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

David Kelsey (STFC-RAL)

WLCG GDB, Taipei (ISGC2017)

8 Mar 2017

Outline

- An update from the HEPiX IPv6 WG
  - Since January 2017 GDB
- Reminder: *WLCG MB approval (20 Sep 2016) of the plan to support IPv6-only CPU from April 2017*
- Recent news
- Current status
  - Are there any known show-stoppers?
  - Are we ready?

HEPiX IPv6 WG meetings

- Meetings held ~monthly
  - All one hour Vidyo meetings
    - 16 March, 12 April 2017
- Last F2F at CERN 2/3 Feb 2017
- Next F2F at CERN 15/16 May 2017

Google IPv6 statistics

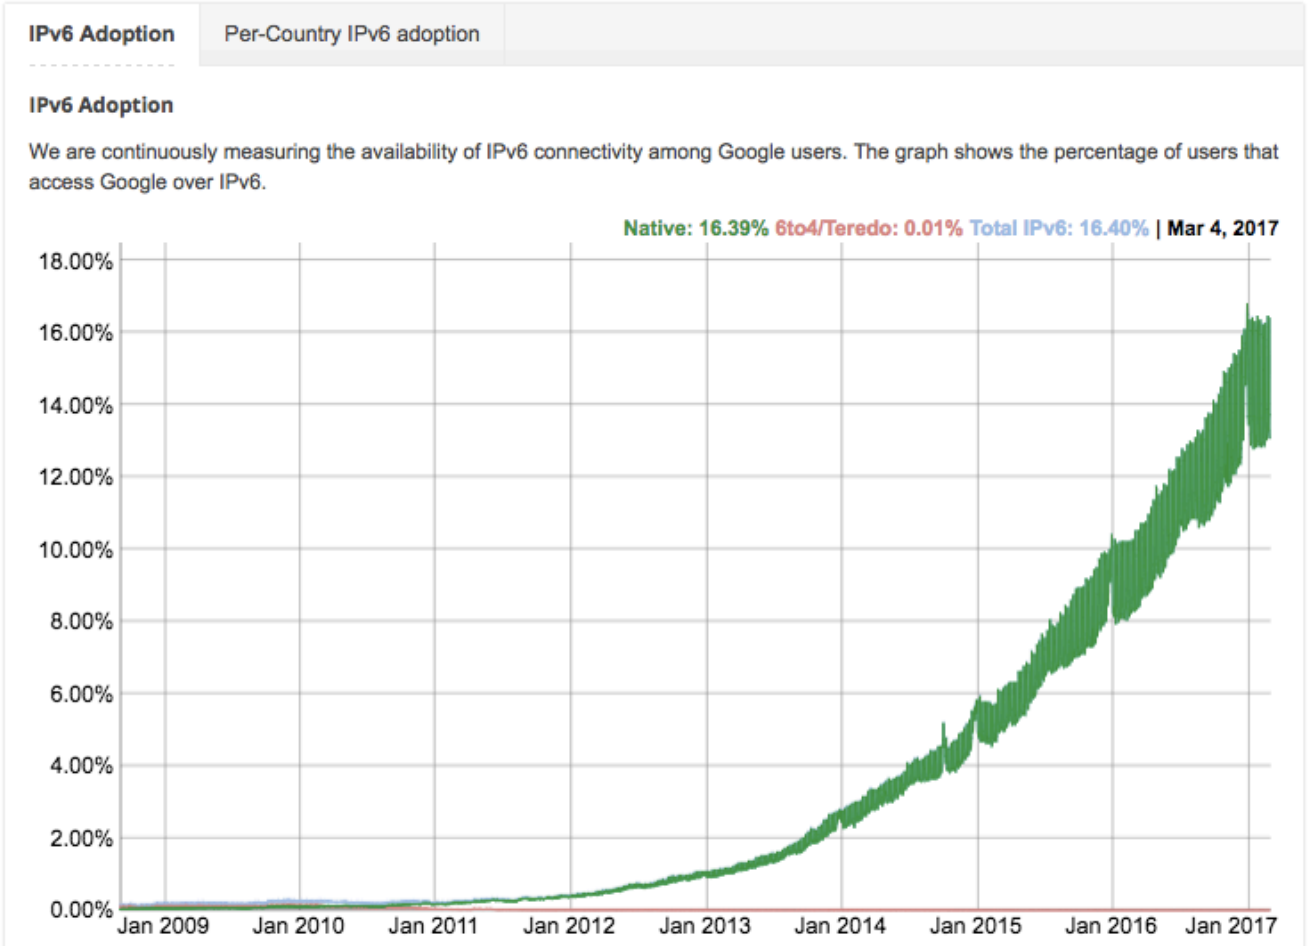

% BDII services (dual-stack)

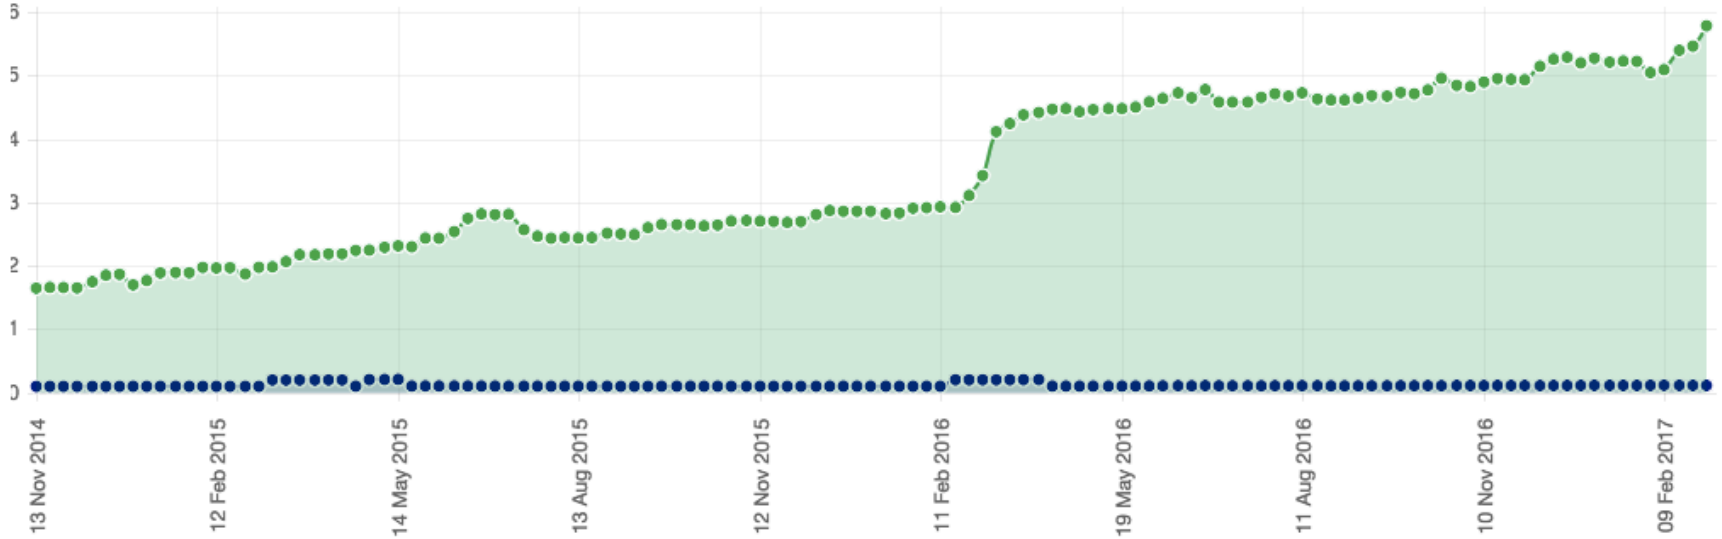

News (experiments)

Update from ATLAS

Input from Alastair Dewhurst (RAL)

- Jobs have successfully run on IPv6-only WN at Brunel using test Panda Server
- Planning on upgrading production Panda Server shortly
- Possible Frontier issue still being investigated

ATLAS (2)

