Network Monitoring for the Grid

Robin Tasker, Daresbury Laboratory

27 March 2002

Daresbury Laboratory
Council for the Central Laboratory of the Research Councils

EU DataGrid

- Objective : to enable next generation scientific exploration emerging in many scientific disciplines including physics, biology, and earth sciences.

- Requirements : intensive computation and analysis of shared large-scaled databases, from hundreds of Terabytes to Petabytes, across widely distributed scientific communities.

- Sharing is made complicated by,
  - distributed communities and resources;
  - the size of the databases; and
  - the limited bandwidth available.

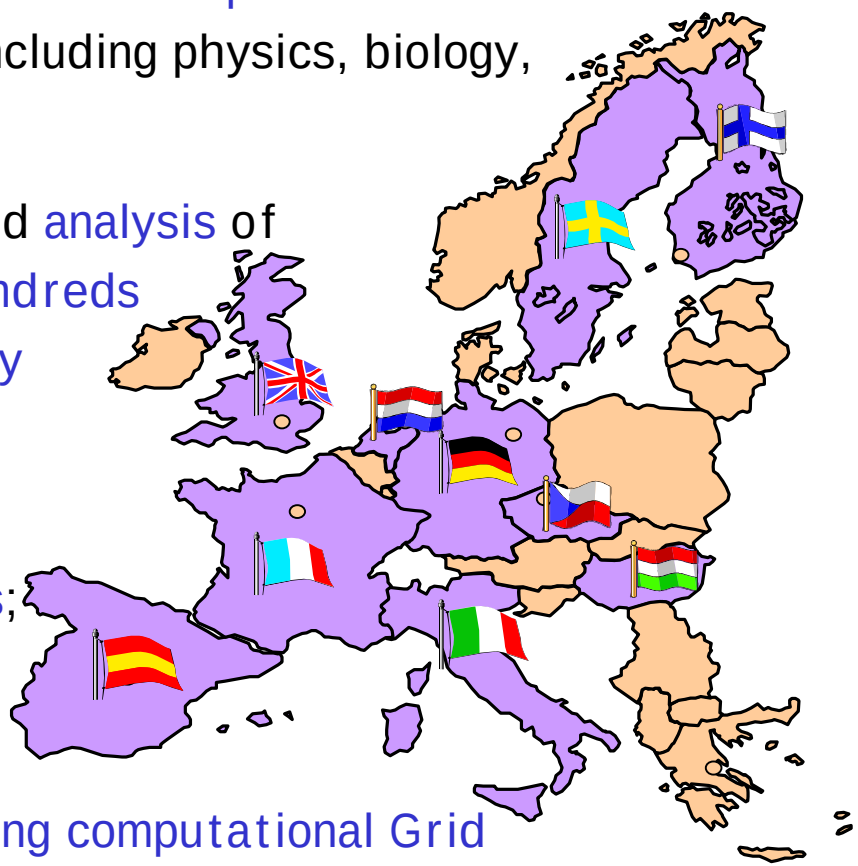

- The EU DataGrid project is using emerging computational Grid technologies to establish a research network to enable the development of the technology components for the implementation of a new worldwide DataGrid on a scale not previously attempted.

What's the aim and the purpose of monitoring for the Grid?

To characterise and quantify network behaviour,

- to inform Grid applications, via the middleware, of the current status of the network, i.e. a publication mechanism is needed
  - to provide source, destination, route choice etc
  - to twiddle TCP parameters, to adjust network QoS parameters

- to assist Grid users,
  - to identify fault conditions in the operation of the Grid
  - to provide input to network configuration and development of the Grid infrastructure.

to understand, and make generally available, the instantaneous, day-to-day, and month-by-month behaviour of the network.

Essential to relate monitoring of network performance to the real world.

Does the monitored view accord with the user/apps experience?

What does the Grid look like?

- Globus Grid architecture contains two layers in which higher-level services can be built on top of core services.

- Globus defines two server types which are used to publish Grid resource information
  - A GRIS contains information about a compute element or storage element.
  - A GIIS is an index of all the other GIISs and GRISs on the Grid.

- Consumers wishing to find out the situation on a particular host on the Grid query the local GIIS and are redirected to where that information resides

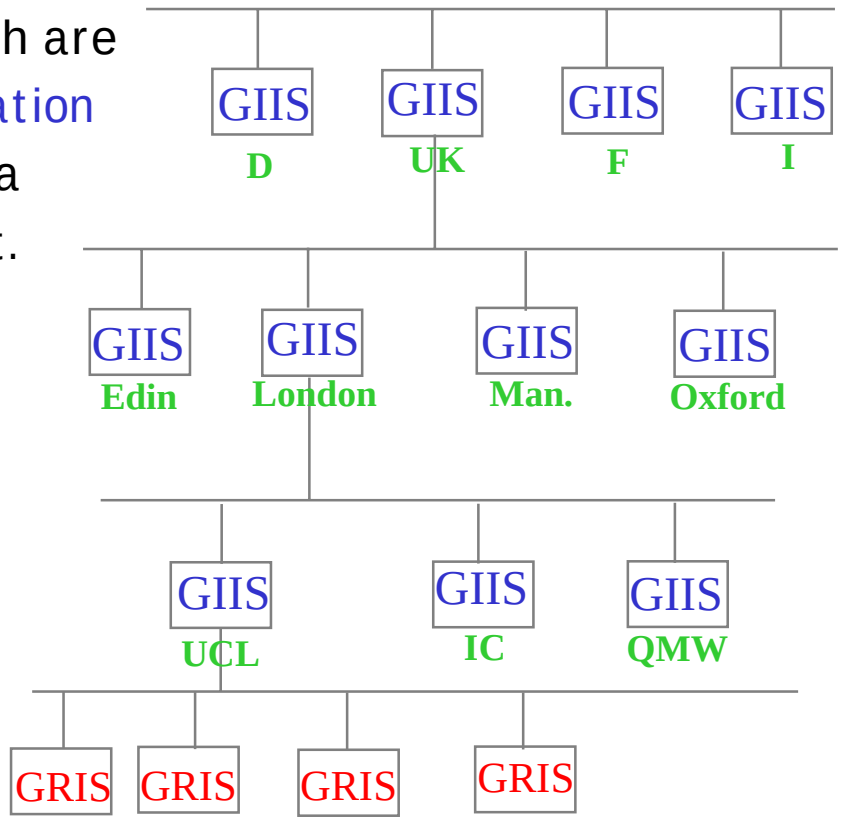

Network Monitoring Architecture

Grid Apps
GridFTP

PingER
(RIPE TTB)
iperf
rTPL
NWS
etc

Local Network Monitoring Store & Analysis of Data (Access)

Backend Information script to fetch metrics
Monitor process to push metrics

Information Schema

local Information Server

Grid Application access via Information Schema to
- monitoring metrics;
- location of monitoring data.

