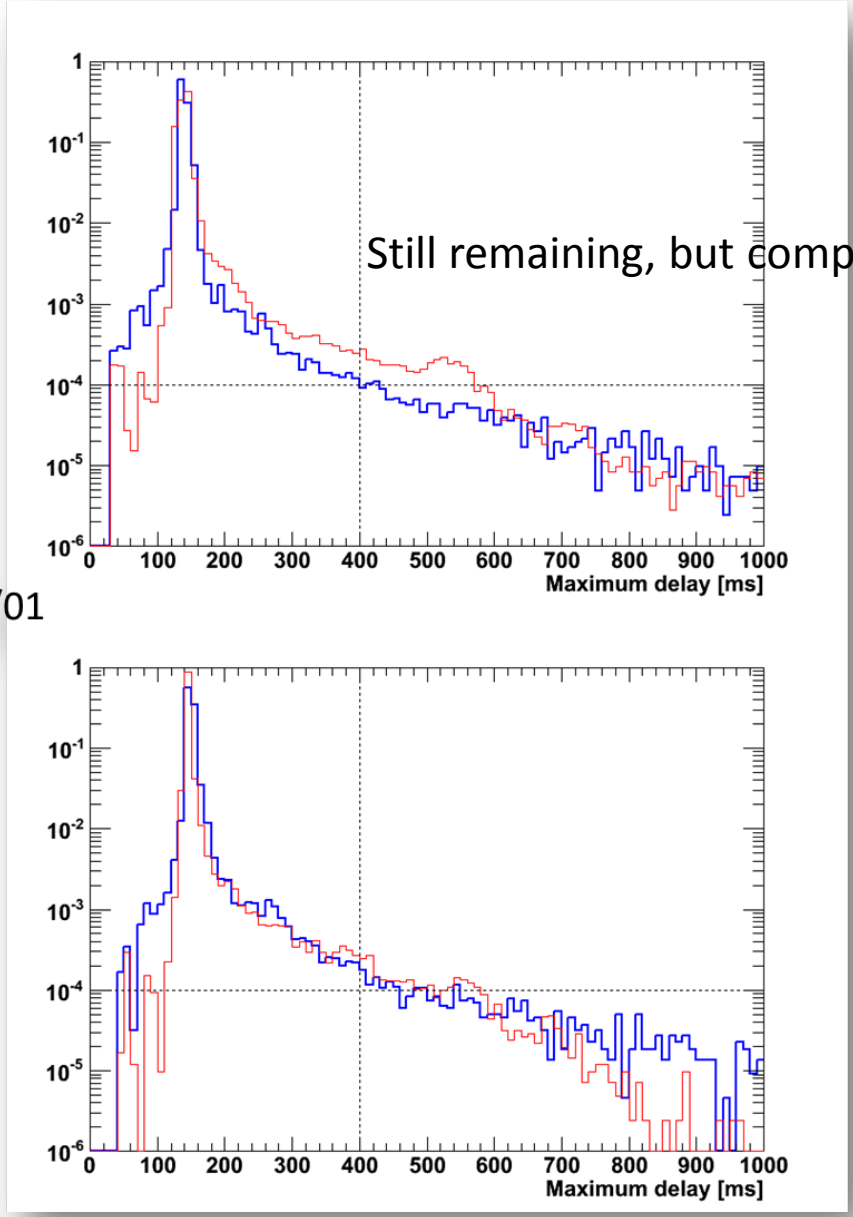

incoming

outgoing

Maximum latency

incoming

outgoing

Bandwidth measurement (CC-IN2P3 and CNAF)

Asymmetric
~38 MB/s (incoming)
~28 MB/s (outgoing)

