

# GridPP

UK Computing for Particle Physics

## perfSONAR and IPv6

Duncan Rand  
Imperial College

- The WLCG has deployed perfSONAR for network monitoring
- Country-based meshes, e.g. UK shown here
- A number of these WLCG production perfSONAR hosts are already dual-stack
- So it makes sense to try to measure IPv4 and IPv6 traffic separately

## UK sites Dashboard

UK sites - UK Cloud BWCTL Mesh Test



PerfSONAR-PS IPv6 status April 19 2011 (Joe Metzger)

## IPv6 in the perfSONAR-PS codebase

(Green check means activity is done. Yellow means in progress. Red means no firm plans to address.)

- perfSONAR-PS software is nominally ready for use with IPv6
- BWCTL measures throughput (iperf)
- OWAMP measures one-way latency and loss
- It is possible to specify IPv4 or IPv6 in a test definition
- If not specified, the tools will prefer IPv6 if available on both endpoints

Tool	IPv6 Development Completed	Limited Deployment	Production Deployment Across Multiple Domains
BWCTL	✓	✓	✓
OWAMP	✓	✓	✓
pSB MA	✓	✓	
Lookup Services	✓		
Topology Service	✓		
SNMP MA	✓	✓	
PinGER	✓	✓	
NDT	✓		
NPAD	✗		
Toolkit Configuration Tools & GUI	✓		



- Set up a mesh for testing dual-stack hosts
- Hosts can optionally have hostnames with -v4 and -v6 suffixes e.g.
  - netmon00.grid.hep.ph.ic.ac.uk has address 146.179.247.74 and IPv6 address 2001:630:12:580:207:43ff:fe11:ffb0
  - netmon00-v4.grid.hep.ph.ic.ac.uk has address 146.179.247.74
  - netmon00-v6.grid.hep.ph.ic.ac.uk has IPv6 address 2001:630:12:580:207:43ff:fe11:ffb0

```
{
  "sites" : [
    {
      "hosts" : [
        {
          "addresses" : [
            "perfsonar03.esc.qmul.ac.uk",
            "perfsonar03-v4.esc.qmul.ac.uk",
            "perfsonar03-v6.esc.qmul.ac.uk"
          ],
          "measurement_archives" : [
            {
              "read_url" : "http://perfsonar03.esc.qmul.ac.uk:8085/perfSONAR_PS/services/pSB",
              "write_url" : "perfsonar03.esc.qmul.ac.uk:8570",
              "type" : "perfsonarbuoy/bwctl"
            },
            {
              "read_url" : "http://perfsonar03.esc.qmul.ac.uk:8086/perfSONAR_PS/services/traceroute_ma",
              "write_url" : "http://perfsonar03.esc.qmul.ac.uk:8086/perfSONAR_PS/services/tracerouteCollector",
              "type" : "traceroute"
            },
            {
              "read_url" : "http://perfsonar03.esc.qmul.ac.uk:8085/perfSONAR_PS/services/pSB",
              "write_url" : "perfsonar03.esc.qmul.ac.uk:8569",
              "type" : "perfsonarbuoy/owamp"
            },
            {
              "read_url" : "http://perfsonar03.esc.qmul.ac.uk:8075/perfSONAR_PS/services/pinger/ma",
              "type" : "pinger"
            }
          ],
          "description" : "UKI-LT2-QMUL Test Host"
        }
      ]
    }
  ],
}
```

### Site definition

- Tests include hostnames with their -v4 and -v6 suffixes
- There is an issue with displaying data on maddash dashboard from hosts that don't have these. Attempted workaround is to use the host's IPv4 or IPv6 IP address explicitly
- Set 'ipv4\_only' or 'ipv6\_only' flags

```
{
  "members" : {
    "members" : [
      "2001:630:10:f00e::1009",
      "ps02-b-v6.ipv6.farm.particle.cz",
      "netmon00-v6.grid.hep.ph.ic.ac.uk",
      "perfsonar03-v6.esc.gmul.ac.uk",
      "t2ps-bandwidth-v6.physics.ox.ac.uk",
      "2001:67c:1148:202::2",
      "2600:900:6:1101:21b:21ff:fe96:a3d4",
      "2001:948:48:1::f502:1"
    ],
    "type" : "mesh"
  },
  "parameters" : {
    "force_bidirectional" : "1",
    "protocol" : "tcp",
    "tool" : "bwctl/iperf",
    "duration" : "30",
    "type" : "perfsonarbuoy/bwctl",
    "interval" : "21600",
    "ipv6_only" : "1"
  },
  "description" : "IPv6 throughput test"
},
```



- Currently eight sites in the mesh
  - FZU
  - T2\_FI\_HIP
  - PIC
  - Brunel, Imperial, QMUL, Oxford
  - UNL
- Measuring:
  - throughput
  - latency and loss
  - traceroute
  - ping

