



# The HEPiX IPv6 Working Group

David Kelsey (STFC-RAL)

WLCG GDB, CERN

12 Feb 2014

# Outline

Update of status and plans

- F2F meeting of IPv6 WG – CERN 23/24 Jan 2014
- HEPiX IPv6 testing news
- Acknowledgement/forgiveness for sites who test dual-stack services in production
- Software and tools survey
  - dCache status
- Survey of sites
- 1<sup>st</sup> Workshop to encourage dual-stack services
  - Pre-GDB at CERN on 10 June 2014

# IPv6 F2F meeting 23/24 Jan

- Agenda included
  - In addition to standard items
  - IPv6, CERN Cloud and OpenStack
    - Luis Fernandez Alvarez
    - OpenStack Havana cannot control IPv6 traffic
    - CERN having to develop its own controls
    - Longer term – OpenStack Neutron – Networking as a service
  - Revisions to CHEP2013 paper
    - Resubmitted and now accepted

# Other news from meeting

- CERN news
  - Progressing well with IPv6 deployment
  - Firewall and DNS ready
  - Aiming to deploy campus-wide in March
  - Still some problems with dhcpv6
- Xrootd V4 still not available (still true?)
- DESY – problems with dhcpv6 fixed
  - Now ready for deployment
- KIT building IPv6 test cluster
- Duncan Rand – news from Imperial – see slides
  - Includes news about PerfSonar (dual stack? separate endpoints?)

# IPv6 file transfer tests

- Tony Wildish (CMS)
- Successful mesh of transfers before CHEP2013
  - GridFTP and UberFTP
- Left unattended
- Lots of sites were broken
  - A variety of unrelated issues – now fixed
  - E.g. globus port range not being set after upgrade
  - Not a problem with IPv6

# ATLAS testing

- Alastair Dewhurst (ATLAS, RAL)
- AGIS
  - Flags added for IPv6 support on storage endpoints
- Start with simple data movement tests
  - `globus_url_copy`, `lcg` commands
- Then data tests using HammerCloud
- Frontier service
  - Squid 2.8 is non IPv6 compatible
    - Squid 3.1 at Tier 2s should be usable

# Next steps in testing

- Once gridftp mesh is working again
- KIT, PIC, FZU, NDGF
  - PhEDEx/SRM/dCache tests (on testbed)
- Then move to use of SRM/FTS/PhEDEx
  - With various **production** endpoints
  - Volunteer sites
- We need clearly defined use case (from expts)
  - CERN Worker Node/Virtual Machine IPv6-only
  - What WLCG services outside CERN need to be dual-stack?

# Forgiveness/Compensation?

- As we move to dual-stack IPv6 testing on production services ...
- Some sites say they will only join if it is accepted by WLCG, by Expts, by Funding agencies that their availability might suffer
  - Either because of service problems
  - Or because sys admin effort is busy on IPv6
- Very similar to middleware testing
- Simone Campana is preparing paper on this for WLCG MB

# Software and Tools survey ...

## IPv6 compliance of WLCG applications

| Software Component     | Type                       | Used by Experiment      | Version | IPv6 Compliance |
|------------------------|----------------------------|-------------------------|---------|-----------------|
| AliEN                  | LHC Experiment Application | ALICE                   |         |                 |
| ARC CE                 | Middleware                 | ATLAS, CMS              |         | YES             |
| ARGUS                  | Middleware                 | ALICE, ATLAS, CMS, LHCb |         | Unknown         |
| BDII                   | Middleware                 | ATLAS, CMS, LHCb        | EMI 2   | YES             |
| BestMAN                | Middleware                 | ATLAS, CMS              |         | Unknown         |
| CASTOR                 | Middleware                 | ALICE, ATLAS, CMS, LHCb |         | NO              |
| cfengine               | Monitoring                 |                         |         | Unknown         |
| CMS Tag Collector      | LHC Experiment Application | CMS                     |         | Unknown         |
| CMSSW                  | LHC Experiment Application | CMS                     |         | Unknown         |
| cmsweb                 | LHC Experiment Application | CMS                     |         | Unknown         |
| CRAB 2                 | LHC Experiment Application | CMS                     |         | Unknown         |
| Cream CE               | Middleware                 | ALICE, ATLAS, CMS, LHCb | 1.16.2  | YES             |
| CVMFS                  | Other Application          | ALICE, ATLAS, CMS, LHCb |         | YES             |
| Dashboard Google Earth | Monitoring                 | ALICE, ATLAS, CMS, LHCb |         | Claimed         |

