JRA4 Network Performance Monitoring: Demonstration

Alistair Phipps (A.Phipps@nesc.ac.uk)
Andrew Jackson (A.Jackson@epcc.ed.ac.uk)
The University of Edinburgh

The need for NPM (1)

- Network performance data is important for:
  - Network Operations Centres (NOCs) and Grid Operations Centres (GOCs) to detect and resolve network problems.
  - Grid middleware to intelligently schedule jobs based on network load and reliability.

The need for NPM (2)

- JRA4/NPM provides uniform access to network performance information from this heterogeneous set of monitoring frameworks.
  - Network monitoring frameworks exist to gather network performance data - both from end-site to end-site and between backbone routers - but no standard mechanism existed to allow data from different frameworks to be utilised.

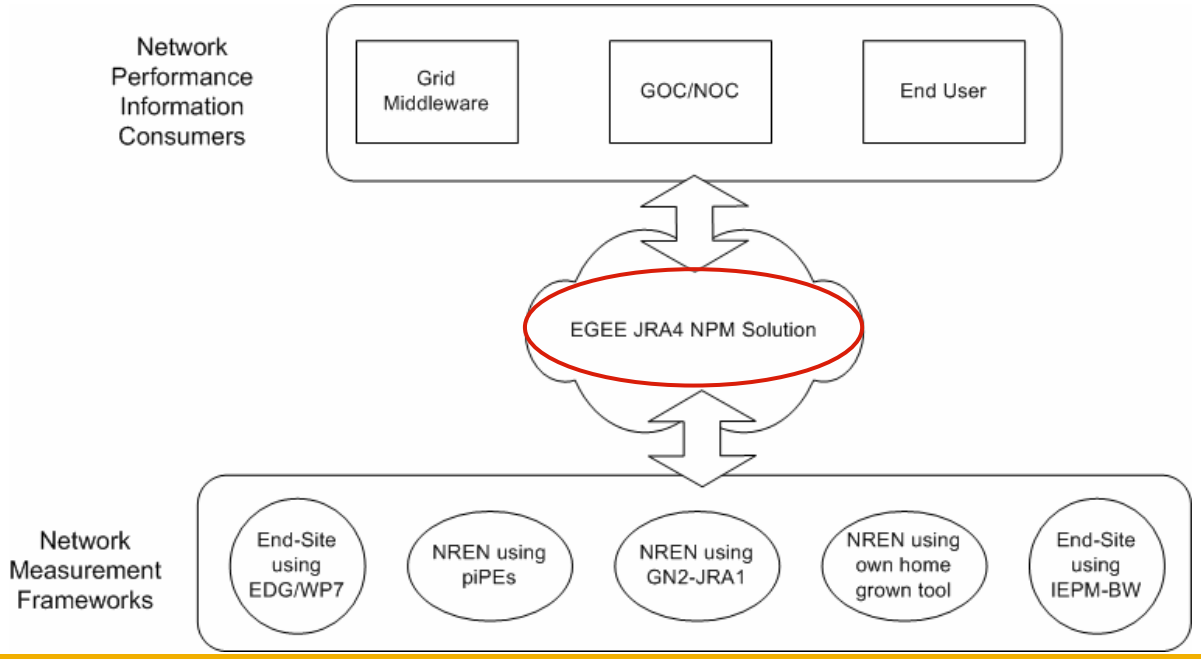

NPM Diagnostic Tool Demonstration

NPM DT Scenario (1)

- **The following slides demonstrate a DT usage scenario.**
- **Scenario:**
  - An end-user has complained to a CIC about intermittent slow GridFTP transfers between UEDIN and CNRS
  - The problem only seems to happen at certain times of day, but the end-user can't give a specific time - though it definitely occurred within the past two days (today is 21/10/2005)
  - The CIC needs to determine if it's a network problem, or something to do with the end host (perhaps it's overloaded or misconfigured)

NPM DT Scenario (2)

- Step 1: Access the NPM Diagnostic Tool.
  - The Diagnostic Tool can be accessed using a standard web browser, which users are individually authorised to use.
  - In the future, we plan to use VOMS for authorisation.

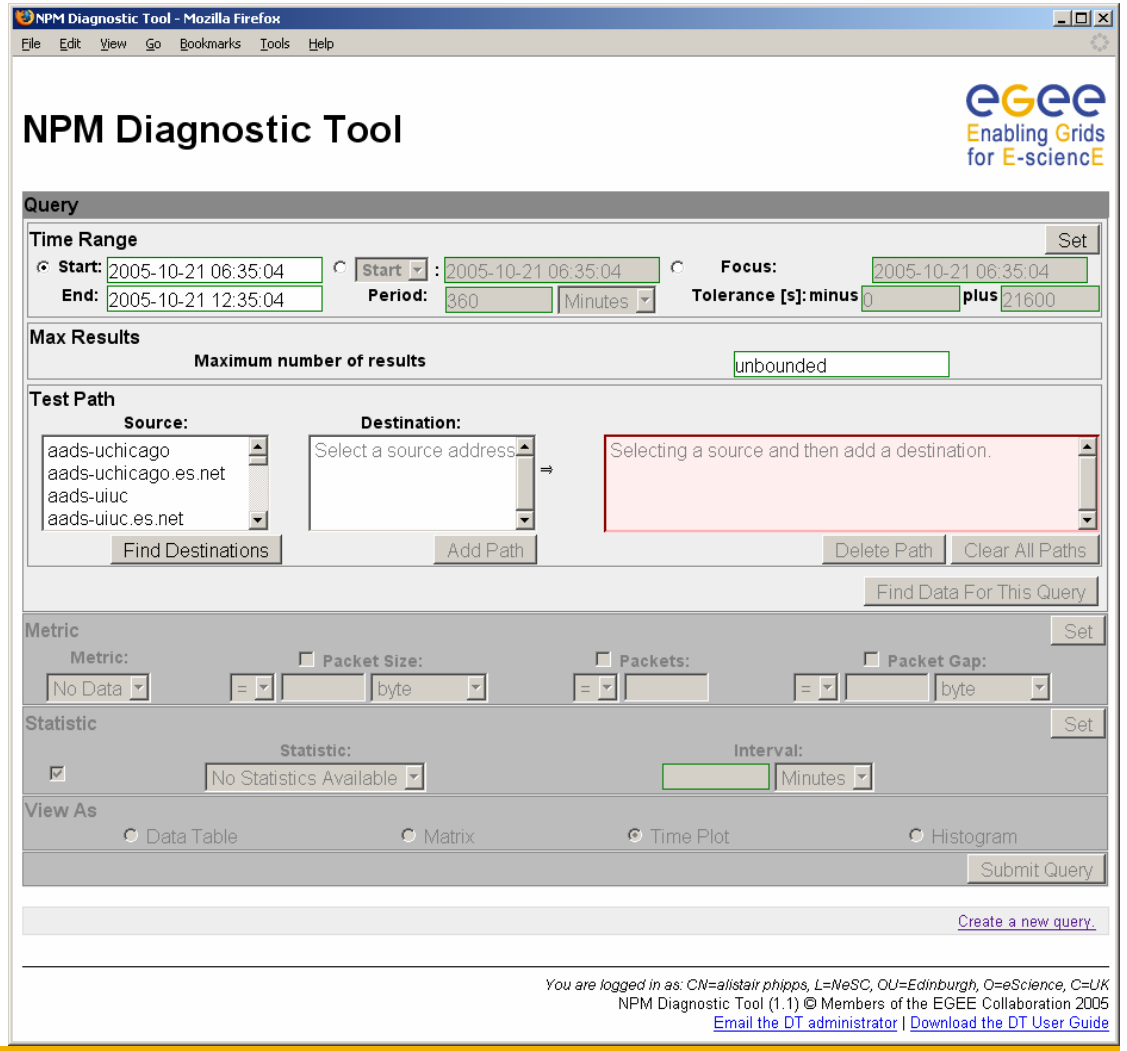

NPM DT Scenario (2)

- Step 2: Select a Time.
  - –The end-user does not have a specific time, but knows the problem occurred within the past two days.
  - –The CIC-user enters the appropriate time range, specifying an End date/time of 2005-10-24 14:30:00 (the current time), and a period of 2 days.
  - –The CIC-user presses the Set button to confirm and the alternate time range representations update.