Access to current and historic data and metrics via the Web, i.e. WP7 NM Pages, access to metric forecasts

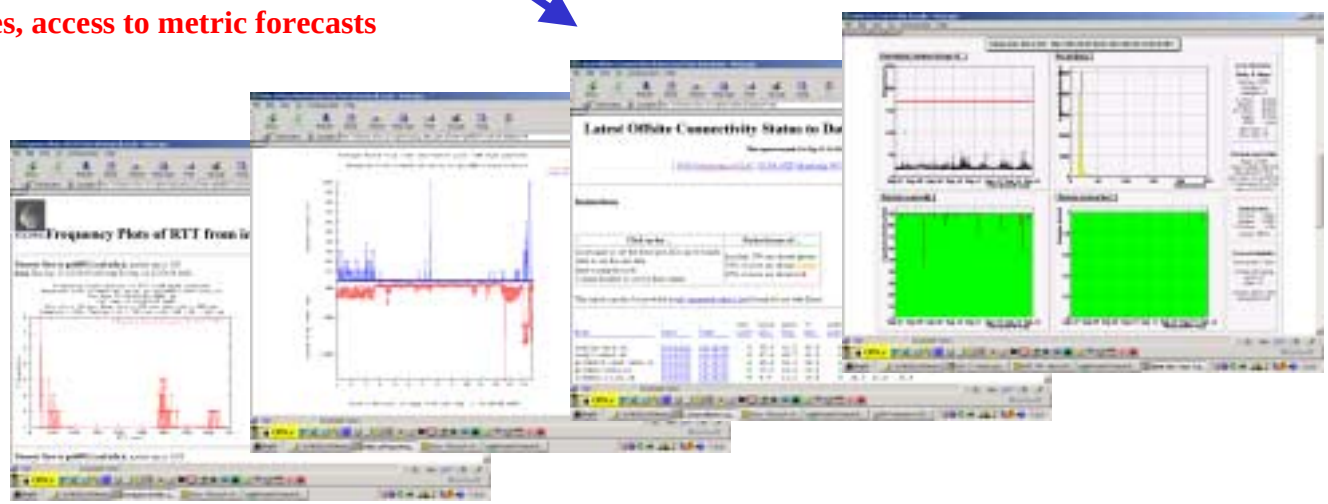

Network Monitoring Architecture

Grid Apps
GridFTP

PingER
(RIPE TTB)
iperf
rTPL
NWS
etc

Local Network Monitoring Store & Analysis of Data (Access)

Backend Information script to fetch metrics
Monitor process to push metrics

Information Schema

local Information Server

Grid Application access via Information Schema to
- monitoring metrics;
- location of monitoring data.

Access to current and historic data and metrics via the Web, i.e. WP7 NM Pages, access to metric forecasts

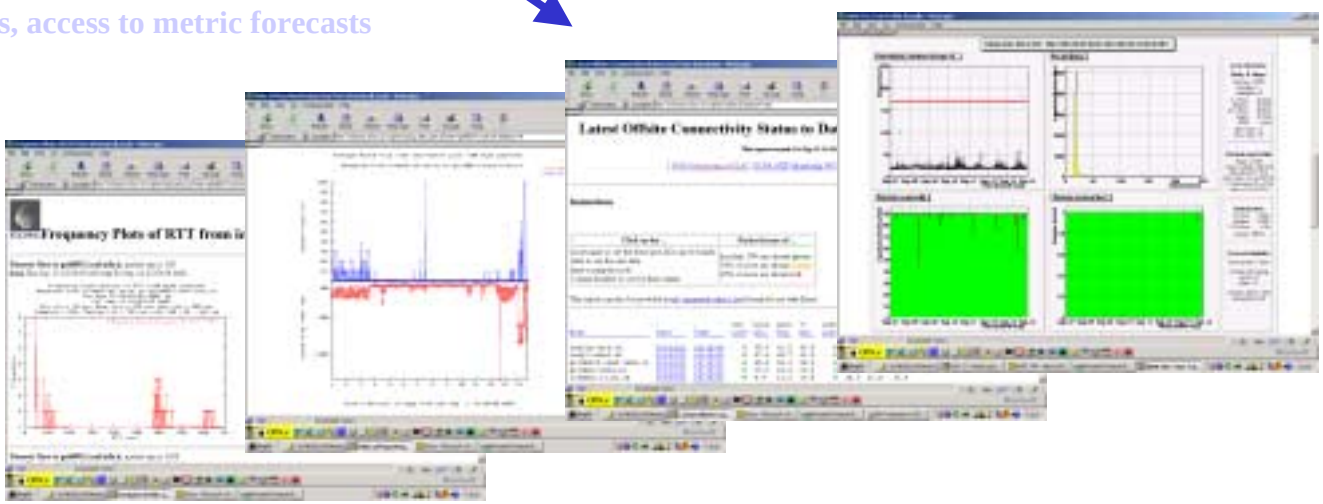

How to characterise a network
What are the metrics of interest?

Starting point of discussion was the IETF IP Performance Monitoring (IPPM) WG output, i.e. RFC 2330, but extended it to include,

Network stuff,

- Packet loss
- One-way delay
- Round Trip Time (RTT), i.e. there and back
- Jitter
- RTT variation (frequency distribution)

Grid stuff,

- Volume, e.g. number of Grid bytes exchanged
- Per flow application throughput
- Aggregate network throughput

What's available for monitoring?

- Basic tools

  traceroute, ping, traceping, pchar, MRTG, netperf, etc..

- RIPE NCC Traffic Test Monitoring box (also Surveyor)

  one-way delay, packet loss, jitter, route information

- PingER Software

  round trip time and frequency, packet loss

  traceroute facility (includes NIKHEF traceroute)

- Throughput measurement tools

  iperf, netperf, rTPL, UDPmon

- Grid applications, e.g. GridFTP, can keep statistics of activities

Network Monitoring Architecture

Grid Apps
GridFTP

PingER
(RIPE TTB)
iperf
rTPL
NWS
etc

Local Network Monitoring Store & Analysis of Data (Access)

Backend LDAP script to fetch metrics
Monitor process to push metrics

LDAP Schema

local LDAP Server

Grid Application access via LDAP Schema to
- monitoring metrics;
- location of monitoring data.

Access to current and historic data and metrics via the Web, i.e. WP7 NM Pages, access to metric forecasts

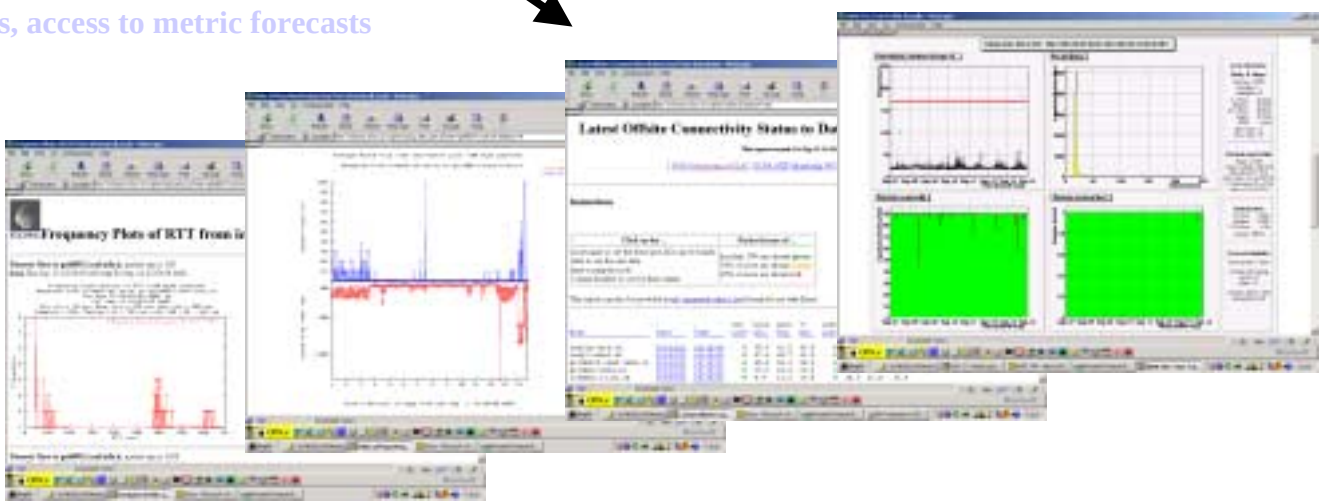

How do the Grid apps access the metrics of network monitoring?

