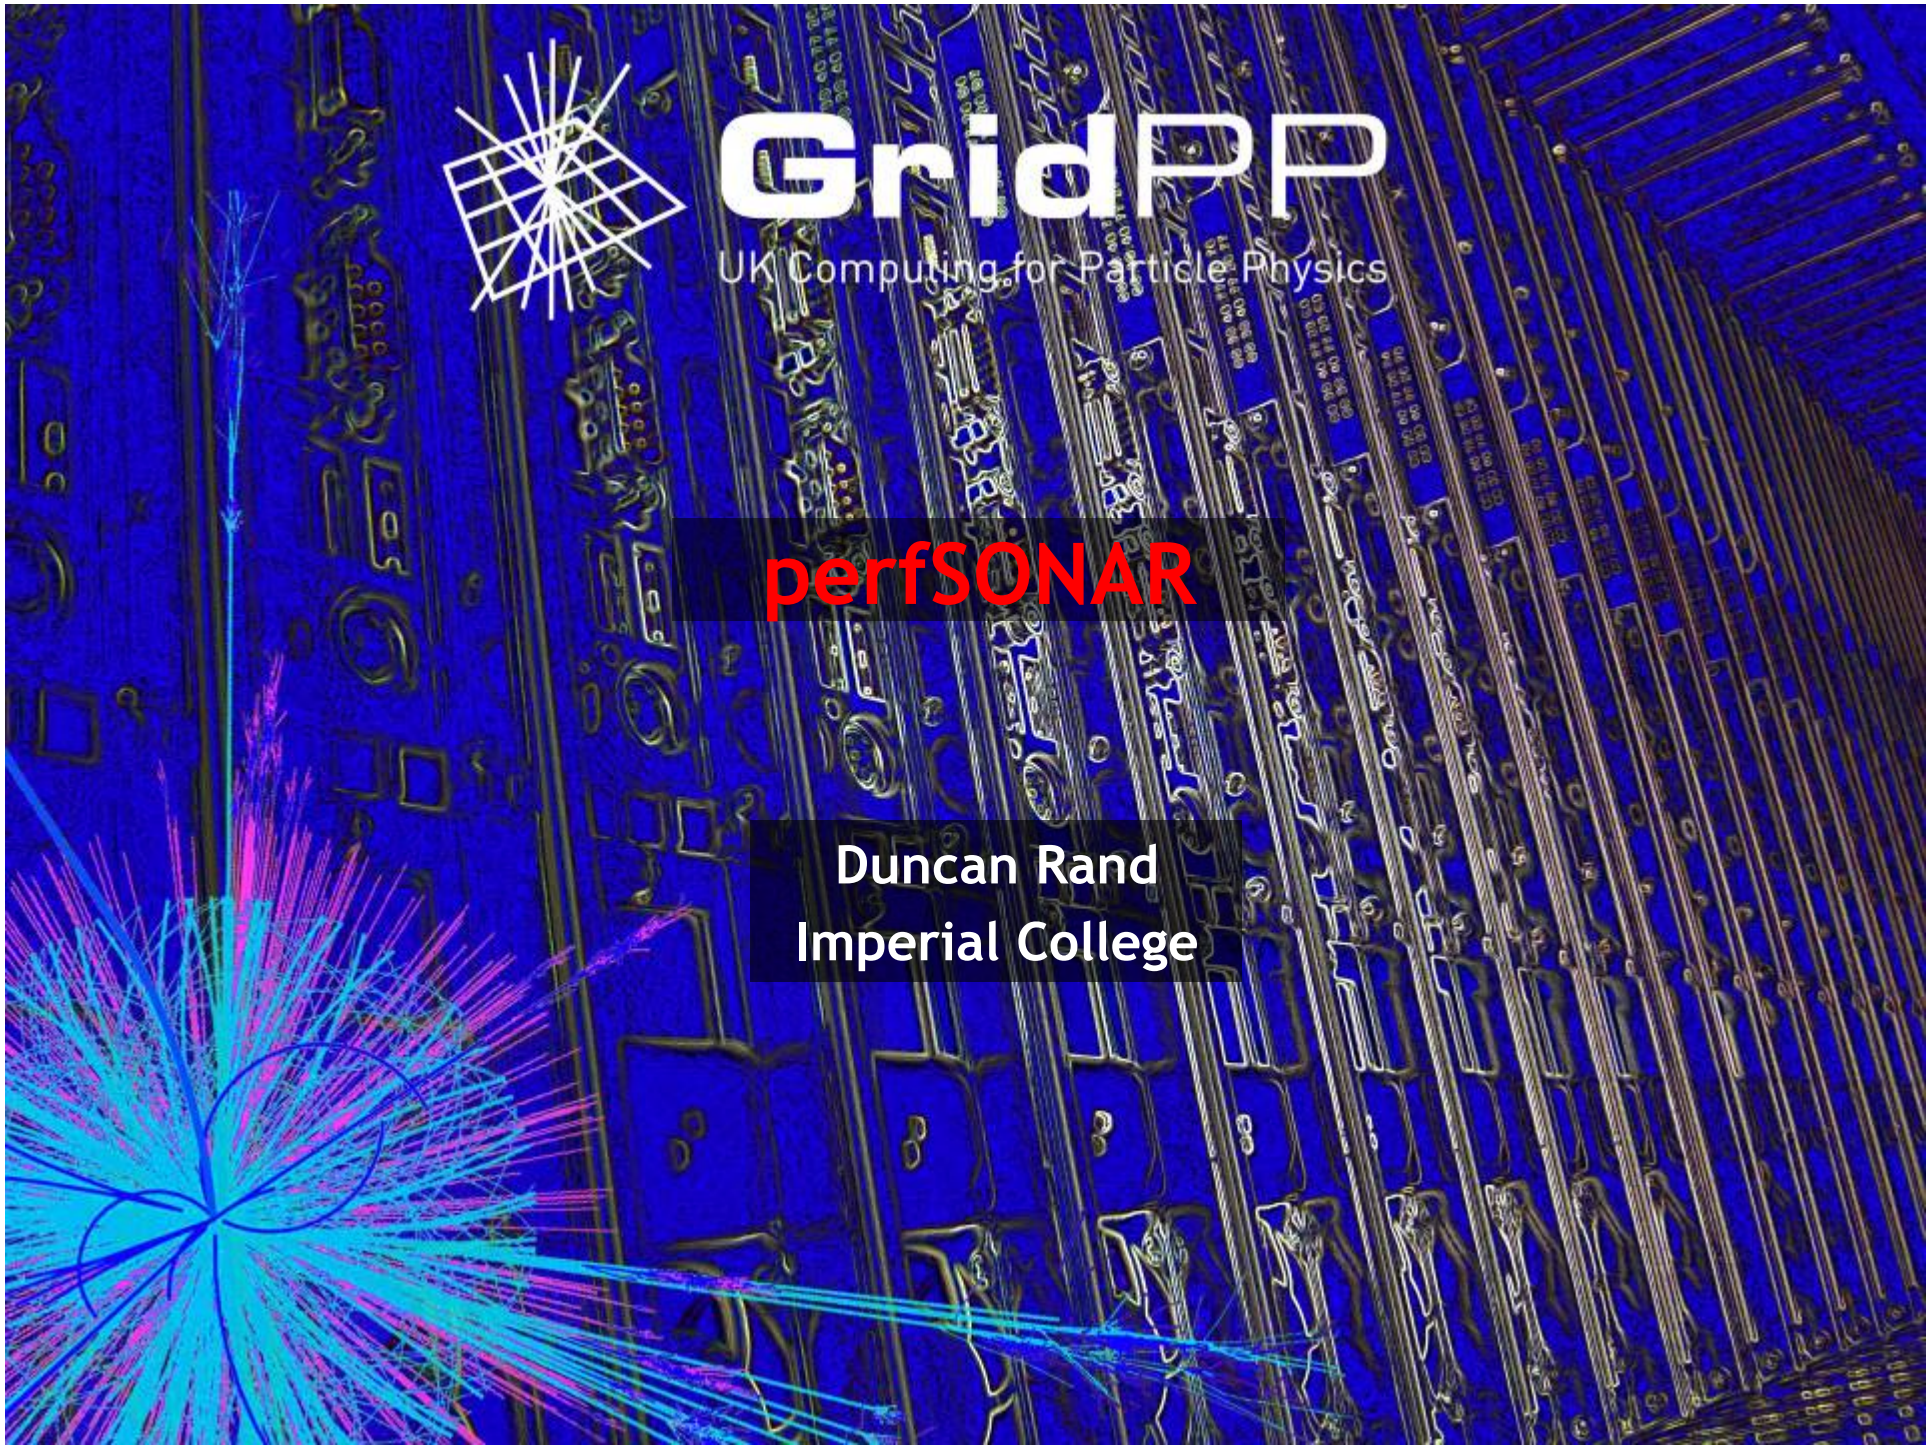
GridPP
UK Computing for Particle Physics
perfSONAR
Duncan Rand
Imperial College