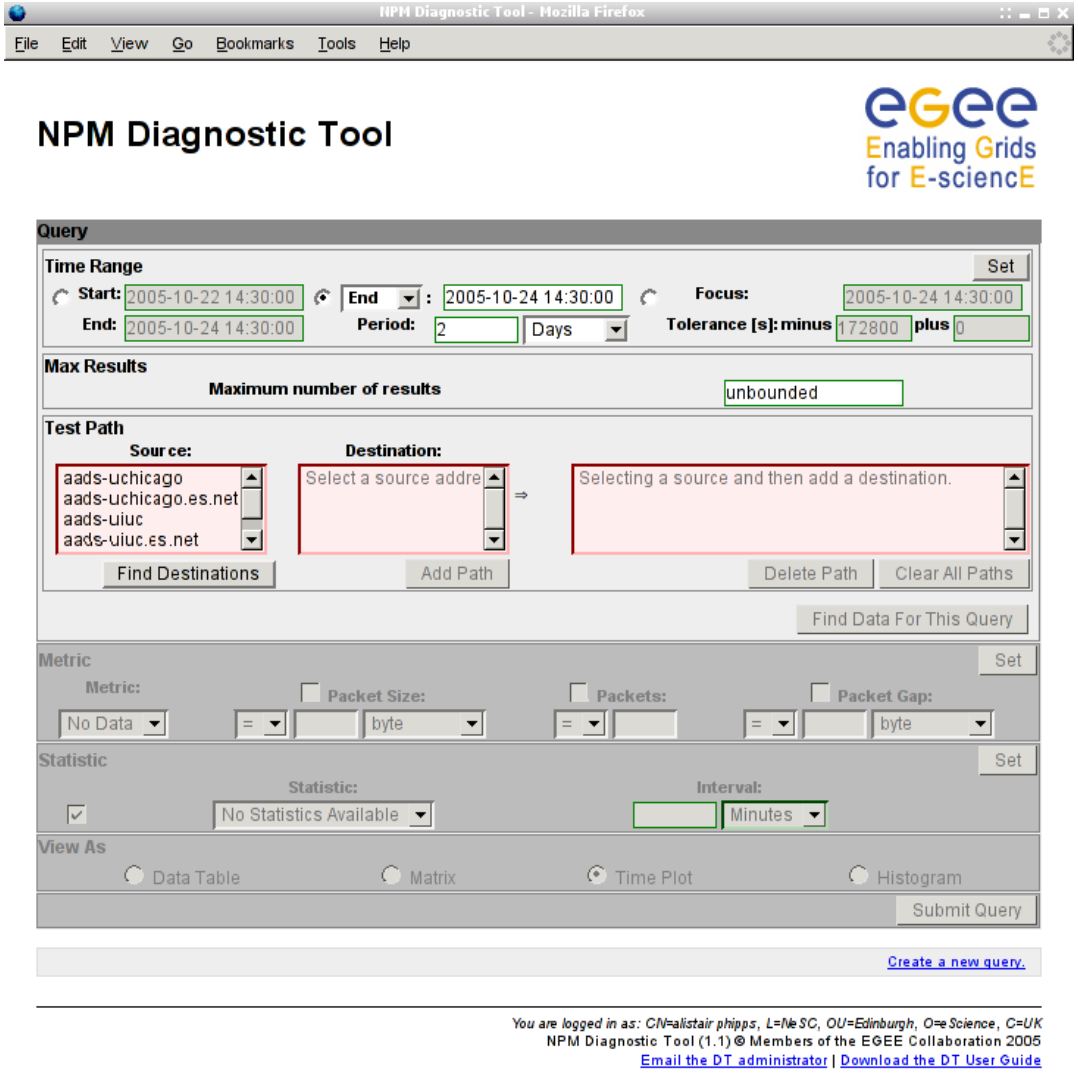

NPM DT Scenario (3)

- Step 3: Select a Path.
  - –The end-user experienced the problem between UEDIN and CNRS.
  - –The CIC-user selects e2emonit sites at UEDIN and CNRS, adds the path and then selects "Find Data For This Query“
    - •E2emonit, formerly known as EDG/WP7, is an end site-to-end site monitoring framework

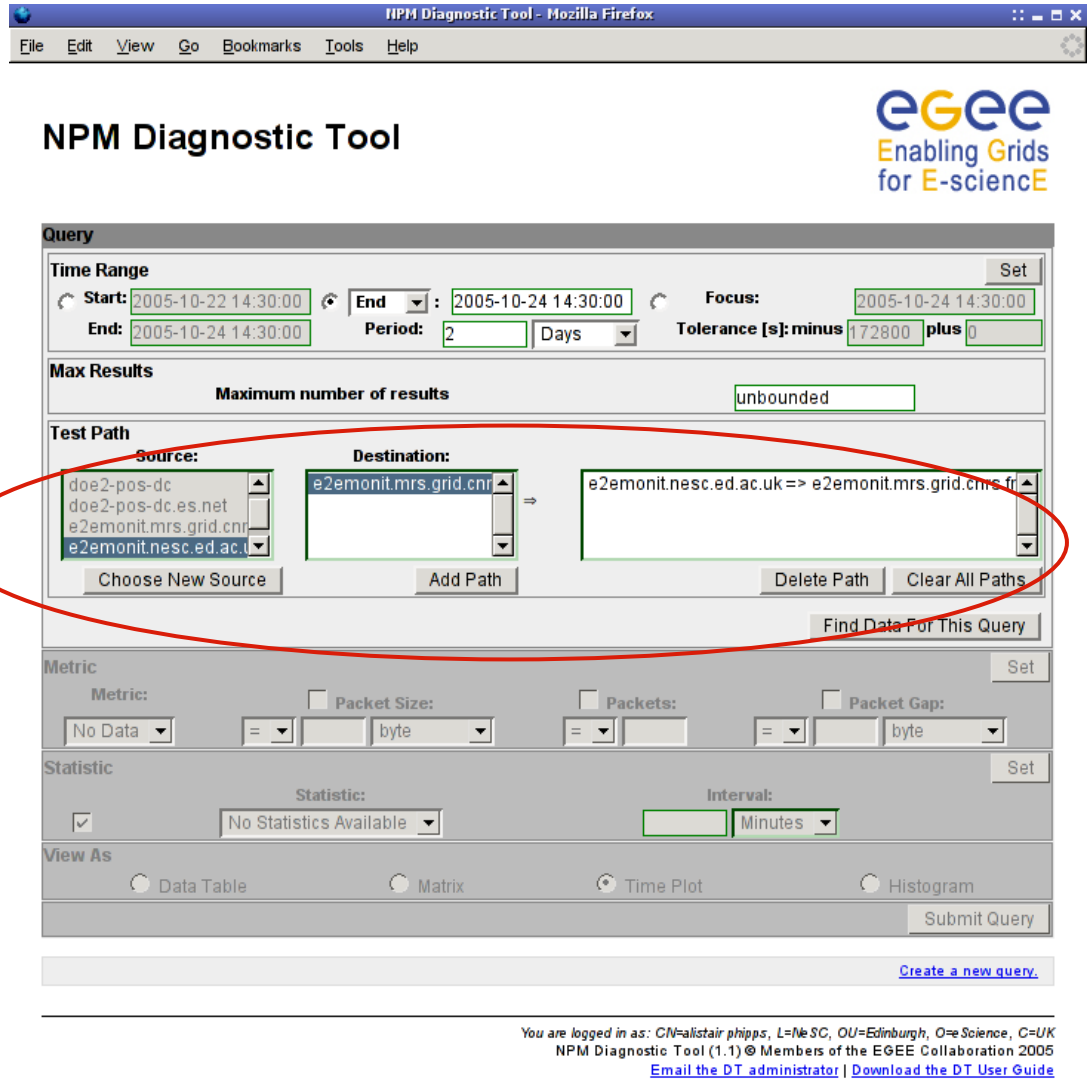

NPM DT Scenario (4)

- Step 4: Select a Metric.
  - –The end-user experienced throughput problems.
  - –Although there are several possibly relevant metrics to choose from (and only those measured are available to select from), the CIC-user decides to look at the Achievable Bandwidth on the path.
  - –Achievable Bandwidth is selected from the Metrics box and the Set button pressed to confirm.

