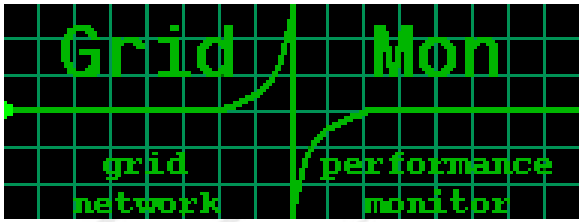

Network Performance Monitoring for the Grid

HEPSYSMAN Meeting
RAL, 9-10th May 2006

Mark Leese

http://gridmon.dl.ac.uk

Mark Leese
Daresbury Laboratory

CCLRC Daresbury Laboratory

Contents

- Why monitor network performance? If it was that important you would have sent yourself not an .mp3, right? ;-)
  - traditional view
  - the Grid as a special case
- Gridmon
  - history
  - current form
  - firewall requirements
- (Grid) Network Performance Monitoring – the wider picture
  - GGF NM-WG
  - EGEE JRA4

There's so much happening that this can only be a taster!

For more info, please see the "more info" slide,
and feel free to get in touch: m.j.leese@dl.ac.uk

Mark Leese
Daresbury Laboratory

A message from our sponsors…

Money by…

Extra intellect by…

Mark Leese
Daresbury Laboratory

Justification (1)

Network performance monitoring has **traditionally** been important to the operation of networks of any significant size:

- fault detection
- determining expected performance, e.g.

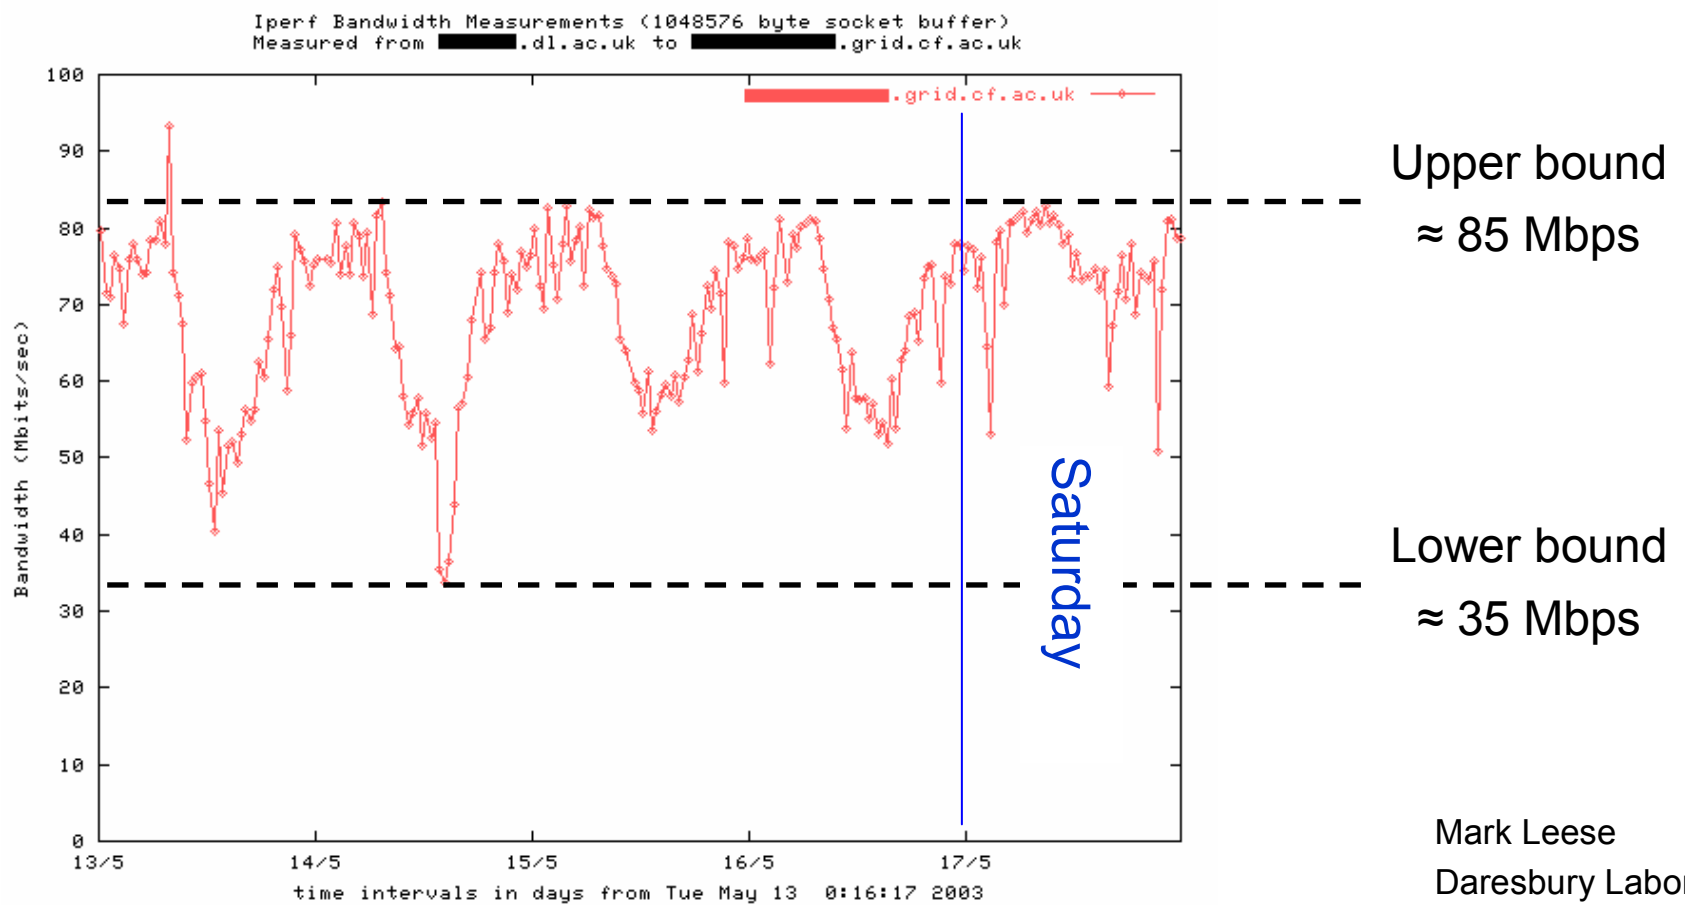

Mark Leese
Daresbury Laboratory

Justification (2)

Fault detection

1. DL's performance to many sites was poor in Dec 2002
2. Yet not man.ac.uk
3. Indicated that problem likely located in NNW or BAR
4. Tracked to router misconfiguration