Networking for WLCG

- LHCONE workshop: networking for WLCG
  - https://indico.cern.ch/event/289679/

- Presentations by LHC experiments on their potential future use of network
- General expectation of increased use of the network
- e.g. ATLAS: "…rely more on the network for just-in-time or real-time data distribution"

- Ian Bird: "Networking is key for the future evolution of WLCG"

- Resources need to be monitored
- perfSONAR provides the monitoring that helps sites and experiments make best use of the network

WLCG perfSONAR Task Force

- WLCG perfSONAR task force led by Simone Campana and Shaun McKee (UK members: Alessandra Forte and Duncan Rand)

- Goals:
  - Find and isolate “network” problems; alerting in a timely way
  - Characterize existing network usage
  - Provide a source of network metrics for higher level services

- The first step is to get the monitoring in place so as to be able to record the current situation

- Next step to monitor the network and alert sites when a problem occurs

- “perfSONAR’s purpose is to aid in network diagnosis by allowing users to characterize and isolate problems. It provides measurements of network performance metrics over time as well as “on-demand” tests”

WLCG deployment

- DRI finance gave UK a great opportunity with which to equip GridPP sites with perfSONAR hosts

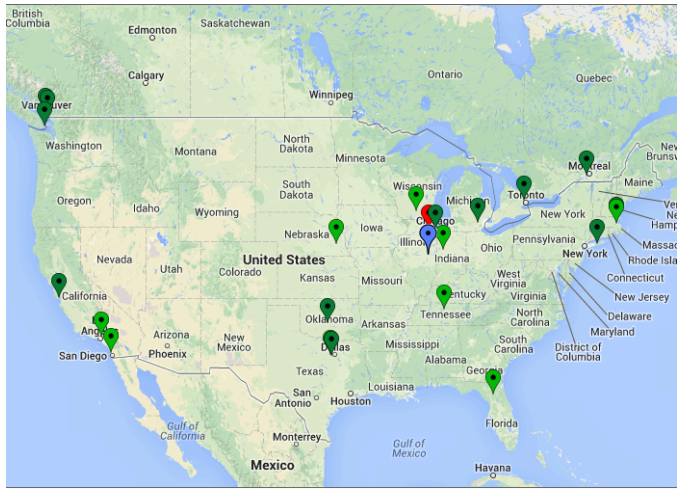

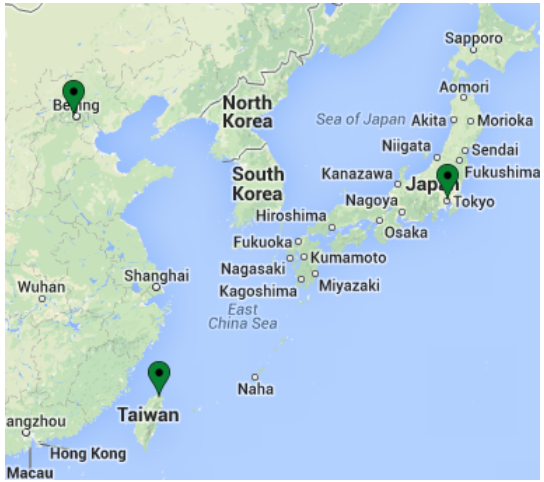

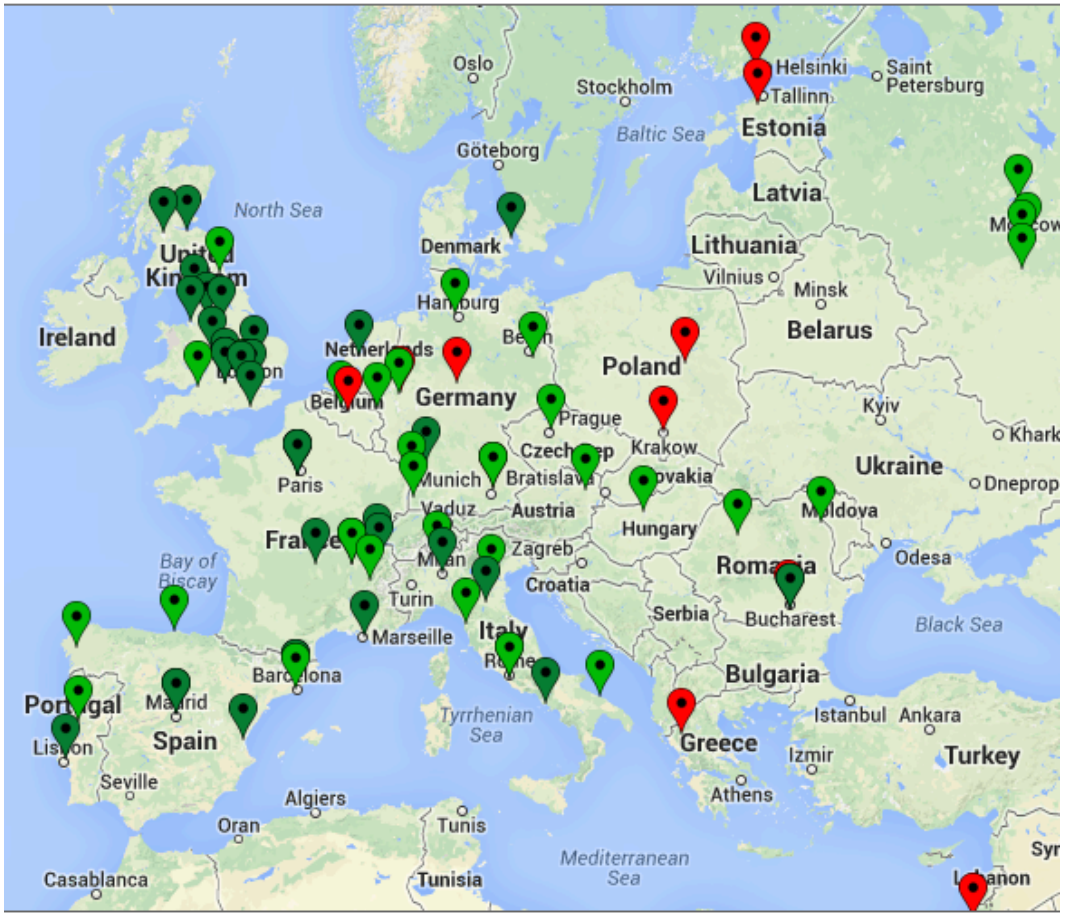

WLCG Deployment

- About 85% of WLCG sites have perfSONAR installed
- Issues include
  - Firewalls blocking services
  - Sites not using the mesh configuration
  - Versions are too old or not fully configured
  - Nodes are down/crashed
- Release 3.3.2 of perfSONAR released 3rd February 2014
  - Improvements in security, minor bug fixes, other improvements
  - this release is now the baseline
- April 1st 2014 is the deadline for all sites to have perfSONAR deployed, configured, registered and accessible from outside for monitoring

- Bandwidth tests
  - 30s tests: 6 hours intra-region, 12 hours T1-T2 inter-region, 1 week all WLCG sites
- Latency test: 10 Hz packets to each WLCG site
- Traceroute tests between all WLCG sites every hour
- Ping tests between all sites every 20 mins

Dashboards

- Centralised dashboards **really** help with visualisation of perfSONAR data
- Dashboard we previously used at Brookhaven no longer being developed
- Prototype dashboard using MaDDash (Monitoring and Debugging Dashboard) replacing it
  - QMUL deployed a temporary instance for the UK at end of 2013
- Simultaneously, Shawn McKee developed a WLCG-wide MaDDash dashboard
- This has two aspects
  - Open Monitoring Distribution: Nagios monitoring
  - MaDDash: similar to old BNL dashboard
- http://maddash.aglt2.org/maddash-webui/index.cgi?dashboard=UK%20sites

Open Monitoring Distribution (OMD)

- Several different displays available

https://maddash.aglt2.org/WLCGperfSONAR/omd/

OMD - Open Monitoring Distribution

Version: 1.10

This page gives you a central view on the available GUIs in OMD. Just have a look and feel free to choose your favorite GUI. At the bottom of this page you can find short instructions on how to change the default GUI of OMD.

Classic Nagios GUI

The classic Nagios GUI is based on CGI program written in C. It retrieves its status information from `status.dat`. This interface is not longer actively developed and does not perform well in large installations.

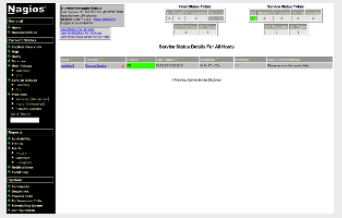

Check_MK Multisite

Check_MK Multisite is a fast and flexibile status GUI written in Python. It supports user definable views and is able to display the status of several sites in one combined view. It uses MK Livestatus for getting the status data from the sites.