Production

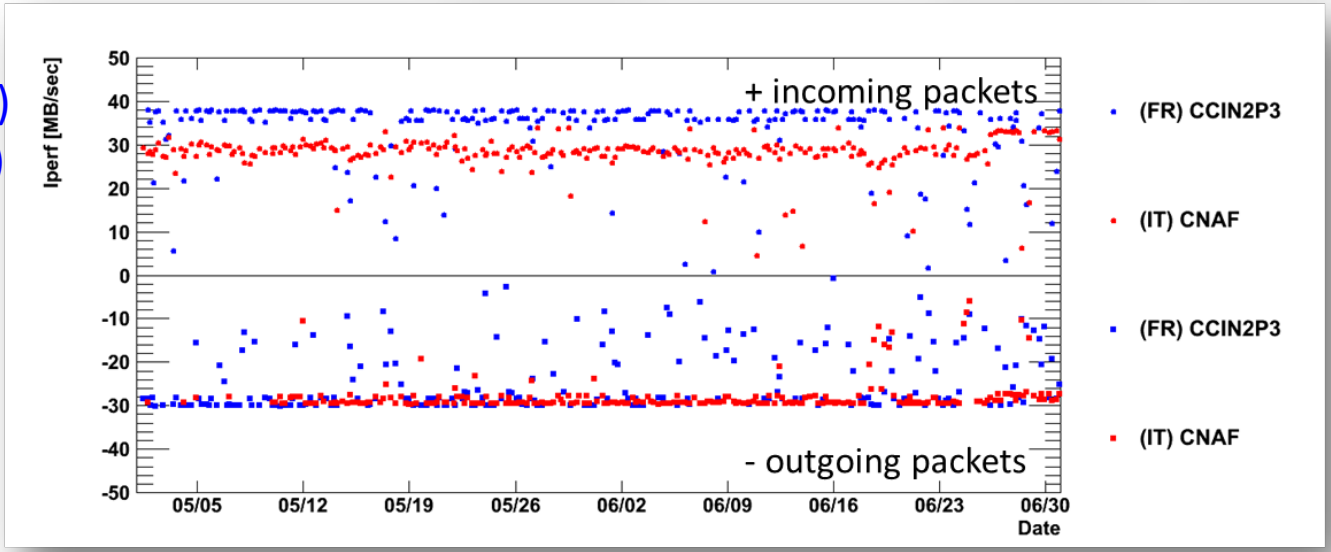

Symmetric, but unstable
~34 MB/s (incoming)
~35 MB/s (outgoing)

LHCONE

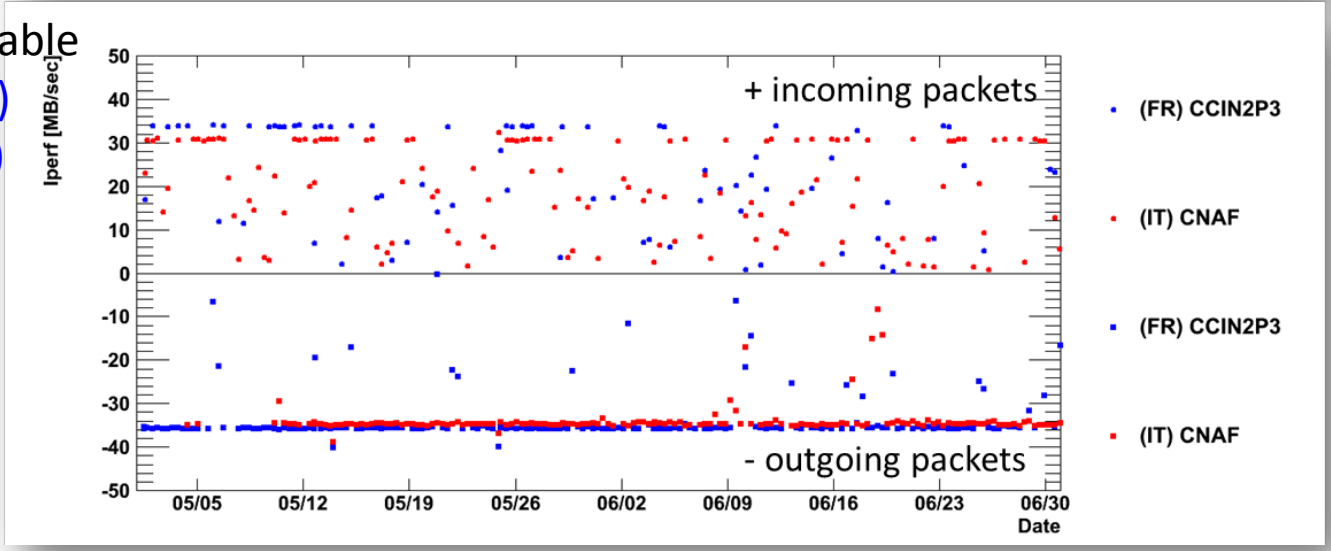

Minimum latency (CC-IN2P3 in 2014)