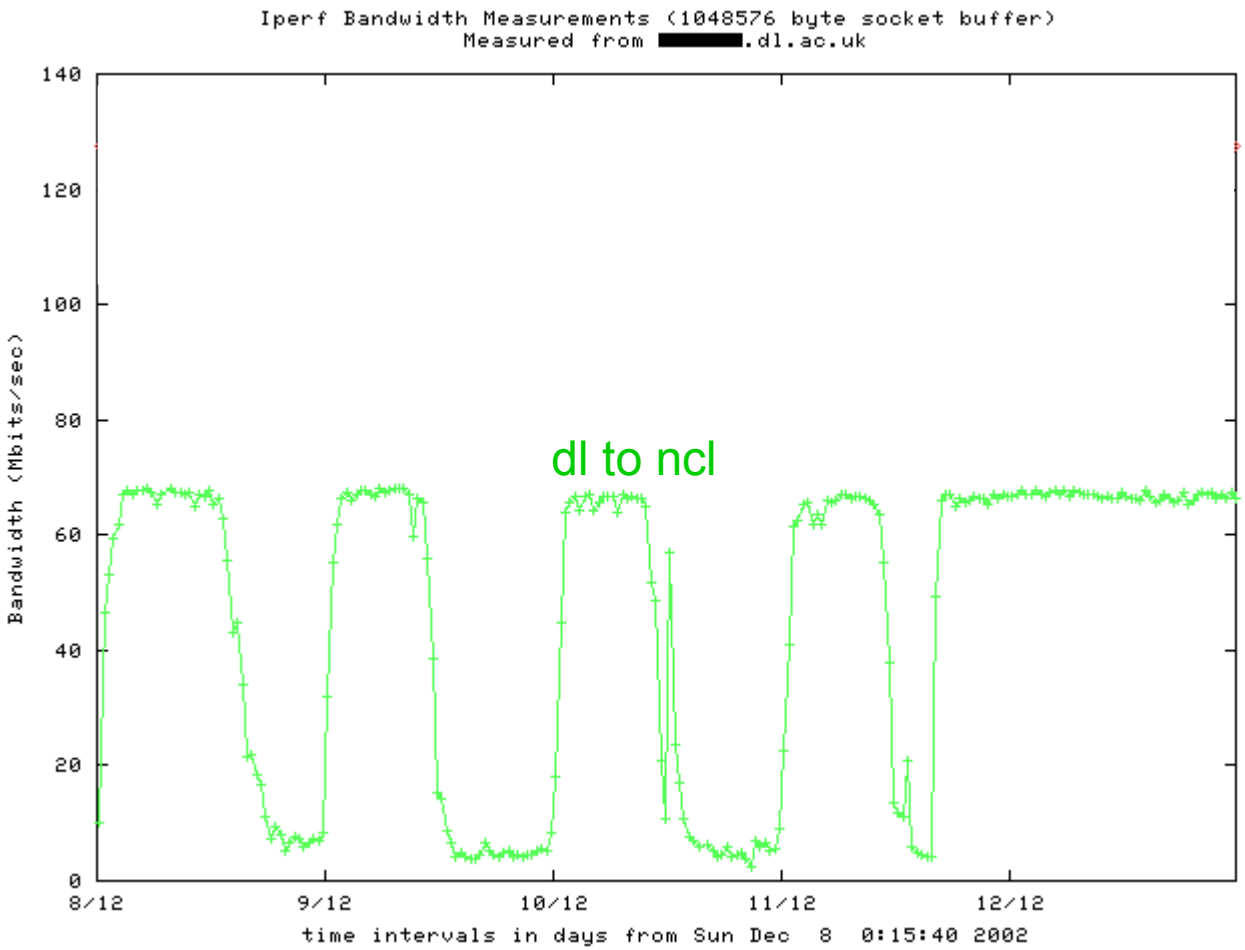

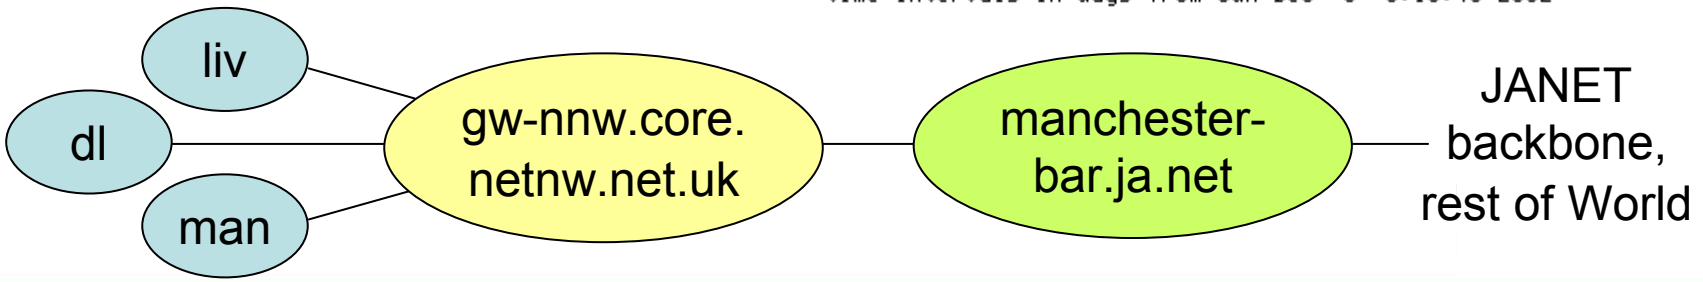

A Grid Use Case

- File replication = proven technique for improving data access
- Distribute multiple copies of same file across the Grid
- A file has Logical File Name (LFN) which maps to 1 or more Physicals
- Replica(tion) Manager responsible for replication issues, such as:
  - maintaining mapping between L and PFNs
  - deciding which replicas should exist, and where (e.g. based on recent use)
- Replica Manager includes Replica Selection Service which uses network performance data (from somewhere) to find the "best" replica

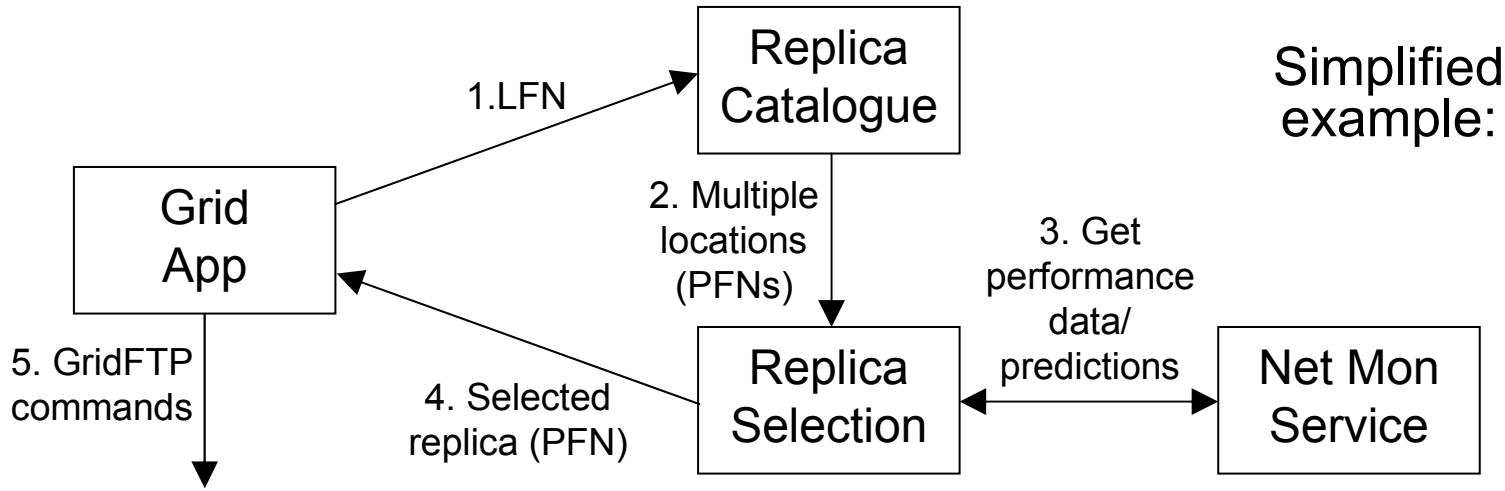

Mark Leese
Daresbury Laboratory

Justification (3)

Q: Okay, so if network performance monitoring has been around for as long as networks themselves, it must be well understood. So why the fuss over the Grid?

A: The Grid is a special case:

- As we've just seen, Grid middleware and applications could use network data to optimise their performance, adapting to changing network conditions = addition of **non-human consumption**
- We're talking about moving and sharing datasets, the sizes of which we haven't seen before. Data intensive applications (e.g. LHC in the PetaByte range, VLBI, RealityGrid…) need networks debugged for efficiency
- The Grid in its "utility computing" guise needs measurable SLAs

Justification (4)

Q: Okay, so why don't we just throw some more bandwidth at the problem? Upgrade the links.