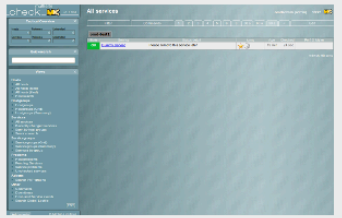

Thruk Monitoring Webinterface

Thruk is a complete rework of the classic interface in Perl. While maintainig the original look and feel it brings lots of improvements and new features. Just as Multisite it uses MK Livestatus as backend and supports the visualization of multiple sites.

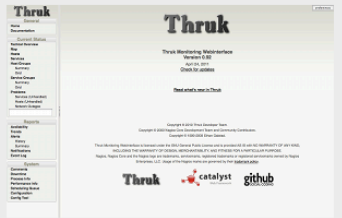

Classic Icinga GUI

Icinga's "classical" GUI is a derivate of the classical Nagios GUI and has been directly evolved from the original CGI programs in C. It has its own look and feel and brings useful improvements. It is not bound to Icinga and can be used with the other monitoring cores as well.

Classic Nagios

UK perfSONAR-PS Toolkit nodes (UK)

Host	Status	Services	Actions
dc2-grid-ps-00.brunel.ac.uk	UP	9 OK	
epgperf.ph.bham.ac.uk	UP	9 OK	
grid-perfsonar.hpc.susx.ac.uk	UP	9 OK	
gridpp-ps-band.ecdf.ed.ac.uk	UP	4 OK 5 CRITICAL	
gridpp-ps-lat.ecdf.ed.ac.uk	UP	3 OK 6 CRITICAL	
heplnx129.pp.rl.ac.uk	DOWN	9 CRITICAL	
heplnx130.pp.rl.ac.uk	DOWN	9 CRITICAL	
hepsonar1.ph.liv.ac.uk	UP	9 OK	
hepsonar2.ph.liv.ac.uk	UP	9 OK	
lcg-sonar01.hep.ucl.ac.uk	UP	9 CRITICAL	
lcgnetmon.phy.bris.ac.uk	UP	9 OK	
lcgnetmon02.phy.bris.ac.uk	UP	9 OK	
lcgperf.shef.ac.uk	UP	5 OK 4 CRITICAL	
netmon00.grid.hep.ph.ic.ac.uk	UP	2 OK 7 CRITICAL	
netmon02.grid.hep.ph.ic.ac.uk	UP	3 OK 6 CRITICAL	
perfmon.dur.scotgrid.ac.uk	UP	9 OK	
perfson1.ppgrid1.rhul.ac.uk	UP	9 OK	
perfsonar-bandwidth.esc.qmul.ac.uk	UP	9 OK	
perfsonar-bw.tier2.hep.manchester.ac.uk	UP	9 OK	
perfsonar-latency.esc.qmul.ac.uk	UP	9 OK	
perfsonar-lt.tier2.hep.manchester.ac.uk	UP	9 OK	
ps001.gla.scotgrid.ac.uk	UP	9 OK	
ps002.gla.scotgrid.ac.uk	UP	9 OK	
pygrid-sonar1.lancs.ac.uk	UP	9 OK	
pygrid-sonar2.lancs.ac.uk	UP	9 OK	
serv04.hep.phy.cam.ac.uk	UP	9 OK	
t2ps-bandwidth.physics.ox.ac.uk	UP	9 OK	
t2ps-latency.physics.ox.ac.uk	UP	9 OK	

Check_MK: WLCG sites

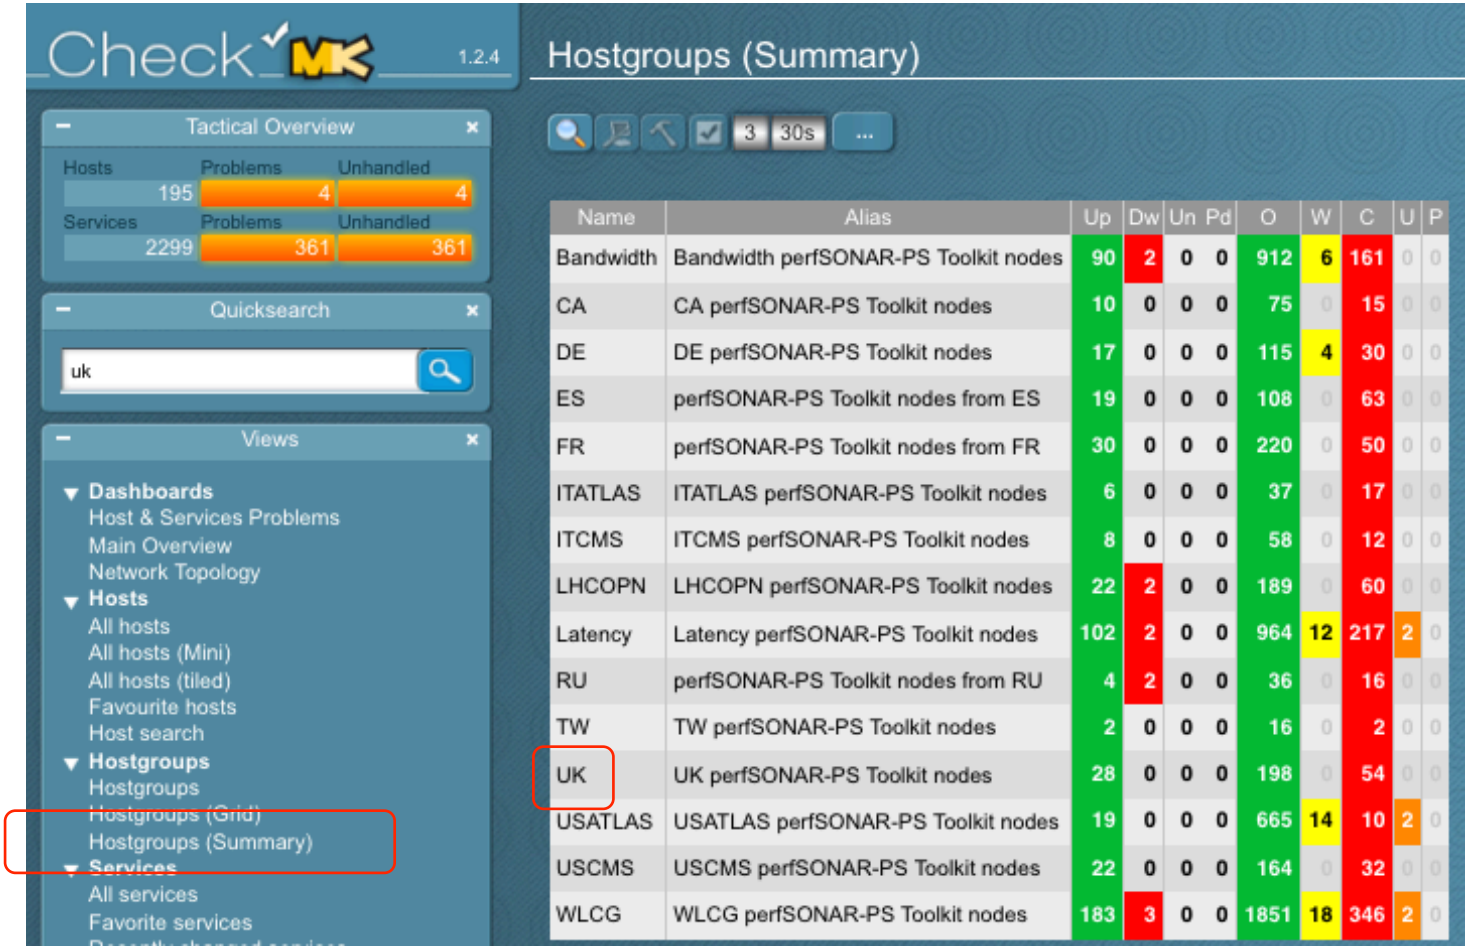

Name	Alias	Up	Dw	Un	Pd	O	W	C	U	P
Bandwidth	Bandwidth perfSONAR-PS Toolkit nodes	90	2	0	0	912	6	161	0	0
CA	CA perfSONAR-PS Toolkit nodes	10	0	0	0	75	0	15	0	0
DE	DE perfSONAR-PS Toolkit nodes	17	0	0	0	115	4	30	0	0
ES	perfSONAR-PS Toolkit nodes from ES	19	0	0	0	108	0	63	0	0
FR	perfSONAR-PS Toolkit nodes from FR	30	0	0	0	220	0	50	0	0
ITATLAS	ITATLAS perfSONAR-PS Toolkit nodes	6	0	0	0	37	0	17	0	0
ITCMS	ITCMS perfSONAR-PS Toolkit nodes	8	0	0	0	58	0	12	0	0
LHCOPN	LHCOPN perfSONAR-PS Toolkit nodes	22	2	0	0	189	0	60	0	0
Latency	Latency perfSONAR-PS Toolkit nodes	102	2	0	0	964	12	217	2	0
RU	perfSONAR-PS Toolkit nodes from RU	4	2	0	0	36	0	16	0	0
TW	TW perfSONAR-PS Toolkit nodes	2	0	0	0	16	0	2	0	0
UK	UK perfSONAR-PS Toolkit nodes	28	0	0	0	198	0	54	0	0
USATLAS	USATLAS perfSONAR-PS Toolkit nodes	19	0	0	0	665	14	10	2	0
USCMS	USCMS perfSONAR-PS Toolkit nodes	22	0	0	0	164	0	32	0	0
WLCG	WLCG perfSONAR-PS Toolkit nodes	183	3	0	0	1851	18	346	2	0