- ATLAS will expect Tier 2s to provide dual stack storage by end of Run 2 (Jan 2019)
- Sites that can't will only be used for secondary data
- More information at ATLAS Computing and Software week (13th -17th March 2017)

LHCb

Input from Raja Nandakumar

- IPv6 is no longer something new or an "issue”
- We can and do run in the same way whether it is ipv4 / ipv6 / dual-stack
- Any issues will be dealt with by the production team as a normal incident with tickets opened through the usual means, like GGUS
- look forward for more dual-stack storage site(s)
  - PIC is only one for now

LHCb (2)

- Look forward to having **100%** SEs that support LHCb running in dual-stack mode!
- If any site wants to provide cpu in dual-stack or IPv6-only, that too should be fine with us
  - though notice needed for pure-IPv6-only WNs
  - until the majority of our SEs (esp. CERN) are dual-stack

News
(Tier 0 & 1s)

News CERN

Input from Edoardo Martelli

- CERN has started deploying RFC 6939 (Client Link-Layer Address Option in DHCPv6) on its Campus and Datacentre routers
- The deployment of the option has suddenly become urgent because the Network Manager of the latest CERN Centos 7 (CC7) distribution doesn't allow to set a DUID that contains the MAC address (-LLT or -LL)
- fact that hampers the lease of addresses to CC7 machines

Tier 1 & IPv6

- Good IPv6 adoption
  - 9 Tier1s and the Tier0 peering over IPv6
  - dual-stack perfSONAR installed in all of them
- LHCOPN IPv6 still missing from:
  - UK-T1-RAL (IPv6 Tier 1 on 8 March 2017)
  - KR-KISTI (new router hardware needed - by June 2016?)
  - RRC-KI-T1 KIAE (IPv6 deployment started)
  - RRC-KI-T1 JINR (will follow KIAE)

Tier 1 - PIC

Input from Fernando Lopez

- All production squid servers now dual-stack
- Aiming for 100% of WNs before summer
- A lot of recent work with enabling dual-stack worker nodes, updated stats are:
  - Nodes: 36.2%
  - Slots: 34.7%
  - HS06: 38.0%
- *Second Tier 1 with dual-stack WNs (after NDGF)*

News
(Tier 2s)

Brunel Tier 2

Input from Raul Lopes

- The news from the Brunel IPv6-only node
- Atlas can jobs on pure IPv6 node
  - It's even listed in the Atlas bigpanda website
- LHCB can also run jobs on the pure node
- CMS has been using IPv6-only node for several months

Other News

IGTF CA – CRLs & IPv6

- Ulf Tigerstedt has been monitoring status for 1 year
- *http://cvmfs-6.ndgf.org/ipv6/overview.php*
- Lowest: 31 working, 8 "has AAAA-record but the network does not work" and 56 IPv4 only CRL servers to
- 1st March 2017: 42/1/52
- So 11 more dual-stacked CAs during one year
- IGTF is currently pushing for all!

“How to” deploy IPv6 - Tier 2

Andrea Sciaba

- Started work collecting knowledge and advice
- https://hepix-ipv6.web.cern.ch/content/how-deploy-ipv6-wlcg-tier-2-site
- Tutorial in WLCG Workshop?
  - Manchester – June 2017

CHEP papers IPv6 & security

- Two submitted to CHEP2016 proceedings
  - "Deployment of IPv6-only CPU resources at WLCG sites"
  - "IPv6 Security"
- IPv6 Security checklist for Sites at:
  - *http://hepix-ipv6.web.cern.ch/content/ipv6-security-checklist-wlcg-site-system-administrators-and-networking-teams*

Monitoring Data Transfers

- https://monit.cern.ch/
- FTS Transfer plots
- Add “data.ipv6:true” to any FTS monitor page for IPv6 only traffic