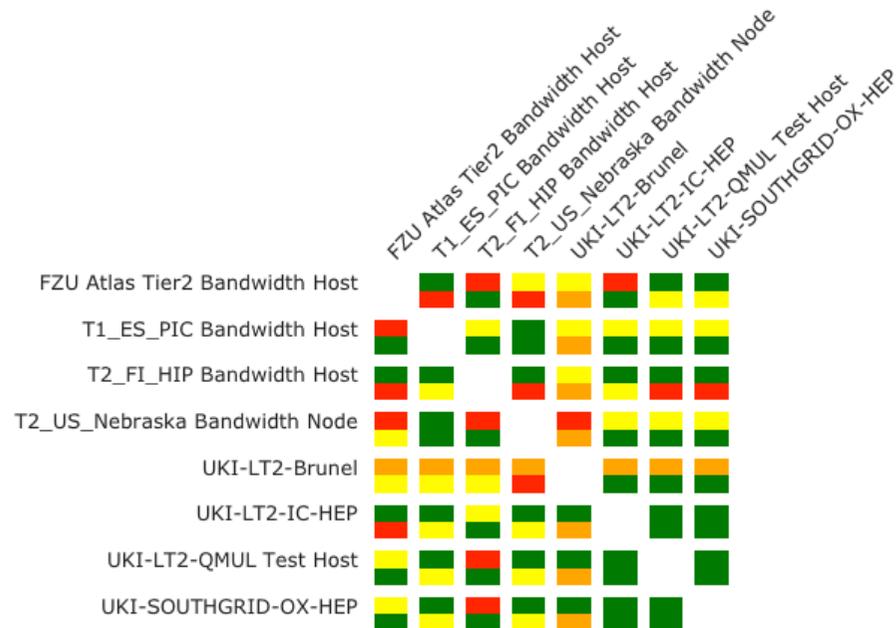
## perfSONAR Dashboard

Dashboards

### HEPiX IPv6 sites Dashboard

HEPiX IPv6 sites - IPv4 throughput test

■ Throughput  $\geq$  400Mbps   
 ■ Throughput < 400Mbps   
 ■ Throughput  $\leq$  100Mbps   
 ■ Unable to retrieve data



List of dashboards



## perfSONAR Dashboard

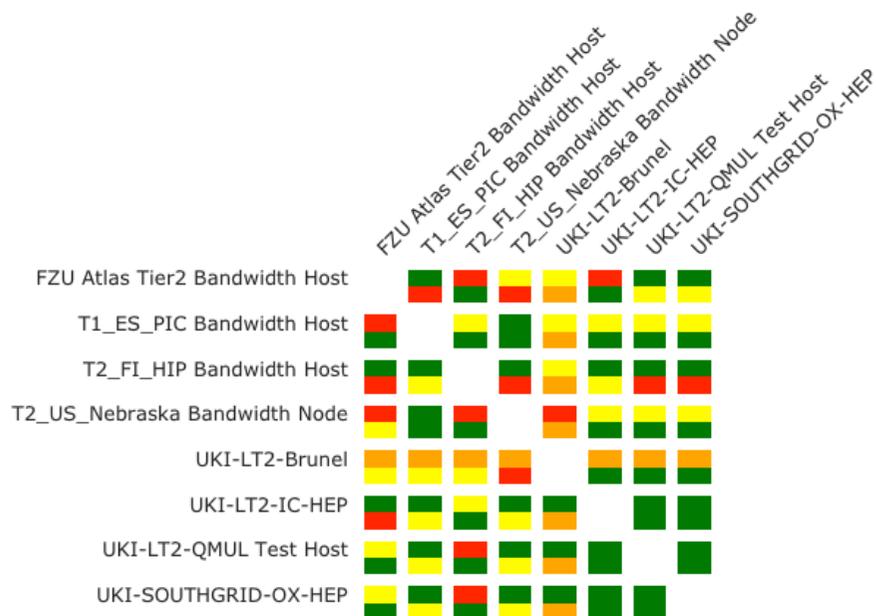
Dashboards

**HEPiX IPv6 sites Dashboard**

UK sites - IPv4 throughput test

All Grids

■ Throughput  $\geq$  400Mbps   
 ■ Throughput  $<$  400Mbps   
 ■ Throughput  $\leq$  100Mbps   
 ■ Unable to retrieve data



Rate thresholds

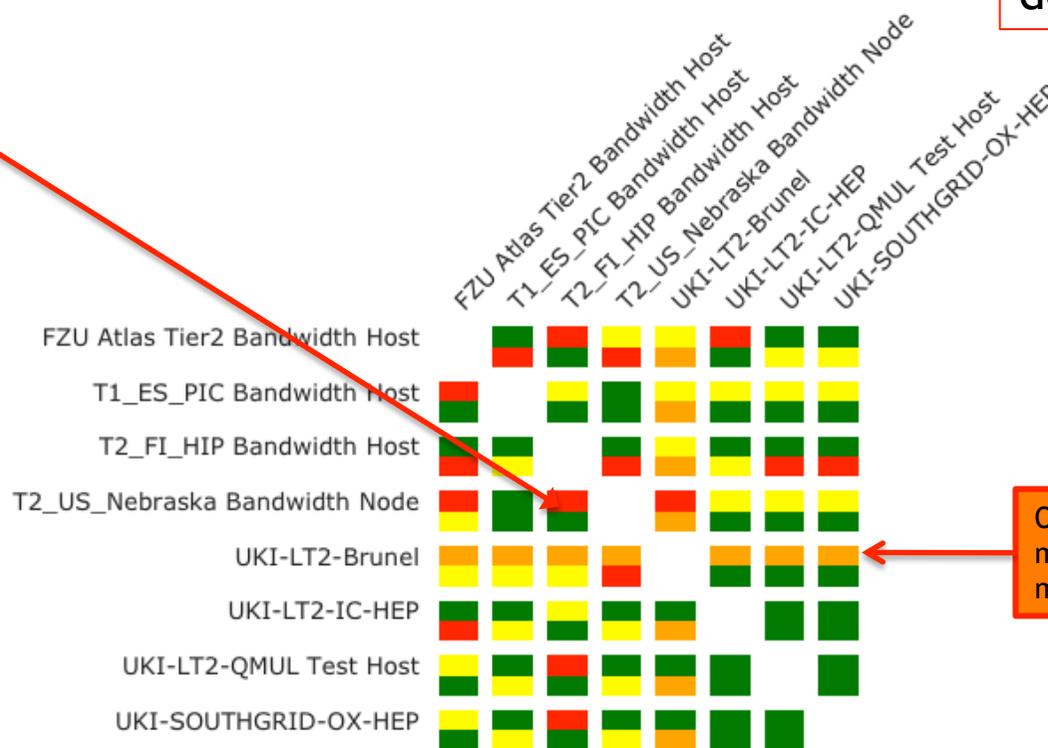


Hover to get rates from both measurement archives of sites involved

destination

Click to drill down (next slide)

source



Orange square means "no measurement"