UK sites

Hostgroup UK perfSONAR-PS Toolkit nodes

2 30s Availability ...

state	Host	Icons	Alias	OK	Wa	Un	Cr	Pd
UP	dc2-grid-ps-00.brunel.ac.uk		dc2-grid-ps-00.brunel.ac.uk	9	0	0	0	0
UP	epgperf.ph.bham.ac.uk		epgperf.ph.bham.ac.uk	9	0	0	0	0
UP	grid-perfsonar.hpc.susx.ac.uk		grid-perfsonar.hpc.susx.ac.uk	9	0	0	0	0
UP	gridpp-ps-band.ecdf.ed.ac.uk		gridpp-ps-band.ecdf.ed.ac.uk	4	0	0	5	0
UP	gridpp-ps-lat.ecdf.ed.ac.uk		gridpp-ps-lat.ecdf.ed.ac.uk	3	0	0	6	0
UP	heplnx129.pp.rl.ac.uk		heplnx129.pp.rl.ac.uk	0	0	0	9	0
UP	heplnx130.pp.rl.ac.uk		heplnx130.pp.rl.ac.uk	1	0	0	8	0
UP	hepsonar1.ph.liv.ac.uk		hepsonar1.ph.liv.ac.uk	9	0	0	0	0
UP	hepsonar2.ph.liv.ac.uk		hepsonar2.ph.liv.ac.uk	9	0	0	0	0
UP	lcg-sonar01.hep.ucl.ac.uk		lcg-sonar01.hep.ucl.ac.uk	0	0	0	9	0
UP	lcgnetmon.phy.bris.ac.uk		lcgnetmon.phy.bris.ac.uk	9	0	0	0	0
UP	lcgnetmon02.phy.bris.ac.uk		lcgnetmon02.phy.bris.ac.uk	9	0	0	0	0
UP	lcgperf.shef.ac.uk		lcgperf.shef.ac.uk	5	0	0	4	0
UP	netmon00.grid.hep.ph.ic.ac.uk		netmon00.grid.hep.ph.ic.ac.uk	2	0	0	7	0
UP	netmon02.grid.hep.ph.ic.ac.uk		netmon02.grid.hep.ph.ic.ac.uk	3	0	0	6	0
UP	perfmon.dur.scotgrid.ac.uk		perfmon.dur.scotgrid.ac.uk	9	0	0	0	0
UP	perfson1.ppgrid1.rhul.ac.uk		perfson1.ppgrid1.rhul.ac.uk	9	0	0	0	0
UP	perfsonar-bandwidth.esc.qmul.ac.uk		perfsonar-bandwidth.esc.qmul.ac.uk	9	0	0	0	0
UP	perfsonar-bw.tier2.hep.manchester.ac.uk		perfsonar-bw.tier2.hep.manchester.ac.uk	9	0	0	0	0
UP	perfsonar-latency.esc.qmul.ac.uk		perfsonar-latency.esc.qmul.ac.uk	9	0	0	0	0
UP	perfsonar-lt.tier2.hep.manchester.ac.uk		perfsonar-lt.tier2.hep.manchester.ac.uk	9	0	0	0	0
UP	ps001.gla.scotgrid.ac.uk		ps001.gla.scotgrid.ac.uk	9	0	0	0	0
UP	ps002.gla.scotgrid.ac.uk		ps002.gla.scotgrid.ac.uk	9	0	0	0	0
UP	pygrid-sonar1.lancs.ac.uk		pygrid-sonar1.lancs.ac.uk	9	0	0	0	0
UP	pygrid-sonar2.lancs.ac.uk		pygrid-sonar2.lancs.ac.uk	9	0	0	0	0
UP	serv04.hep.phy.cam.ac.uk		serv04.hep.phy.cam.ac.uk	9	0	0	0	0
UP	t2ps-bandwidth.physics.ox.ac.uk		t2ps-bandwidth.physics.ox.ac.uk	9	0	0	0	0
UP	t2ps-latency.physics.ox.ac.uk		t2ps-latency.physics.ox.ac.uk	9	0	0	0	0

Individual sites

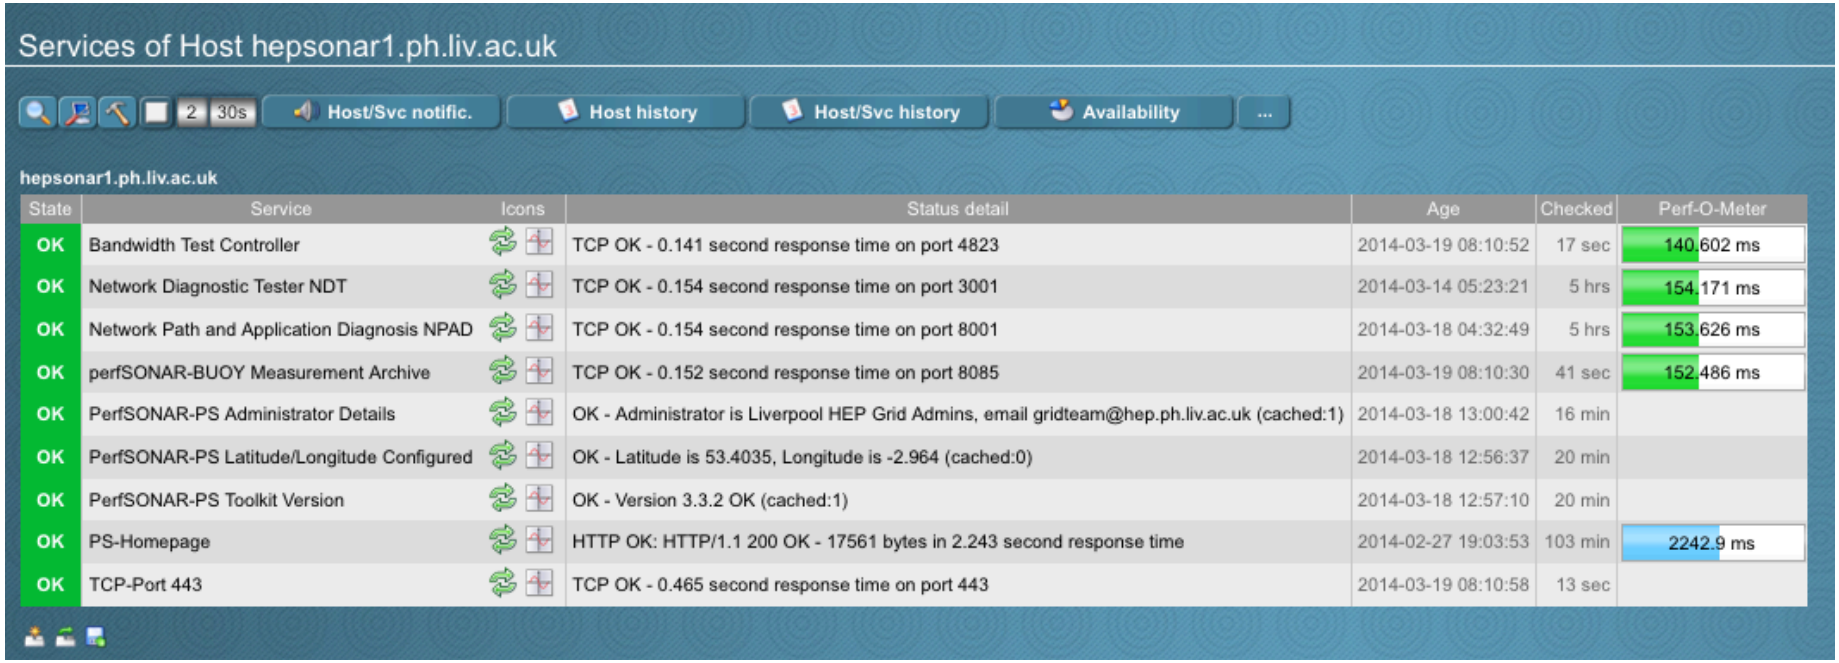

Services of Host hepsonar1.ph.liv.ac.uk

2 30s | Host/Svc notific. | Host history | Host/Svc history | Availability | ...