What is the RTT viewed from UCL to RAL and QMW?

Query an LDAP server that makes use of an LDAP Schema containing the ObjectClass RTT to find out.

How?

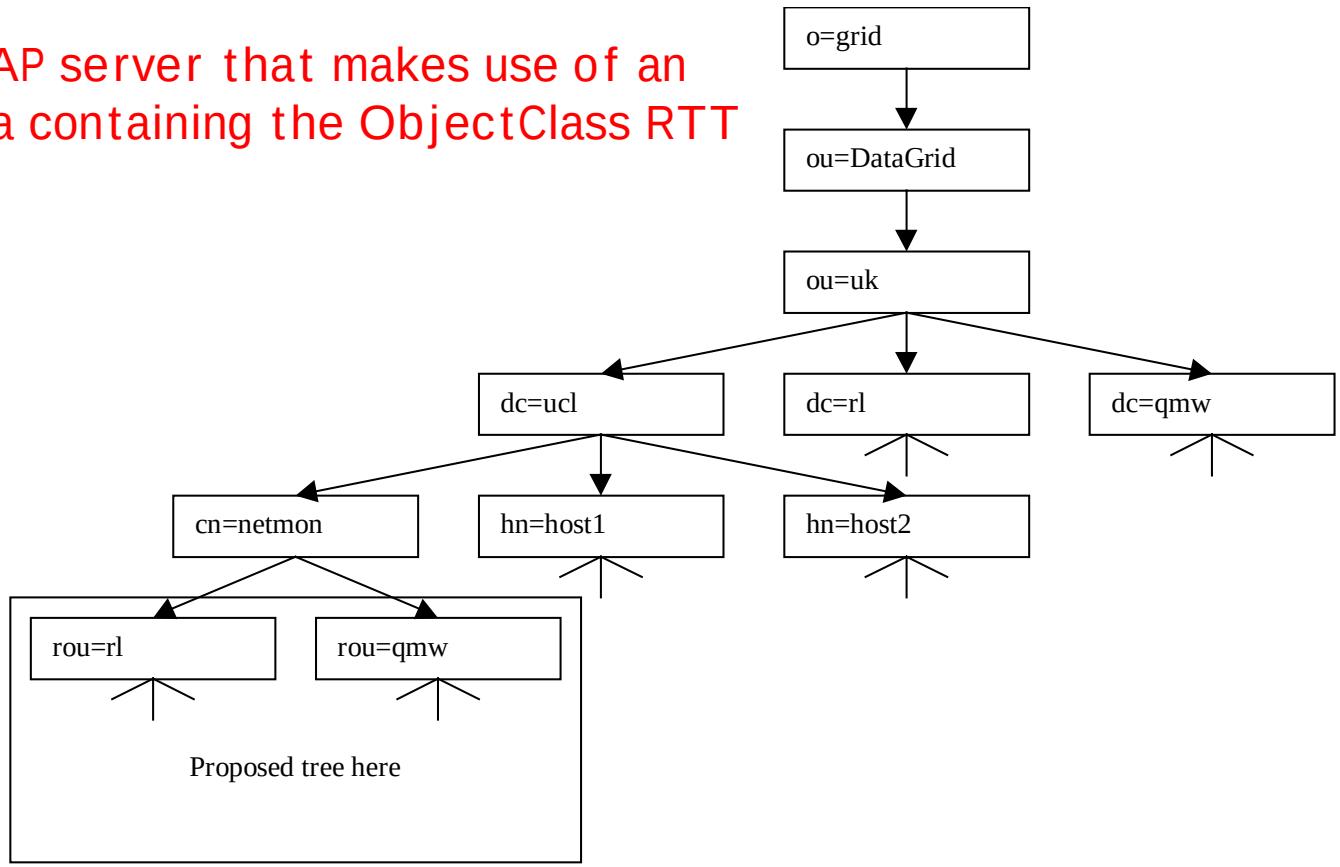

and find the details from the LDAP Schema for publication of the RTT metric

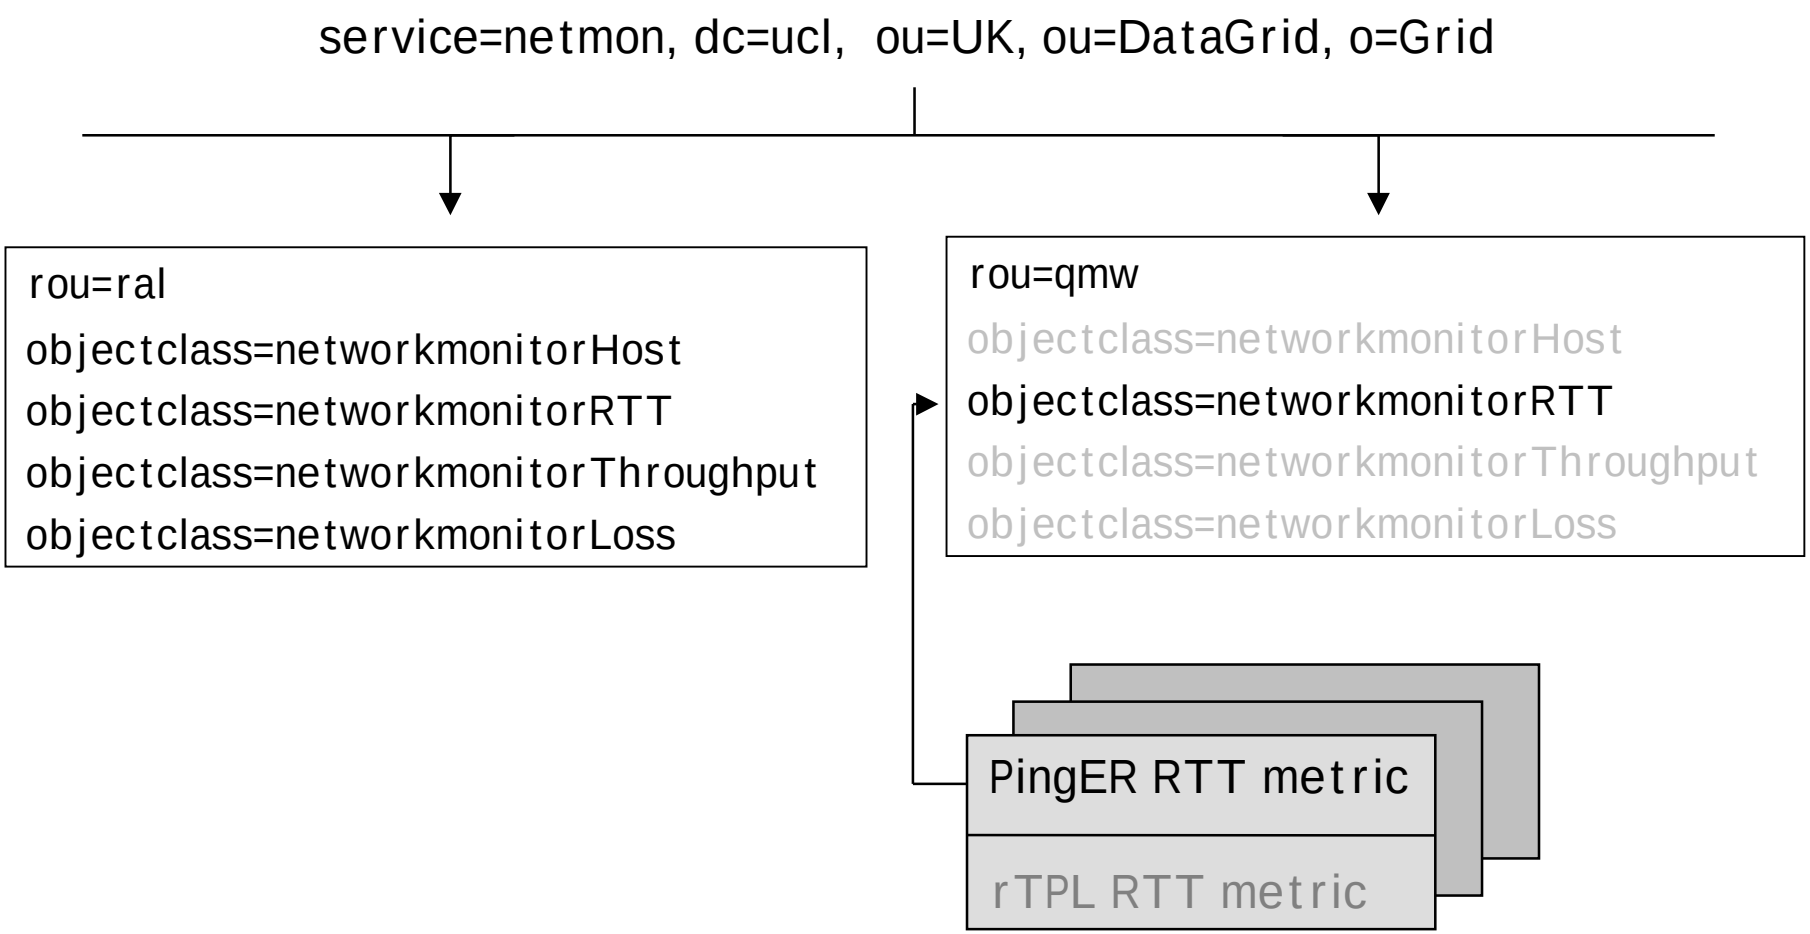

And the mechanics of it all