## netmon00-v6.grid.hep.ph.ic.ac.uk to perfsonar03-v6.esc.qmul.ac.uk (Throughput)

Status: **OK** Last Checked: June 05, 2014 05:14:02 AM BST Next Check: June 05, 2014 13:14:02 PM BST

Summary History Check Details

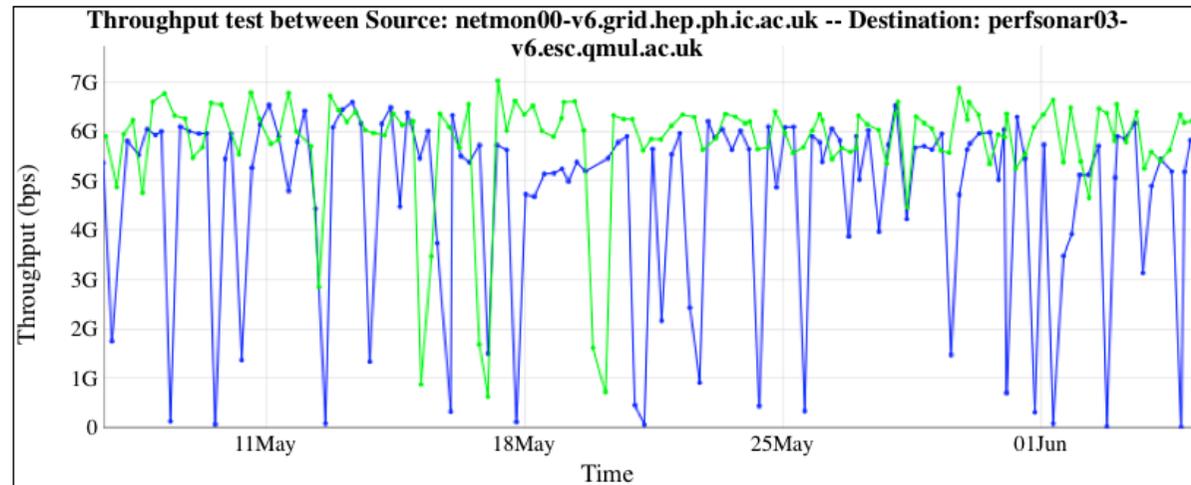
▼ Current Results

**Current Status: OK**  
**Result of last check: OK**  
**Message For Current Status:** Average throughput is 4.341Gbps

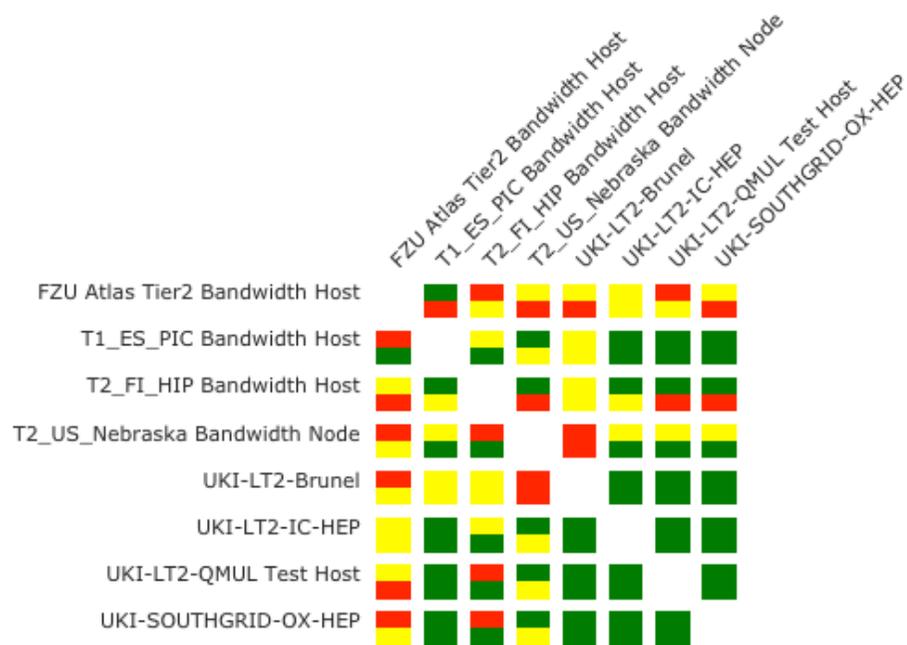
▶ Statistics

▼ Graph

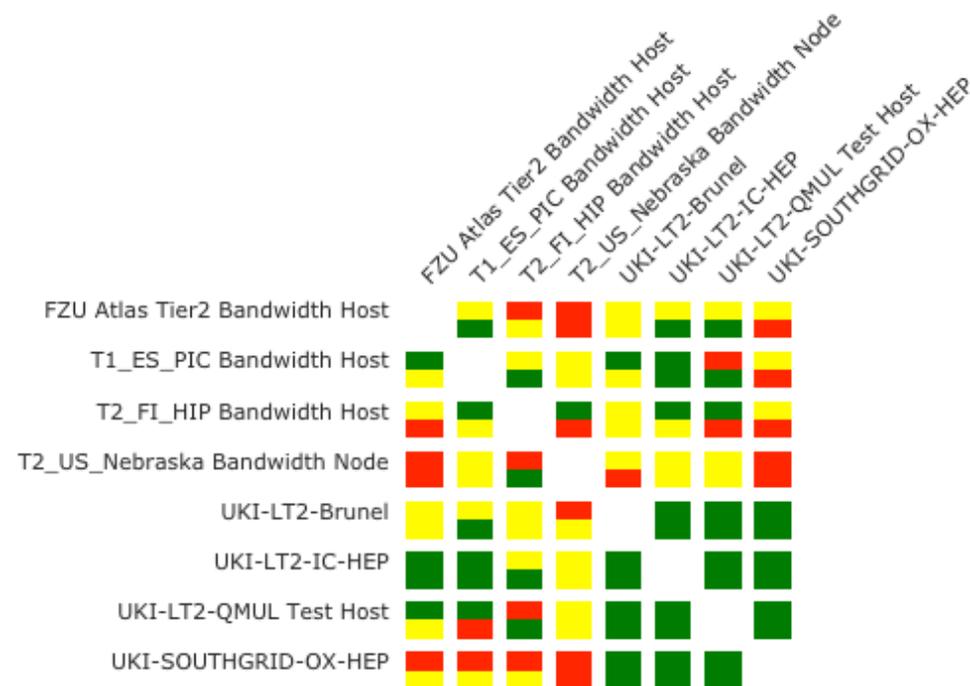
perfSONAR BWCTL Graph



## IPv4



## IPv6



■ Throughput  $\geq$  400Mbps    
 ■ Throughput  $<$  400Mbps    
 ■ Throughput  $\leq$  100Mbps    
 ■ Unable to retrieve data    
 ■ Check has not yet run