hepsonar1.ph.liv.ac.uk

State	Service	Icons	Status detail	Age	Checked	Perf-O-Meter
OK	Bandwidth Test Controller		TCP OK - 0.141 second response time on port 4823	2014-03-19 08:10:52	17 sec	140.602 ms
OK	Network Diagnostic Tester NDT		TCP OK - 0.154 second response time on port 3001	2014-03-14 05:23:21	5 hrs	154.171 ms
OK	Network Path and Application Diagnosis NPAD		TCP OK - 0.154 second response time on port 8001	2014-03-18 04:32:49	5 hrs	153.626 ms
OK	perfSONAR-BUOY Measurement Archive		TCP OK - 0.152 second response time on port 8085	2014-03-19 08:10:30	41 sec	152.486 ms
OK	PerfSONAR-PS Administrator Details		OK - Administrator is Liverpool HEP Grid Admins, email gridteam@hep.ph.liv.ac.uk (cached:1)	2014-03-18 13:00:42	16 min	
OK	PerfSONAR-PS Latitude/Longitude Configured		OK - Latitude is 53.4035, Longitude is -2.964 (cached:0)	2014-03-18 12:56:37	20 min	
OK	PerfSONAR-PS Toolkit Version		OK - Version 3.3.2 OK (cached:1)	2014-03-18 12:57:10	20 min	
OK	PS-Homepage		HTTP OK: HTTP/1.1 200 OK - 17561 bytes in 2.243 second response time	2014-02-27 19:03:53	103 min	2242.9 ms
OK	TCP-Port 443		TCP OK - 0.465 second response time on port 443	2014-03-19 08:10:58	13 sec	

Imperial College dashboard

- We have also installed our own UK instance which we can adapt (e.g. add IPv6 hosts)

List of dashboards

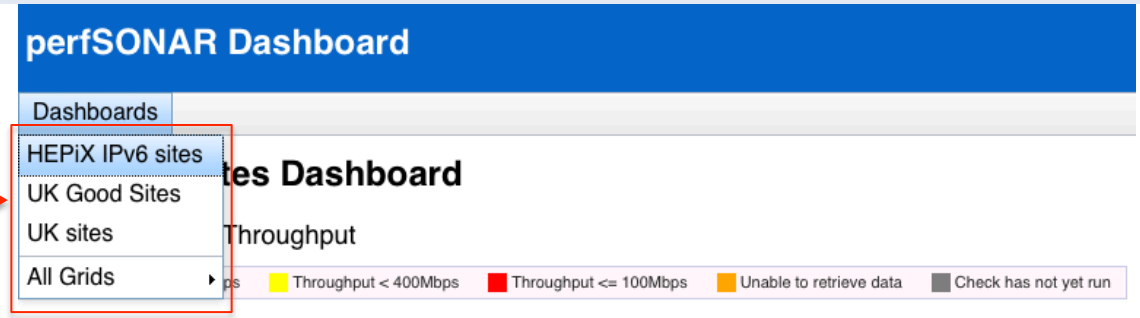

Aside: note red boxes for Durham <100Mbps

No measurement

http://netmon02.grid.hep.ph.ic.ac.uk:8080/maddash-webui

Rate thresholds. Hard to define sensible values for sites of largely varied connectivity (i.e. >10 Gbps or not). Might split into two sections.

destination

Hover to get rates from both measurement archives of sites involved

Click to drill down (next slide)

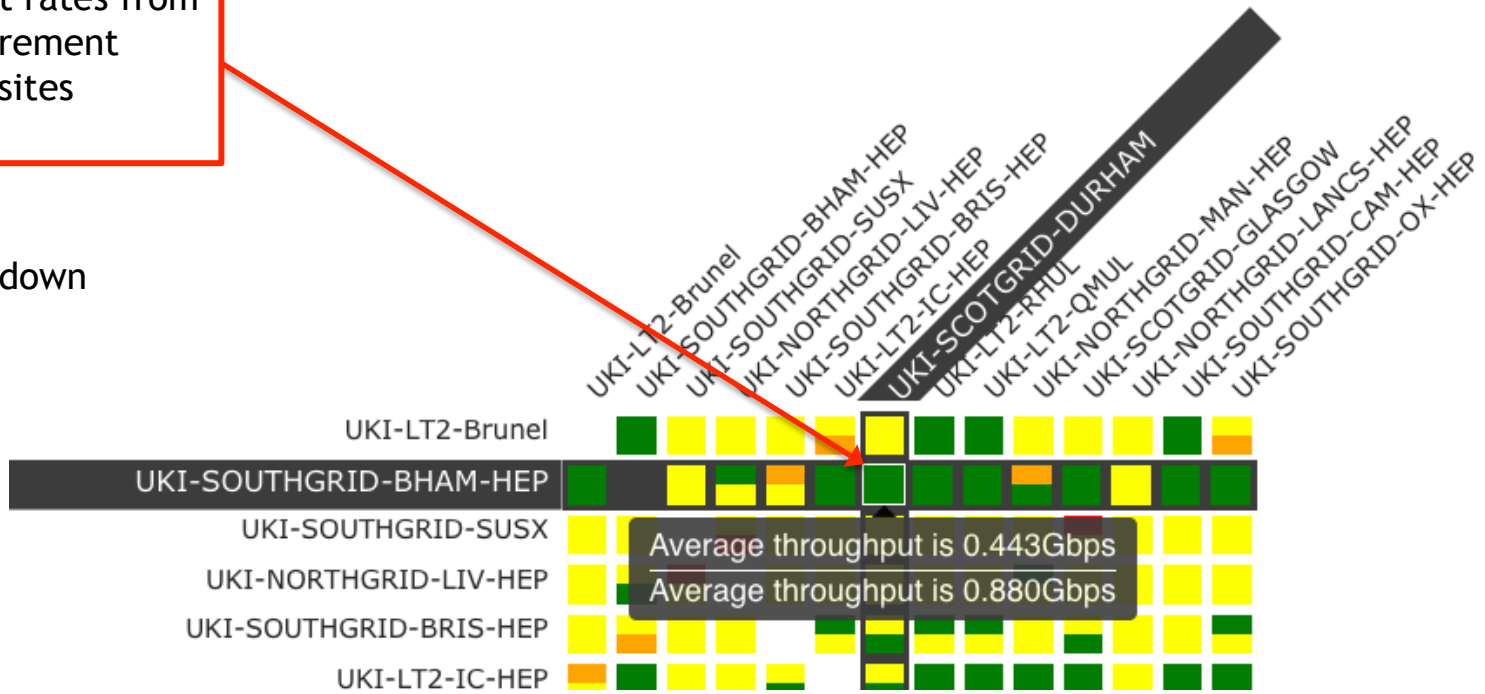

source

Drill down to time series graphs

epgperf.ph.bham.ac.uk to perfmon.dur.scotgrid.ac.uk (Throughput)

Status: OK **Last Checked:** March 19, 2014 05:22:52 AM GMT **Next Check:** March 19, 2014 13:22:52 PM GMT

Summary | History | Check Details

▾ Current Results

Current Status: OK
Result of last check: OK
Message For Current Status: Average throughput is 0.443Gbps