A: For want of a more interesting term to make sure you're still paying attention, this is what I call the Heroin Effect…

- You start off with a little, but that's not really doing it for you; it's not solving the problem. So you keep increasing the dose, yet it's never as good as you thought it would be.
- By analogy you keep buying more and more bandwidth to take you to new highs but it's never quite as good as you thought it would be
- Simple over-provisioning is not sufficient
- Doesn't address the key issue of **end-to-end** performance
  - Network backbone in most cases is genuinely not the source of the problem
  - Last mile (campus network→end-user system→your app) often cause of the problem: firewall, wiring, hard disc, **application** and many more potential culprits

Justification (5)

Q: Okay, so why don't we use dedicated optical fibre everywhere?

A: Costs are still prohibitive. LHC will have 19 Tier-2 sites in the UK.

Q: Okay then smarty pants, what if we share existing fibre, and used circuit-switched lightpaths? That's dedicated bandwidth, but without the cost of dedicated fibre.

A: Good idea in theory, and we can see the benefits from the UKLight infrastructure and ESLEA*, but this still doesn't address the end-to-end issue. Take a **real-life** ESLEA example (thanks to ESLEA for the figures)…

- The UK wanted to transfer data from FermiLab to UCL for analysis, before returning the results
- datasets **currently** 1-50TB
- 50TB would take > 6 mths on production net, or one week at 700Mbps
- As a result, a 1 Gbps circuit-switched light path was provisioned
- Result = disc-to-disc transfers @ 250Mbps, just 1/4 of theoretical network max
- Tests revealed a problem at an end site

* Exploitation of Switched Lightpaths for e-Science Applications

Previous UK Work

- "...design and deploy an infrastructure for network performance monitoring within the UK e-Science community" – June 2002
- MPs (Monitoring Points) at each e-Science Centre
- Full mesh of tests
- Human access (Web interface) to monitor performance and find faults
- Based on EDG WP7 work, using pingER, iperfER and udpmon(ER)

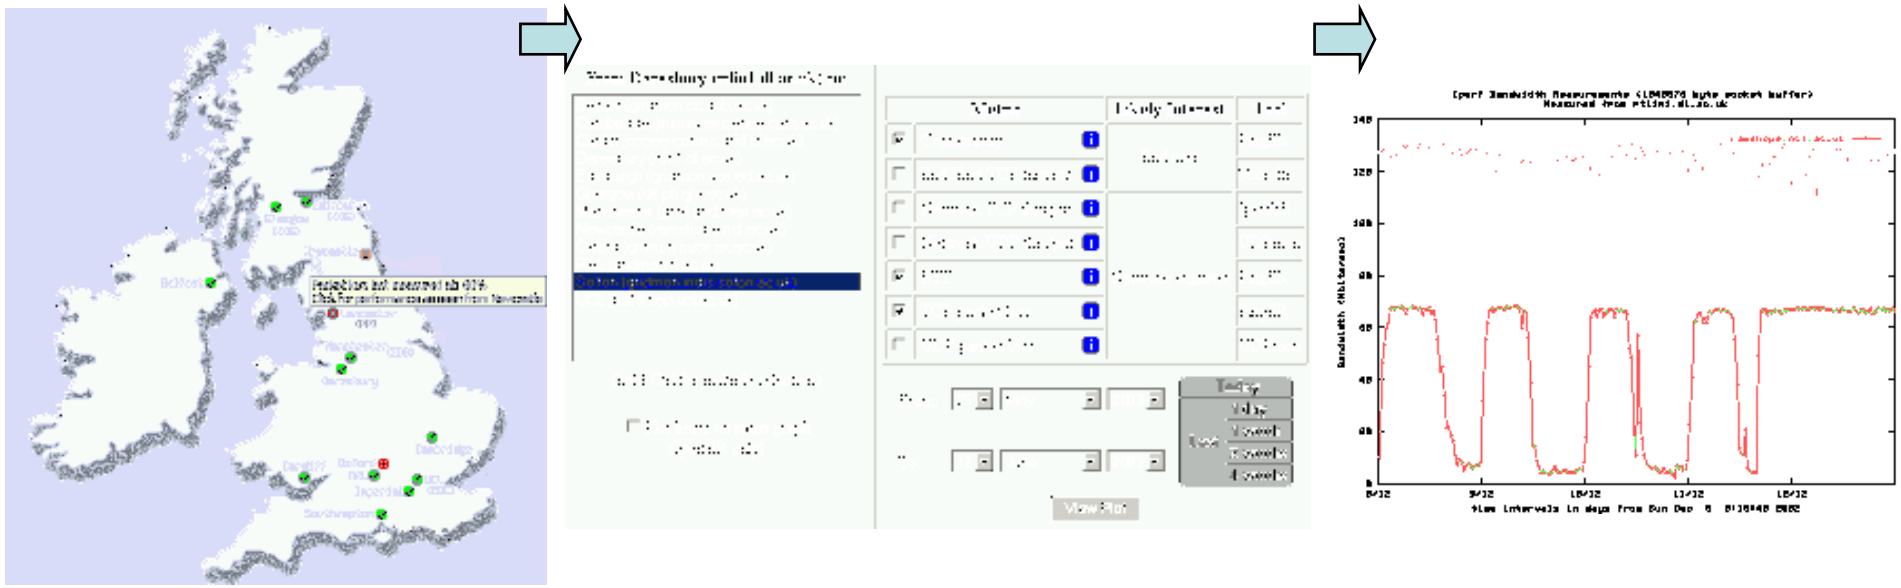

Current UK Work (1)

- Rik Tyer also working on Gridmon V2 (r.p.tyer@dl.ac.uk, http://www.e-science.clrc.ac.uk/web/staff/richard_tyer)

- V1 well received and grew interest (e.g. within UK HEP/PP community), however…
- Version 1 infrastructure proved to be unsustainable
  - many institutions were very helpful, but…
  - varying spec of machines, flavours of Linux, security rules etc.

- V1 MP:
  - Ran tests
  - Stored data locally in simple text files
  - Served data to human users using Web server running on the MP
  - Would have provided Web Services interface into collected data using Tomcat running locally

- By comparison V2…

Current UK Work (2)

