

WLCG and support for IPv6-only CPU

David Kelsey (STFC-RAL)
WLCG GDB, CERN
8 Jun 2016



Outline


- Some background
- Yesterday's IPv6 pre-GDB workshop

Google IPv6 statistics

- <https://www.google.com/intl/en/ipv6/statistics.html>

IPv6 Adoption

We are continuously measuring the availability of IPv6 connectivity among Google users. The graph shows the percentage of users that access Google over IPv6.



Global ~ 12%
Belgium ~ 43%
USA ~ 27%
Germany ~ 24%

Background

- Oct 2015
 - A request from ATLAS in Canada
 - planning deployments of new computing systems in 2016-2017 which will support our Tier-1 and Tier-2
 - can we go with a pure IPv6 scheme for the compute nodes?
- IPv6-only opportunistic resources may also become available
 - At least one Cloud provider offers discount for use of IPv6

HEPiX IPv6 WG - history

- Many years ago
 - Started to investigate compliance of our software
- Lots of testbed activity and work with storage developers
 - Lots of good progress here
 - Most storage now works in dual-stack mode
- In 2015
 - Good engagement of LHC experiments
 - Pushing for sites to implement IPv6
 - And dual-stack perfSONAR

IPv6 WG in 2015-16


- Moved away from use of the IPv6 testbed
- Deploy dual-stack storage wherever possible
 - In production
 - But members of the working group
 - Close monitoring of behaviour
- *From testbed to production deployment*
- Also work has been done on Monitoring

CERN IPv6 traffic stats

IPv6: Internet

ALL Daily Weekly Monthly Yearly

Yearly



Aims of pre-GDB IPv6 workshop

- Present experiences of sites running dual-stack services
- Look at other developments and best practice
 - Monitoring, IPv6 security, ...
- What is needed to support IPv6-only CPU?
 - Experiment status and requests at the GDB
- Some fun along the way with quiz, demos and hands-on

IPv6 workshop – 7 June 2016

- <https://indico.cern.ch/event/394830/>
- ~ 20 people in the room plus ~ 15-20 over Vidyo
- Agenda (see slides for all the details!)
 - IPv6 intro and get your laptop working over IPv6
 - Intro and aims of the day
 - Experiences from sites particularly with storage
 - dCache, XRootD, DPM, EOS, FTS, ...
 - IPv6 online quiz
 - Monitoring
 - Dashboards, ETF, perfSONAR
 - LHCOPN/LHCONE and IPv6
 - IPv6 security
 - Update from OSG
 - An online game over IPv6