▸ Statistics

▾ Graph

perfSONAR BWCTL Graph

perfSONAR

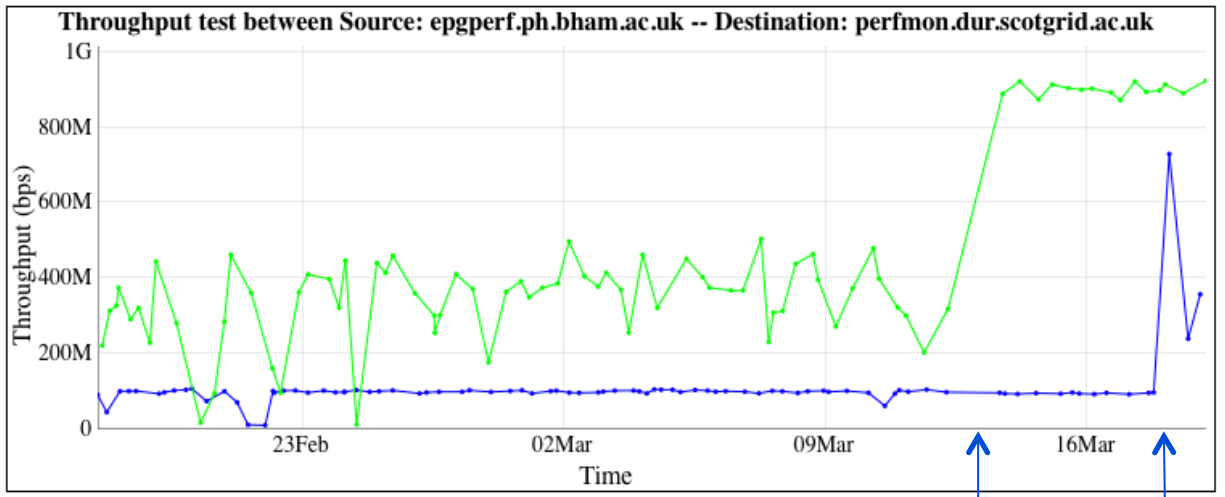

Moved physical location of Durham host

Replaced motherboard of Durham host

UK Deployment

Orange pollution

All UK sites

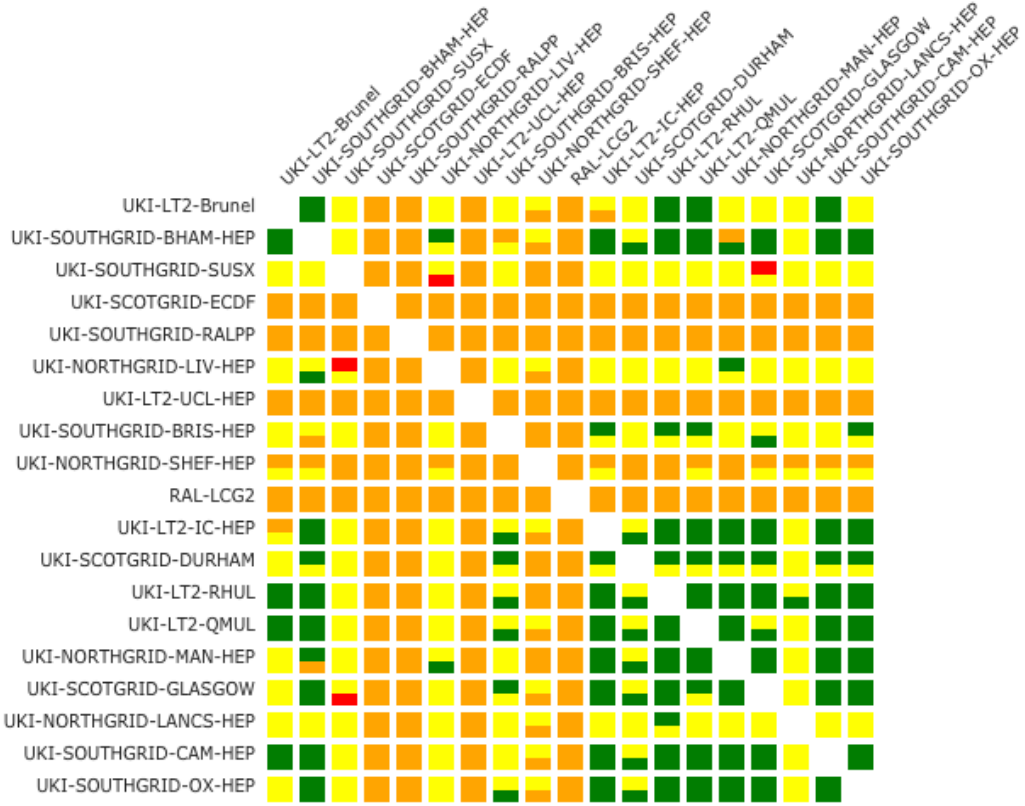

'Good' UK sites

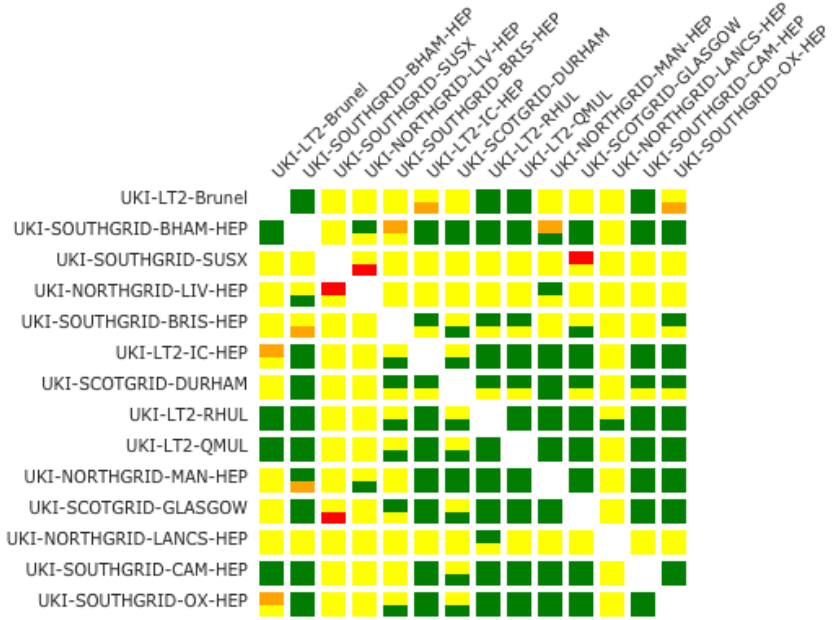

UK Deployment

Orange pollution

All UK sites

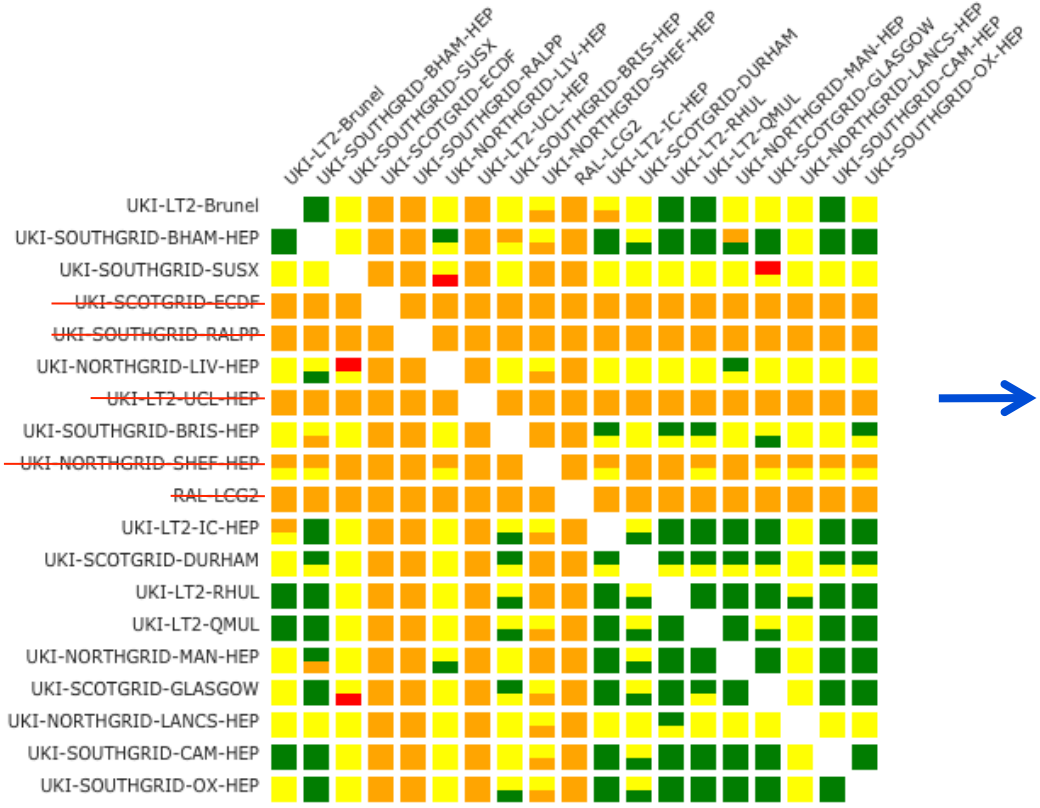

‘Good’ UK sites

5/19 sites have problems: ECDF, RALPP, UCL, SHEF, RAL

The Lancaster conundrum

GridPP
UK Computing for Particle Physics

UKI-NORTHGRID-LANCS-HEP

Throughput >= 400Mbps Throughput < 400Mbps

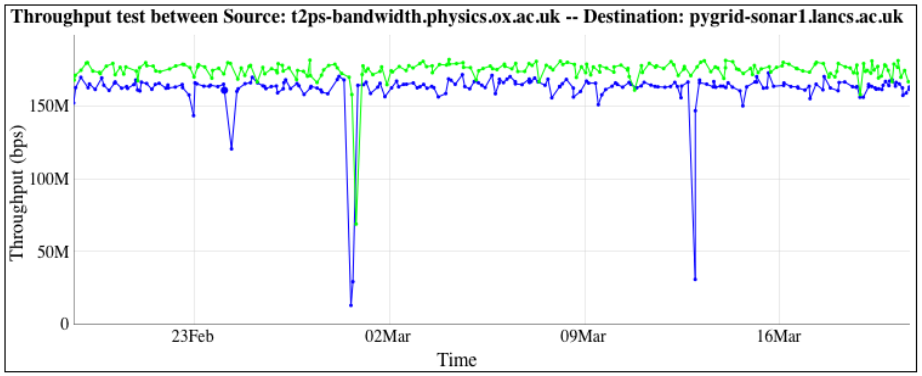

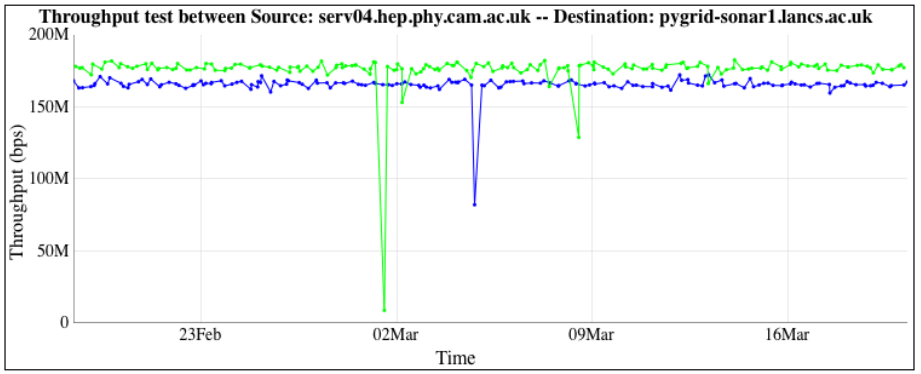

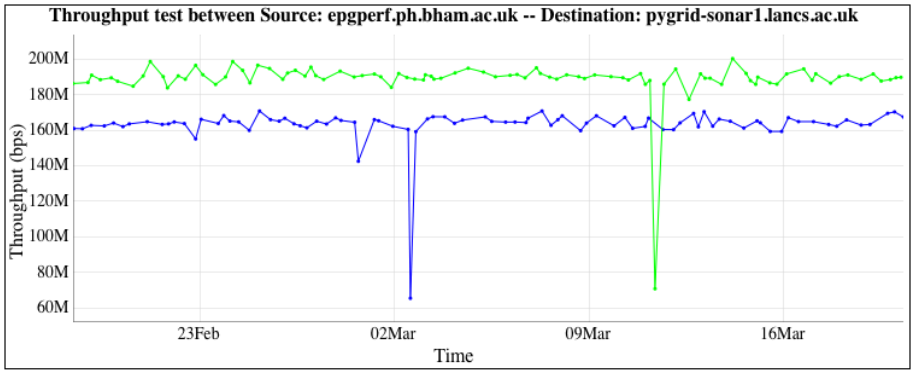

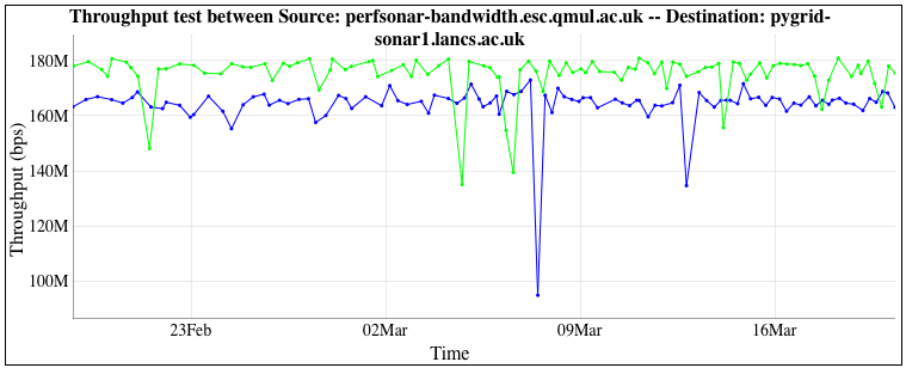

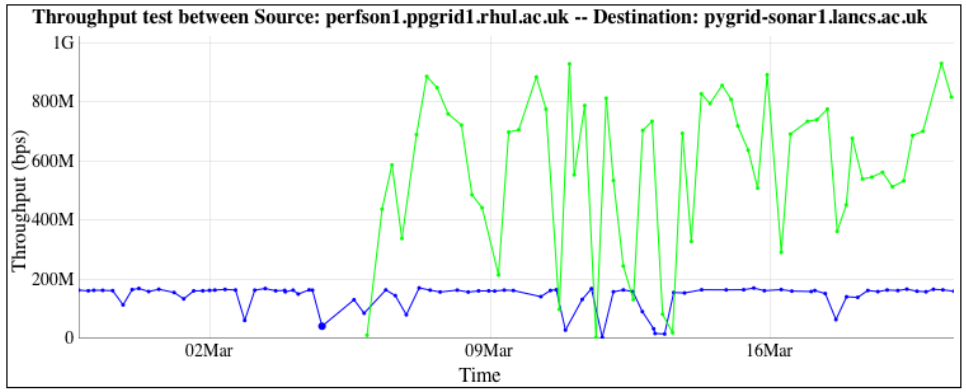

Graph Key

Src-Dst throughput

Dst-Src throughput

- Ian Bird's talk

perfSONAR deployment

- WLCG agreed on perfSONAR as the core toolkit for network monitoring in the infrastructure
  - Strong push came from experiments
- Deployment of perfSONAR has been (and still is) sometimes problematic
  - Some sites refuse to install it at all
  - Some sites still run very old versions
- perfSONAR needs to be treated as any other service in WLCG
  - Including the level of commitment in installing, configuring, operating it.

Improving deployment

- UK marching closely in step with the WLCG
- Best practice wiki
  - http://www.usatlas.bnl.gov/twiki/bin/view/Projects/LHCperfSONAR
- Aim to improve resiliency and use OMD to alert when services fail
- Disentangle problems with perfSONAR and problems with the network (eg Durham and maybe Lancaster?)
- Development roadmap
  - https://code.google.com/p/perfsonar-ps/wiki/RoadMap

QMUL helping test perfSONAR releases

WLCG perfSONAR Release Test Mesh Dashboard

WLCG perfSONAR Release Test Mesh - TCP BWCTL Test Between USATLAS Bandwidth Hosts