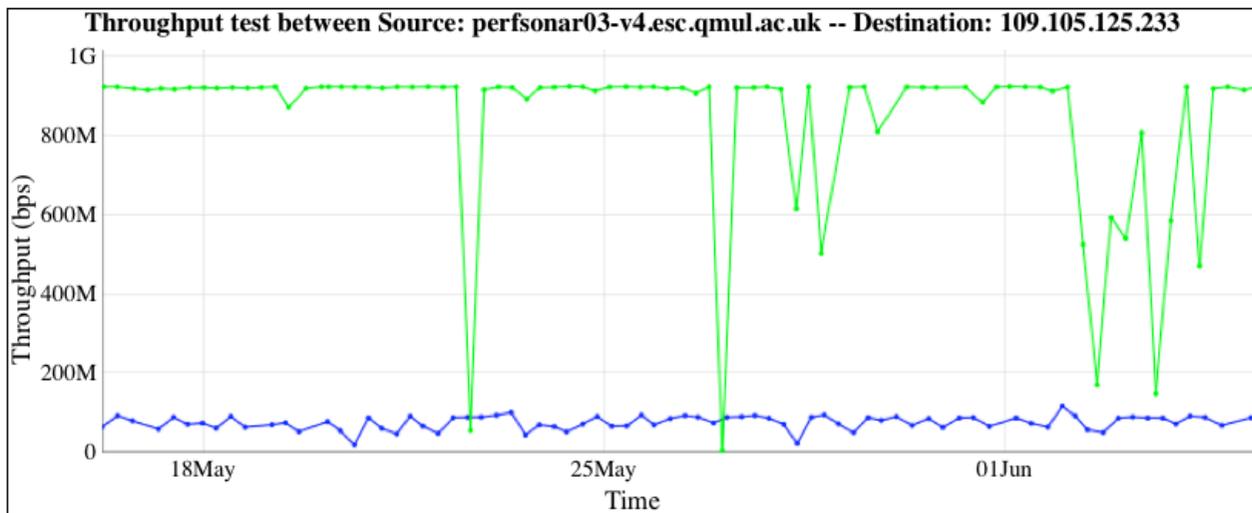


### Graph Key

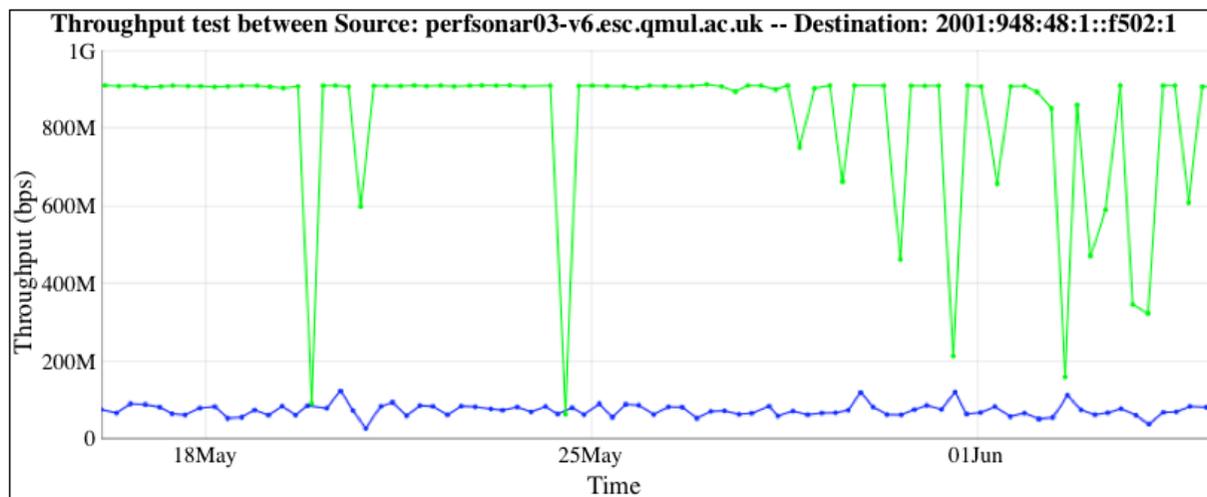
- Src-Dst throughput
- Dst-Src throughput

IPv4

QMUL to  
T2\_FI\_HIP  
slower over  
both IPv4  
and IPv6...