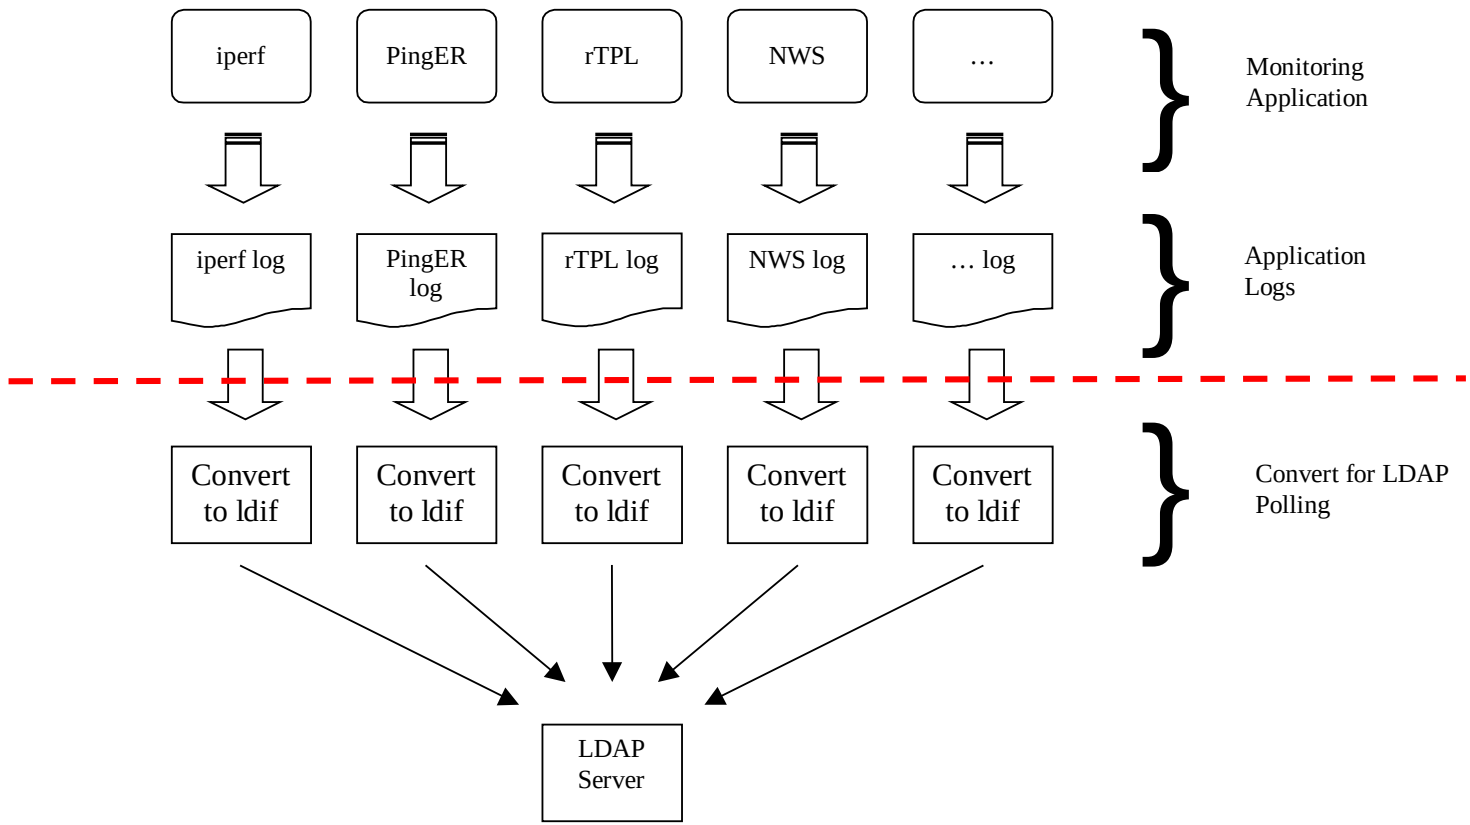

to allow a request to return a valid metric

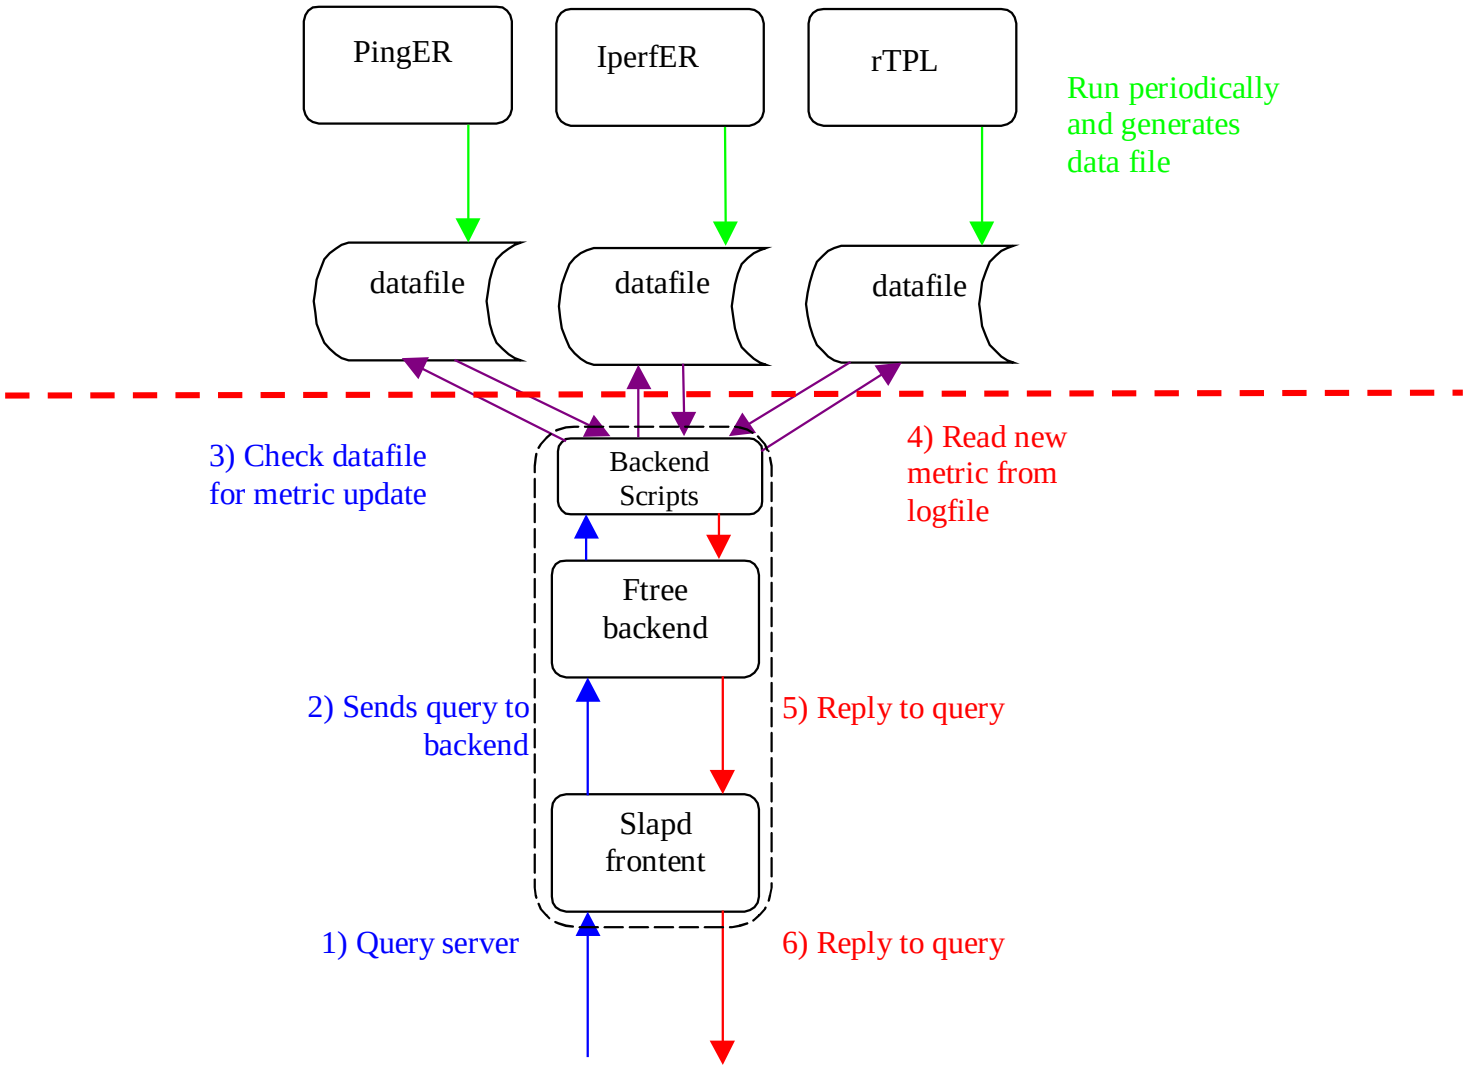

Network Monitoring Architecture

Grid Apps
GridFTP

PingER
(RIPE TTB)
iperf
rTPL
NWS
etc

Local Network Monitoring Store & Analysis of Data (Access)

Backend LDAP script to fetch metrics
Monitor process to push metrics