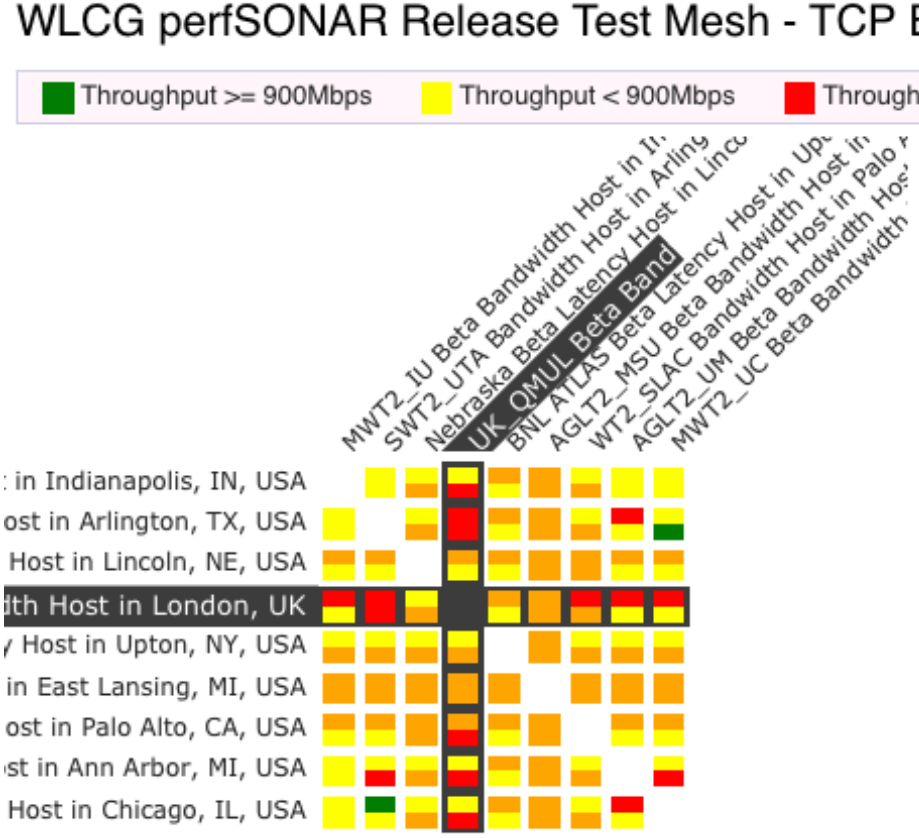

QMUL's test host has joined US beta testing group

Tier-1s

- Many Tier-1s not engaged it seems - perhaps only testing LHCOPN sites

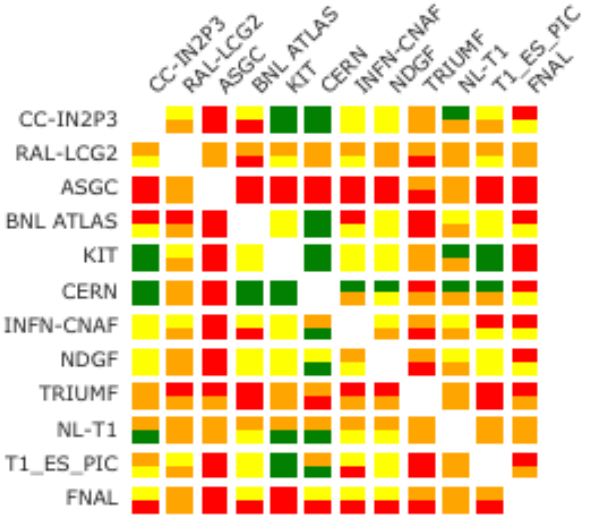

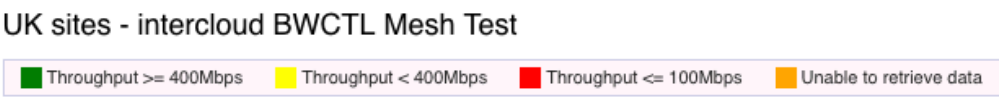

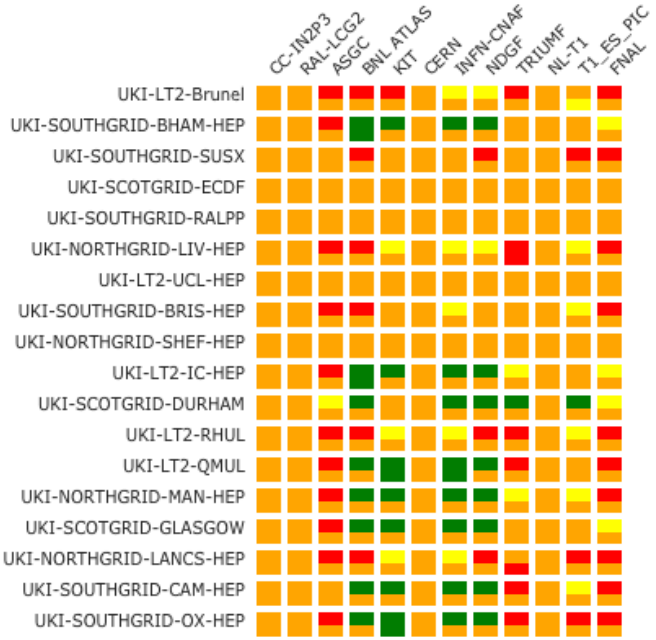

WLCG sites dashboard

- Each site tests each other site once a week
- A lot of orange
- Still some work to do…
- Sites currently sorted by hostname rather than site name which would group all UK sites together

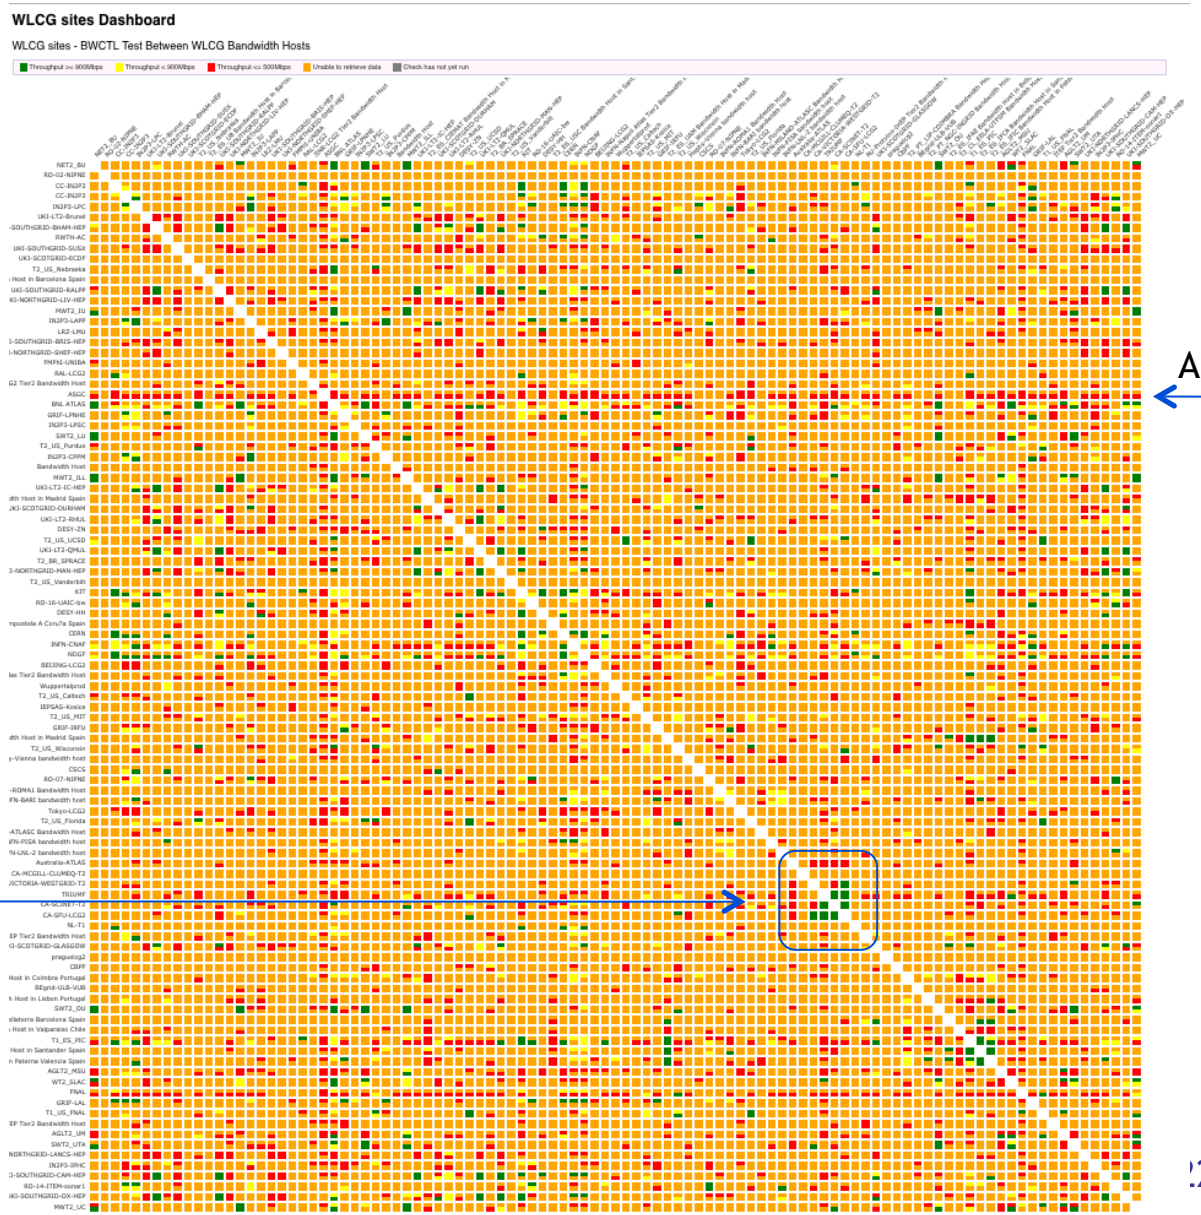

Future Use of Network Metrics

WLCG
Worldwide LHC Computing Grid

- Once we have a source of network metrics being acquired we need to understand how best to incorporate those metrics into our facility operations.
- Some possibilities:
  - Characterizing paths with "costs" to better optimize decisions in workflow and data management (underway in ANSE)
  - Noting when paths change and providing appropriate notification
  - Optimizing data-access or data-distribution based upon a better understanding of the network between sites
  - Identifying structural bottlenecks in need of remediation
  - Aiding network problem diagnosis and speeding repairs
  - In general, incorporating knowledge of the network into our processes
- **We will require testing and iteration to better understand when and where the network metrics are useful.**

Acknowledgement: Shawn McKee's GDB talk
http://tinyurl.com/p8l4xm

Conclusion

- The UK was fortunate to be able to equip our sites with perfSONAR hosts
- The new MaDDash dashboards have made the perfSONAR instances much more useable
- The new OMD nagios testing should make the perfSONAR instances more maintainable
- We should play our part within WLCG by running the service well
- We need to build up experience of running perfSONAR as a service
- The LHC VOs are planning to use perfSONAR…