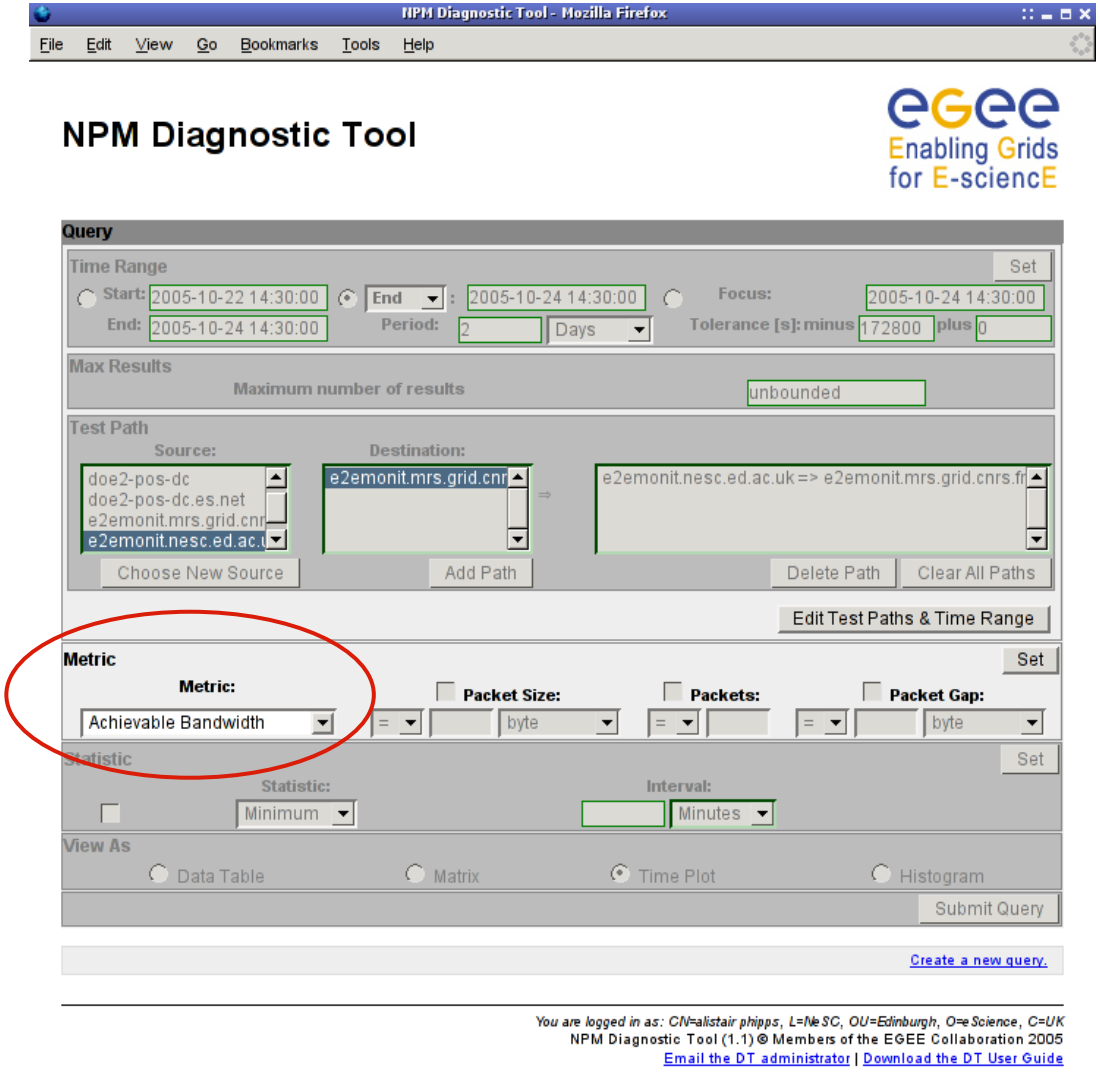

NPM DT Scenario (5)

- Step 5: Select a Statistic.
  - –Several types of statistical data are available, such as Minimum, Maximum, Mean.
  - –A particular interval can be applied to each, to provide, for example, an hourly mean over the past two days.
  - –The CIC-user just wants a general overview of measurements and elects to retrieve raw data (Statistic check-box not checked).

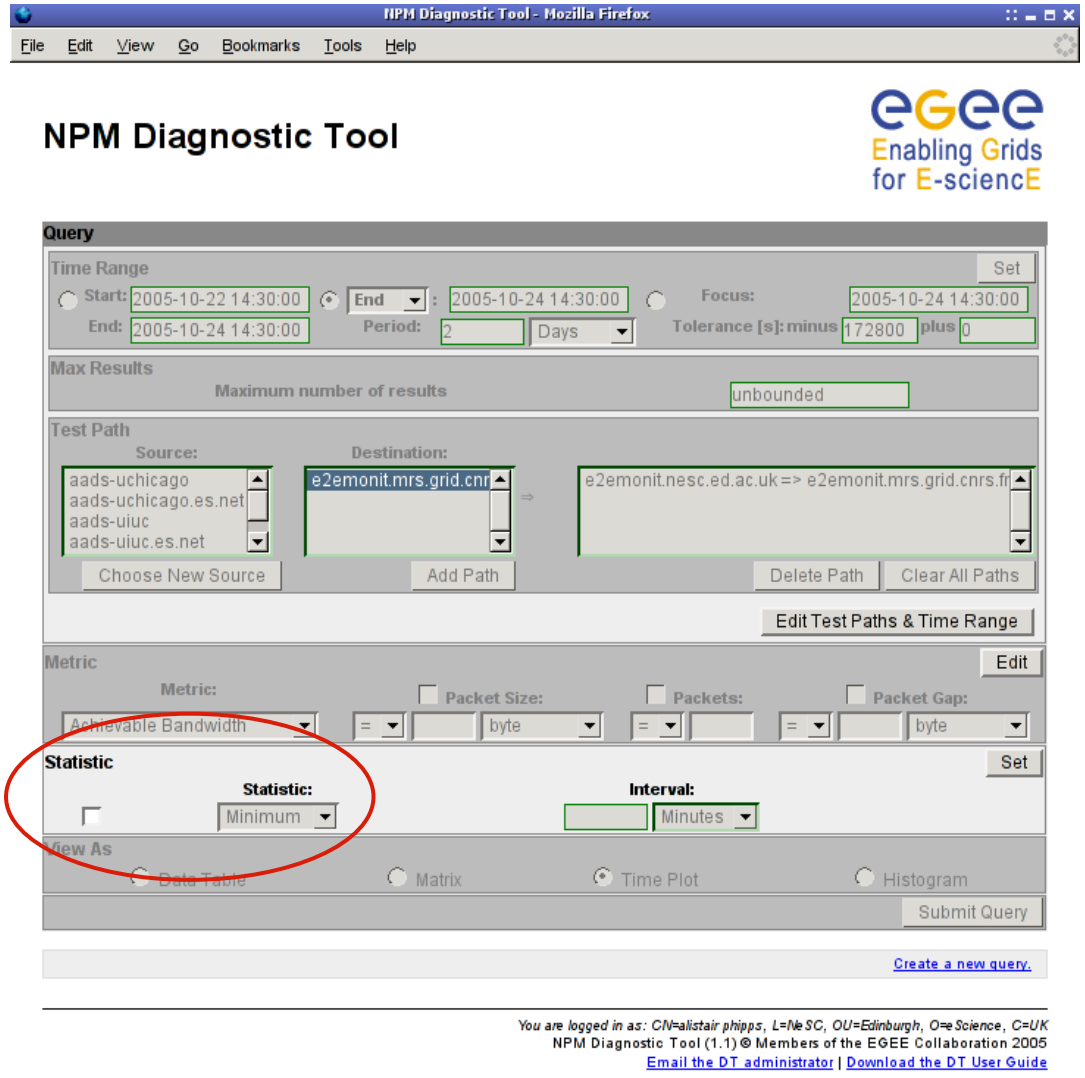

NPM DT Scenario (6)

- Step 6: Select a View.
  - –Currently Data Table and Time Plot views are available.
  - –The CIC-user wants an overview of how the Achievable Bandwidth has changed over time, so selects the Time Plot.
  - –The Query entry is complete, and the CIC-user selects Submit Query.

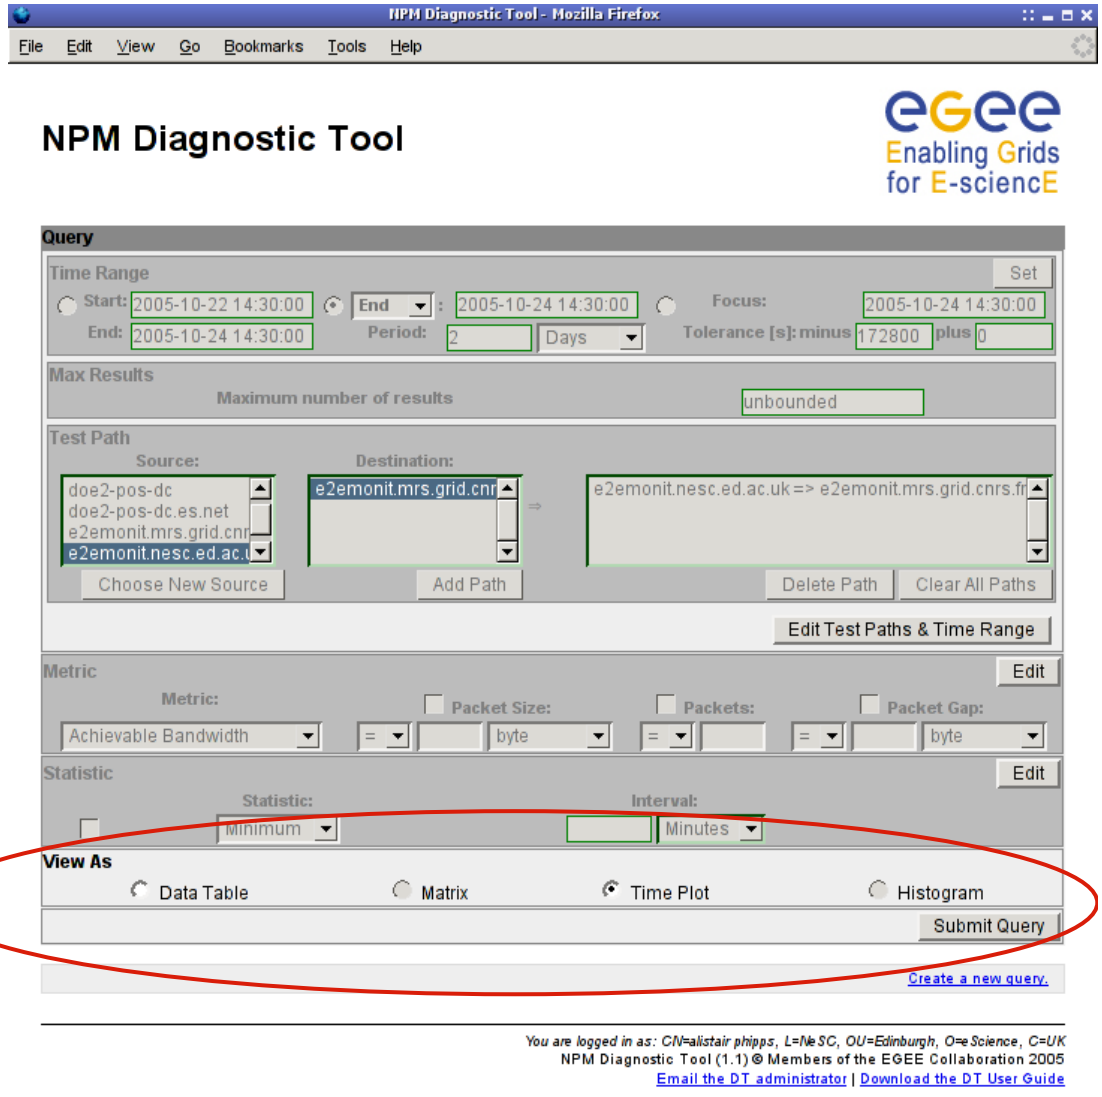

NPM DT Scenario (7)

