



GridPP

UK Computing for Particle Physics

Network Monitoring: Continuing Progress & Status Report

Robin Tasker (r.tasker@dl.ac.uk)
CCLRC, Daresbury Laboratory

28 June 2006



Network Performance Monitoring (thanks to Mark Leese and Rik Tyer¹ at Daresbury)

1. Rik Tyer funded by the JISC through UKERNA to embed the Gridmon network performance monitoring into the national network infrastructure in support of Grid computing. Rik works for the ESC Centre at Daresbury

[Background and Theory](#)

[Implementation Status](#)



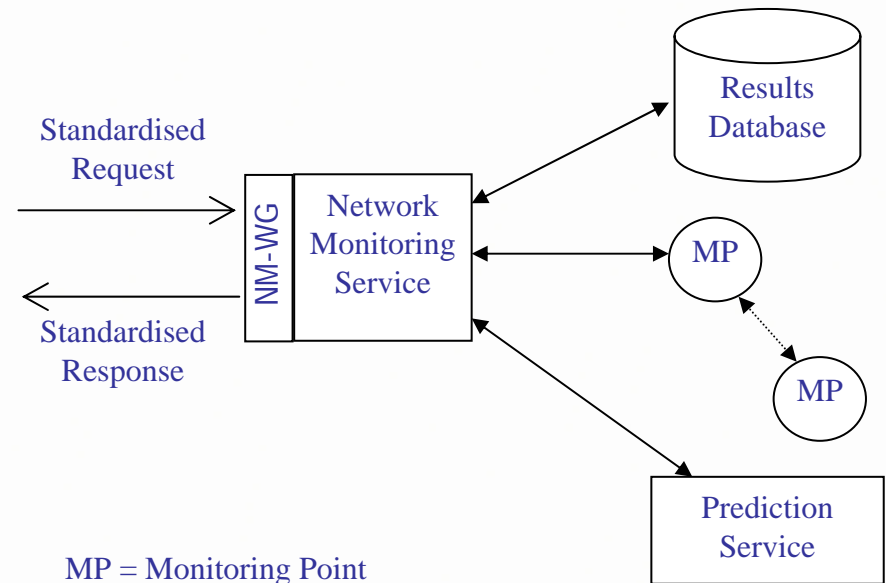
Request / Response Approach

Interactive systems clients will be able to request

- historic data,
- future or on-demand tests
- predictions (i.e. NWS style).

All request and result messages are formatted using standardised schemas developed within the **GGF NMWG**.

This allows heterogeneous monitoring systems to interact providing that they use the same schemas.





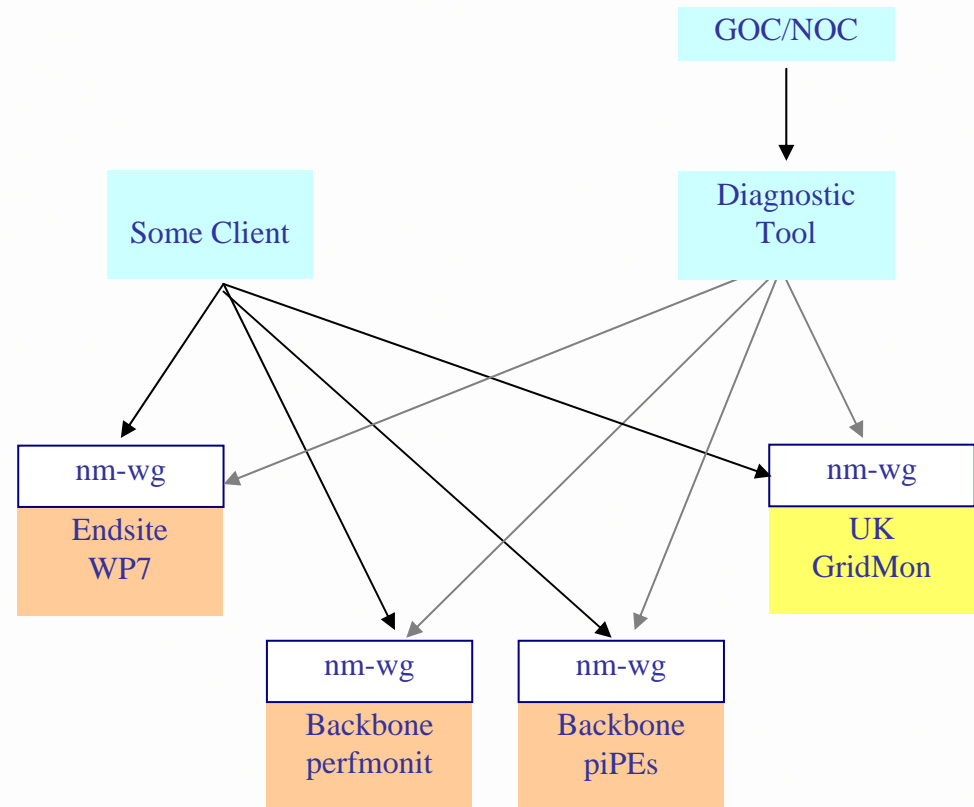
If a Client, whatever that may be, wishes