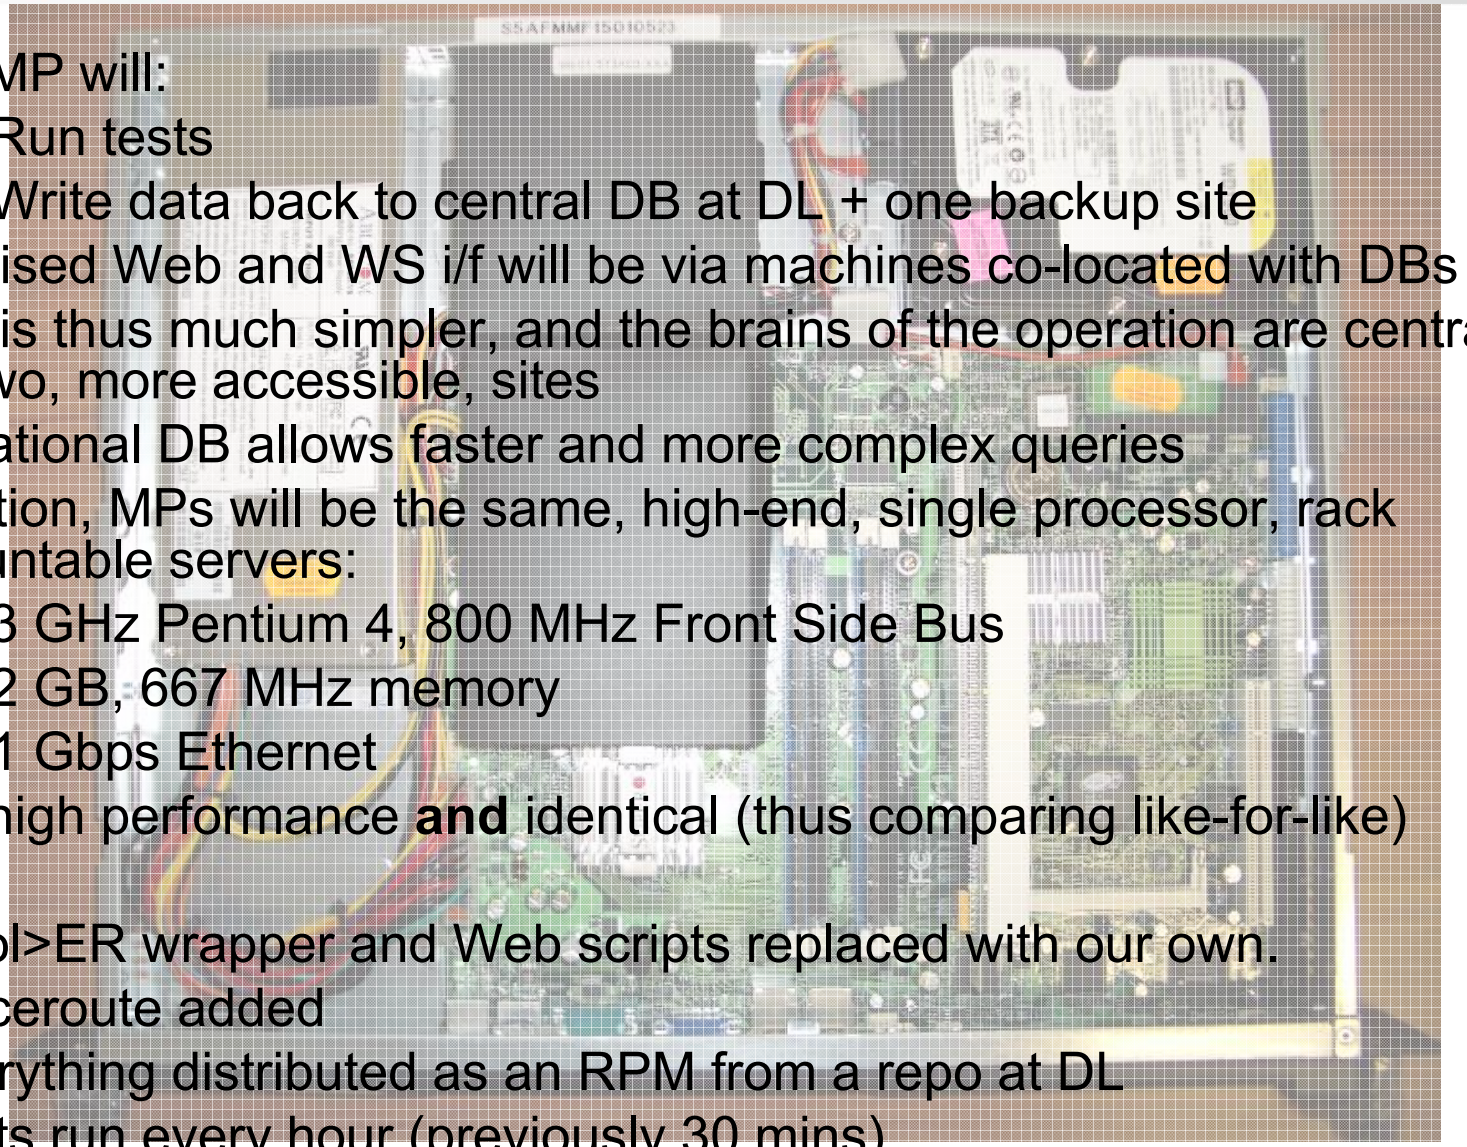

- V2 MP will:
  - Run tests
  - Write data back to central DB at DL + one backup site
- Revised Web and WS i/f will be via machines co-located with DBs
- MP is thus much simpler, and the brains of the operation are centralised at two, more accessible, sites
- Relational DB allows faster and more complex queries

In addition, MPs will be the same, high-end, single processor, rack mountable servers:

  - 3 GHz Pentium 4, 800 MHz Front Side Bus
  - 2 GB, 667 MHz memory
  - 1 Gbps Ethernet
- So high performance **and** identical (thus comparing like-for-like)

- <tool>ER wrapper and Web scripts replaced with our own.
- Traceroute added
- Everything distributed as an RPM from a repo at DL
- Tests run every hour (previously 30 mins)

Test Topology

"Full mesh" testing does not scale:

- As you add hosts it becomes more and more difficult to avoid contention between tests
- In this particular case, LHC aids us by using a topology of a central star and several mini-meshes for its data flows
- Each site only tests to/from the Tier-1 and other sites within their Tier-2
- Work is underway within EGEE to resurrect the EDG pcpd software (token passing control of tests)

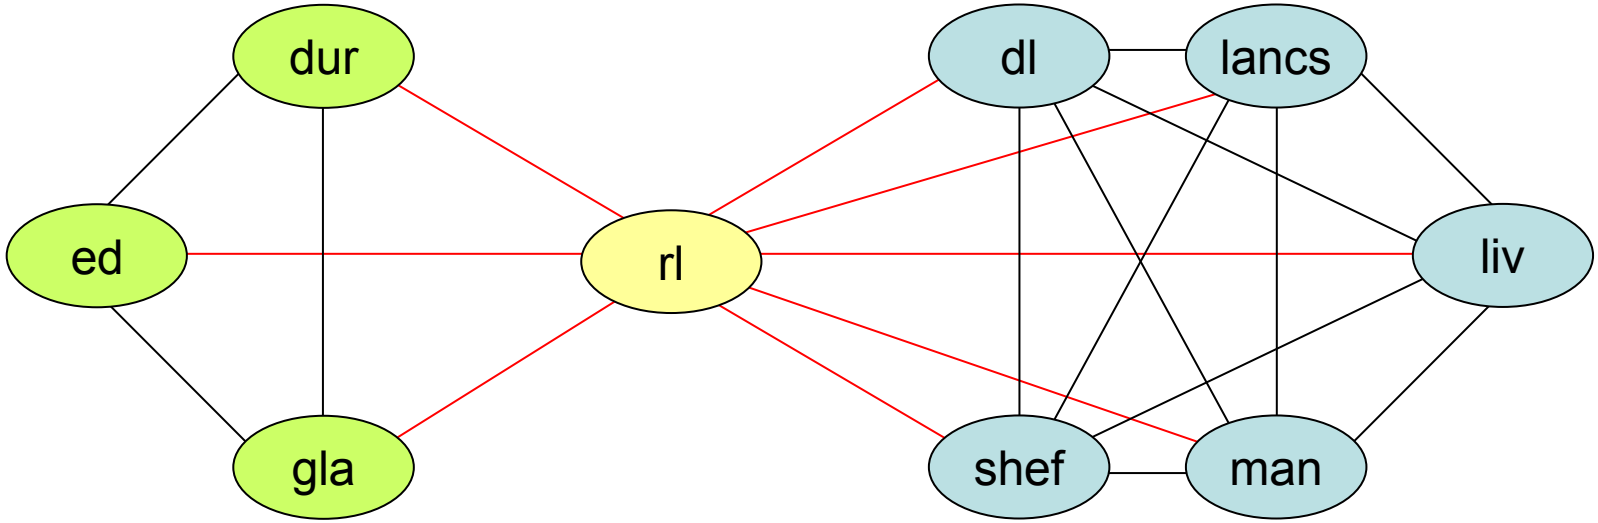

Mark Leese
Daresbury Laboratory

CCLRC Daresbury Laboratory

Monitoring Tools (1)

Tool	Origin	Metrics	Relevance
ping	Mike Muuss, 1983	Round Trip Time (RTT)	TCP is a send-acknowledge protocol. It sends a "block" of data and waits for an acknowledgement that it has arrived correctly before sending the next "block". If the ack takes a long time to arrive, the sender is sitting idle, wasting the available bandwidth between the sender and receiver.
		TCP Packet Loss	We use ping (ICMP) packet loss to approximate TCP packet loss. With TCP, if the receiver detects missing (or damaged) packets all data currently in transit is thrown away, and a re-transmit of the current "block" is requested. This is mega-wasteful, and causes a huge performance hit.
		Connectivity	The reverse of packet loss, e.g. 10% pkt loss = 90% connectivity