local
LDAP
Server

LDAP
Schema

Grid Application access via
LDAP Schema to
- monitoring metrics;
- location of monitoring data.

Access to current and historic data and metrics via the Web, i.e. WP7 NM Pages, access to metric forecasts

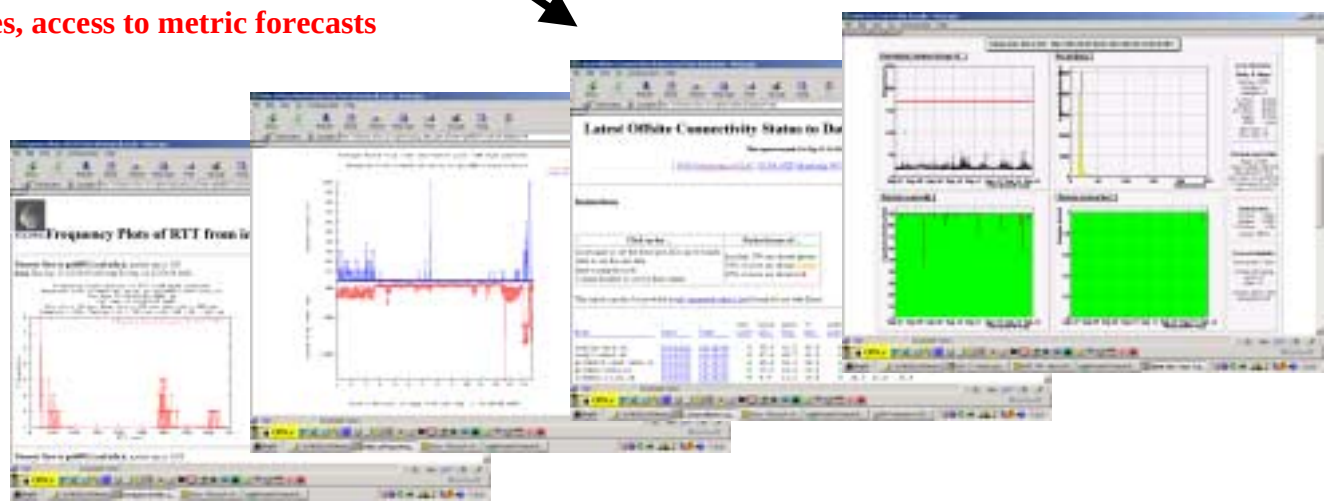

So how does a user access the metrics of network monitoring?