IPv6



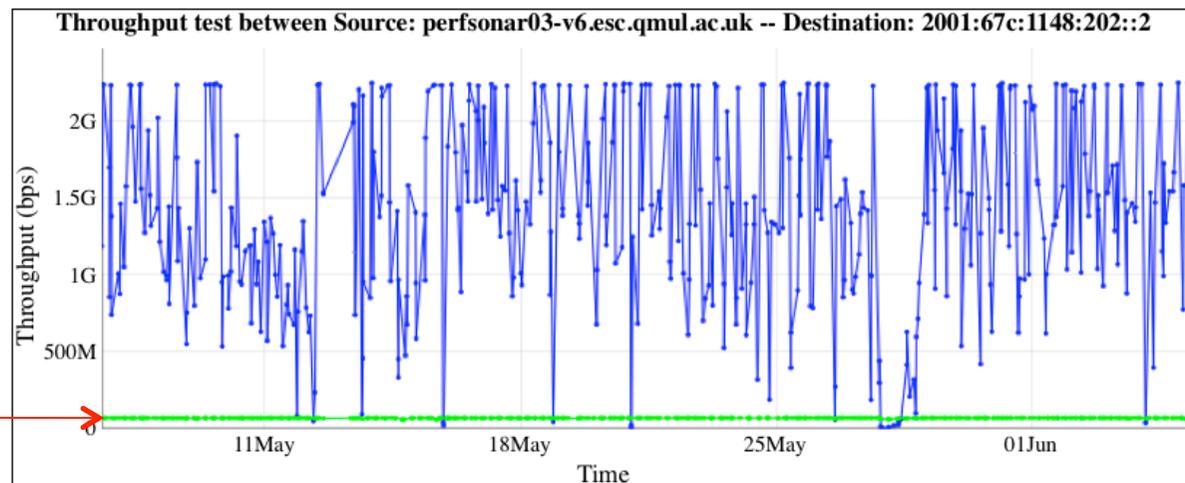
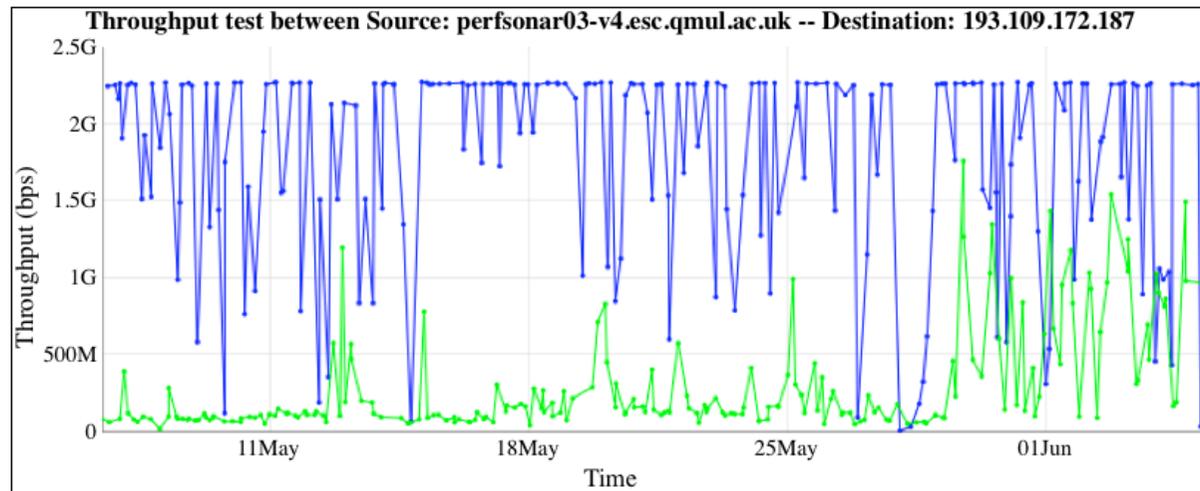
**Graph Key**

- Src-Dst throughput
- Dst-Src throughput

IPv4

PIC to QMUL  
slower over  
IPv6. Both  
using jumbo  
frames....

