

Network Performance Monitoring for the Grid

ABSTRACT – Networkshop 2006

Suggested subject area: Monitoring and Measurement

Mark Leese, Rik Tyer and Robin Tasker
CCLRC, Daresbury Laboratory, Warrington, WA4 4AD
m.j.leese@dl.ac.uk, r.tyer@dl.ac.uk, r.tasker@dl.ac.uk,

Kostas Kavoussanakis, Alistair Phipps and Andrew Jackson
EPCC, The University of Edinburgh, Mayfield Road, Edinburgh, EH9 3JZ
kavousan@epcc.ed.ac.uk, alistair@epcc.ed.ac.uk, andyj@epcc.ed.ac.uk

Network performance monitoring has traditionally been important to the operation of networks of any significant size as an aid to fault detection and for determining expected performance.

The need for such monitoring is enhanced in the Grid computing arena. Without it:

- Grid middleware and applications cannot optimise their performance by adapting to changing network conditions.
- networks cannot be debugged for efficiency, an essential step in supporting data intensive applications because simple over-provisioning is not sufficient.
- the Grid cannot be “self healing”, or apply the SLAs associated with the “utility computing” model.

The talk will outline UK and European efforts to provide the required user-focused monitoring capabilities. These efforts predate but closely complement similar work underway within UKERNA.

From inception in 2002, the UK e-Science funded GridMon project has aimed to create a basic, UK-wide network monitoring infrastructure. This was achieved by establishing a monitoring presence at each of the original e-Science Centres with access to measured network performance data provided via localised Web interfaces.

The infrastructure was well received and grew interest in such work, while its use in debugging networks demonstrated the benefits of end-to-end monitoring; that is performance seen over a complete network path, including the end systems. End-to-end monitoring is a key aspect of GridMon. Network backbones are in the main efficient and well run, leaving the “last mile” (campus networks, end-user systems and even end-user software applications) as the most frequent cause of performance problems.

Despite its successes the chosen architecture proved an unsustainable administrative burden. The complexity required to make each monitoring point self-sufficient combined with the plethora of machines made available as monitoring nodes resulted in each installation requiring customisation and frequent maintenance. Further work funded by GridPP and the JISC is now establishing a second version of the infrastructure with many aspects centralised to minimise the operational effort required. A revised web interface is also being developed.

A key aspect of both GridMon and the EGEE-JRA4 work described below is machine (Web Services) access to collected performance data. The Network Measurements-Working Group (NM-WG) of the Grid Global Forum standards body, co-chaired by Daresbury, are defining XML schemas for communicating with network monitoring systems via Web Services technology. In suggesting a standard by which software (nominally monitoring systems) can request and share data, it is hoped that true end-to-end performance monitoring and testing will become possible, spanning both administrative and geographical boundaries.

JRA4 are the group responsible for development of network services within the EGEE (Enabling Grids for E-scienceE) project. Building on the work of NM-WG, JRA4 have produced network performance intermediary software, the Mediator. This software provides a single interface from which Grid applications, EGEE middleware and GOC/NOCs can request network performance data. Requests are forwarded to relevant underlying monitoring infrastructures, whether they are operated by a large

network operator such as DANTE, a local NREN or a project such as EGEE, hiding this complexity from the requestor. It is hoped that the software will develop to liaise with different administrative domains over the running of on-demand and scheduled tests, again making a powerful abstraction for the client. Indeed, the JRA4 Diagnostic Tool, essentially a graphical interface for the Mediator, has already been used to show the Mediator retrieving performance data from the Abilene and ESNet US research and academic backbone networks, their pan-European companion (GÉANT2), and the EGEE e2emonit end-to-end monitoring infrastructure (formerly known as EDG-WP7). This was an important step, demonstrating the retrieval of data from multiple-administrative domains using a unified interface.

JRA4 are currently investigating the publication of performance data to Grid middleware. This is the first stage in providing middleware and Grid applications that can adapt to changing network conditions.

Short demonstrations will also be possible.