Network Monitoring of  **sites for WP7 from Daresbury**

- View connectivity to the DataGRID test sites for today or for the previous 7 days
- Plot **Today's** Connectivity from **Daresbury** to | Bologna | CERN | Lyon | Rutherford | SARA |
- Plot Previous **Fortnight's** Connectivity from **Daresbury** to | Bologna | CERN | Lyon | Rutherford | SARA |
- Traceroute from Daresbury back to your machine
- Review iperf measurements from UCL to WP7 sites here
- View the rTPL monitoring of WP7 sites here - (Beware! the Back button on your browser won't navigate you back to this page!)
- Visit the other DataGrid WP7 Network Monitoring sites at | Bologna | CERN | Lyon | Rutherford | SARA |

Return to Daresbury Network Monitoring Page

Last updated on 2 October 2001, Robin Tasker

Here's some PingER output

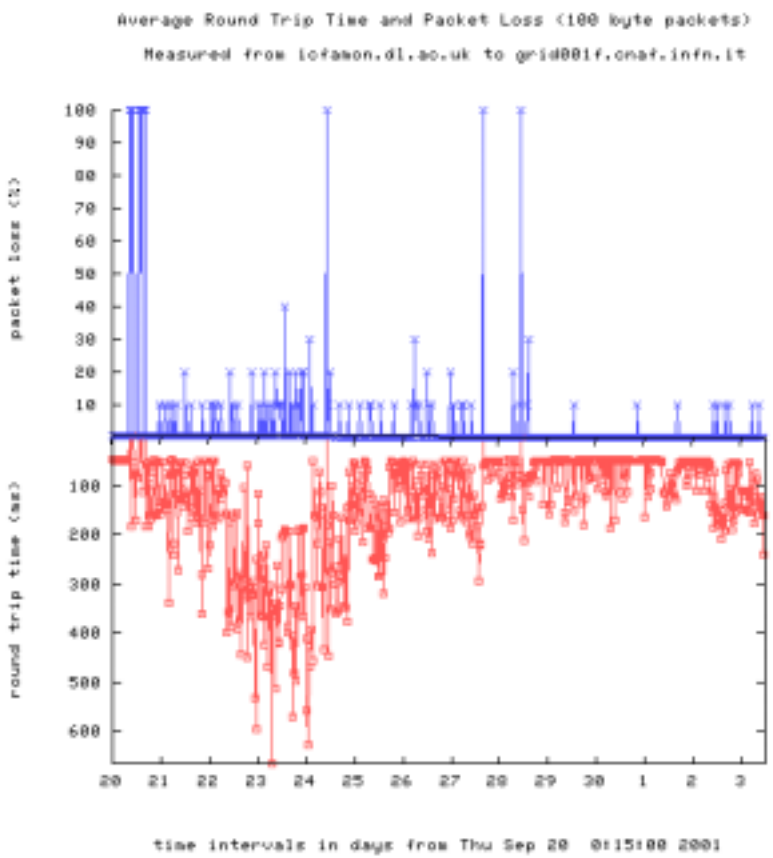

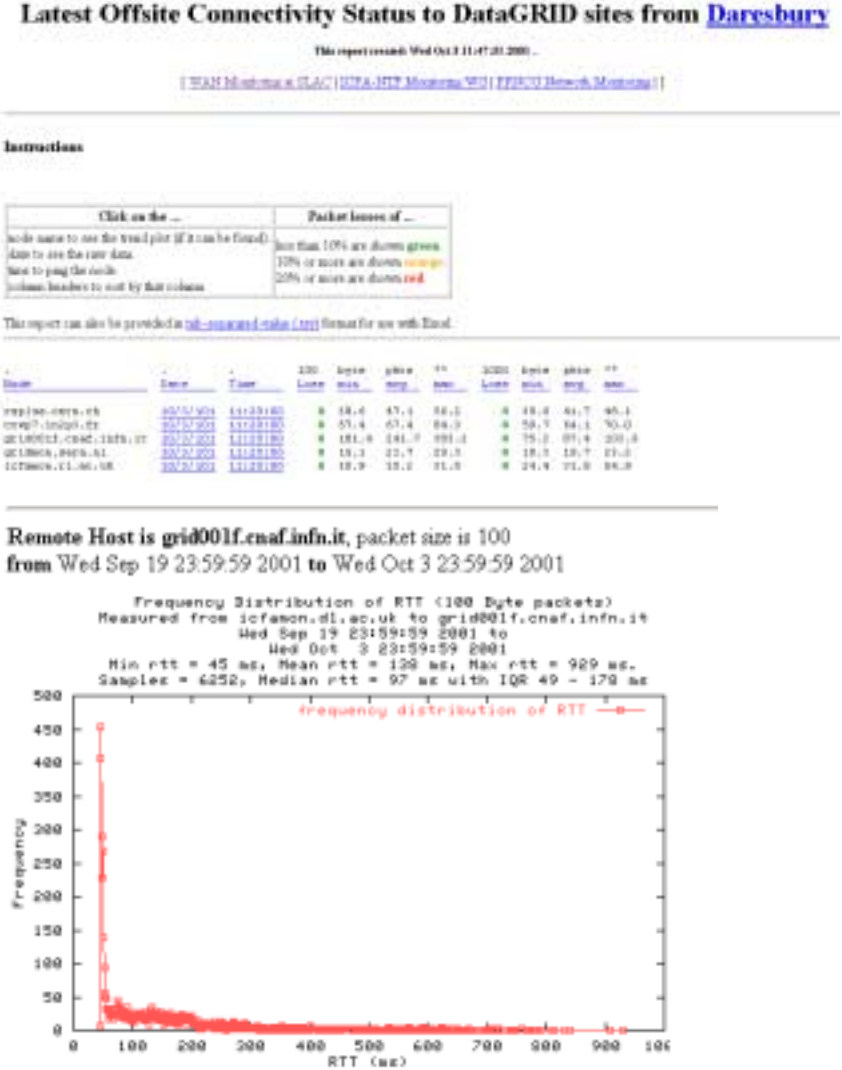

.....and the same from rTPL.....

Ping Min [ms]
(row >> column)

	CLRC	UCL	ENS	CERN	CNAF	SARA
CLRC	---	7.200	90.000	71.900	81.300	53.100
UCL	7.184	---	---	---	---	45.216
ENS	0.000	---	---	19.800	35.500	40.900
CERN	72.000	---	19.000	---	16.000	21.000
CNAF	85.700	---	36.100	16.100	---	34.800
SARA	49.900	44.700	41.000	21.400	35.200	---

Throughput [Mbit/s]
(row >> column)

	CLRC	UCL	ENS	CERN	CNAF	SARA
CLRC	---	17.45	0.02	.83	.39	2.14
UCL	22.86	---	---	---	---	7.97
ENS	.28	---	---	22.93	4.41	7.76
CERN	0.01	---	23.04	---	5.19	11.14
CNAF	.31	---	7.05	6.71	---	7.18
SARA	0.06	.17	10.4	18.06	4.09	---