IPv6 at PIC Tier-1

*Pre-GDB, IPv6 Workshop at CERN
7th June 2016*

F. López *

On behalf of the PIC Tier1 Network, Storage, Batch System teams and LHC Office

V. Acín, E. Acción, C. Acosta, J. Casals, J. Flix, A. Pérez-Calero and A. Vedae

** PIC Tier-1 network manager*



EXCELENCIA
SEVERO
OCHOA

Institut de Física
d'Altes Energies



Ciemat

Centro de Investigaciones
Energéticas, Medioambientales
y Tecnológicas

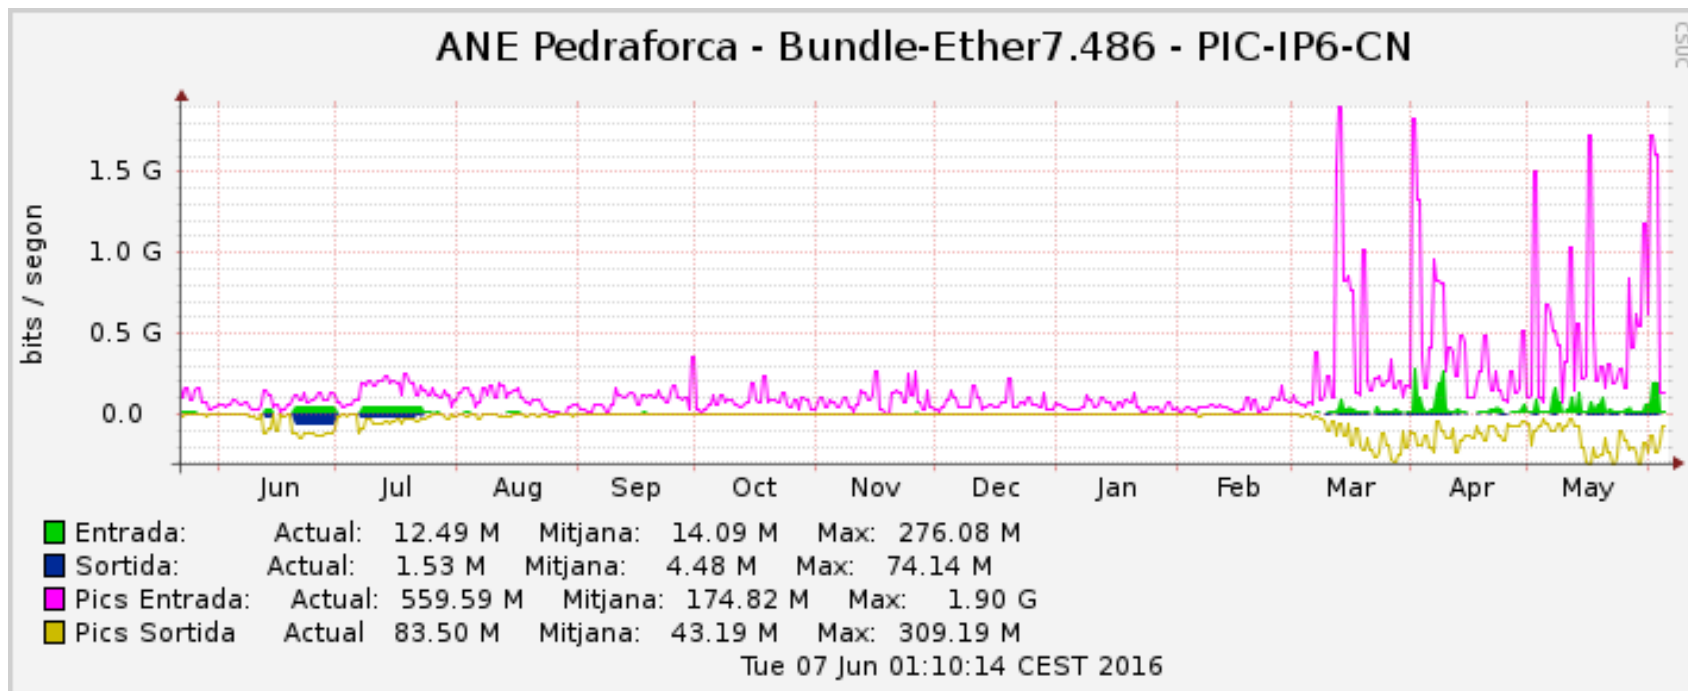


EXCELENCIA
MARÍA
DE MAEZTU

Production dCache

- Current production dCache version 2.13.28+
 - Mar/16: SRM + transfers protocols (gridFTP, dcap|gsidcap and xrootd) + CMS pools
 - Apr/16: ATLAS (Tier-1 and Tier-2) + LHCb pools
- dCache MWR instance for CMS is running 2.15.5