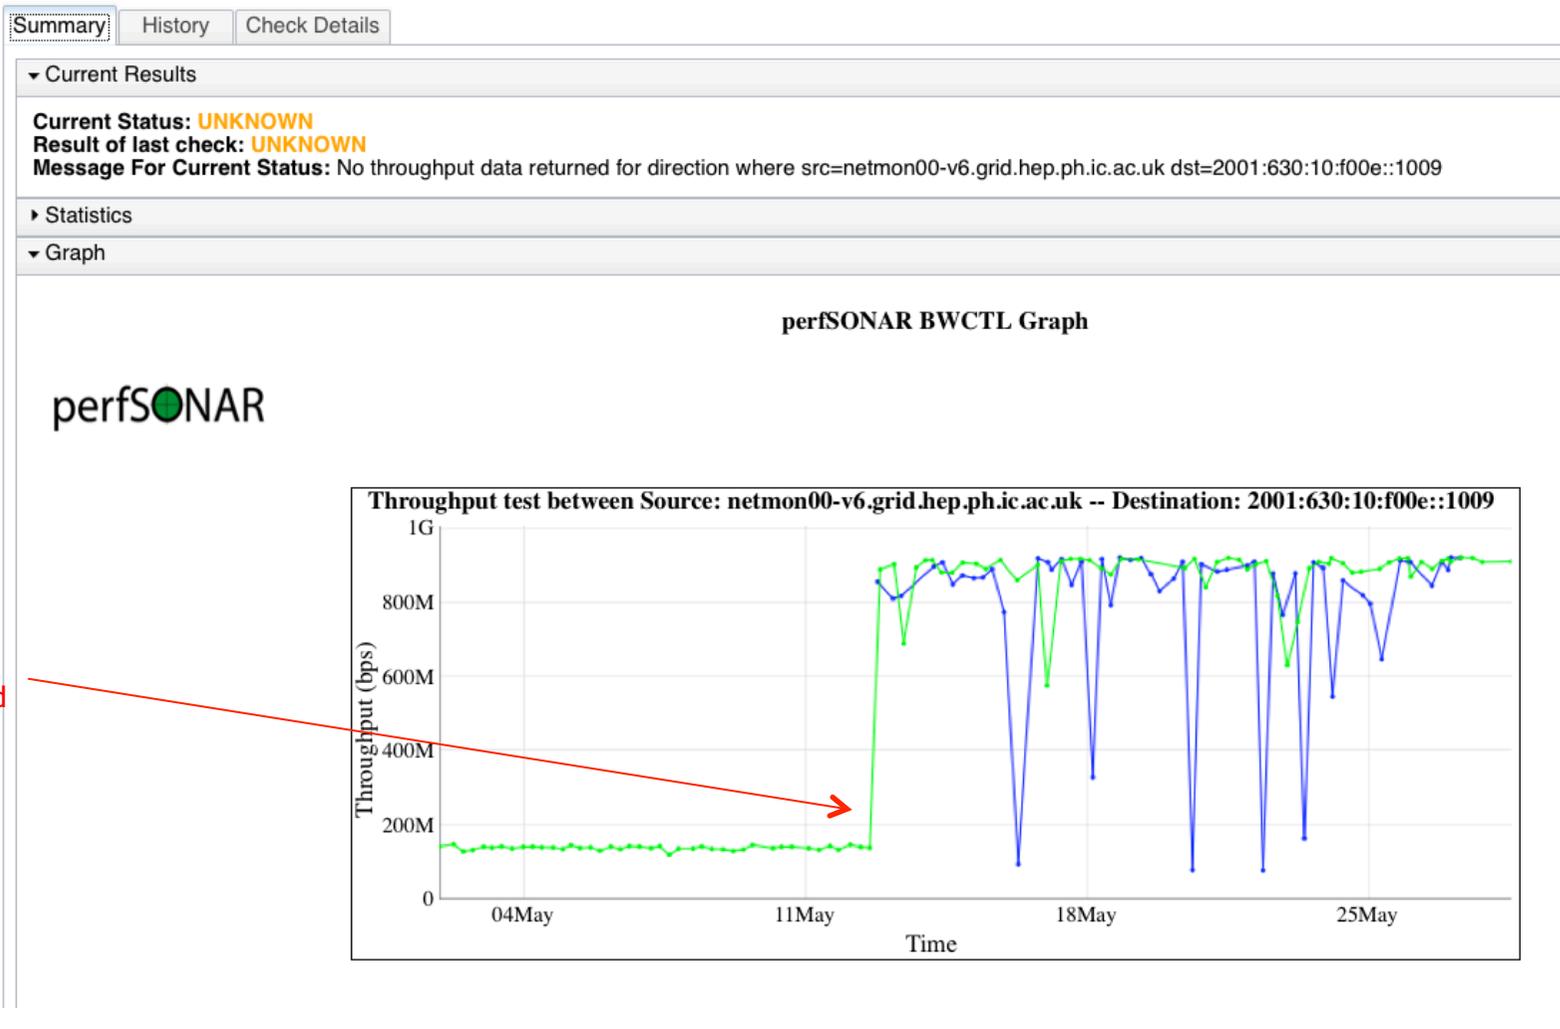
IPv6



0.07 Gbps

## netmon00-v6.grid.hep.ph.ic.ac.uk to 2001:630:10:f00e::1009 (Throughput)

Status: **UNKNOWN** Last Checked: May 28, 2014 13:51:07 PM BST Next Check: May 28, 2014 21:51:07 PM BST