to receive detailed information concerning a network path;

or

to concatenate data from several backbone networks to produce a picture of the full backbone path.

Then the Client has to make several requests, and if necessary, perform its own data aggregation.

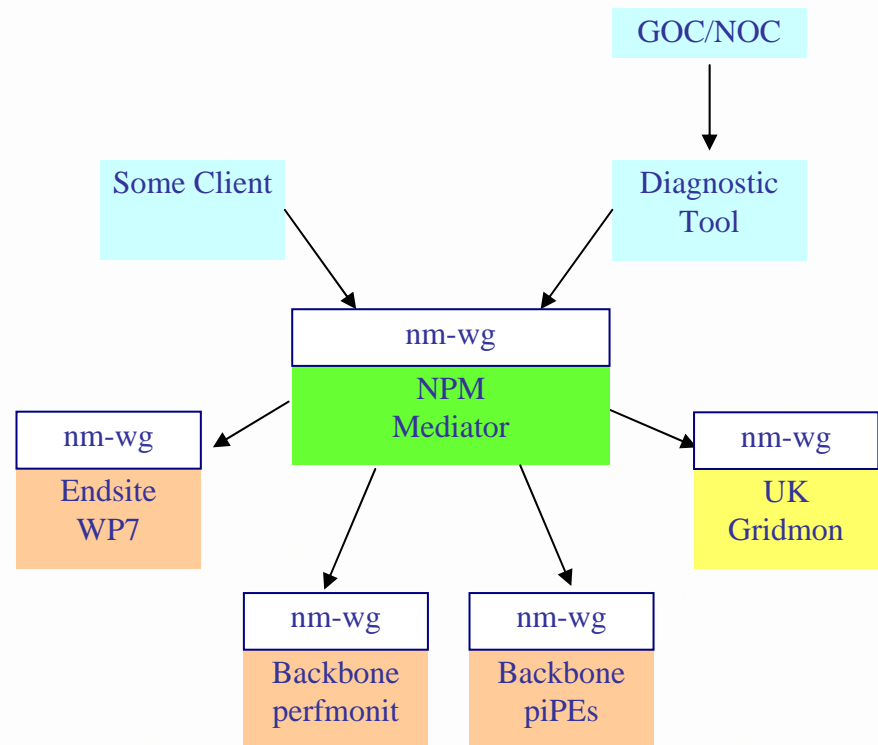




EGEE JRA4 is responsible for “Development of Network Services” .