# Software survey (2)

|                       |                            |                         |        |                  |
|-----------------------|----------------------------|-------------------------|--------|------------------|
| dCache                | Middleware                 | ALICE, ATLAS, CMS, LHCb | 2.6.19 | Claimed          |
| dCache                | Middleware                 | ALICE, ATLAS, CMS, LHCb | 1.9.12 | NO               |
| DIRAC                 | LHC Experiment Application | LHCb                    |        | Unknown          |
| DPM                   | Middleware                 | ALICE, ATLAS, CMS, LHCb |        | YES with caveats |
| EGI Accounting Portal | Monitoring                 | ALICE, ATLAS, CMS, LHCb |        | Unknown          |
| EOS                   | Middleware                 | ALICE, ATLAS, CMS, LHCb |        | NO               |
| Experiment Dashboards | Monitoring                 | ALICE, ATLAS, CMS, LHCb |        | Unknown          |
| Frontier              | LHC Experiment Application | ATLAS, CMS              |        | NO               |
| FTS                   | Middleware                 | ATLAS, CMS, LHCb        | FTS 3  | YES              |
| Ganglia               | Monitoring                 |                         |        | YES              |
| GFAL/lcg_util         | Middleware                 | ALICE, ATLAS, CMS, LHCb |        | Claimed          |
| glideinWMS            | Middleware                 | CMS                     |        | Claimed          |
| globus toolkit        | Middleware                 | ALICE, ATLAS, CMS, LHCb | 5.2.5  | YES              |
| GOCdb                 | System Management Tool     | ALICE, ATLAS, CMS, LHCb |        | Unknown          |
| Gratia Accounting     | Monitoring                 |                         |        | Unknown          |
| Gridsite              | Middleware                 |                         |        | Claimed          |
| Gstat                 | Monitoring                 |                         |        | Unknown          |

# Software survey (3)

|               |                            |                         |       |                  |
|---------------|----------------------------|-------------------------|-------|------------------|
| GUMS          | Middleware                 | ATLAS, CMS              |       | Unknown          |
| Hadoop        | Middleware                 | ATLAS, CMS              |       | NO               |
| HammerCloud   | LHC Experiment Application | ATLAS, CMS, LHCb        |       | Unknown          |
| HTCondor      | Middleware                 | ATLAS, CMS              |       | YES with caveats |
| iCMS          | LHC Experiment Application | CMS                     |       | Unknown          |
| LFC           | Middleware                 | ATLAS, LHCb             | 1.8.7 | YES              |
| MonALISA      | Monitoring                 |                         |       | Claimed          |
| MyOSG         | Monitoring                 |                         |       | Unknown          |
| MyProxy       | Middleware                 |                         | EMI 1 | Unknown          |
| MyWLCG        | Monitoring                 | ALICE, ATLAS, CMS, LHCb |       | Unknown          |
| Nagios        | Monitoring                 | ALICE, ATLAS, CMS, LHCb |       | Claimed          |
| OpenAFS       | Other Application          | ALICE, ATLAS, CMS, LHCb | Last  | NO               |
| PanDA         | LHC Experiment Application | ATLAS, CMS              |       | Unknown          |
| perfSONAR     | Monitoring                 | ATLAS, CMS              |       | YES              |
| PhEDEx agents | LHC Experiment Application | CMS                     |       | YES              |
| Puppet        | System Management Tool     | ALICE, ATLAS, CMS, LHCb |       | YES              |
| REBUS         | Monitoring                 | ALICE, ATLAS, CMS, LHCb |       | Unknown          |