Production dCache



Basic functionality tests

- IPv6-only client → xrootd diskserver (EOS fst) dual-stack

```
[root@p05151113432846 /]# netstat -natup | grep xrootd
tcp      0      0 :::1094          :::*              LISTEN     53524/xrootd
tcp      0      0 :::1095          :::*              LISTEN     54118/xrootd
```

```
[root@p05151113432846 /]# ps -ef | grep xrootd
xrootd  53524  1 1 15:08 7 00:00:08 /usr/bin/xrootd -f /var/log/xrootd/xrootd.log -c /etc/xrootd/xrootd-clustered.cfg -k fto -f v6 -b -s
```

```
xavierc3@xplus0046:~$ env XRD_NETWORKSTACK=IPv6 xrdcopy -f -d1 /tmp/2GB root://p05151113432846.ipv6.cern.ch:/tmp/2GB
```

```
[dxc1f00@root://p05151113432846.ipv6.cern.ch:1094/tmp/2GB?oss.asize=2097152000] Sending a close command for handle 0x0 to p05151113432846.ipv6.cern.ch:1094
[2016-05-25 15:19:54.043396 +0200][Debug ][File ][dxc1f00@root://p05151113432846.ipv6.cern.ch:1094/tmp/2GB?oss.asize=2097152000] Close returned from
p05151113432846.ipv6.cern.ch:1094 with: [SUCCESS]
[1.953GB/1.953GB][100%]-----[500MB/s]
[2016-05-25 15:19:54.058971 +0200][Debug ][TaskMgr ][Requesting unregistration of: "TickGeneratorTask for: p05151113432846.ipv6.cern.ch:1094"
[2016-05-25 15:19:54.058999 +0200][Debug ][AsyncSock ][p05151113432846.ipv6.cern.ch:1094 #0.0] Closing the socket
[2016-05-25 15:19:54.058922 +0200][Debug ][Poller ][<[2001:1458:201:a8:100:333]:58006><-><[2001:1458:301:34:100:13]:1094> Removing socket from the poller
[2016-05-25 15:19:54.059037 +0200][Debug ][PostMaster ][p05151113432846.ipv6.cern.ch:1094 #0] Destroying stream
[2016-05-25 15:19:54.059064 +0200][Debug ][AsyncSock ][p05151113432846.ipv6.cern.ch:1094 #0.0] Closing the socket
```

- IPv6-only client → EOS dual-stack:

```
[root@p05088916c58033 ~]# eos space ls
# type # name #w(fs) #r(fs-rw) #sum(usedbytes) #sum(capacity)
spaceview ipv6 11 11 452.77 M 38.99 T
```

Standard EOS put from IPv6-only client:

```
08:53:53:xavierc3@xplus0046:~$ eos cp /etc/group /eos/ipv6/test.xavi$RANDOM
[eos-cp] going to copy 1 files and 862 B
```

```
[eos-cp] group Total 0.00 MB [==] 100.00 % [0.0 MB/s]
[eos-cp] copied 1/1 files and 862 B in 0.43 seconds with 1985 B/s
```

- Next: PoC with an IPv6-only diskserver pool on EOS + dual-stack setup for permanent IPv6-only clients testing (via an specific namespace 'endpoint')

