Support for IPv6-only CPU – an update from the HEPiX IPv6 WG

David Kelsey (STFC-RAL)

WLCG GDB, CERN

11 Jan 2017

Outline

- An update from the HEPiX IPv6 WG
  - Since Sep 2016 GDB
- Reminder: WLCG MB approval (20 Sep 2016) of the plan to support IPv6-only CPU from April 2017
- Brief review of current status
  - Experiments, Tier 0/1/2, central services?
  - Are there any known show-stoppers?

HEPiX IPv6 WG meetings

- Meetings held monthly
  - All one hour Vidyo meetings
  - 6 Oct, 3 Nov, 15 Dec 2016
  - Next F2F at CERN 2/3 Feb 2017

We need to track …

- Experiment services
- Tier 1 status
  - LHCOPN peering over IPv6
  - Dual-stack perfSONAR
  - Dual-stack storage
  - Other (central) services – dual stack
- Tier 2 status
  - LHCONE peering over IPv6
  - Dual-stack perfSONAR
  - Dual-stack storage
- Central services
  - Dual-stack

Update from ATLAS

- Input from Alastair Dewhurst (RAL)
- QMUL and Brunel have IPv6-only CPU already (testing)
  - QMUL use NAT64 to handle communication with IPv4 only hosts
  - QMUL periodically provides list of hosts that have been accessed via NAT64 (so easy to check if it is working)
- Several sites have dual stack storage
- Rucio
  - At the end of last year, Rucio team migrated all nodes to CC7
  - Enabled IPv6 at the same time
  - All required nodes are now done!
  - Rucio UI web front (rucio-ui.cern.ch) also made dual stack

ATLAS (2)

FTS Transfers

5

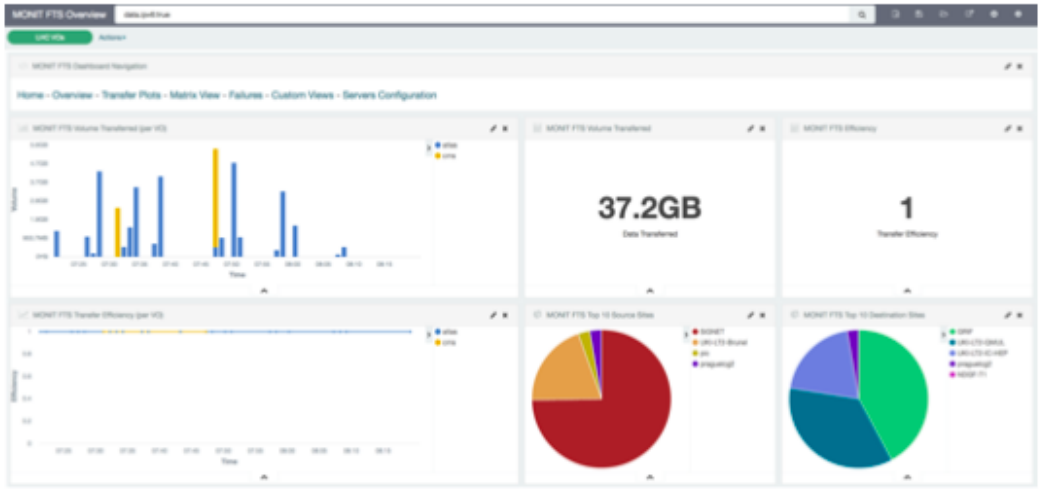

- Dual stack storage appears to be working.
- Add "data.ipv6:true" to any FTS monitor page for IPv6 only traffic.

https://monit.cern.ch/goto/99ae5597ed958e334ca46b8cb535be1f

Alastair Dewhurst, 9th January 2017

ATLAS (3)

- PANDA
  - Slower progress
  - Production nodes still IPv4-only
  - Delays caused by availability of pilot code on dev nodes
  - Waiting on sending jobs to QMUL IPv6-only
  - Plan for progress during next week or so
- Other services
  - Not critical for April 2017
  - AGIS not started
  - APF – CERN done not others yet
  - Frontier not started