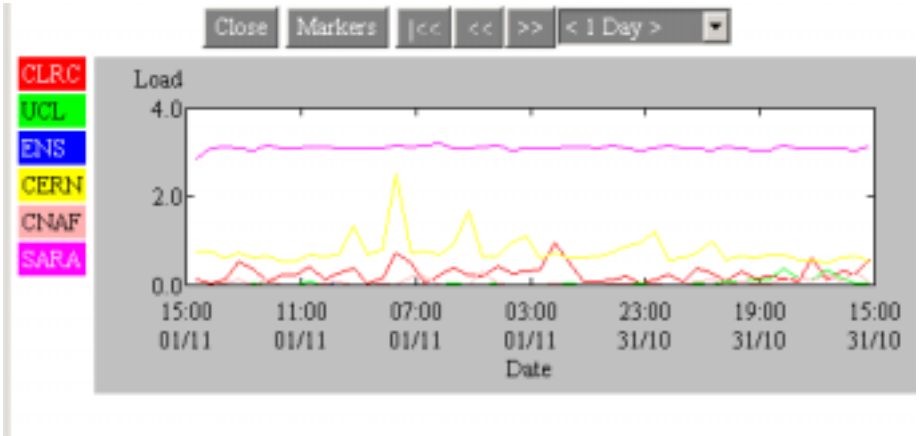

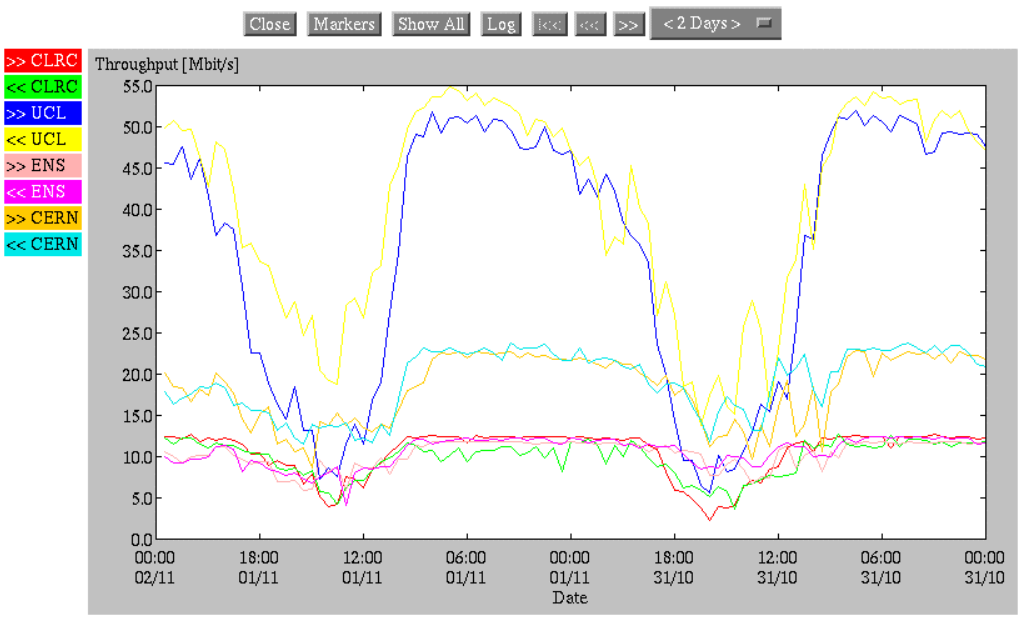

Use mouse to zoom in.

...and the RIPE NCC TTM box output...

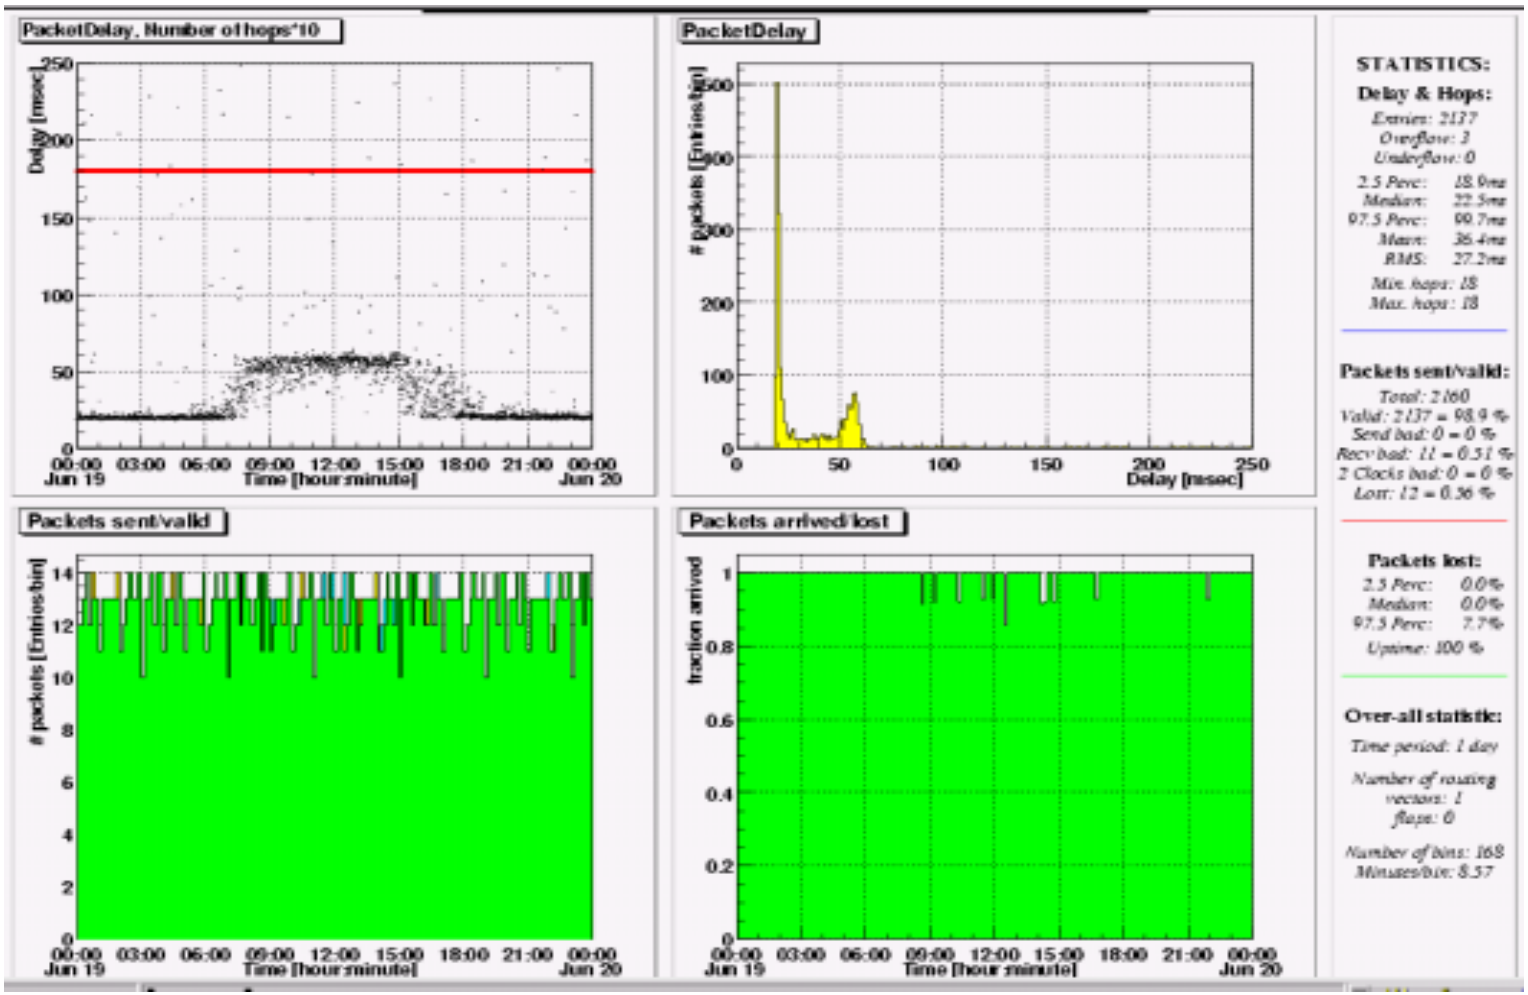

...and IperfER output...