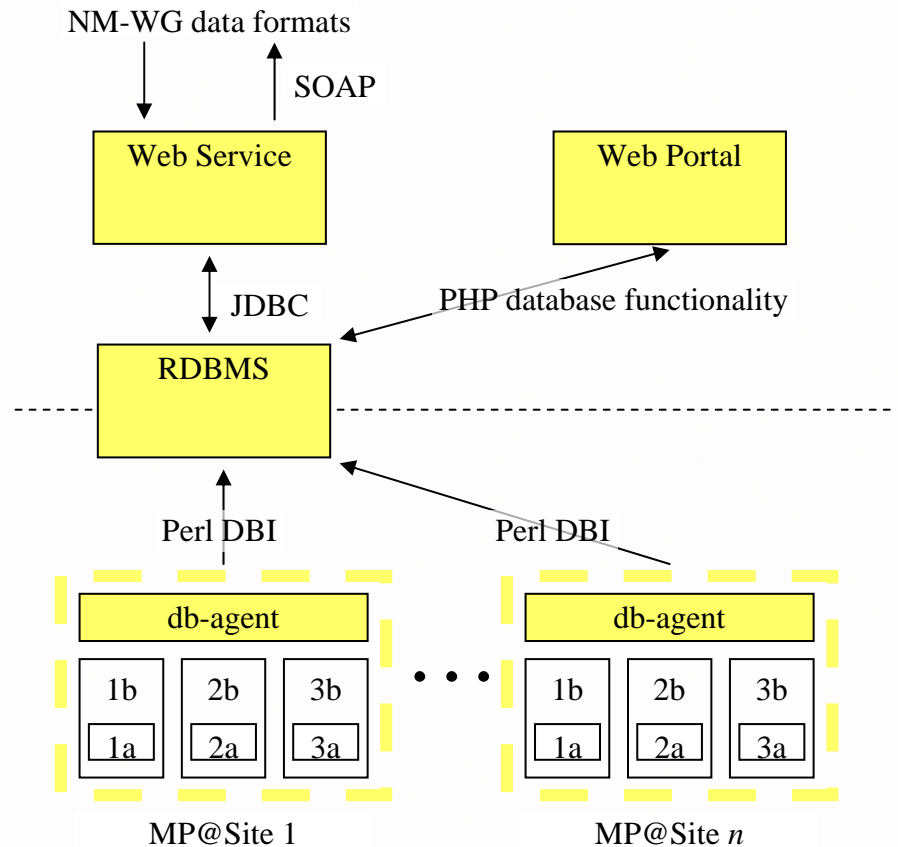
Within JRA4’s Network Performance Monitoring (NPM) activity, the “Mediator” software has been developed.

This will greatly simplify this process by unifying access to network performance measurement data.






1. An infrastructure in which each site (MP) will store its test data in a central database.
2. Web Services and human (web) access to the data will also be via services running centrally.
3. Storing and providing access to the data from a central location reduces complexity of the individual monitoring nodes
4. Critically, we can also move to a relational database model, speeding up access to the data and allowing considerably more advanced queries to be made