Other news

- LHCb starting to test IPv6-only WN at Brunel
- Testing by CMS of ETF on IPv6-only WN
- Imperial team reports successful Dirac job submission to same IPv6 WN for some VOs
  - GridPP, Pheno, CERNatSchool
- dCache and transfer tests started to IHEP/CN

LHCOPN & Tier 1s

- Update yesterday from Edoardo Martelli
- Rising IPv6 adoption
  - 9 Tier1s and the Tier0 peering over IPv6
  - dual-stack perfSONAR installed in all of them
- LHCOPN IPv6 still missing from:
  - KR-KISTI (by end of 2016?)
  - RRC-KI-T1 JINR (will follow KIAE)
  - RRC-KI-T1 KIAE (IPv6 deployment started?)
  - UK-T1-RAL
    - RAL now making good progress – infrastructure ready by mid March 2017

Need to keep this up to date

Sites IPv6 connectivity

http://hepix-ipv6.web.cern.ch/sites-connectivity

Title	Type	LHCOPN IPv6 peering	LHCONE IPv6 peering	LHCONE IPv6 peers	NREN IPv6 peers	IPv6 LAN	dualstack perfSONAR	dualstack storage	Network Statistics
NDGF	Tier1	Yes	Yes	NORDUnet	NORDUnet	Yes	Yes	Yes	
INFN CNAF	Tier1	Yes	Yes	GARR	GARR	Yes	Yes	No	https://gins.garr.it/Statistics/viewer.php?stroke_ipv6=on&target%5B%5D=l...
FR-CCIN2P3	Tier1	Yes	Yes	RENATER	RENATER	Yes	Yes	No	N/A
ES-PIC	Tier1	Yes	Yes	RedIRIS	RedIRIS	Yes	Yes	Yes	N/A
DE-KIT/GridKa	Tier1	Yes	Yes	2a00:1398:104::/46	DFN	Yes	Yes	Yes	N/A
CH-CERN	Tier0	Yes	Yes	GEANT, ESnet, CERNlight	GEANT, SWITCH, RENATER, SURFnet, NORDUnet, ASGCnet, KREOnet, Internet2, CANARIE	Yes	Yes	No	https://netstat.cern.ch/monitoring/network-statistics/ext/?q=IPv6&p=EXT&...

ETF and IPv6

- Input from Marian Babik
- ETF IPv6 instance provides dual-stack testing support for SAM

EFT IPv6 (2)

- Works for all experiments (though only ATLAS and CMS are configured now)
- Using experiments production topologies
  - custom hosts/services can be added manually
- Does not publish data to SAM3
  - so the test results it takes are not part of the official reports (yet)
- Aim is to help sites understand status/availability of their IPv6 resources as compared to IPv4
- Some additional remarks:
  - ETF IPv6 test only services that have IPv6 address - it parses a list of CEs/SEs from the experiments feeds and only monitor those that have an IPv6 entry
  - Uses the exact same plugins and configuration we currently run in production and will thus receive all the updates (wrt. topology, metrics, updated tests, etc.)
  - It groups services to sites, accessible via host groups
  - Custom host groups and tests can be defined, such as e.g. ATLAS central services to check DNS/TCP reachability of the central services
  - Not auto-reloaded since central services are not part of the experiments feeds, but can be extended via API or manually from a static list of hosts/ports
- This instance can be added to the ETF central, which provides an overview of site services across all experiments
  - so it can be used to compare how site services perform (wrt IPv4 vs IPv6)

perfSONAR & IPv6

- http://hepix-ipv6.web.cern.ch/perfsonar-ps
- Another table to keep up to date
- The online version contains more sites