- Step 7: Examine results.
  - The results are plotted, with Time on the x-axis and Achievable Bandwidth on the y-axis.
  - The parameters used to gather measurements are shown - here, showing that the iperf tool was used to gather the achievable bandwidth information.
  - These parameters can be useful in interpreting the results.

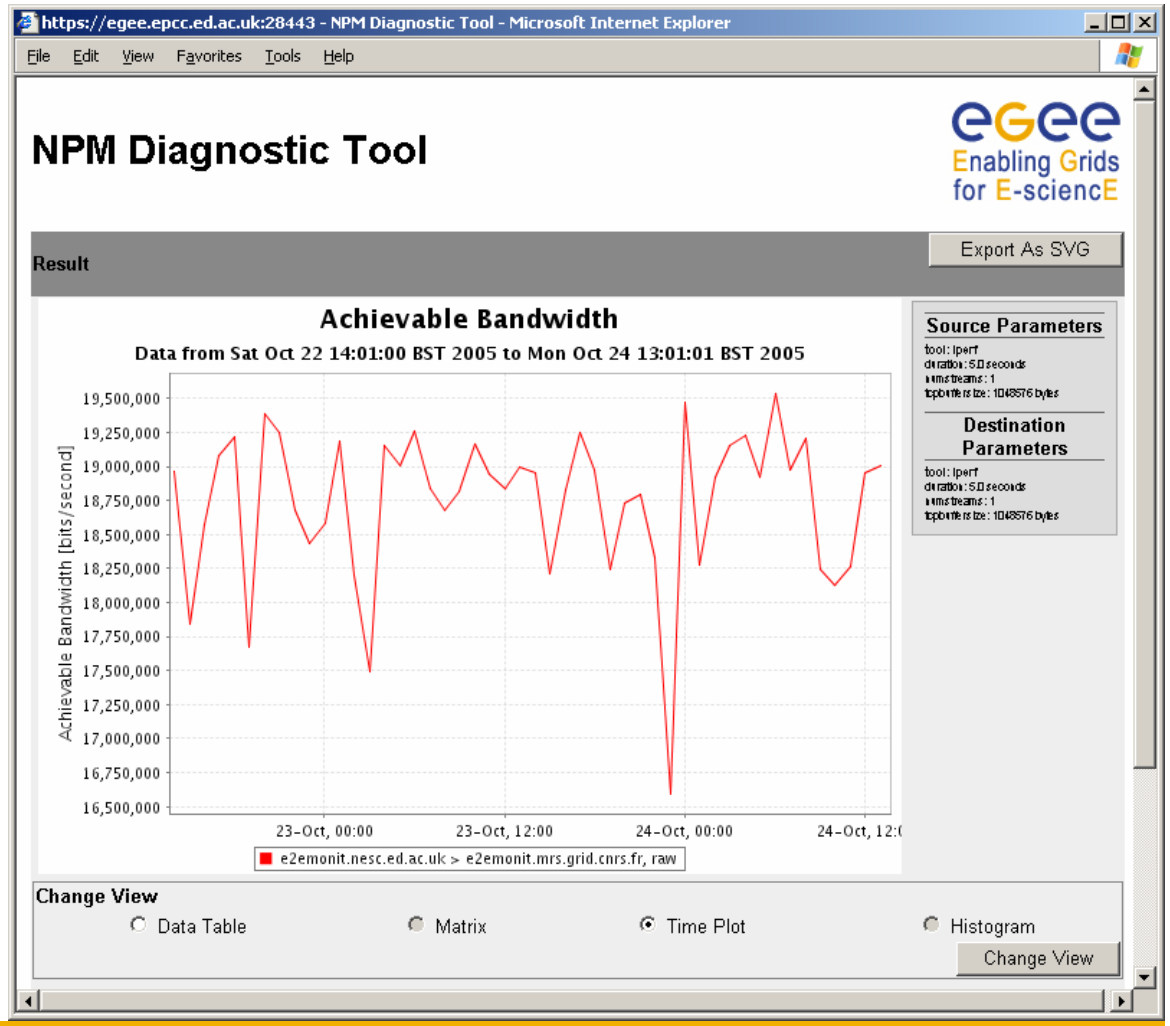

NPM DT Scenario (8)

- Step 7: Examine results (continued).
  - Our user was reporting intermittent poor throughput between these two sites.
  - The results show that the achievable bandwidth is fluctuating, though perhaps no more than should be expected.

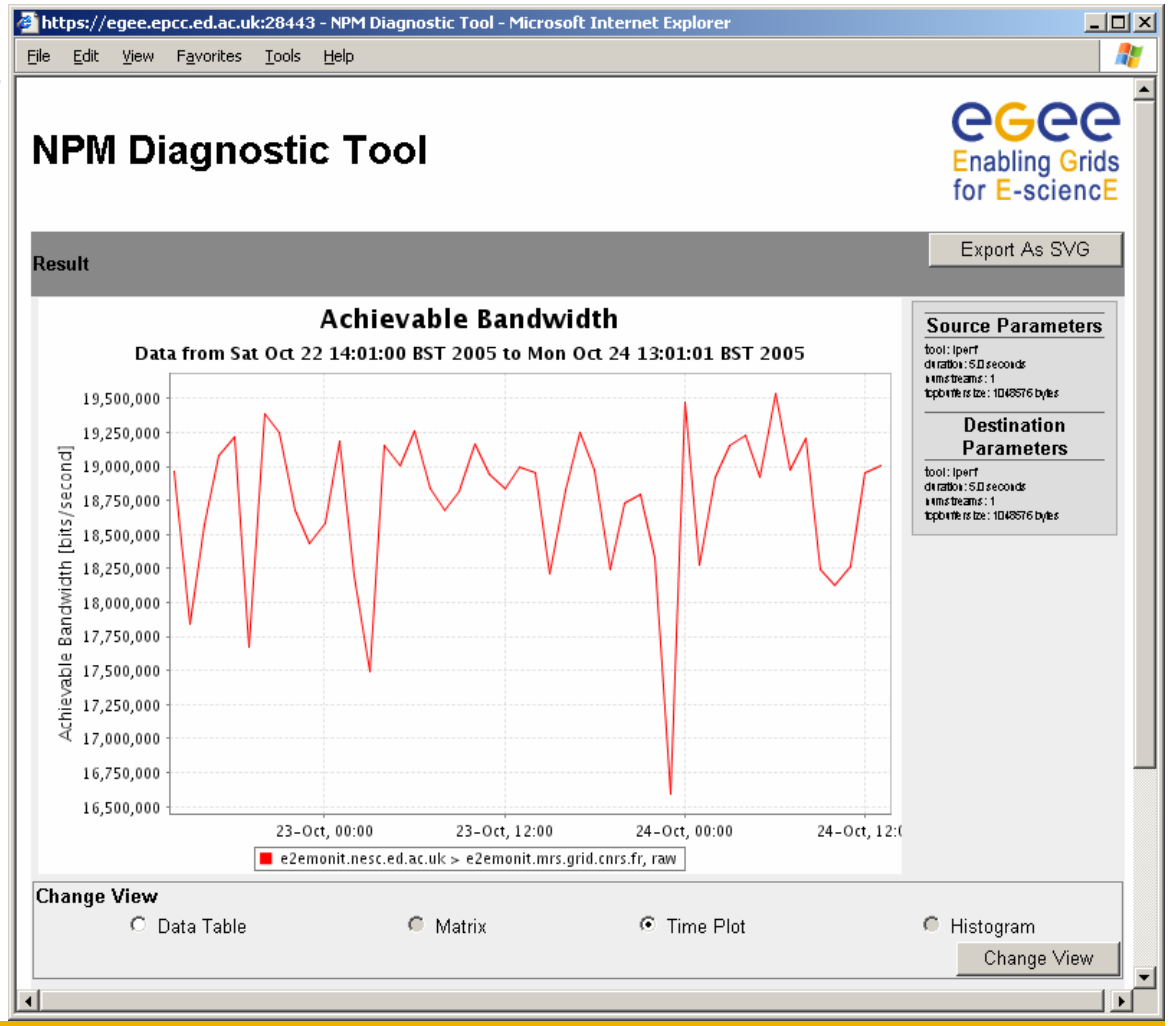

NPM DT Scenario (9)

- Step 7: Examine results (continued).
  - –The CIC-user may wish to examine further by:
    - –Checking tabulated achievable bandwidth statistics
    - –Examining another metric for comparison
    - –Comparing results for other sites
  - –All these actions are possible with the DT.