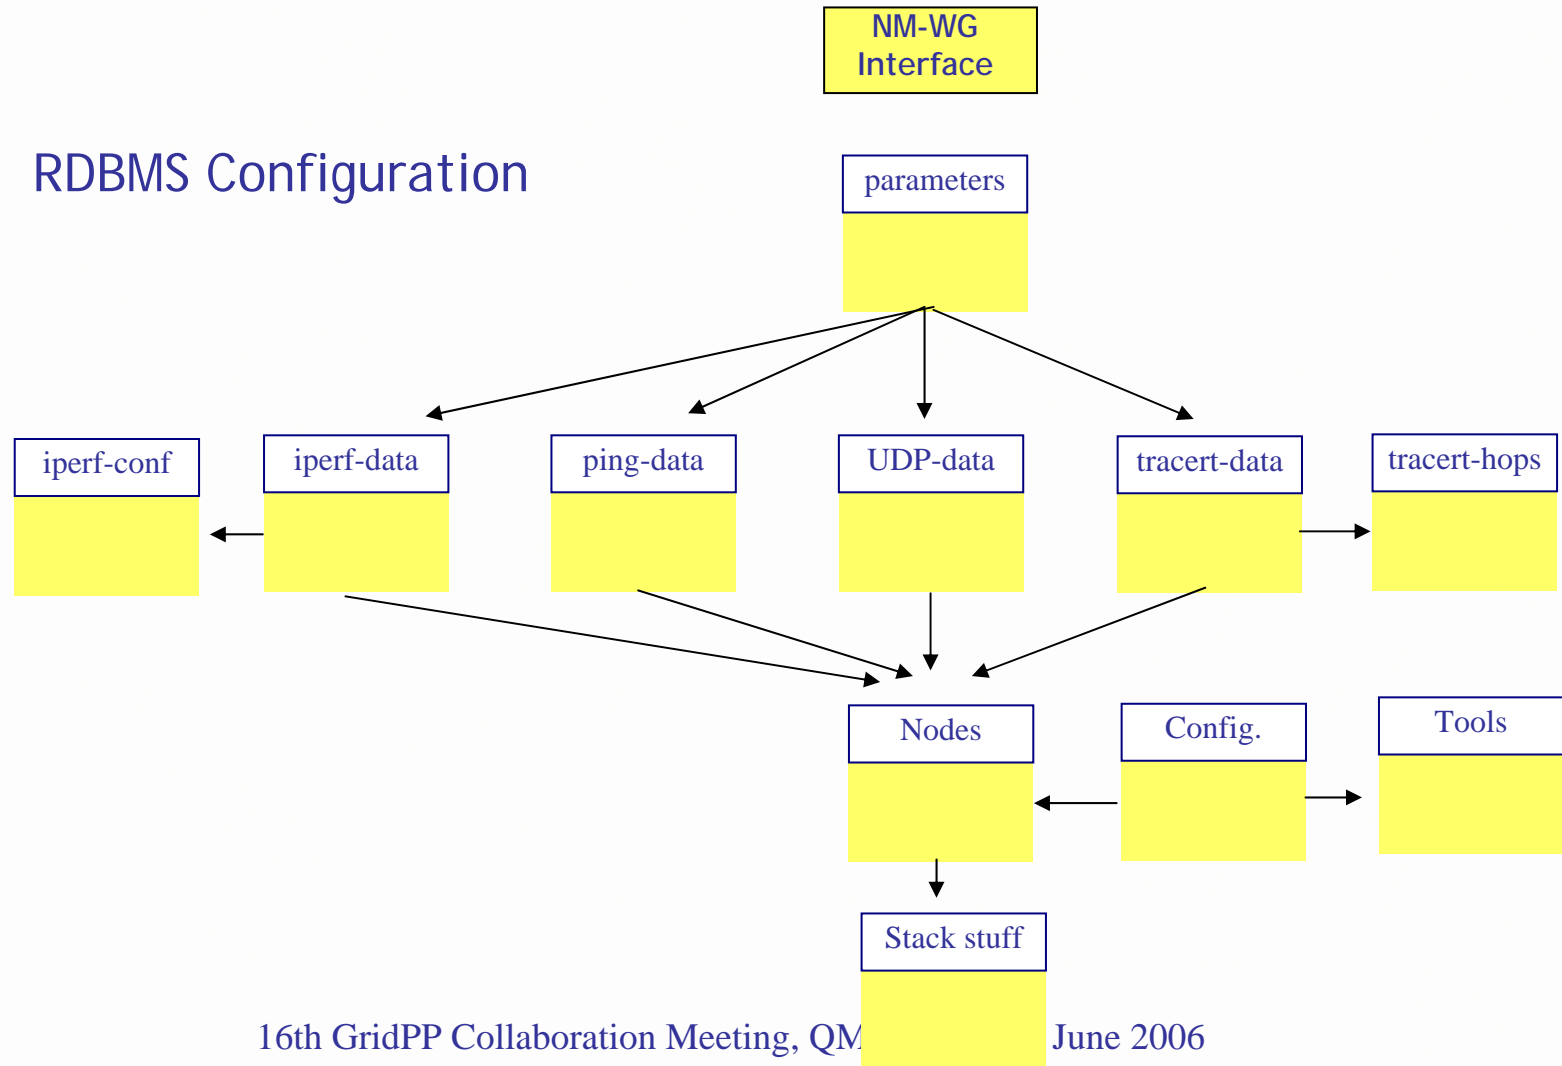
Monitor Point Configuration



MP@ GridPP Tier 2 Site



RDBMS Configuration





1. Map view:

- snapshot view of site's connectivity within the last 30 mins
- looking at adding representation of thru'put (achievable bandwidth)
- of interest to end-users, GOC and NOC