Xavier Espinal (IT/ST) pre-GDB IPv6 Workshop CERN 7/06/2016




GridPP

UK Computing for Particle Physics


Recent Imperial experiences with FTS3

Duncan Rand

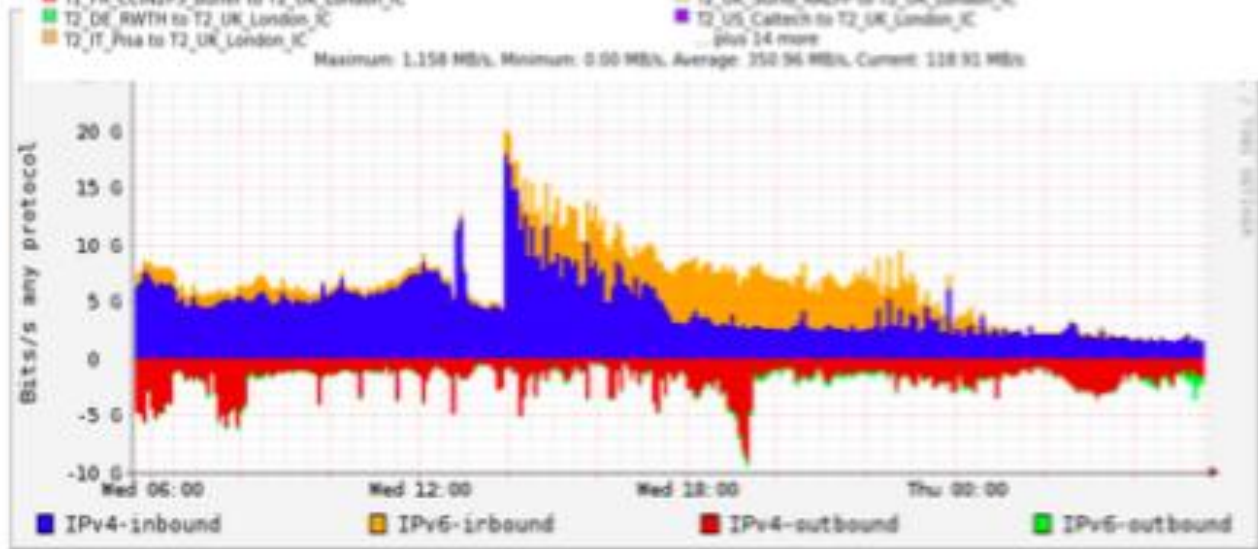
Imperial College London



Large transfer to Imperial
 Much of it from PIC
 PIC and Imperial both dual-stack storage



Orange shows IPv6 inbound transfers



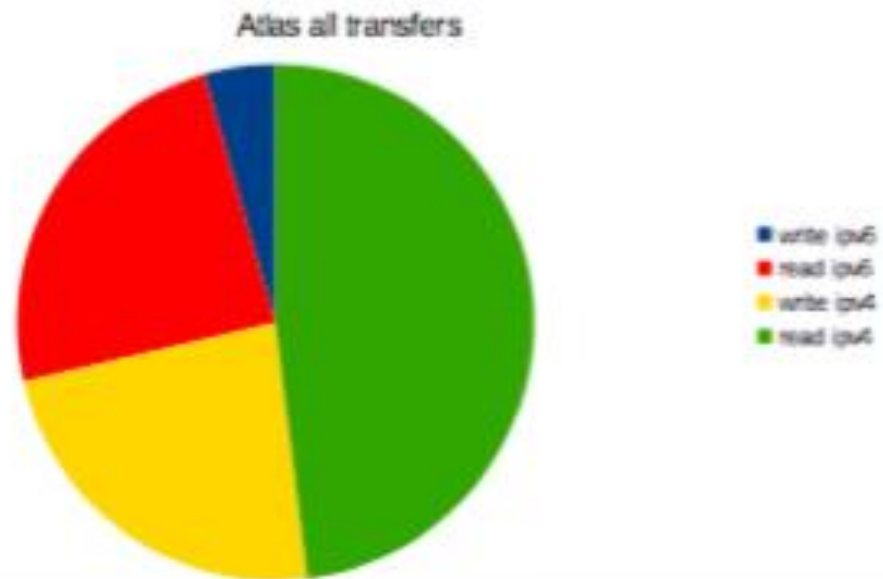


NDGF-T1 transfer statistics

Pre GDB
7.6.2016

Ulf Tigerstedt
<ulf.tigerstedt@csc.fi>

Atlas all transfers



406

2206

2129

4374

3

WLCG Experiments Test Framework (ETF)

Marian Babik, Andrea Sciabà

IPv6 pre-GDB
7th June 2016



IPv6

- Production ETF should eventually be dual-stack
 - No ill effects on any site, transparent to both sites and experiments
 - Results used for the official WLCG reports
 - Timescale? April 2017? Before?
- A dedicated ETF service for IPv6 testing
 - **Prototype running since December** but needs to be integrated with Puppet management
 - Only CMS tests running on the IPv6 WG testbed nodes
 - **Other experiments interested in running their probes?**
 - **Do we want to have an IPv6-only ETF host**, given that there are ways (e.g. via env variables) to force IPv6?
 - IPv6-only not officially supported at CERN

- etf-ipv6-dev.cern.ch demo


perfSONAR

Marian Babik (IT/CM)

IPv6 pre-GDB
7th June 2016



Parts of this presentation are a result of work by the perfSONAR Project and are licensed under CC BY-SA 4.0 (<https://creativecommons.org/licenses/by-sa/4.0/>).