CCLRC Daresbury Laboratory

Monitoring Tools (2)

Tool	Origin	Metrics	Relevance
iperf	NLANR/ DAST	TCP achievable bandwidth	Your TCP speedo
udpmon	Richard Hughes-Jones, man.ac.uk HEP	UDP achievable bandwidth	Your UDP speedo
		UDP packet loss	Unlike TCP, UDP has no in-built mechanisms for automatically recovering from packet loss or reordering – it's the responsibility of the application. Their importance is largely dependant on the specific needs of your application (e.g. there is no point re-transmitting lost VC packets) but in all cases they provide general indicators of UDP "quality".
		UDP packet reordering	

CCLRC Daresbury Laboratory

Monitoring Tools (2)

Tool	Origin	Metrics	Relevance
traceroute	van Jacobson	Hoplist	The nodes (routers) visited along a test path. Can be used to monitor route changes. Changes may not be beneficial, or may indicate other problems (e.g. failure of your site's main link has caused it to route traffic via a resilient link, provided by another route) Can identify cases of asymmetric routing (e.g route A→B is not the reverse of B→A)
		Hop delays	Provides three RTT measurements for each hop along the path. Can be used to indentify bottlenecks or problem hops, e.g. if one hop's RTT is significantly larger than the previous hop's, and there is no obvious reason for this, such as the hop has taken you across the Atlantic.

Firewall Requirements (1)

- This is the **minimum** required for the monitoring infrastructure to be worthwhile. It would be very useful if your firewall was opened to the other UK test machines, so that when required, we can run manual tests for locating and diagnosing problems. A list of machines is available on request.

SSH

Allow inbound access from:

- gridmon.dl.ac.uk (193.62.119.20) for Mark
- rikster.dl.ac.uk (193.62.113.31) for Rik

Yum (for patching)

- Allow port 80 outbound for ebro.dl.ac.uk (193.62.125.80) for Gridmon updates. Alternatively, provide details of your web proxy.
- Allow outbound FTP access to ftp.scientificlinux.org (131.225.110.41) for access to the Scientific Linux repositories. This is client access only - the Gridmon boxes do not run an ftp server.

Ping & Traceroute

- Allow inbound and outbound pings and ICMP (not UDP) traceroute to the other test machines in your Tier-2 and the RAL Tier-1 test machine.

Mark Leese
Daresbury Laboratory

Firewall Requirements (2)

Allow inbound Iperf and Udpmon test

1. Allow inbound access to <IperfPort/udpmonPort> from port > 1024 on the other test machines in your Tier-2 and the RAL Tier-1 test machine, thus allowing machines to run tests to your iperf/udpmon server.
2. Allow port <IperfPort/udpmonPort> outbound access to port > 1024 on the other test machines in your Tier-2 and the RAL Tier-1 test machine, thus allowing your iperf/udpmon server to reply to machines running tests to it.

Allow outbound Iperf and Udpmon tests

3. Allow port > 1024 outbound access to <IperfPort/udpmonPort> on the other test machines in your Tier-2 and the RAL Tier-1 test machine.
4. Allow port <IperfPort/udpmonPort> on the other test machines in your Tier-2 and the RAL Tier-1 test machine inbound access to port > 1024

- Steps 2 and 4 can be discounted if have a stateful firewall, i.e. one which logs allow inbound and outbound connections that you've made and automatically allows traffic in the return direction.
- This presentation may go on the web, so I haven't revealed the port numbers or machine names. Please email me for the details.

The machines are also firewalled locally to this affect!

Example (1)

- We've had test nodes in place at DL, NeSC and RAL for some time.
- Similar plots will be available via the new Web interface.
- Period covers March's LHC Service Challenge: http://www.gridpp.ac.uk/news/-1143610450.137009.wlg

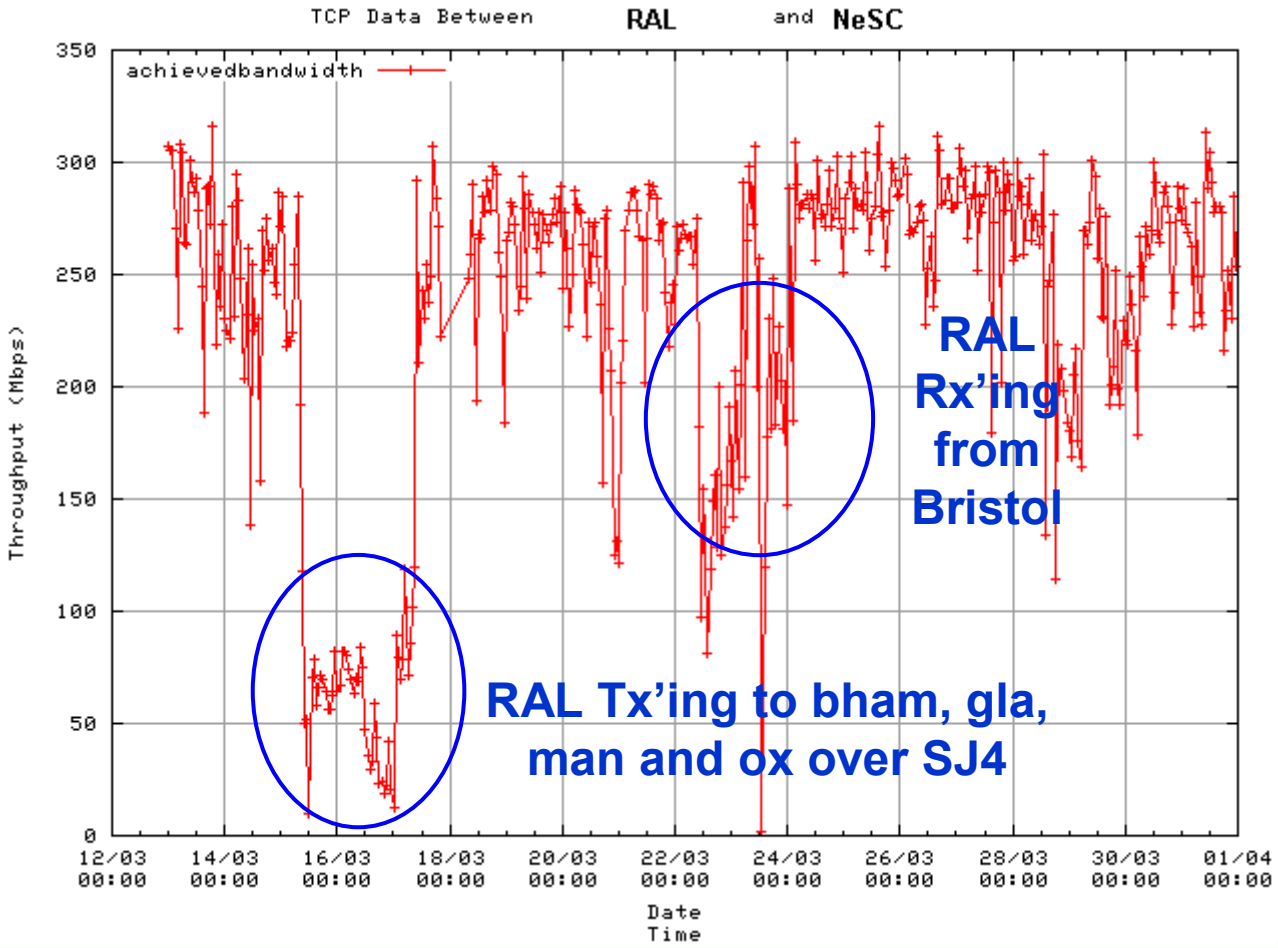

Very spikey, but general level around 250-300Mbps. Over commodity IP network that's a good start.

Note that RAL Rx'ing from Bristol **appears** to impair RAL's ability to transmit.