Remote Host is ccwp7.in2p3.fr
buffersize is 262144
from Mon Sep 24 14:37:16 2001 **to** Tue Oct 2 15:00:01 2001
There is a time interval every day

View bandwidth frequency plot of this data here

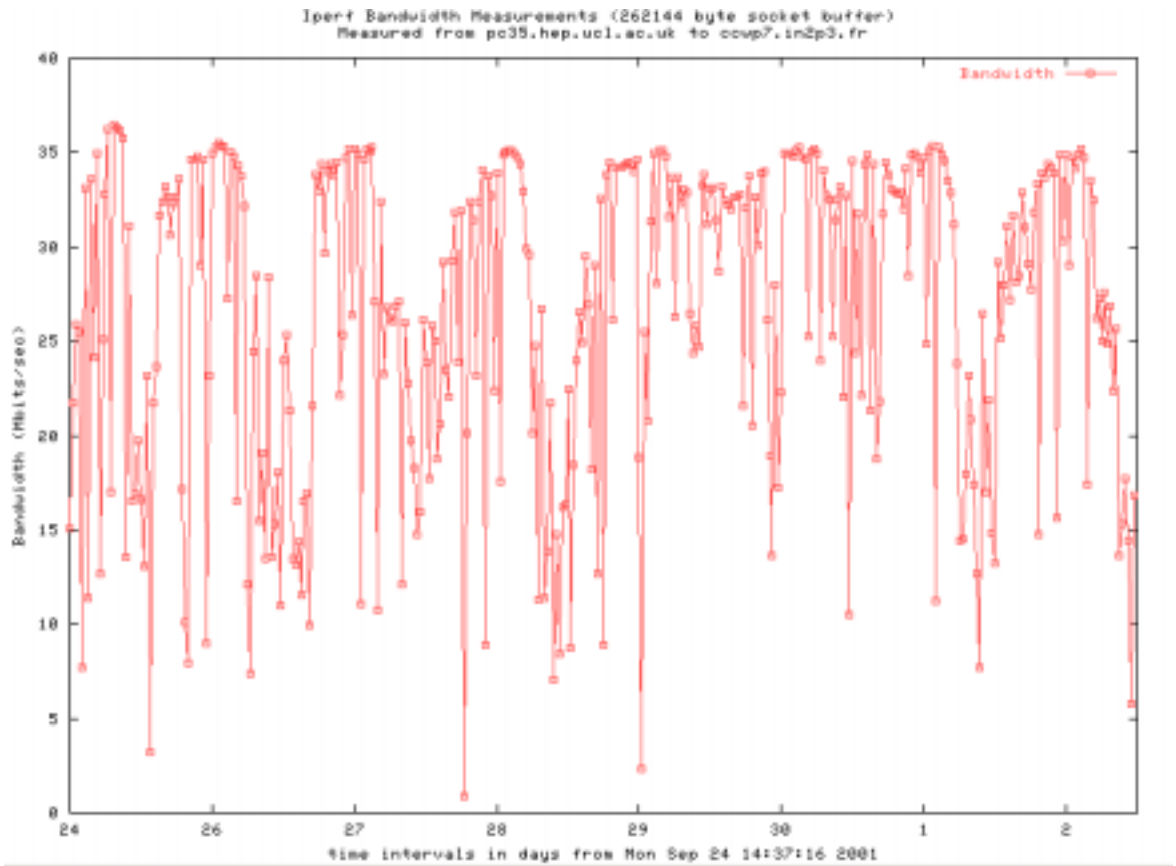

...and finally UDPmon.

UDPmon packet loss
Man – SARA
300 * 1400 byte frames

UDPmon throughput Mbit/s
Man – SARA
300 * 1400 byte frames

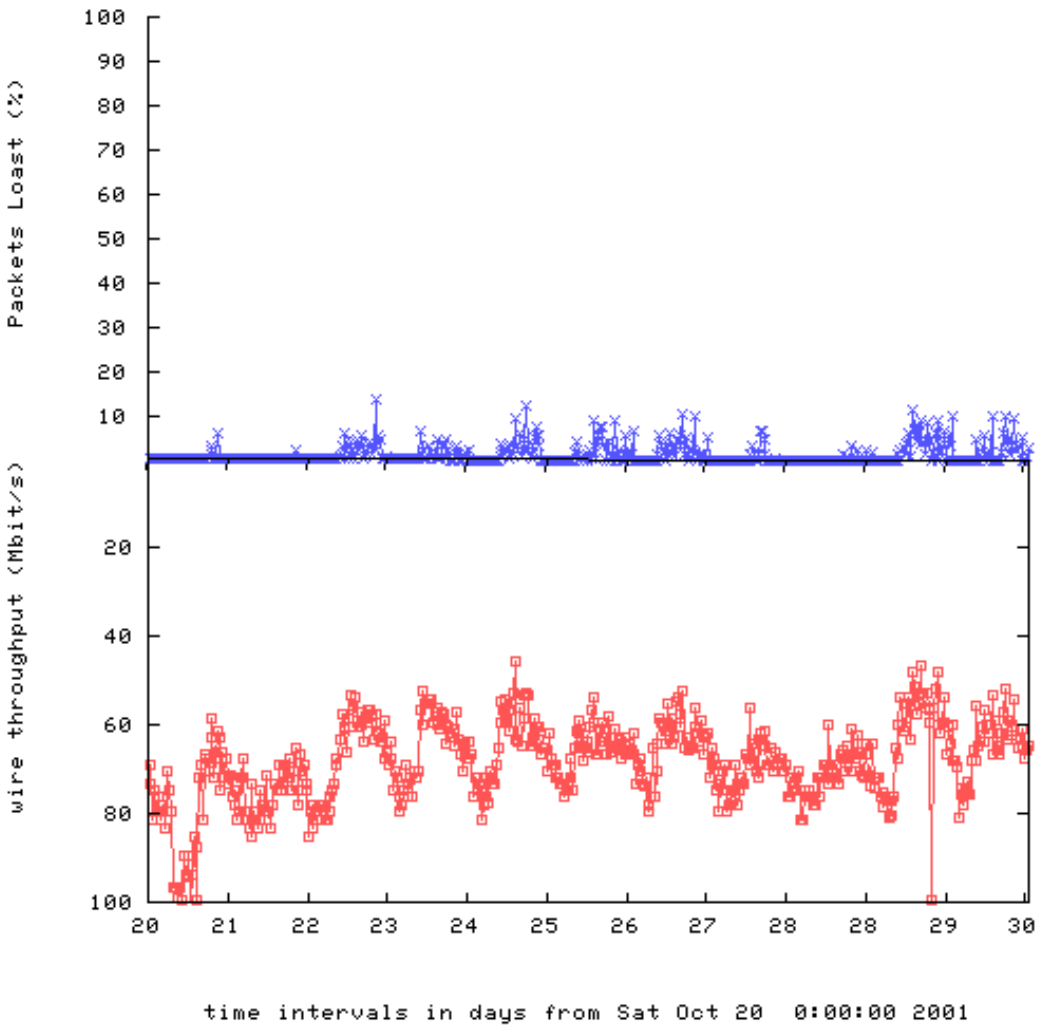

And the synthesis of all this monitoring http://ccwp7.in2p3.fr/