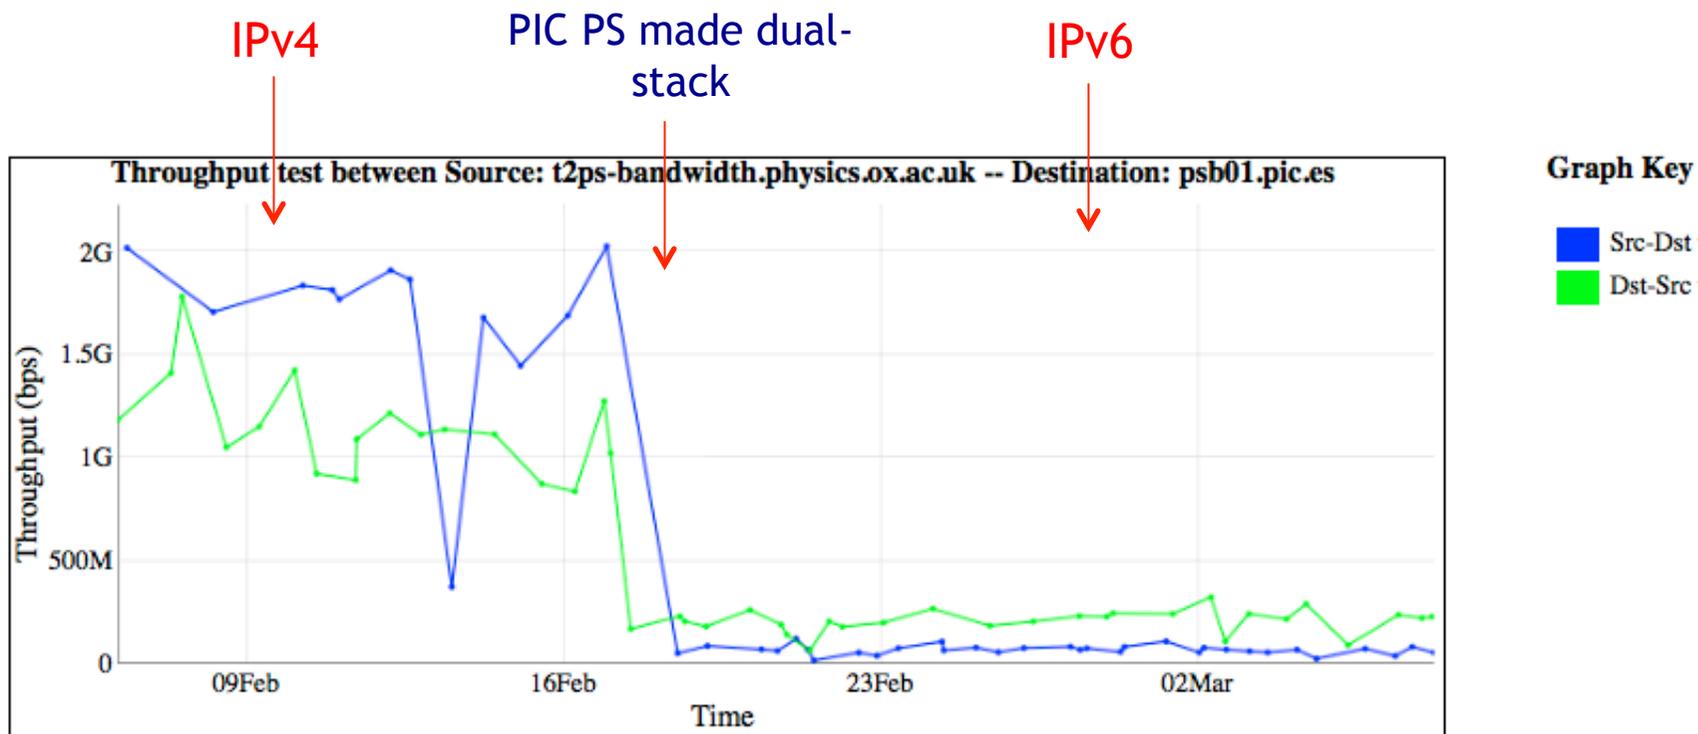


No IC to Brunel IPv6 results until ICMPv6 unblocked

ICMPv6 unblocked



- If the 'ipv4\_only' or 'ipv6\_only' flags are unset then dual-stack perfSONAR hosts favour IPv6 by default
- For example Oxford to PIC suffered a 'drop' in measured rate as PIC switched their perfSONAR host from IPv4-only to dual-stack and IPv6 rather than IPv4 rate is shown in the graph

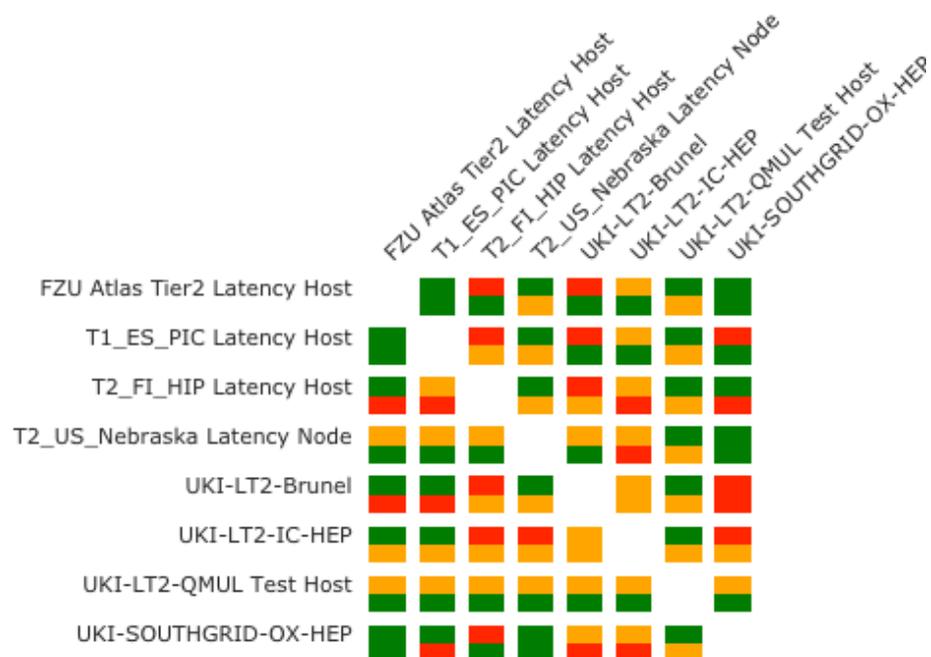


[<- 1 month](#)

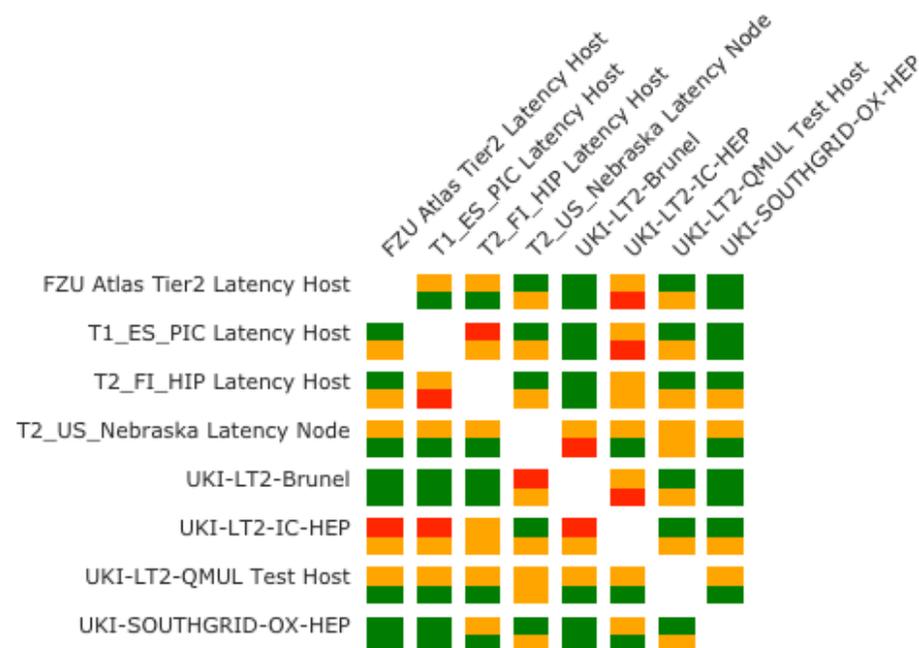
[1 month ->](#)

Timezone: GMT+0000 (GMT)

## IPv4



## IPv6



■ Loss rate is  $\leq 0$     
 ■ Loss rate is  $\geq 0$     
 ■ Loss rate is  $\geq 0.01$     
 ■ Unable to retrieve data    
 ■ Check has not yet run



## ps01-l-v6.ipv6.farm.particle.cz to t2ps-latency-v6.physics.ox.ac.uk (Loss)

Status: **OK** Last Checked: June 09, 2014 15:05:13 PM CEST Next Check: June 09, 2014 15:35:13 PM CEST