perfSONAR-ps

Title	Location	Site	Url	Notes
pship0[12].csc.fi	FI_HIP_T2	Tier 2	pship0[12].csc.fi	GPI
ps0[1-l/2-b].farm.particle.cz	praguelcg2	Tier 2	ps0[1-l/2-b].farm.particle.cz	GPI
perfsonar-[bandwidth/latency].esc.qmul.ac.uk	UKI-LT2-QMUL	Tier 2	perfsonar-[bandwidth/latency].esc.qmul.ac.uk	GPI
perfsonar-ps-0[12].desy.de	DESY	Tier 2	http://perfsonar-ps-0[12].desy.de/toolkit/	LHCONE / Gerneral Purpose Internet (GPI)
perfmon.dur.scotgrid.ac.uk	UKI-SCOTGRID-DURHAM	Tier 2	perfmon.dur.scotgrid.ac.uk	GPI
netmon00.grid.hep.ph.ic.ac.uk	UKI-LT2-IC-HEP	Tier 2	netmon00.grid.hep.ph.ic.ac.uk	GPI
icgperf.shef.ac.uk	UKI-NORTHGRID-SHEF-HEP	Tier 2	icgperf.shef.ac.uk	GPI
hcc-ps0[12].unl.edu	University of Nebraska-Lincoln	Tier 2	http://hcc-ps0[12].unl.edu/	GPI
dc2-grid-ps-00.brunel.ac.uk	UKI-LT2-Brunel	Tier 2	dc2-grid-ps-00.brunel.ac.uk	GPI
ps[bl]01.pic.es	PIC	Tier 1	ps[bl]01.pic.es	LHCOPN / GPI
perfsonar-ps[2].ndgf.org	NDGF	Tier 1	perfsonar-ps[2].ndgf.org	LHCOPN + GPI

Dual-stack mesh

- http://psmad.grid.iu.edu/maddash-webui/index.cgi?dashboard=Dual-Stack%20Mesh%20Config