Production (mask)

RMS: 7.1 (incoming)

RMS: 7.1 (outgoing)

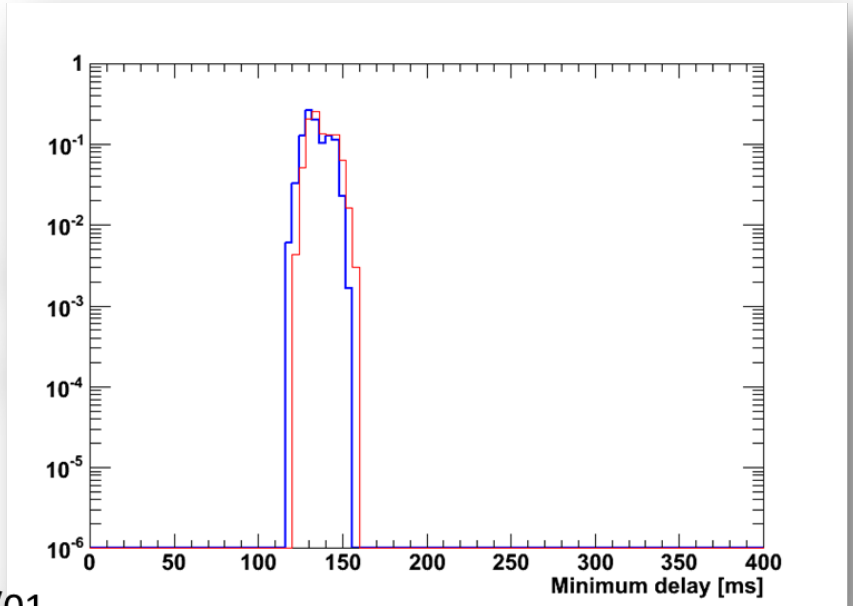

2014/01/01 - 2014/03/01

Minimum latency

LHCONE

RMS: 4.2 (incoming)

RMS: 3.8 (outgoing)

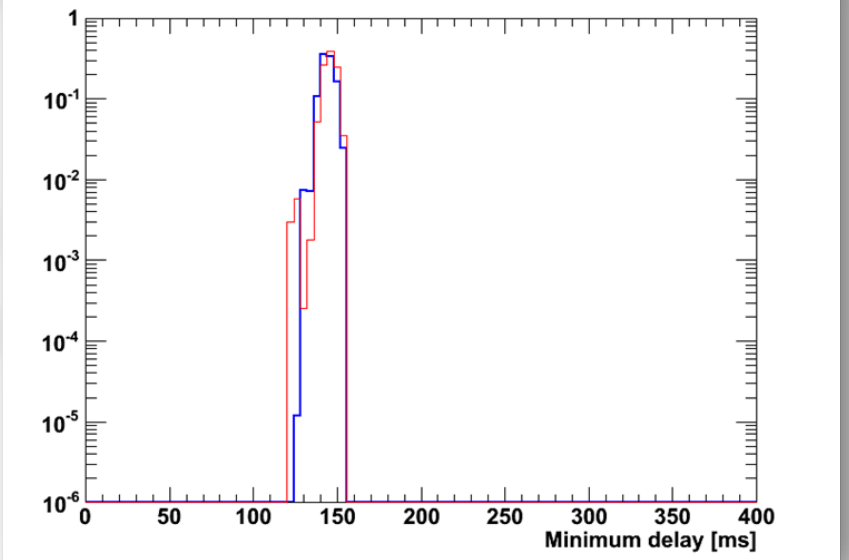

Minimum latency (CC-IN2P3 in 2014)

Production (mask)

RMS: 7.1 (incoming)

RMS: 7.1 (outgoing)