Summary History Check Details

Current Results

Current Status: **OK**  
Result of last check: **OK**  
Message For Current Status: Loss is 0.000%

Statistics

Graph

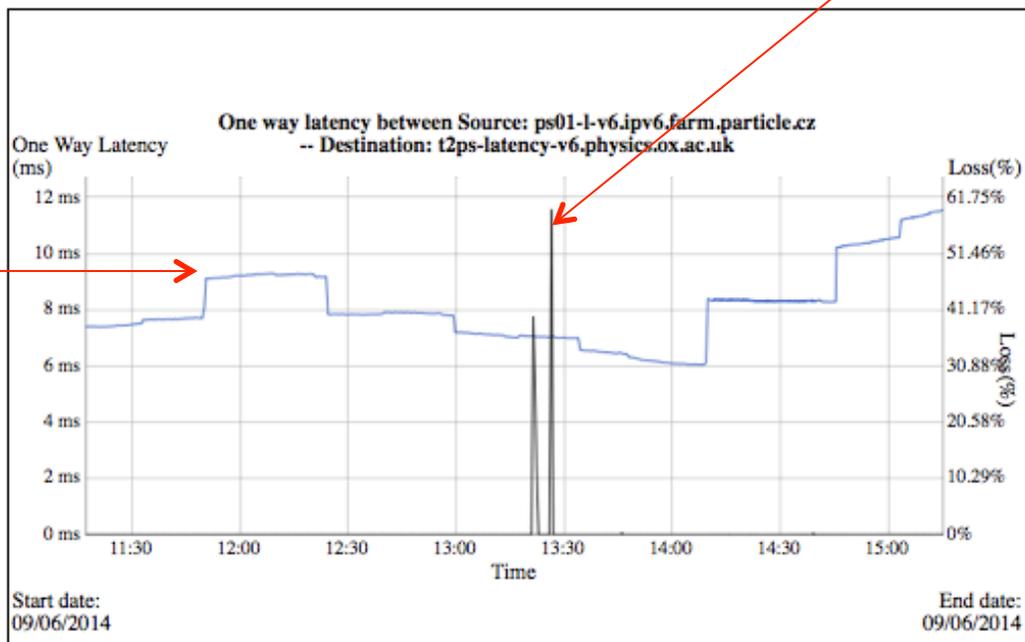
### perfSONAR One Way Latency



Scale Y axis from 0  Show Reverse Direction Data

Graph Key (Src-Dst)

- Max delay
- Min delay
- Loss



latency

loss

<- 4 hours

Timezone: GMT+0200 (CEST)

- Also collecting IPv4 and IPv6 traceroute data
- Can be useful for identifying asymmetric routing
- Routes between Imperial College and Nebraska are shown below

Select endpoints available on [http://localhost:8086/perfSONAR\\_PS/services/traceroute](http://localhost:8086/perfSONAR_PS/services/traceroute)

netmon02-v4.grid.hep.ph.ic.ac.uk (146.179.247.184) ----> hcc-ps01.unl.edu (129.93.239.148)

Do not de-duplicate results

Select endpoints available on [http://localhost:8086/perfSONAR\\_PS/services/tracerouteMA](http://localhost:8086/perfSONAR_PS/services/tracerouteMA)

netmon02-v6.grid.hep.ph.ic.ac.uk (2001:630:12:580:216:3eff:fe7f:153) ----> hcc-ps01.unl.edu (2600:900

Do not de-duplicate results

**Topology beginning at Tue May 27 00:24:00 2014 (UTC +1)**

Hop	Router	IP
1	hep-outside-hsrp1-slave.net.ic.ac.uk	146.179.246.253
2	te0-4-0-1.londsh-rbr1.ja.net	146.97.137.157
3	be2.londic-rbr1.ja.net	146.97.66.34
4	ae24.londpg-sbr1.ja.net	146.97.37.197
5	ae29.londhx-sbr1.ja.net	146.97.33.1
6	janet.mx1.lon.uk.geant.net	62.40.124.197
7	ae0.mx1.ams.nl.geant.net	62.40.98.81
8	ae0.mx1.fra.de.geant.net	62.40.98.128
9	abilene-wash-gw.mx1.fra.de.geant.net	62.40.125.18
10	et-9-0-0.115.rtr.chic.net.internet2.edu	198.71.45.56
11	ks-96-xe-11-3-3-100.greatplains.net	164.113.255.249
12	164.113.255.10	164.113.255.10
13	hcc-ps01.unl.edu	129.93.239.148

IPv4

**Topology beginning at Tue May 27 00:26:35 2014 (UTC +1)**

Hop	Router	IP
1	2001:630:12:580::2	2001:630:12:580::2
2	2001:630:0:9001::199	2001:630:0:9001::199
3	2001:630:0:1001:10::1	2001:630:0:1001:10::1
4	ae24.londtw-sbr1.ja.net	2001:630:0:10::23d
5	ae29.londtn-sbr1.ja.net	2001:630:0:10::1d2
6	ae0.lond-gw-ixp4.ja.net	2001:630:0:10::152
7	40ge1-3.core1.lon2.he.net	2001:7f8:4::1b1b:1
8	100ge1-1.core1.nyc4.he.net	2001:470:0:2cf::2
9	100ge7-2.core1.chi1.he.net	2001:470:0:298::1
10	wiscnet.v219.core1.chi1.he.net	2001:470:0:f2::2
11	2001:4e0:0:300c::2	2001:4e0:0:300c::2
12	2600:900:0:1::2	2600:900:0:1::2
13	hcc-ps01.unl.edu	2600:900:6:1101:7a2b:cbff:fe68:6238

IPv6

- Prototype perfSONAR dashboard allows differences between IPv4 and IPv6 network performance to be identified
- For sites wanting to get started with IPv6, perfSONAR hosts are a good place to start: let me know and I can add your hosts to the HEPiX IPv6 mesh