# Software survey (4)

|                                  |                            |                         |        |                  |
|----------------------------------|----------------------------|-------------------------|--------|------------------|
| REBUS                            | Monitoring                 | ALICE, ATLAS, CMS, LHCb |        | Unknown          |
| SAM                              | Monitoring                 | ALICE, ATLAS, CMS, LHCb |        | Unknown          |
| <a href="#">Scientific Linux</a> | Operating System           | ALICE, ATLAS, CMS, LHCb | 5.7    | YES              |
| SDT IB and QA pages              | LHC Experiment Application | CMS                     |        | Unknown          |
| Squid                            | Other Application          | ALICE, ATLAS, CMS, LHCb | 2.8    | NO               |
| Squid                            | Other Application          | ALICE, ATLAS, CMS, LHCb | 3.1    | YES              |
| StoRM                            | Middleware                 | ALICE, ATLAS, CMS, LHCb | 1.11.2 | YES with caveats |
| Ticket system -- GGUS            | System Management Tool     | ALICE, ATLAS, CMS, LHCb |        | Unknown          |
| various D web tools              | LHC Experiment Application | CMS                     |        | Unknown          |
| VOMS                             | Middleware                 | ALICE, ATLAS, CMS, LHCb | 2.0.11 | YES              |
| VOMS-Admin                       | Middleware                 | ALICE, ATLAS, CMS, LHCb | 2.0.11 | YES with caveats |
| WLCG squid monitoring            | Monitoring                 | ATLAS, CMS              |        | Claimed          |
| WMAgent                          | LHC Experiment Application | CMS                     |        | Unknown          |
| WMS                              | Middleware                 | ALICE, ATLAS, CMS, LHCb | EMI 2  | Unknown          |
| xrootd                           | Middleware                 | ALICE, ATLAS, CMS       | 3.1.0  | NO               |
| xrootd                           | Middleware                 |                         | 4.0.0  | Claimed          |

# What has/will change(d)?

- Descriptions of yes/no etc needed on web
  - *Yes* = we have tested
  - *Yes with caveats* = needs configuration or works only in a subset of protocols/configs
  - *Claimed* = by developers but not yet tested
  - *Unknown* = ticket submitted to developers
  - *Blank* = not even started investigation
- Click on product for details
  - See example on next slide

# Example: status of Frontier

**Name of Supplier and/or Developer:**

Dave Dykstra

**Does developer claim IPv6 compliance?:**

No

**Did you try running on dual stack system?:**

No

**Does it bind to both stacks?:**

No

**Is IPv6 preferred by default?:**

No

**Can it be configured to prefer IPv6 or IPv4?:**

No

**More information and/or link(s) to ticket(s):**

Frontier is not IPv6 compliant and there's a pretty big obstacle in its way. squid3 is IPv6 compliant, but it is missing two important features so we can't yet use it. Apparently one of them (collapsed forwarding) has been implemented but not yet released, but the other (If-Modified-Since support, the famous squid bug #7) is top on their priority list of features that are not yet staffed. I am expecting that if they don't get to it by early next year, I'm going to have to put significant time into contributing the implementation of that feature to the squid project. In addition, I'm pretty sure the frontier client is not completely IPv6 compliant, but I don't expect that to be a significant amount of work to fix. The frontier tomcat server doesn't matter because it is only accessed via localhost from squid. It does access the DB on its backend but I am assuming that since that's always on a LAN that an IPv4 address will be available if JDBC has a problem with IPv6. Googling indicates that IPv6 is supported starting with Oracle 11gR2 along with the associated JDBC driver from them, so it should be fine in any case. (Dave Dykstra, 7/11/2013)

-

# dCache & IPv6

- Pre-release of dCache 2.8.0 tested at NDGF
  - IPv6 fixes will not be backported to V2.6
- Seems to work
  - BUT all Gridftp-1, thus proxied via the FTP door instead of the data being sent from client directly to the pool
- dCache uses the response to the PASV FTP command to steer data connections to a pool of servers for load balancing
  - And obeying RFC2428 which says protocol and host fields of the EPSV response "MUST be blank" (for NAT)
- But Globus, jGlobus, Uberftp clients \*do\* parse these fields
- Discussions ongoing with dCache team

# Next steps

- Survey of all WLCG sites
  - When will they run out of IPv4 addresses
  - When will they be capable of IPv6 support
- We mentioned summer IPv6 workshop
- We prefer a one-day pre-GDB meeting
  - 10 June 2014 at CERN
  - Aimed primarily at Tier 1 sites (others welcome)
  - Encourage IPv6 support
  - Move to dual-stack WLCG services
  - Testing needs to move to production services

# Next IPv6 meetings

- 13 Feb 2014 - Vidyo
- 13 Mar 2014 - Vidyo
- 10/11 April 2014 - F2F at CERN
- Pre-GDB (full day) on 10 June at CERN

# Further info

- HEPiX IPv6 web

<http://hepixonv6.web.cern.ch>

- HEPiX IPv6 wiki

<https://w3.hepixonv6.org/ipv6-bis/>

- Working group meetings

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

- WLCG Operations IPv6 Task Force

<http://hepixonv6.web.cern.ch/content/wlcg-ipv6-task-force-0>

# Questions?