2. Graph view:

- will offer more flexible graphing (e.g. multiple metrics and/or multiple site on the same graph...up to a limit!)
- of likely interest to NOC, GOC and end-users (in that order)

3. Text view:

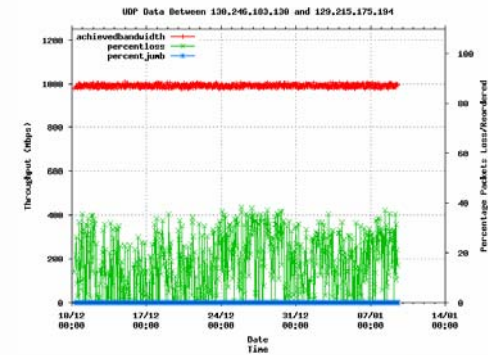
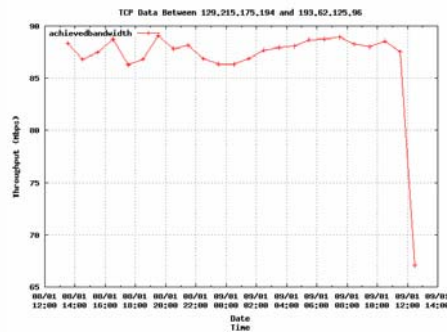

- will allow users to make specific queries for textual data, e.g. min/max values for TCP thru'put over last day
- to prevent abuse will initially will be tied down to specific queries
- of likely interest and provides more detail and finer control of queries for NOCs or local Sys and NetAdmins



1. GridMon V1 (for UK e-Science) ran tests in a **full mesh** (between all nodes) for the 12 original e-Science centres and 1-2 others
2. The LHC model is different - RAL (as Tier-1) distributes data to NorthGrid, Southgrid *etc.* where data is processed within that Regional Tier-2.
3. Proposed that Full mesh therefore replaced with:
 - **star** from RAL (Tier-1) to all sites (Tier-2s)
 - **full mesh** within each Regional Tier-2
 - e.g. Edinburgh runs tests to Glasgow, Durham and RAL but not Cambridge or Manchester or UCL or...
4. Could be changed if desired, but test contention is an issue - meshes do not scale, and GridPP has approx 20 sites



1. MP and RDBMS software up and running with prototype Gridmon V2 nodes located at Daresbury, RAL and Edinburgh collecting data...



2. MP machines purchased and delivered, about to be configured and shipped to all sites by February/March timeframe.



SuperMicro SuperServer 5015M-MF

3. Portal implementation will be visible through January...design meetings will be organised with interested parties



1. Machines shipped to site at the end of March (except for the one left under the desk in Stores at Daresbury 😊)
2. Up and running: Birmingham, Bristol, Brunel, Daresbury, Durham, Edinburgh, Glasgow, Lancaster, Liverpool, Oxford, RAL Tier-1, RAL Tier-2, RHUL
3. Caveats: 4-5 sites are failing to run of some tests, e.g. currently trying to agree with firewall people that ICMP (e.g. ping) traffic is okay
4. Manchester and Sheffield have the machines but they need to be switched on 😊
5. Awaiting info: Cambridge, Imperial, QMUL, UCL
6. **MANY, MANY THANKS** to the local site contacts for their help in bringing these machines into service



	Inbound				Outbound						NIC (Mbps)
	Ping	Tracert	Iperf	UDPmon	Ping	Tracert	Iperf	UDPmon	NTP	yum	
Birmingham	N	N	Y	Y	N	N	Y	Y	Y	Y	1000
Bristol	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	1000
Brunel	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	1000
Cambridge	no info										
Daresbury	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	1000
Durham	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	1000
Edinburgh	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	1000
Glasgow	N	N	Y	Y	N	N	Y	Y	Y	Y	1000
Imperial	no info										
Lancaster	Y	Y	Y	?	Y	Y	Y	?	Y	Y	1000
Liverpool	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	1000
Manchester	no power										
Oxford	Y	Y	Y/N	Y	Y	Y	Y/N	Y	Y	Y	1000
QMUL	no info										
RAL Tier-1	?										
RAL Tier-2	Y	Y	N	N	Y	Y	Y	Y	N	Y	1000
RHUL	Y	Y	Y/N	Y	Y	Y	Y	Y	Y	Y	1000
Sheffield	no power										
UCL	no info										