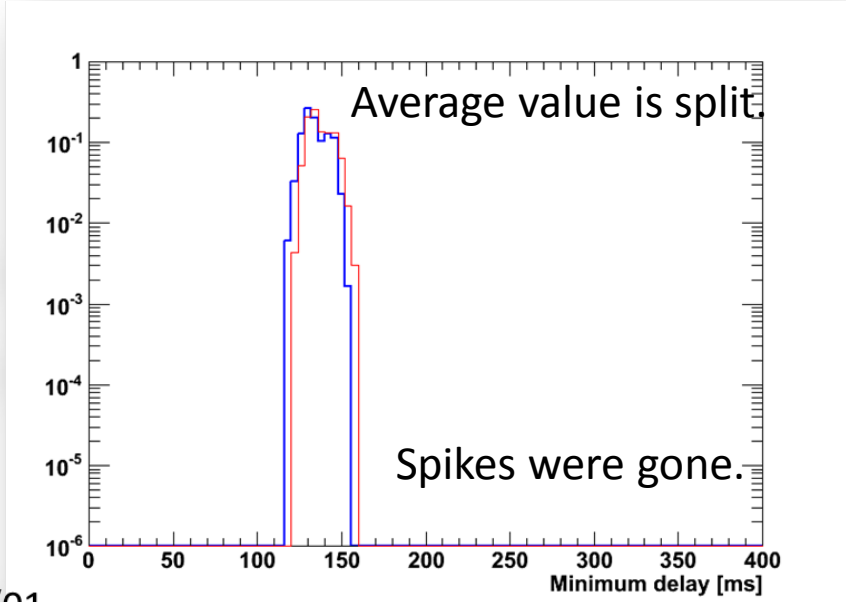

2014/01/01 - 2014/03/01

Minimum latency

LHCONE

RMS: 4.2 (incoming)

RMS: 3.8 (outgoing)

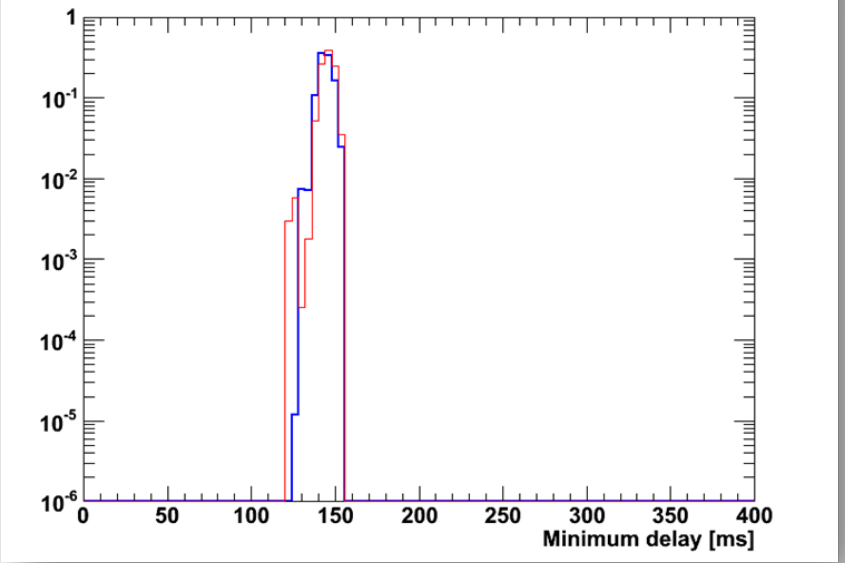

Latency in one day (CC-IN2P3)

Load balancing somewhere in NY or GEANT?