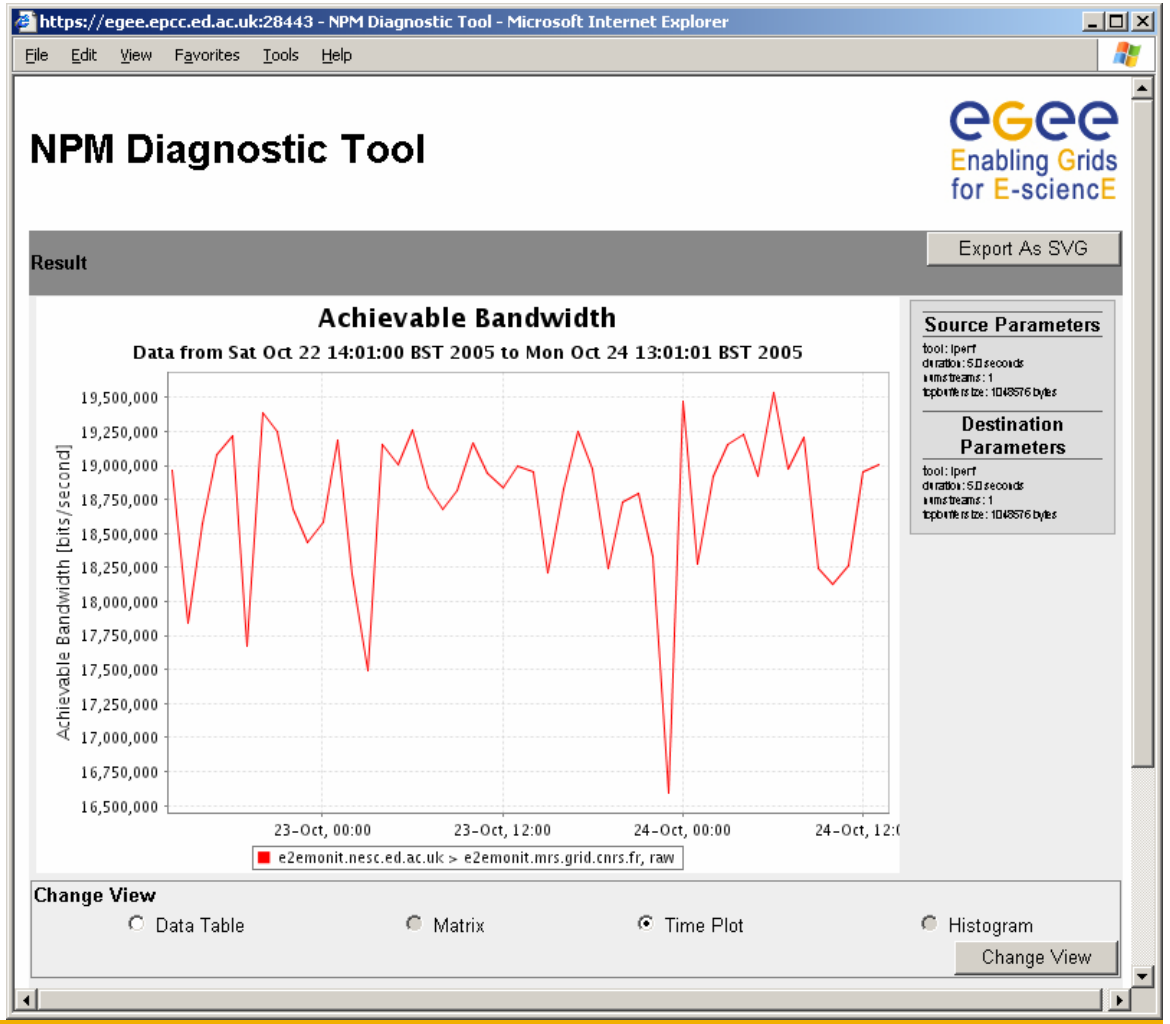

NPM Architecture

NPM Demo Architecture

- The currently deployed NPM client, services and accessible frameworks are shown below:

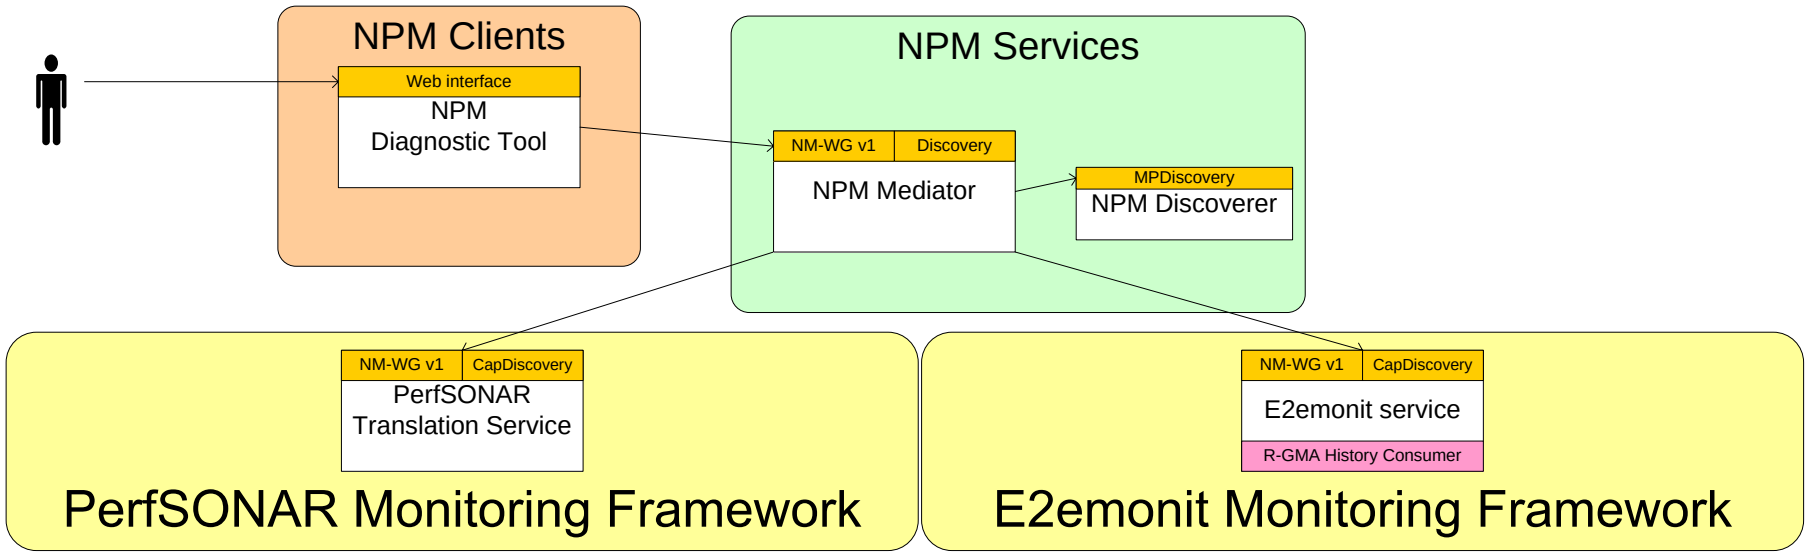

NPM Services (1)

- The central component is the NPM Mediator.
- It acts as a single point of contact for clients:
  - Receiving requests for network monitoring data.
  - Locating a monitoring point capable of providing an answer.
  - Relaying the request to that monitoring point.
  - Returning the result to the client.