Mark Leese
Daresbury Laboratory

Example (2)

- The opposite direction (into RAL) is very spikey. This is most likely contention with other traffic - the Tier-1 does have quite substantial flows inbound onto the subnet on which the Gridmon box resides.
- Without prior knowledge of the other traffic flows this is hard to avoid, though provisioning extra capacity can help

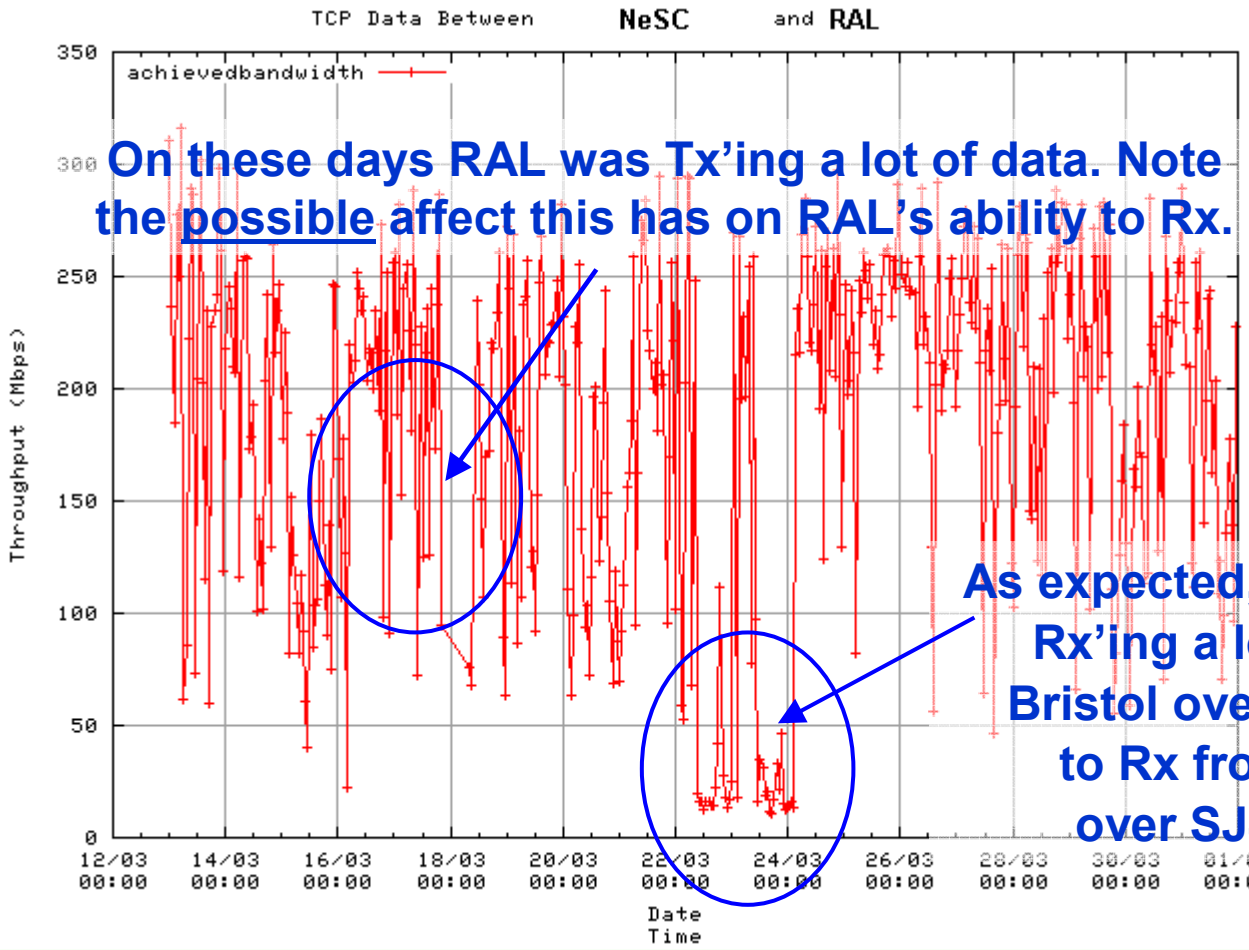

On these days RAL was Tx'ing a lot of data. Note the possible affect this has on RAL's ability to Rx.

As expected, when RAL starts Rx'ing a lot of data from Bristol over SJ4, its ability to Rx from other sites over SJ4 is affected.

GGF NM-WG in one slide

- Grid Global Forum Grid computing standards body
- Mark co-chairs the Network Measurements-Working Group

- In the period 2002-'05 we produced unified XML schemas for requesting and subsequently publishing network performance data (e.g. results of bandwidth tests)
- In a sentence, the schemas: **provide a unified interface for network operators, and Grid middleware/apps to share performance data** electronically (using Web Services)

- V1 schemas:
  - Monolithic, "one size fits all" solution
  - Just two schemas: less to maintain, less to worry about parsing etc.
  - Test deployments identified some gaps in group thinking, but were crucial in selling the idea to people
- V2 schemas:
  - Successful trial implementations happening (inc. Internet2 and ESNet in the US, and DANTE in Europe)

Mark Leese
Daresbury Laboratory

EGEE JRA4

- **European Grid project, the successor to EDG, completed on 31st March 2006:**
  - EGEE Joint Research Activity JRA4 = group responsible for “Development of Network Services”, inc. Network Performance Monitoring (NPM)
  - Some work continues in EGEE II, under Service Activity SA1 (European Grid Operations, Support and Management)

- **The work was about standardising access to NPM data across multiple domains and using it:**
  - NM-WG schemas as the selected basis for standardisation

- **Outputs included:**
  - Mediator: standardising access to NPM data
  - Diagnostic Tool: Web interface presenting the data to Network and Grid Operations Centres (NOCs and GOCs)

Architecture (1)