1. First prototype available here:
<http://gridmon3.dl.ac.uk/gridmon/>
2. Provides access to plots of all the data while extra functionality added
3. Please let us know when you find any bugs ☹
4. Planned functionality (happy to take any good ideas)
 - multiple path plots on same graph (same src or dst, or arbitrary, e.g. dl->ral with ed->gla)
 - possibly multiple metrics on same graph, e.g. UDP bandwidth with UDP packet loss
 - frequency distributions
 - status map (left hand tab)
 - textual access to data (right hand tab) such as min/max over a time period, traceroute results...
5. Web Services interfaces work continues, via GGF NM-WG



Map View (Disabled)

Graph View

Text View (Disabled)

Reset Form

This page can be used to plot raw data and statistical summaries for the monitored network connections.

From	To
Tier-1 RAL (x.gridpp.rl.ac.uk)	Tier-1 RAL
London Brunel (x.brunel.ac.uk) Imperial (x.ic.ac.uk) Queen Mary (x.qmul.ac.uk) Royal Holloway (x.pp.rhul.ac.uk) UCL (x.ucl.ac.uk)	London Brunel Imperial Queen Mary Royal Holloway UCL
NorthGrid Daresbury (x.dl.ac.uk) Lancaster (x.lancs.ac.uk) Liverpool (x.ph.liv.ac.uk) Manchester (x.man.ac.uk) Sheffield (x.shef.ac.uk)	NorthGrid Daresbury Lancaster Liverpool Manchester Sheffield
ScotGrid Durham (x.dur.scotgrid.ac.uk) Edinburgh1 (x.epcc.ed.ac.uk) Edinburgh2 (x.nesc.ed.ac.uk) Glasgow (x-gla.scotgrid.ac.uk)	ScotGrid Durham Edinburgh1 Edinburgh2 Glasgow
SouthGrid Birmingham (x.ph.bham.ac.uk) Bristol (x.phy.bris.ac.uk) Cambridge (x.cam.ac.uk) Oxford (x.physics.ox.ac.uk) RAL (x.pp.ral.ac.uk)	SouthGrid Birmingham Bristol Cambridge Oxford RAL

Metric	Likely Interest	Tool
<input type="checkbox"/> Connectivity	1	ping
<input type="checkbox"/> Network TCP thru'put	1	iperf
<input checked="" type="checkbox"/> Network UDP thru'put	1	udpmom
<input type="checkbox"/> RIT	1	ping
<input type="checkbox"/> TCP (≠ICMP) packet loss	1	GOC, NOC ping
<input type="checkbox"/> UDP packet loss	1	udpmom
<input type="checkbox"/> UDP packet reordering	1	udpmom



Auto-scale Y axis Mark data points

Data Type: Raw data

From: 11 June 2006 00:00
To: 25 June 2006 00:44

Today
1 day
Last 1 week
2 weeks
4 weeks

View Plot





1. Gridmon V1 (for UK e-Science) ran tests in a **full mesh** (between all nodes) for the 12 original e-Science centres and 1-2 others
2. To reflect the LHC T1/T2 operational model, the full mesh therefore replaced with:
 - star from RAL (Tier-1) to all sites (Tier-2s)
 - full mesh within each Regional Tier-2
e.g. Edinburgh runs tests to Glasgow, Durham and RAL but not Cambridge or Manchester or UCL or...
3. Could be changed **but** test contention is an issue
 - meshes do not scale, and
 - GridPP has approx 20 sites



GridPP

UK Computing for Particle Physics

Questions?

16th GridPP Collaboration Meeting, QMUL 27 - 29 June 2006