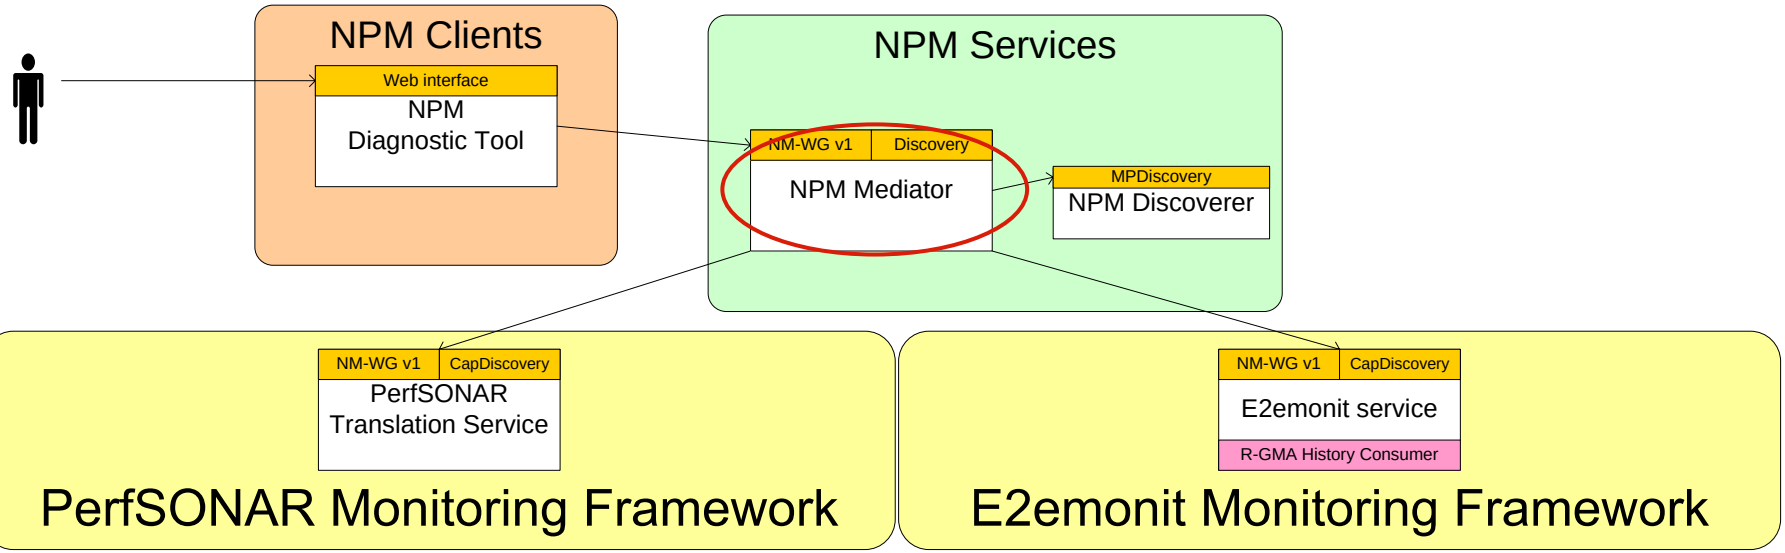

NPM Services (2)

- Location data is provided by the NPM Discoverer.
- This is currently statically configured with a list of available measurement paths within each monitoring framework
  - The architecture permits dynamic configuration.

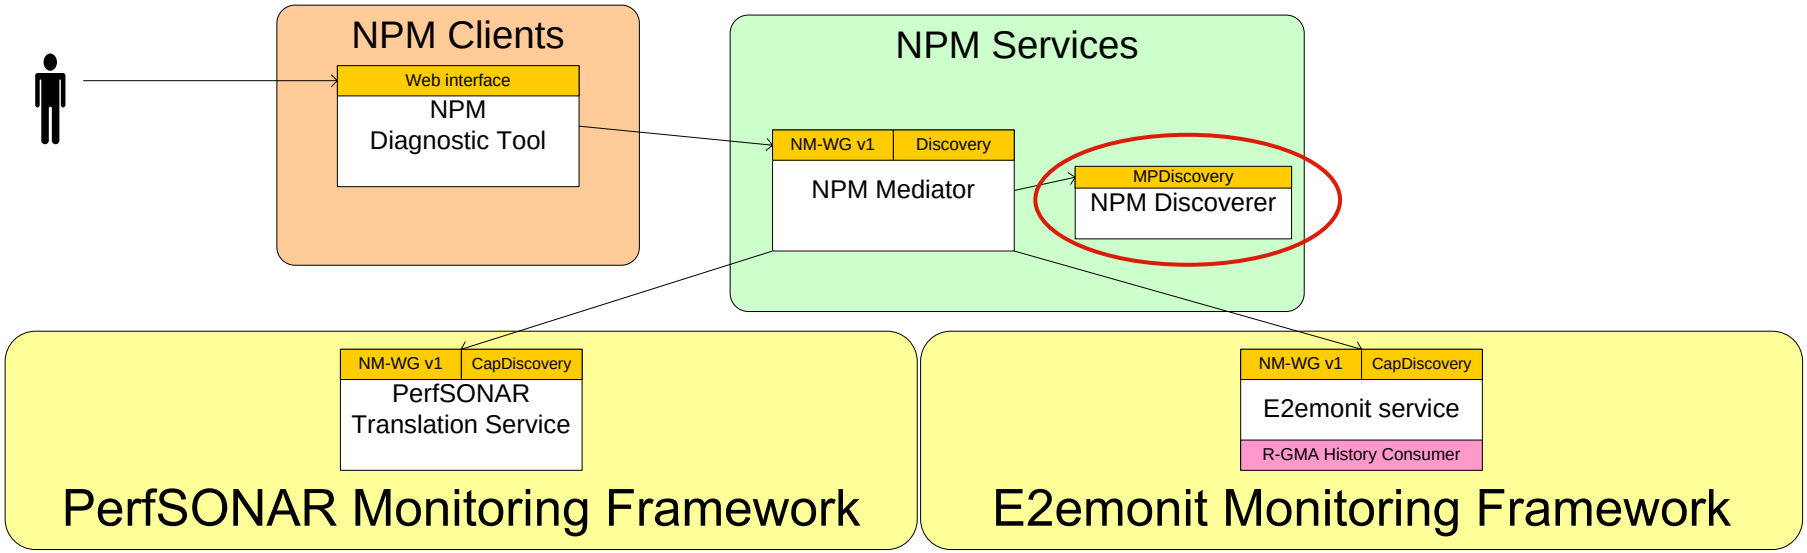

NPM Services (3)

- Each network monitoring framework exposes one or more web services complying with the GGF NM-WG v1 XML schema, allowing the monitoring frameworks to be accessed uniformly by the Mediator.

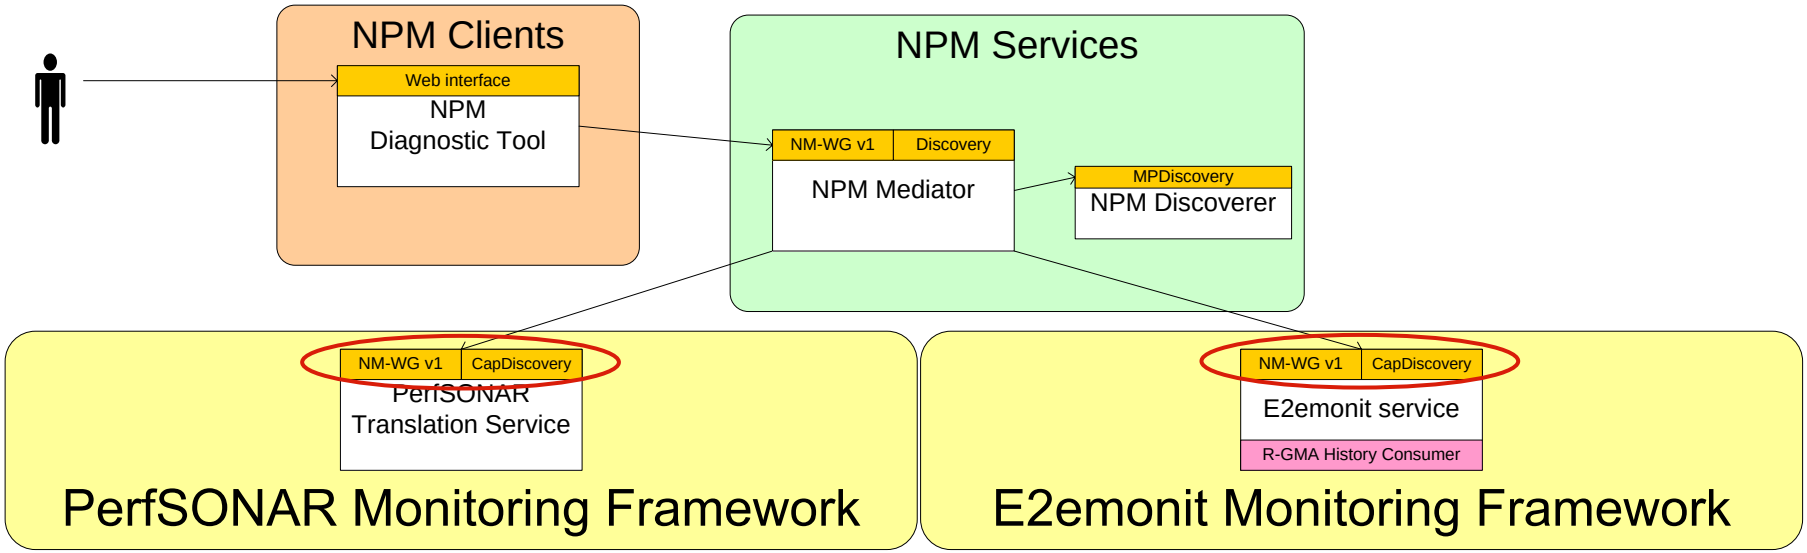

NPM Services (4)

- The Mediator exposes to clients the same NM-WG interface adhered to by frameworks, allowing its services to be used by any NM-WG v1 compliant client.