Incoming

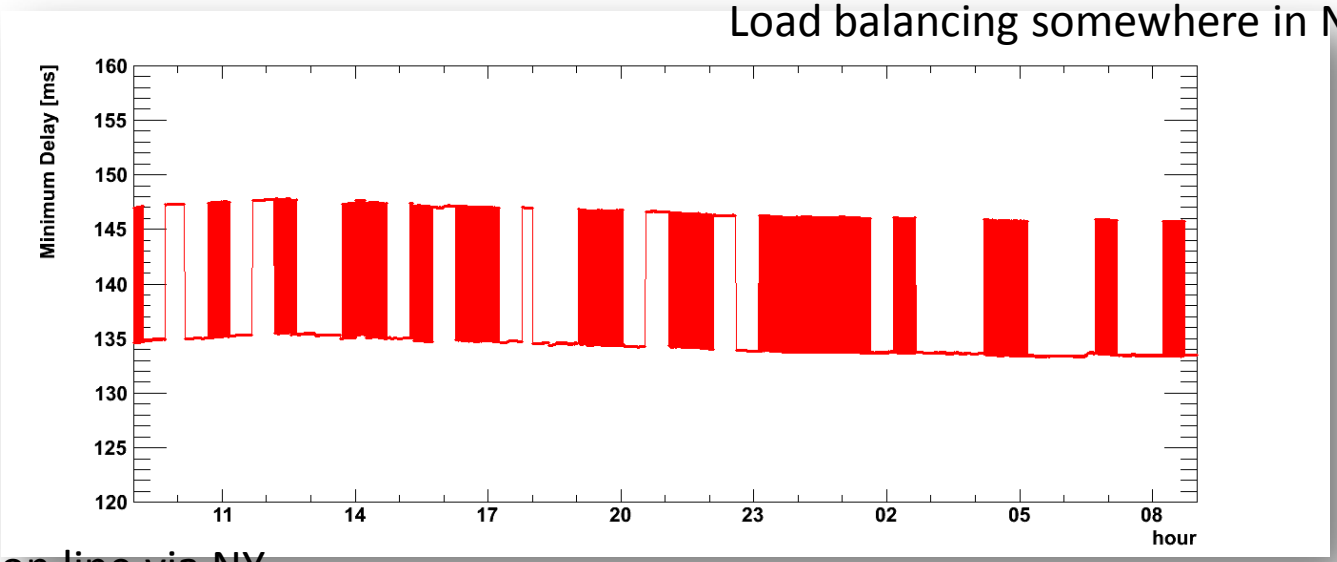

Both production line via NY

Outgoing

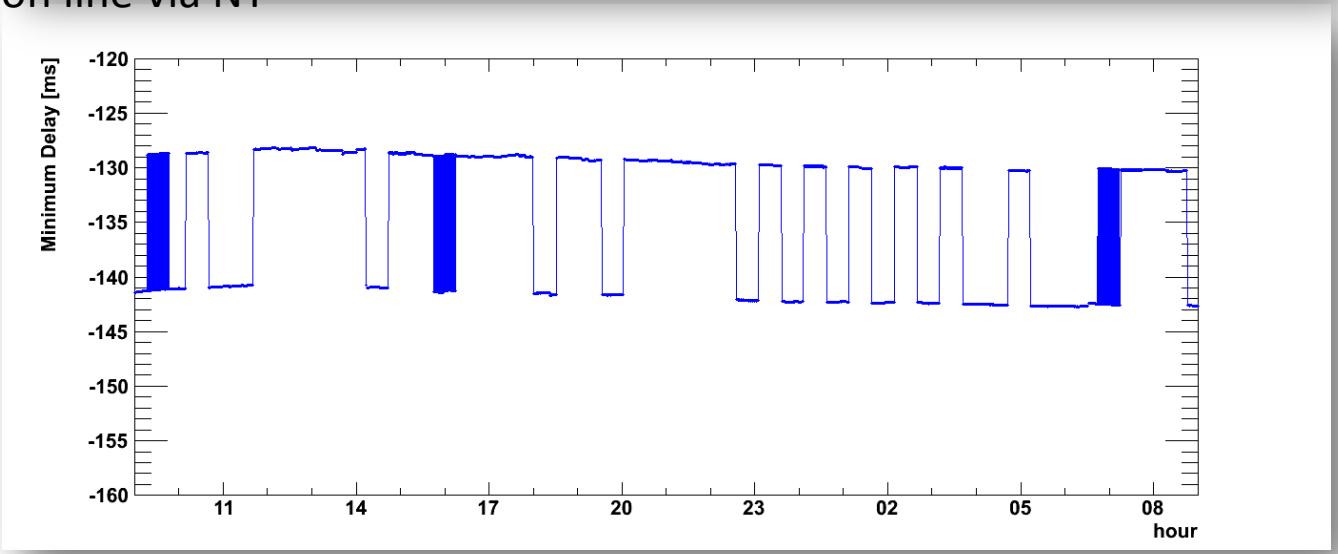

Maximum latency (CC-IN2P3, 2014)

Production
(mask)