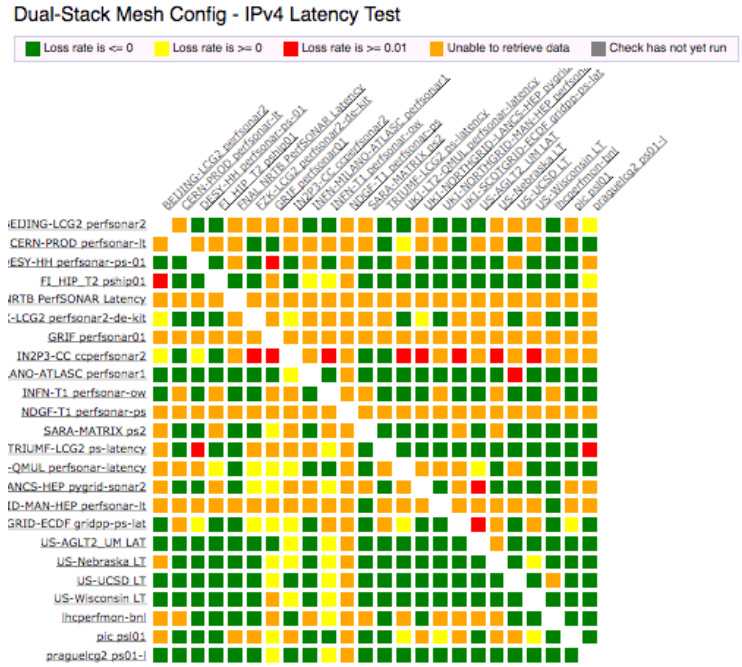

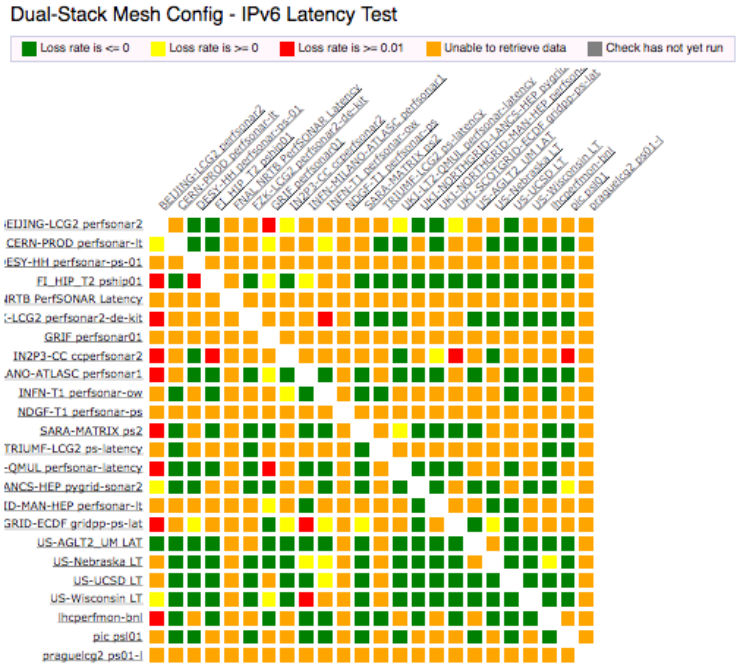

Dual-stack services

http://orsone.mi.infn.it/~prelz/ipv6_bdii/

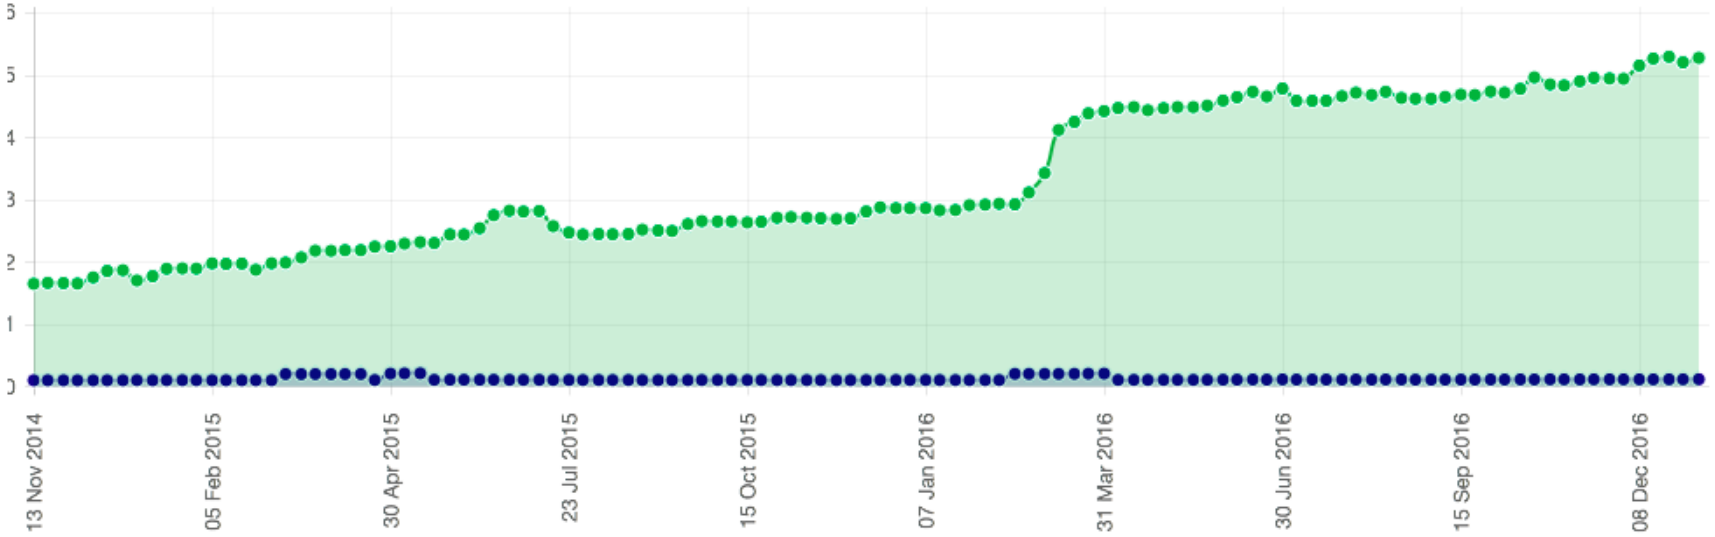

Summary

- Much improved engagement by Tier 1s
  - Most are ready
  - BUT – still limited dual-stack storage
    - Except at WG sites who already had this
- A good number of Tier 2s run dual-stack
  - BUT *MANY* do not
- No show-stoppers identified to date
- Still a lot of work ahead of us
- Need to decide how best to track/urge/encourage/support the Tier 2's
  - A task for WLCG Operations?

Links

- HEPiX IPv6 web

http://hepix-ipv6.web.cern.ch

- Working group meetings

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

- WLCG Operations IPv6 Task Force

http://hepix-ipv6.web.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*

2016 – *two papers to be submitted*

Questions?