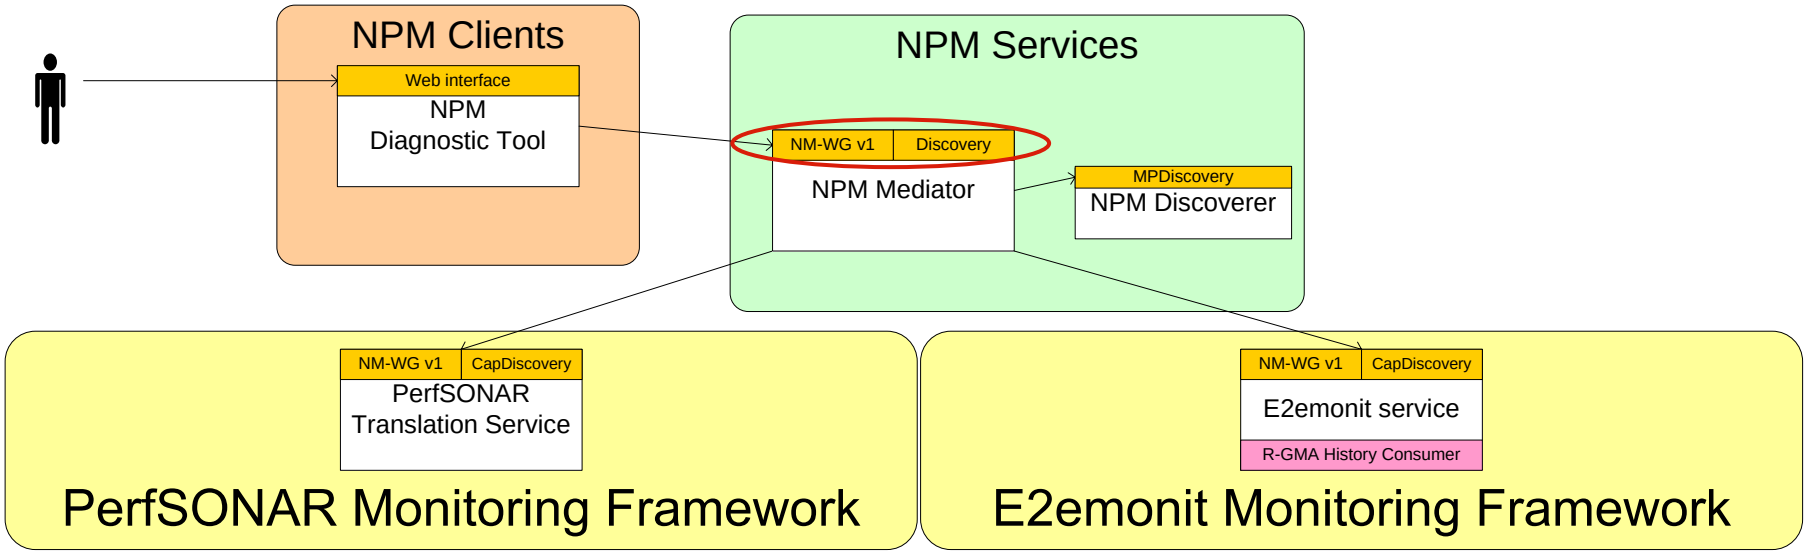

NPM Diagnostic Tool

- The NPM Diagnostic Tool (DT) is a client of the NPM Mediator.
- It is intended to be used by NOCs and GOCs (including CICs) to find and examine network monitoring data.
- The DT is a web-based application that can be accessed with any modern web browser.
- Since it is a client of the Mediator, it allows access to data from any monitoring framework accessible by the Mediator.

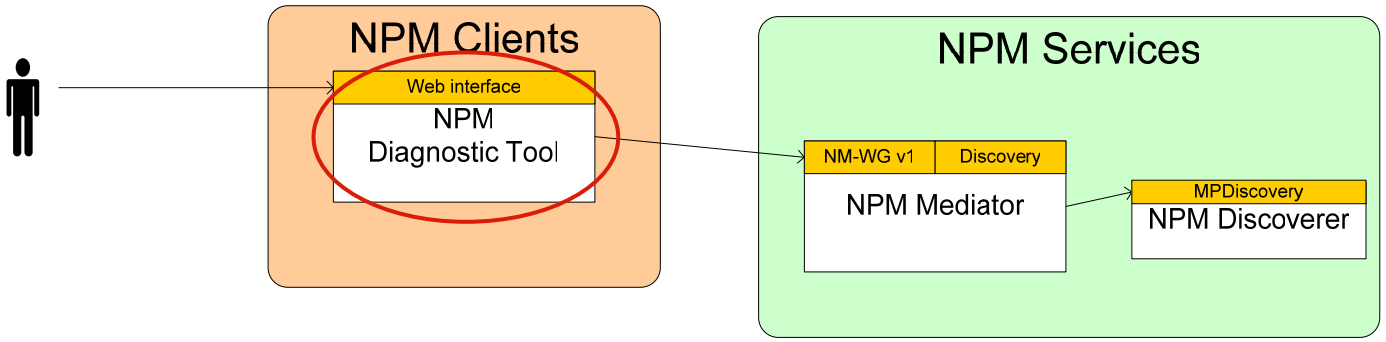

NPM Frameworks

- NPM clients, such as the DT, get their data from the Mediator - but it ultimately gets its data from monitoring frameworks.
- NM-WG v1 web services have been developed by JRA4 for PerfSONAR and e2emonit.

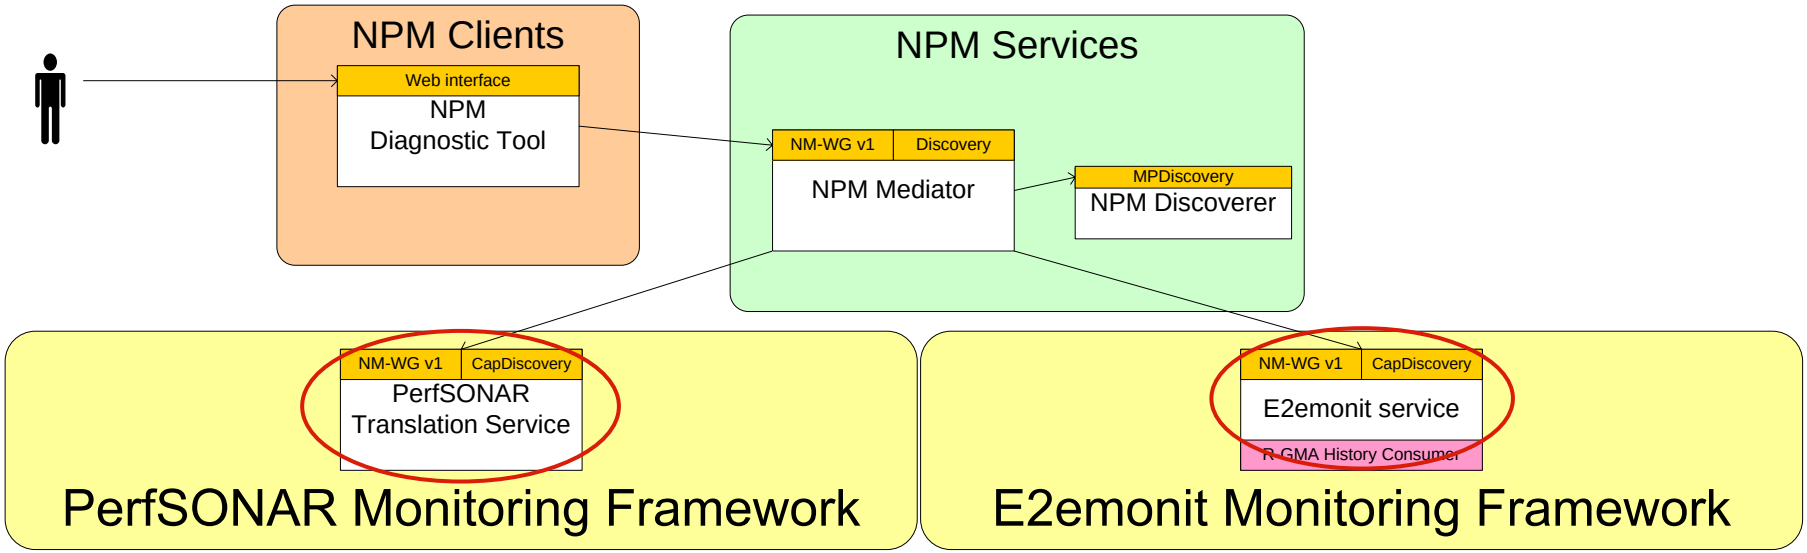

NPM Frameworks: PerfSONAR (1)

- PerfSONAR is a monitoring framework providing information for core backbones and NRENs.
- It is under active development by a global collaboration of networking partners.