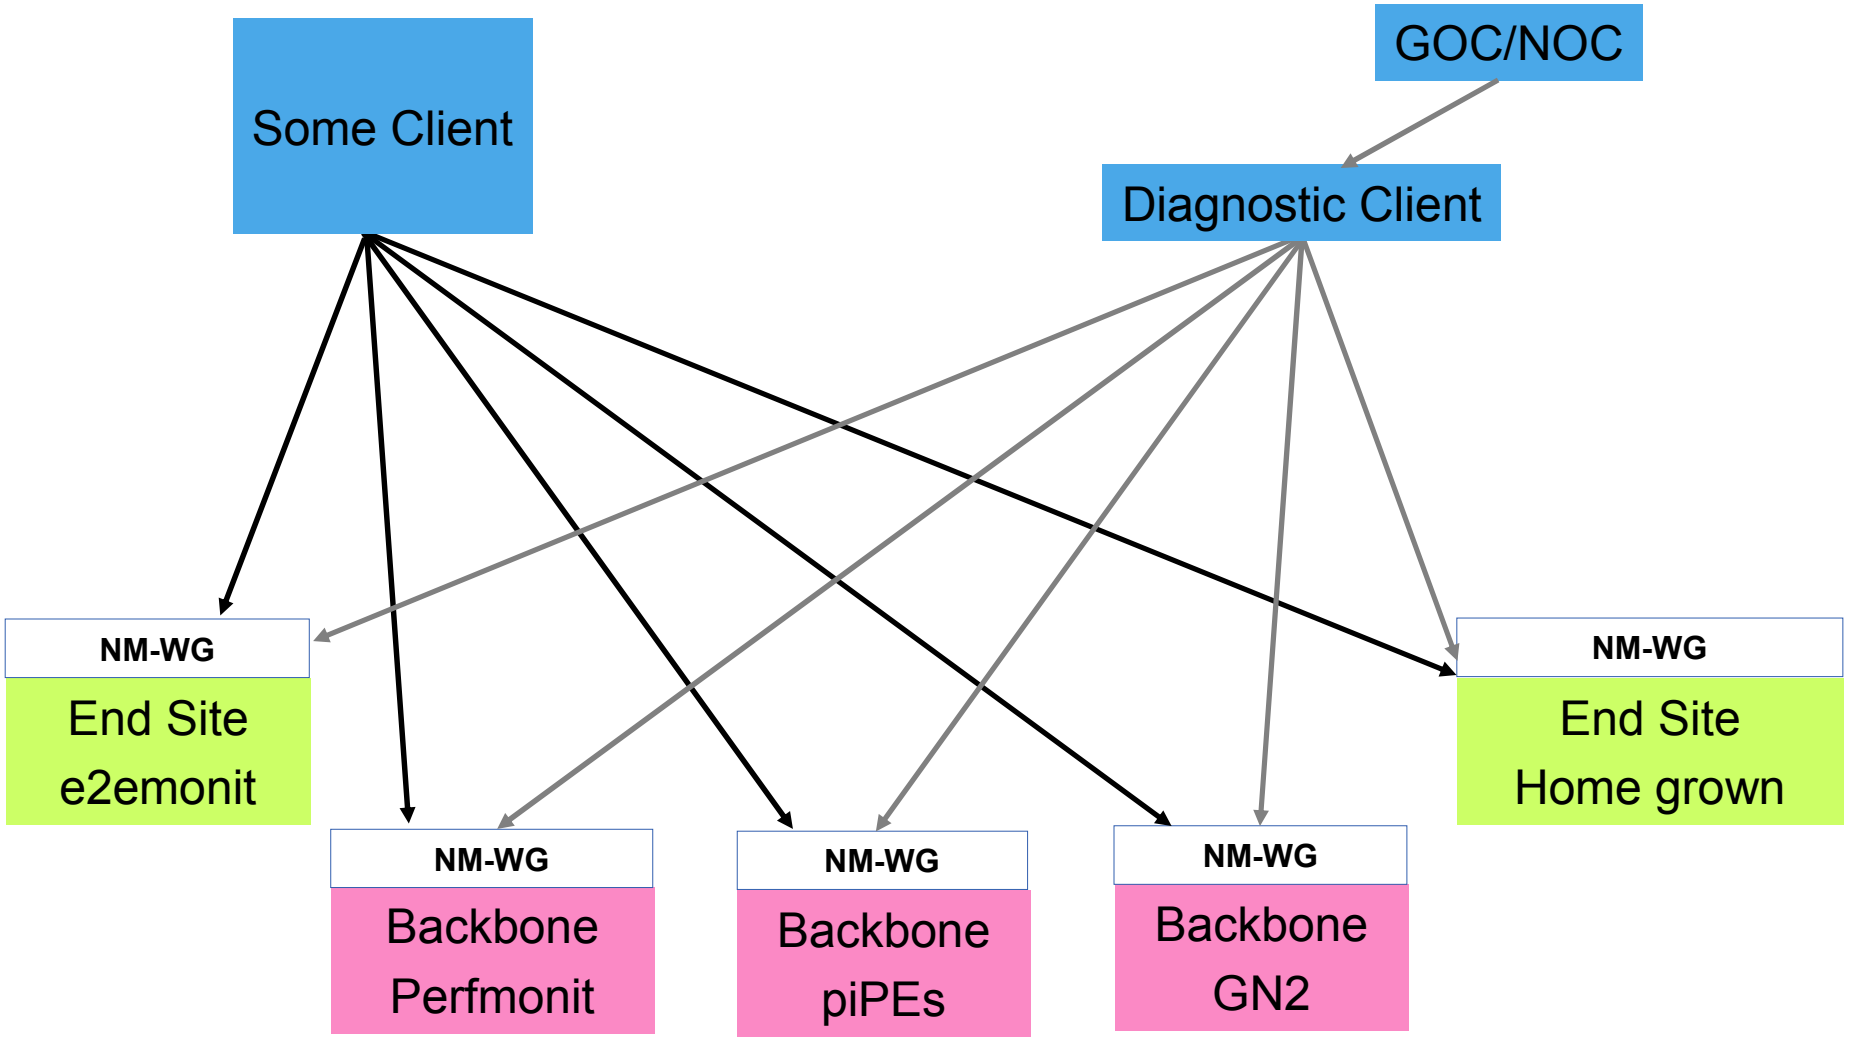

Architecture (2)