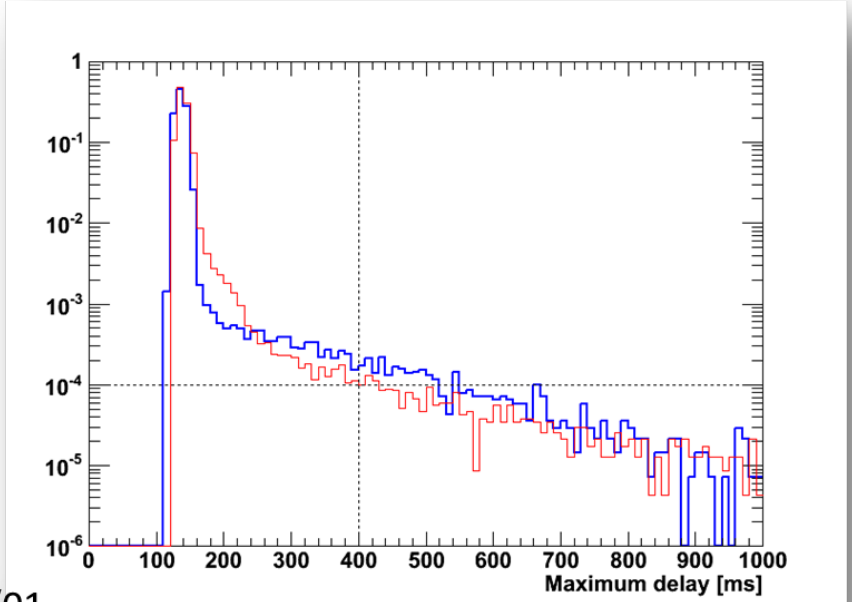

incoming

outgoing

2014/01/01 - 2014/03/01

Maximum latency

LHCONE

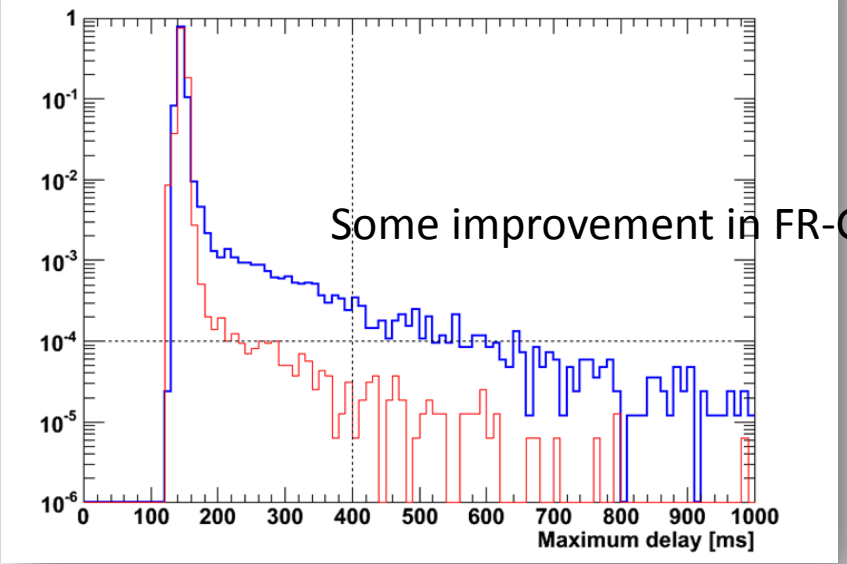

incoming

outgoing

Some improvement in FR-Geneva?

Bandwidth measurement (latest data)

Still asymmetric
~35 MB/s (incoming)
~24 MB/s (outgoing)