FTS transfers (last week) Total

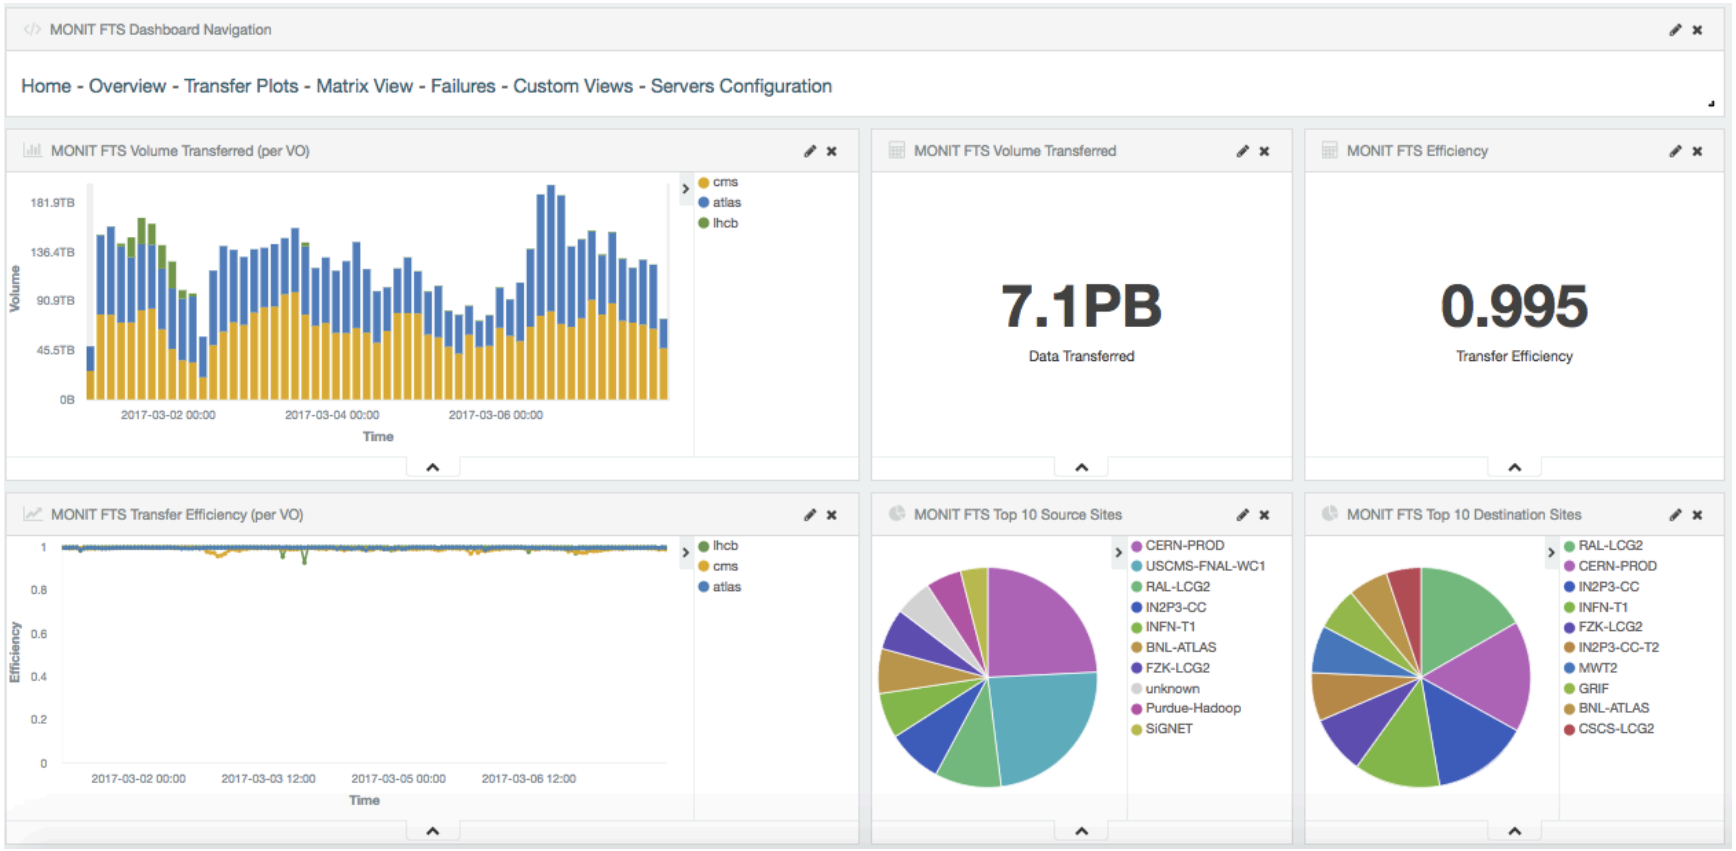

FTS transfers (last week) IPv6

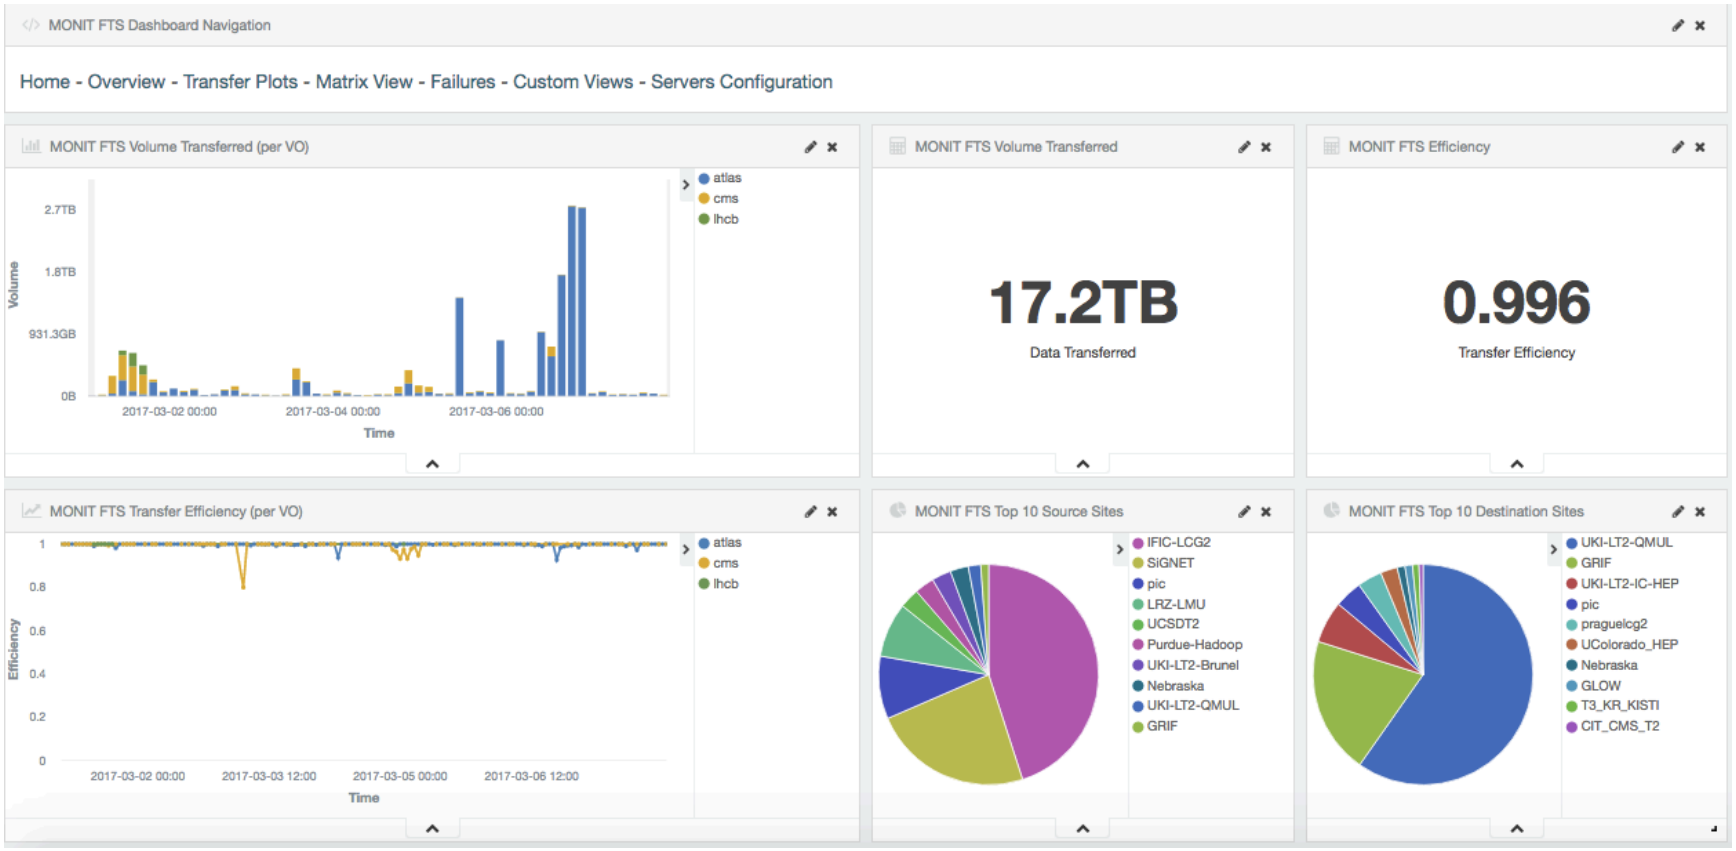

Need to keep this up to date

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

perfSONAR-ps

Title	Location	Site	Url	Notes
pship0[12].csc.fi	FI_HIP_T2	Tier 2	pship0[12].csc.fi	GPI
ps0[1-|2-b].farm.particle.cz	praguelcg2	Tier 2	ps0[1-|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 / General 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
lcgperf.shef.ac.uk	UKI-NORTHGRID-SHEF-HEP	Tier 2	lcgperf.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
perfsonar-de-kit.gridka.de	DE-KIT	Tier 1	http://perfsonar-de-kit.gridka.de/toolkit/	LHC[OPN/ONE] / GPI
perfsonar-bw.cern.ch	CERN	Tier 1	https://perfsonar-bw.cern.ch/	LHC[OPN/ONE] / GPI
lhcperfmon.bnl.gov	BNL	Tier 1	lhcperfmon.bnl.gov	latency node
lhcmon.bnl.gov.	BNL	Tier 1	lhcmon.bnl.gov.	throughput node
ccperfsonar[12].in2p3.fr	FR-CCIN2P3	Tier 1	ccperfsonar[12].in2p3.fr	LHC[OPN/ONE] + GPI

perfSONAR – dual-stack mesh

Production version

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

Test version

- *http://maddash.aglt2.org/maddash-webui/index.cgi?dashboard=Dual-Stack%20Mesh%20Config*