Summary

- perfSONAR is the best tool to assess network performance and debug issues
 - For both IPv4 and IPv6
 - All sites can easily enable dual-stack and use command line tools to test their performance
 - Growing network of dual-stack perfSONARs
 - Also outside of WLCG
- WLCG perfSONAR network
 - Runs regular testing activities to baseline the performance
 - Currently focused on IPv4, we need to plan how we gradually introduce IPv6
 - Dual-stack mesh already available
- WLCG perfSONAR SU is here to help with deployment

LHC[OPN/ONE] → IPv6 → status +

Bruno Hoeft / DE-KIT

STEINBUCH CENTRE FOR COMPUTING - SCC

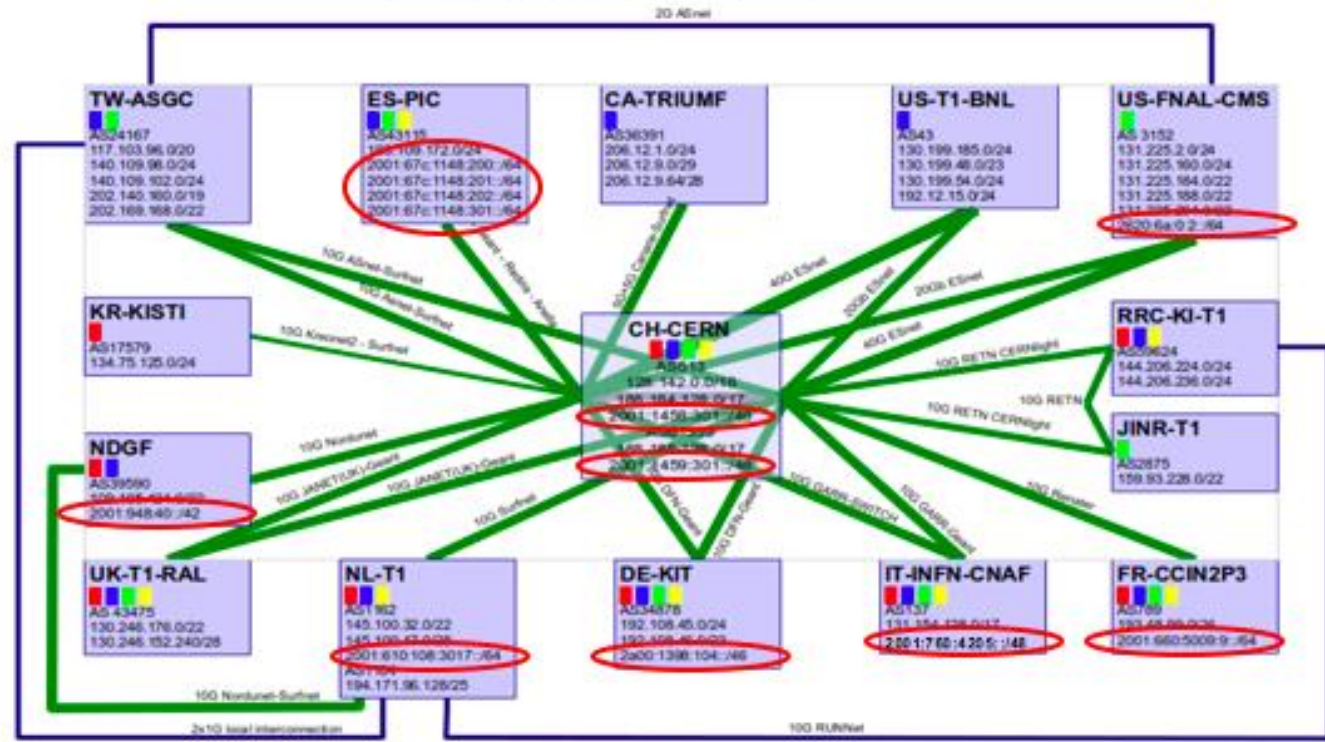


KIT – Universität des Landes Baden-Württemberg und
nationales Forschungszentrum in der Helmholtz-Gemeinschaft

www.kit.edu

Click to add title

LHCOPN



— T0-T1 and T1-T1 traffic
— T1-T1 traffic only
--- Not deployed yet
 (thick) == 100Gbps
 (thin) == 10Gbps