Production

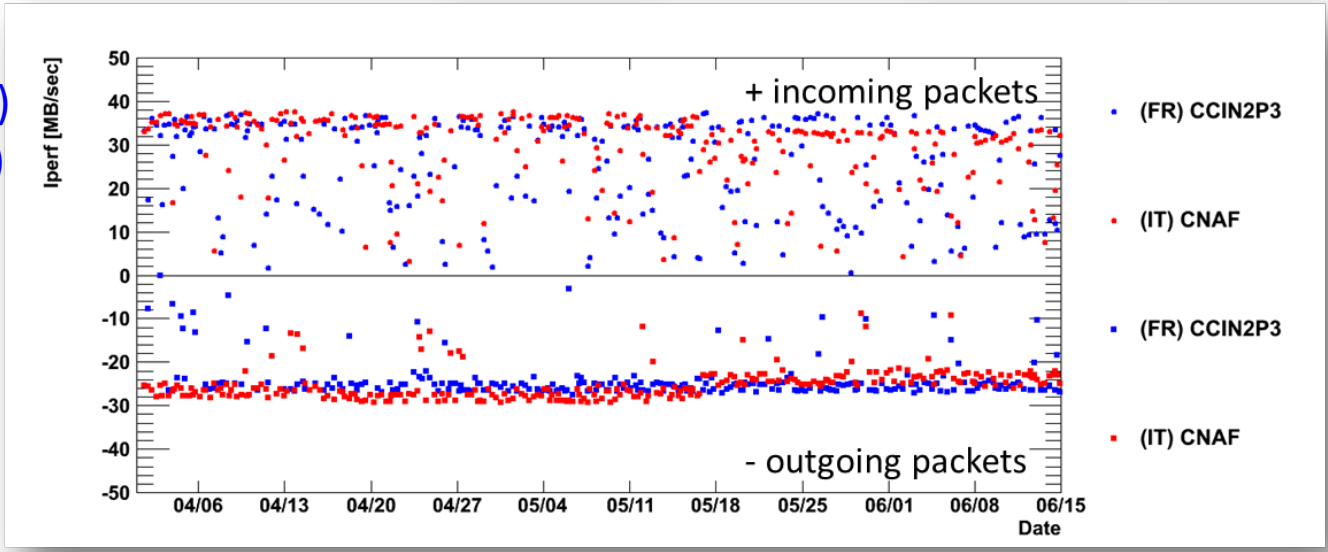

Symmetric, and very stable
~32 MB/s (incoming)
~30 MB/s (outgoing)

LHCONE

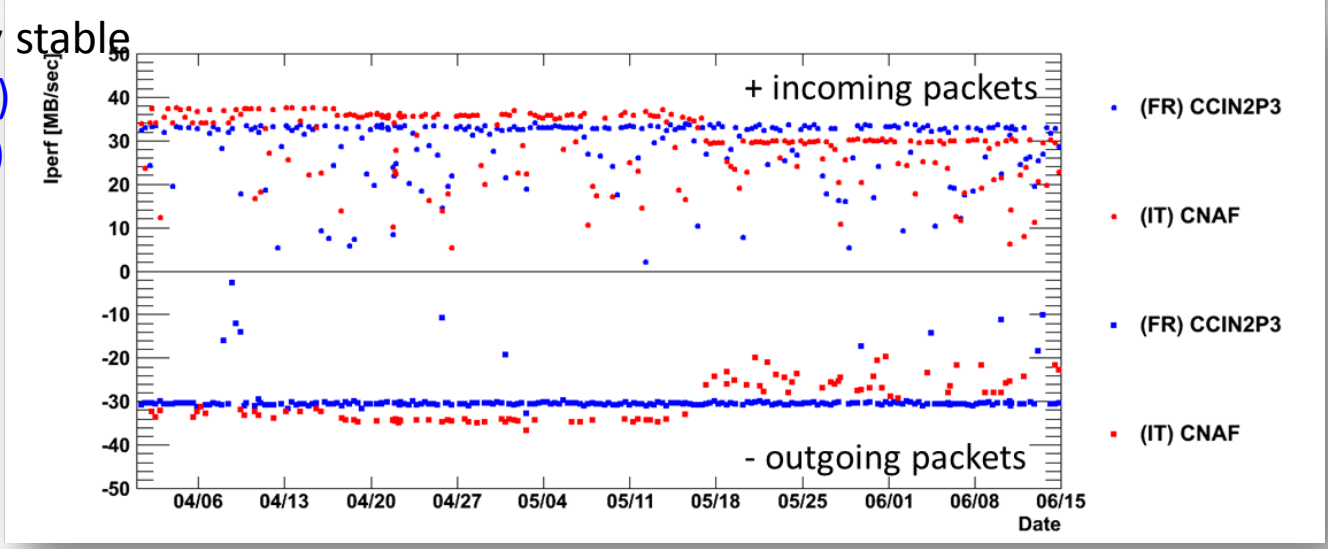

Summary

- Finally, we could have a consensus to ride on LHCONE routing for EU sites,
  (Tokyo-Osaka-Washington-Geneva)
  and it will be applied for all of our production instance.

- We have requested to switch our routing to IT division in University of Tokyo. They have already started some work on detailed configuration and arrangement of appropriate date to switch.

- We can show the performance of data transfer throughput with new line at the next time.

- We will start work to ride on LHCONE for US sites via NY and LA line (internet2).

- In more future, we expect the next generation of SINET5 (from 2016) provide more bandwidth for us (100Gbps to US ?).