Dual-stack mesh

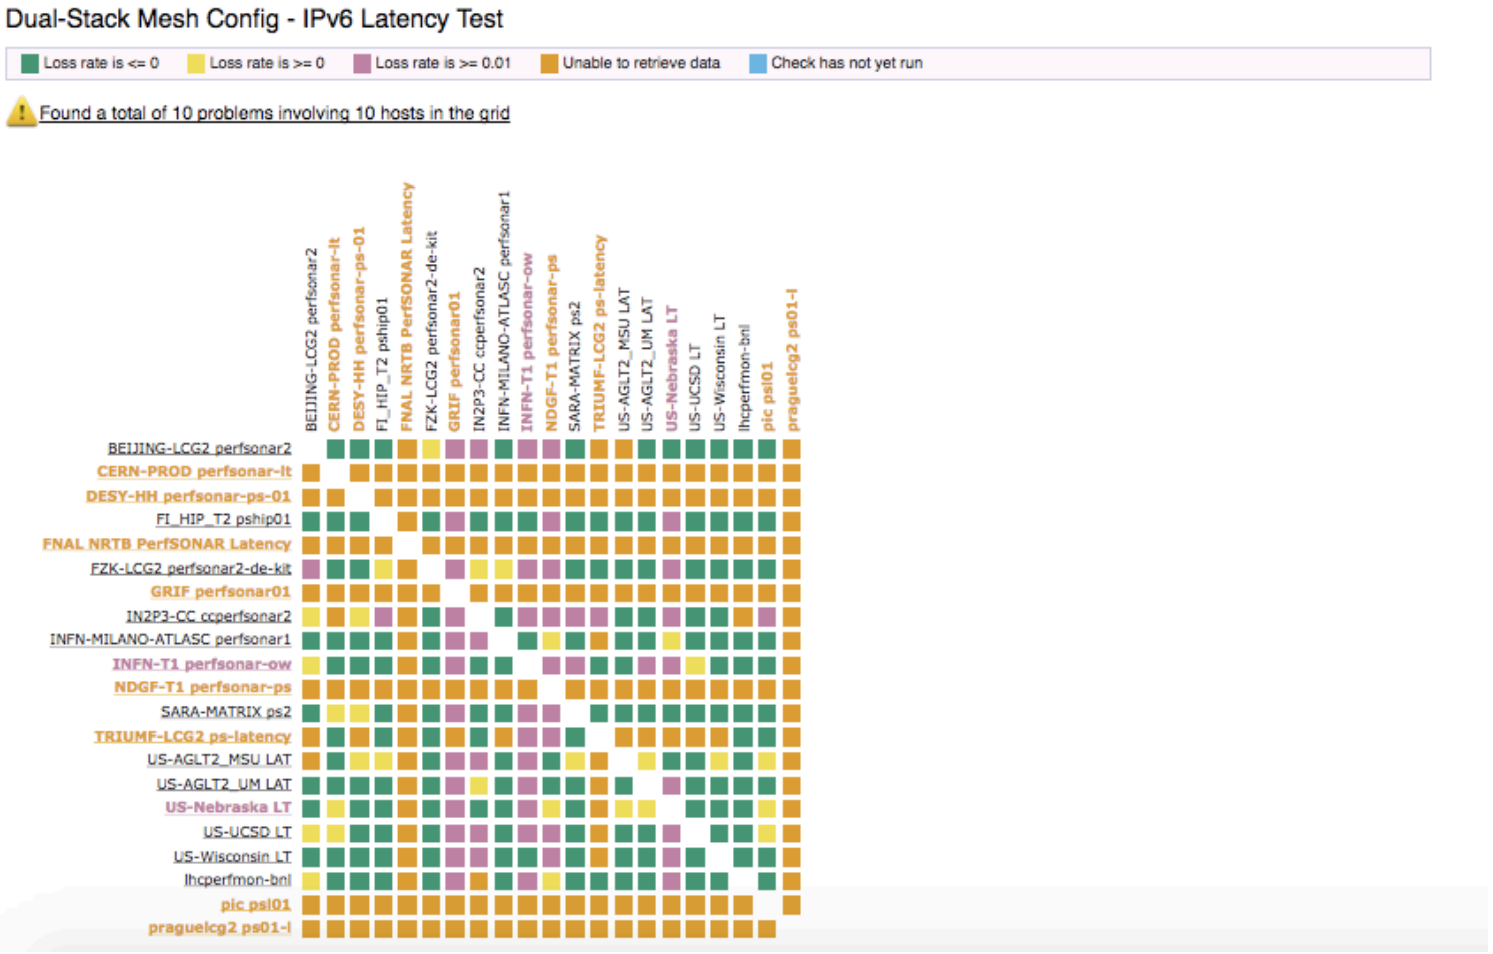

Summary

- Much improved engagement by Tier 1s
  - Most are ready
- BUT – still limited dual-stack storage
    - Except at working group sites who already had this
- A good number of Tier 2s run dual-stack
  - BUT *MANY* do not!
- WLCG Tier 2s must start planning NOW
- No show-stoppers identified to date
- Still a lot of work ahead of us!
- 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 submitted*

Questions?

Backup slides

ALICE

- ALICE central services have been dual stack for more than a year
  - Storage is fully federated
    - Any site can access data from any site
  - To support IPv6-only resources, all data must be available on some IPv6-enabled storage
  - ALICE will not be able to support IPv6-only CPU resources by April 1st 2017
  - ALICE can support IPv6-only CPU resources as soon as enough sites have upgraded their storage to dual stack
    - Goal of the end of Run II