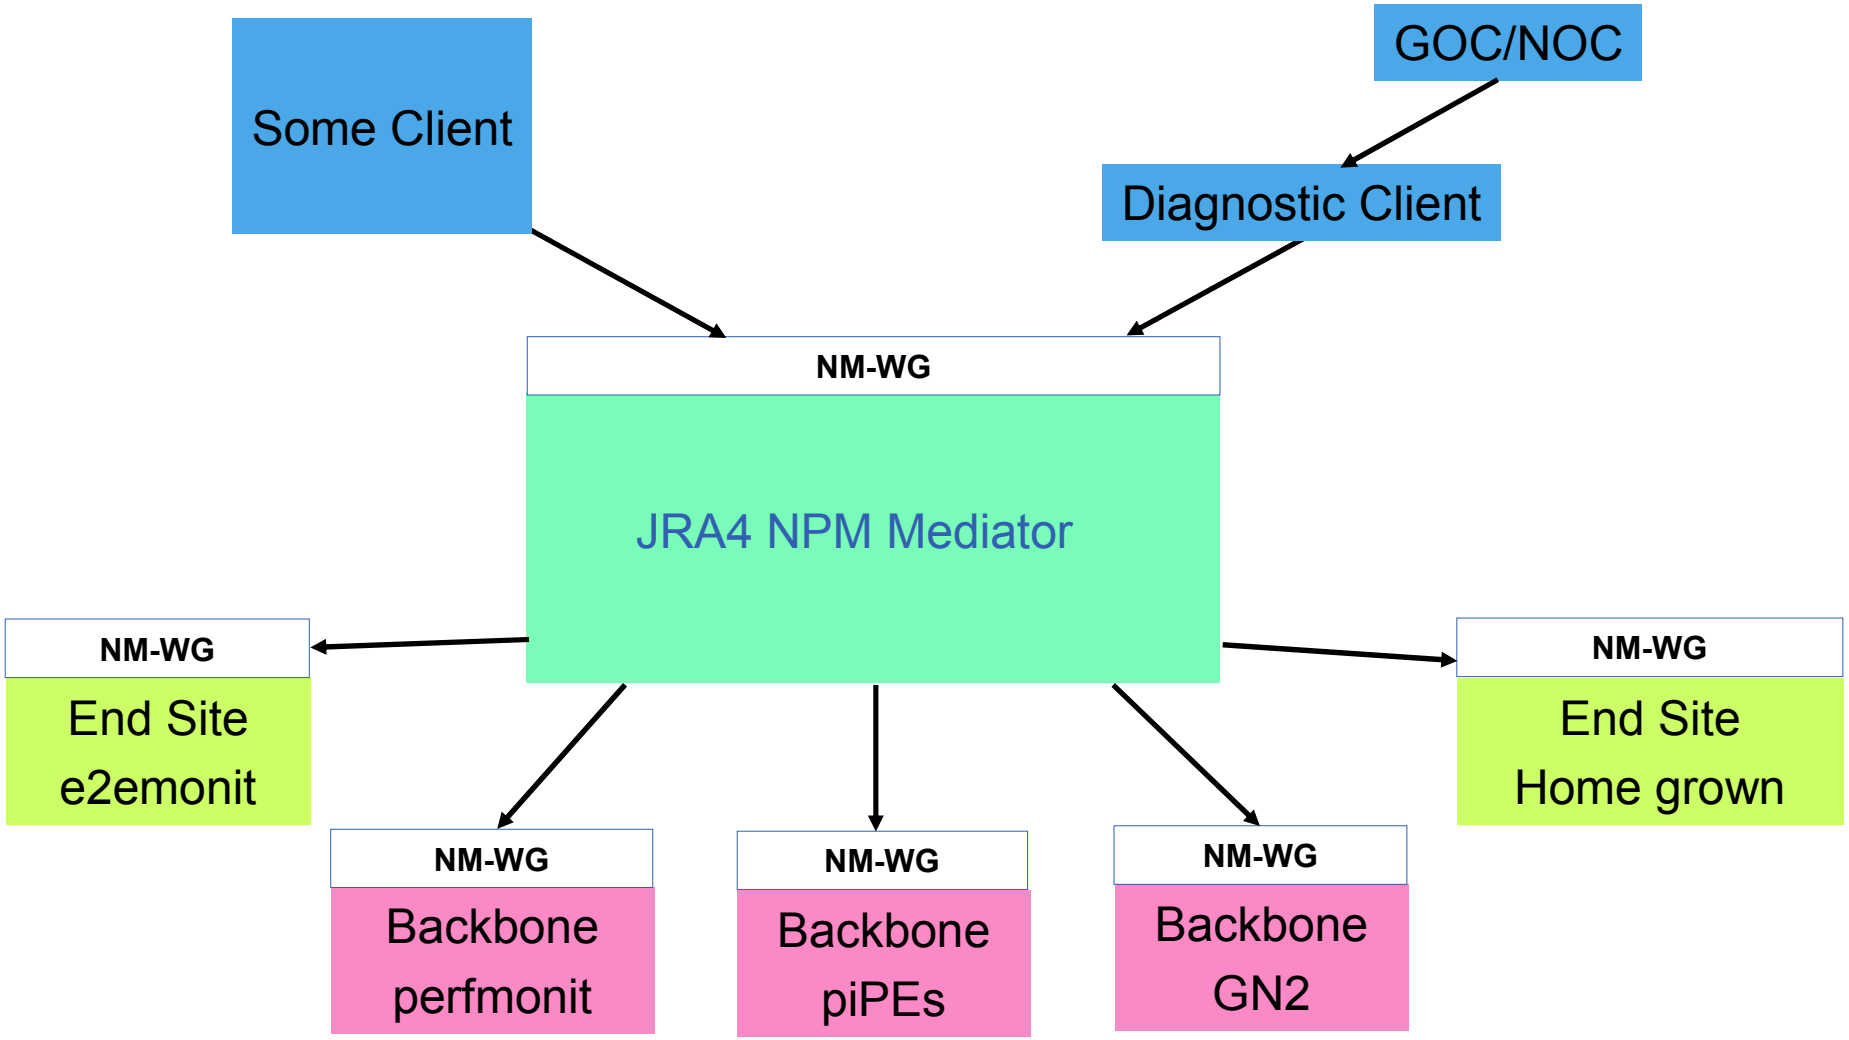

Diagnostic Tool (1)

- **Diagnostic Tool (DT) provides Web interface access to any network data accessible via the Mediator, i.e. any data that the Mediator can access via the unified NM-WG interface – a lot :)**
  - Must have a valid X.509 certificate to gain access
  - https://edms.cern.ch/file/653967/1/EGEE-JRA4-TEC-653967-DTUserGuide-v1-3.pdf
  - Demonstrated at GGF15 and 4th EGEE conference (October 2005) graphing data from Abilene, ESNet, GÉANT2 and e2emonit (JRA4 end-to-end monitoring infrastructure).

- **DT can access lots of data (EGEE, DANTE etc.) BUT must do so through a Web Services interface – not very efficient for graphing.**
- **The Gridmon Web interface can access data more natively using a simple TCP connection via a DB interface such as PerlDBI, BUT is Gridmon/UK only**

- **There are different approaches to deployment & dissemination in use and being developed throughout the World.**
- **People are not necessarily recreating the wheel. There's a need to see what's best for different scenarios.**

Diagnostic Tool (2)

NPM Diagnostic Tool

Query

Time Range — Set

Start: 2006-01-27 08:00:00 | Start : 2006-01-27 08:00:00 | Focus: 2006-01-27 08:30:03

End: 2006-01-27 12:30:00 | Period: 270 Minutes | Tolerance [s]: minus 1803 plus 14397

Max Results

Maximum number of results: unbounded

Test Path

Source: ornl-orau.es.net, ornl-osti, ornl-osti-atm, ornl-osti-atm.es.net

Destination: ornl-osti-atm

sanet.sk1.sk.geant.net => sanet

sanet => sanet.sk1.sk.geant.net

ornl-osti-atm.es.net => ornl-osti-atm

Choose New Source | Add Path

Metric

Metric: Utilization | Packet Size: = byte | Packets =

Statistic

Statistic: Average | Interval: 5 Minutes

View As

Data Table | Matrix | Time Plot

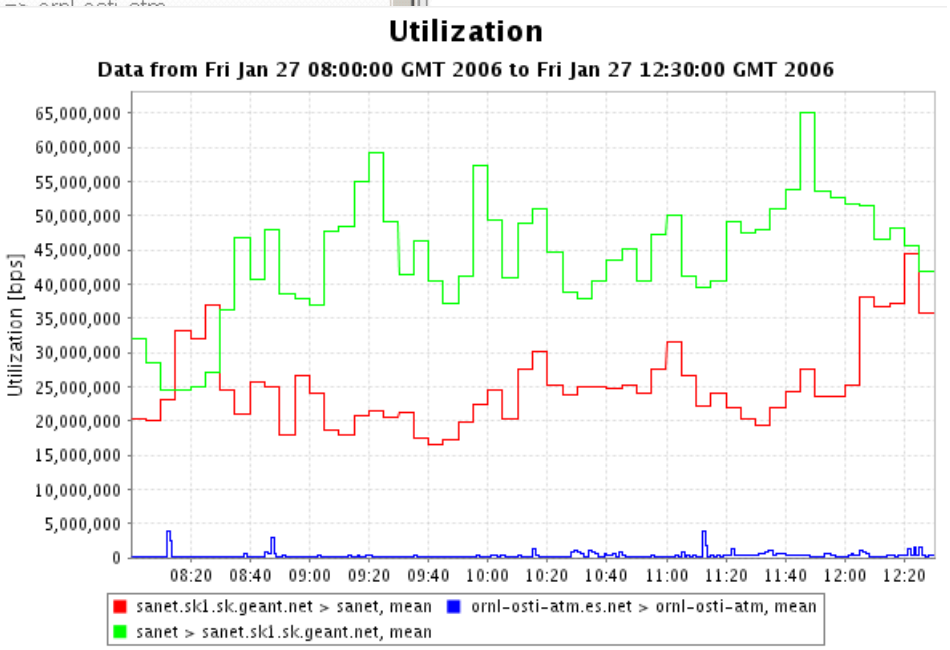

More Info

- **Networks For Non-Networkers (NFNN)** – two workshops organised by Mark, looking at network performance: http://gridmon.dl.ac.uk/nfnn/
- Monitoring use cases - JRA4 Diagnostic Tool use cases: https://edms.cern.ch/document/591777/1
- Gridmon: http://gridmon.dl.ac.uk/ (under re-construction)
  - September 2005 Conference Paper: http://gridmon.dl.ac.uk/~mjl/presentations/LeeseTyerTaskerAHM05pres.pdf
- GGF: http://www.ggf.org, NM-WG: http://nmwg.internet2.edu
- EGEE-JRA4: http://egee-jra4.web.cern.ch/EGEE-JRA4/
  - NPM Final Report: https://edms.cern.ch/file/695235/1/EGEE-DJRA4.7-695235-v1-1.doc
- If you can't find what you're looking for, get in touch: m.j.leese@dl.ac.uk

CCLRC Daresbury Laboratory

Conclusion

Forget the talk (trust me, it's quite easy ;-)
Just take this away with you…

- Network performance monitoring is crucial to the Grid
  - Adaptive behaviour, predicting performance, making the network efficient, fault detection, monitoring SLAs

- We're interested in the **end-to-end** network performance
- Problems are **frequently** not caused by the network:
  - Your hard disc, your application (GridFTP vrs HTTP) etc.
  - If it is the network, look in the last mile, e.g. the firewall

- GridMon is the infrastructure for UK academic network performance monitoring
- There's lots happening. The UK has its fingers in many worthwhile pies, but consensus and development take time.