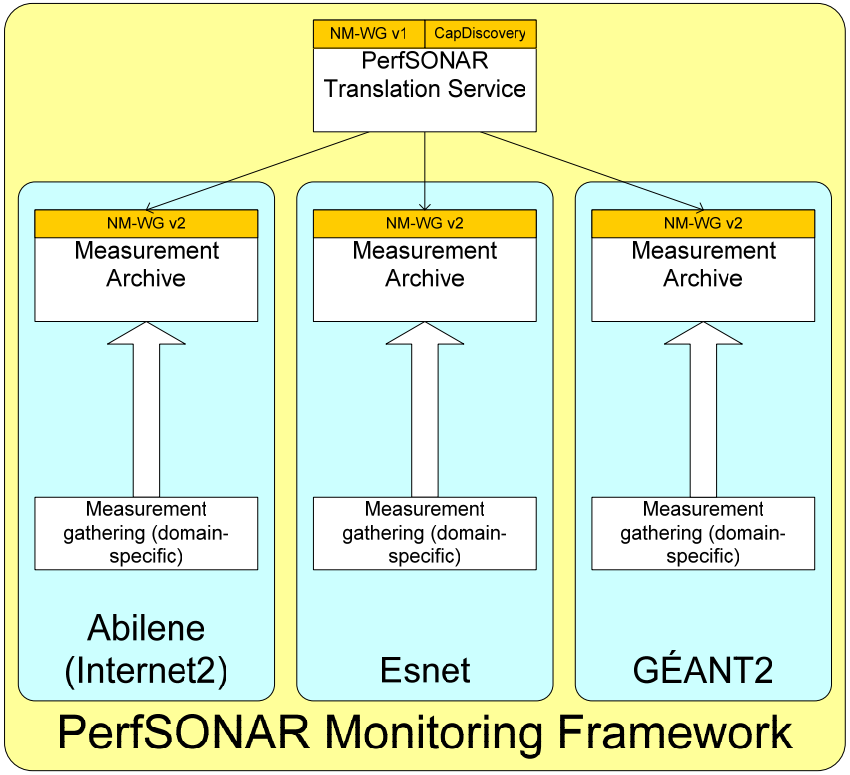

NPM Frameworks: PerfSONAR

- Currently data is available from GÉANT2, Esnet and Abilene (Internet2) - all accessible via the NPM Mediator.
- PerfSONAR natively uses NM-WG v2 schemas; JRA4 has provided a translation service to allow NM-WG v1 requests to be made.

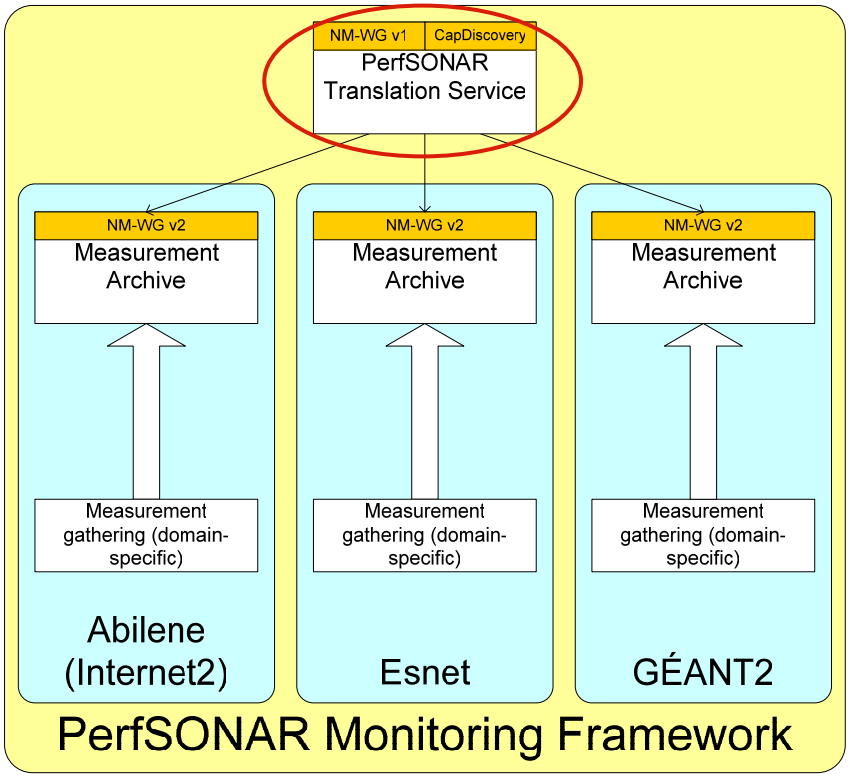

NPM Frameworks: PerfSONAR (2)

- EGEE JRA4's relationship with PerfSONAR has been mutually beneficial:
  - Backbone measurement data is provided for EGEE.
  - EGEE JRA4 - as an early adopter of PerfSONAR - has helped define matters that must be resolved within PerfSONAR, such as data disclosure policies.
- PerfSONAR development is ongoing, and we anticipate continued collaboration between the two groups to support improvements in the framework.

NPM Frameworks: e2emonit (1)

- E2emonit makes active measurements between end-sites, using tools such as iperf, udpmon and ping.
- Data is distributed using the Relational Grid Monitoring Architecture (R-GMA), developed by EGEE JRA1.

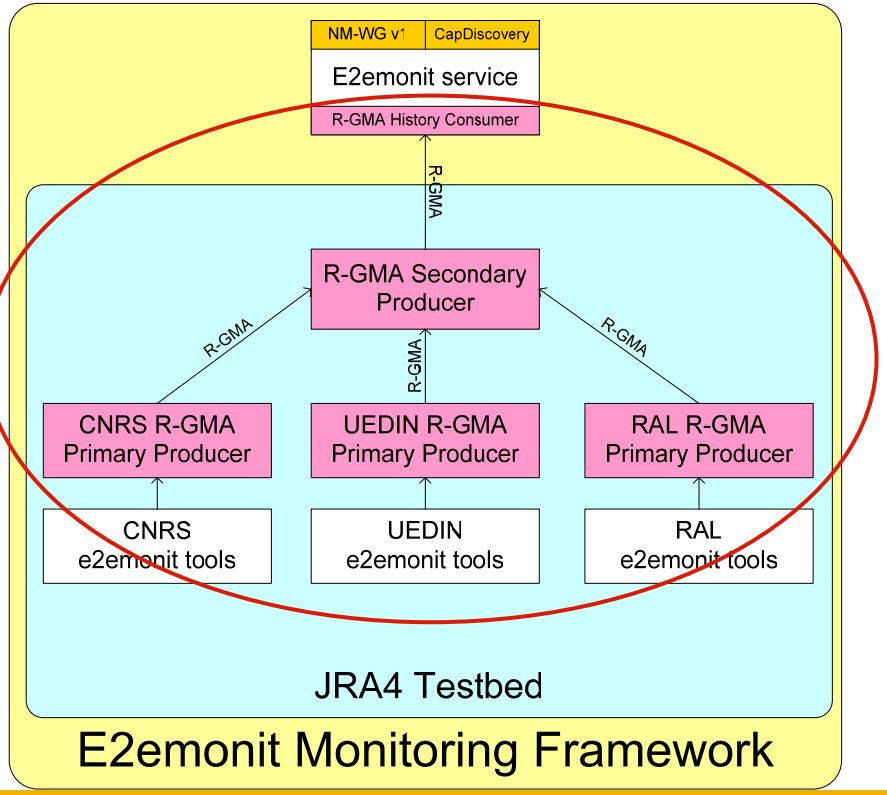

NPM Frameworks: e2emonit (2)

- Collaboration with the JRA1 R-GMA team on e2emonit has been mutually beneficial:
  - JRA4 has benefited from JRA1's R-GMA expertise.
  - JRA1 has benefited from additional user testing.
- JRA4 is now improving the e2emonit framework to support the requirements of middleware and deployments within EGEE.

NPM Current developments (1)

- JRA4 is currently working on a solution for resource-brokering middleware.
- Grid middleware has more stringent latency requirements than can be met with the Mediator:
  - The proposed NPM Publisher pre-caches measurement data and allows requests to be answered almost instantly, whilst also providing a mapping from Compute Element/Storage Element address to Network Monitoring Point address.

NPM Current developments (2)

- EGEE SA1 data challenges is interested in network monitoring between CERN and the LCG Tier 1 sites.
- JRA4 is finalising e2emonit packaging, and will provide support for the SA1 deployment.

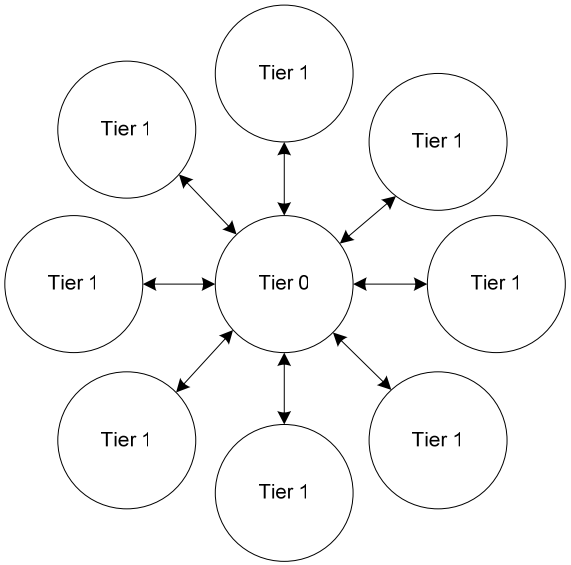

- Additionally, the possibility of accessing GridFTP usage data via the NPM services is under investigation.

Contact JRA4/NPM

- For further information on JRA4/NPM, please see the JRA4 web site:

http://egee-jra4.web.cern.ch/EGEE-JRA4/

- You can also write to EGEE JRA4 via email:

project-eu-egee-jra4@cern.ch