■ = Alice ■ = Atlas
■ = CMS ■ = LHCb

p2p prefix: 192.16.186.0/24 - 2001.1458:302::/48
 eduardo.martelli@cern.ch-20180322

IPv6 tier-1 site readiness ticket

status at Tuesday 07. June

Ticket-ID	Type	VO	Site	Priority	Resp. Unit	Status	Last Update	Subject	
121896		none	BNL-ATLAS	very urgent	OSG(Prod)	assigned	2016-06-06	Tier-1 LHCOPN IPv6 Peering, incl. dualst...	KO
121895		none	USCMS-FNAL-WC1	very urgent	OSG(Prod)	solved	2016-06-03	Tier-1 LHCOPN IPv6 Peering, incl. dualst...	✓
121894		none	RAL-LCG2	top priority	NGI_UK ▶ assigned	in progress	2016-06-01	Tier-1 LHCOPN IPv6 Peering, incl. dualst...	ok
121893		none	Taiwan-LCG2	top priority	ROC_Asia/Pacific	in progress	2016-06-02	Tier-1 LHCOPN IPv6 Peering, incl. dualst...	?
121892		none	SARA-MATRIX	top priority	NGI_NL	in progress	2016-06-01	Tier-1 LHCOPN IPv6 Peering, incl. dualst...	✓
121891		none	RRC-KI-T1	top priority	ROC_Russia ▶ assigned	in progress	2016-06-01	Tier-1 LHCOPN IPv6 Peering, incl. dualst...	?
121890		none	NIKHEF-ELPROD	top priority	NGI_NL	solved	2016-06-01	Tier-1 LHCOPN IPv6 Peering, incl. dualst...	✓
121889		none	NDGF-T1	top priority	NGI_NDGF	solved	2016-06-02	Tier-1 LHCOPN IPv6 Peering, incl. dualst...	✓
121888		none	KR-KISTI-GSDC-01	top priority	ROC_Asia/Pacific	waiting for reply	2016-06-02	Tier-1 LHCOPN IPv6 Peering, incl. dualst...	?
121887		none	JINR-T1	top priority	ROC_Russia	in progress	2016-06-01	Tier-1 LHCOPN IPv6 Peering, incl. dualst...	?
121886		none	INFN-T1	top priority	NGI_IT	solved	2016-06-01	Tier-1 LHCOPN IPv6 Peering, incl. dualst...	✓
121885		none	IN2P3-CC	top priority	NGI_FRANCE	solved	2016-06-01	Tier-1 LHCOPN IPv6 Peering, incl. dualst...	✓
121884		none	pic	top priority	NGI_IBERGRID	solved	2016-06-01	Tier-1 LHCOPN IPv6 Peering, incl. dualst...	✓
121883		none	FZK-LCG2	top priority	NGI_DE	solved	2016-06-01	Tier-1 LHCOPN IPv6 Peering, incl. dualst...	✓
121882		none	CERN-PROD	top priority	ROC_CERN	in progress	2016-06-05	Tier-1 LHCOPN IPv6 Peering, incl. dualst...	✓
121881		none	TRIUMF-LCG2	top priority	ROC_Canada	solved	2016-06-07	Tier-1 LHCOPN IPv6 Peering, incl. dualst...	ok



IPv6 peering + dual stack per sonar

work in progress

no response

Links

- HEPiX IPv6 web

<http://hepixonweb.cern.ch>

- Working group meetings

<http://indico.cern.ch/categoryDisplay.py?categId=3538>

- WLCG Operations IPv6 Task Force

<http://hepixonweb.cern.ch/content/wlcg-ipv6-task-force-0>

- IPv6 working group CHEP papers

2013 - <http://iopscience.iop.org/article/10.1088/1742-6596/513/6/062026>

2015 - <http://iopscience.iop.org/article/10.1088/1742-6596/664/5/052018>

Questions?