

GGF12 Telecom Related Sessions: Telecom Panel, Enterprise Grid Workshop, End-to-End Management of Grids, NMA RG

**Alan J Weissberger
DCT**

Editors Note: While I also attended the two (Grid High Performance Network) GHPN RG meetings, they were too academic for me to combine with the sessions covered in this report.

I. Telecom Panel: Sept 21, 2004

Representatives from **BT, Telecom Italia Lab, and Internet 2** made presentations and took questions from the audience. France Telecom was originally scheduled to present, but they were a no show at the last minute. The panel was chaired by Wolfgang Boch, a representative of the European Commission who is responsible for Grid Technologies (one of two principal CEC Grid initiatives. The other major CEC initiative is on Grid Enabled Infrastructure).

The presentations may be downloaded from:

http://www.ggf.org/Meetings/GGF12/presentations/GridComputing@TelecomItaliaLab-GGF12_v1.pdf

A. John Ward of BT on their plans for Grid Network Services (John is retiring in early November): john.g.ward@bt.com

BT sees the Digital Network Economy taking shape, with grids at its center:

- More complex and robust management software required
- Globus toolkit will kick start grid development by providing open source middleware
- Virtualization business model across the entire IT resource profile, is what grids are all about (but security and control of telco resources are critical issues). Virtualization in the network is already taken place through VPNs, which are finally reaching critical mass.
- Exploiting the fact that grids plays to telco's strengths, e.g. in the management of distributed resources

Grids are important to BT:

- Grid moves customers IT resources into the virtualization world, alongside VPNs in the telco space
- Grids integrated with BTs 21st Century Network= a single unified network to provide multiple types of services

[From Matt Beal of BT: "The 21st Century Network is a unique mix of multiple technologies that no-one else is applying all together in a single program. Whilst many telcos concede these technologies' time will come, none of them are as advanced in implementing them as BT or have had the vision to deploy them all together."]

Here are a few links for **BT's 21st Century Network:**

A press link for BTs 21st Century Network:

<http://www.btplc.com/News/Pressreleasesandarticles/Corporatenewsreleases/2004/nr0445.htm>

BT is now trialing the network:

<http://www.btplc.com/News/Pressreleasesandarticles/Corporatenewsreleases/2004/nr0444.htm>

Sharper focus:

<http://www.btplc.com/Innovation/FutureNetworks/focus/index.htm>

Network benefits:

<http://www.btplc.com/Innovation/FutureNetworks/net/index.htm>

John Ward stated, “Most grids sold today are in a single enterprise (intra- site) environment. Connecting multiple grid computing sites will be a huge opportunity for BT.”

“Resource Management of the Virtual Network will be the key success factor: Service level guarantees must be met!”

BT will define, set, and monitor the SLAs for the underlying network and for the grid services it offers. They do not know of any standards for SLAs and feel that they will be trusted by their customers to monitor SLA compliance.

Network Technology to deliver grid services:

Initially, BT will extend its IP-MPLS VPN (as part of the 21st Century Network initiative) to enter the grid market. However, some customers have complained that performance/ throughput would not be adequate with IP (presumably because packet- by- packet or VPN routing is not fast enough). Hence, BT might choose another network technology for their production 21st Century Network used to interconnect grid sites and deliver grid services to end users. *Either they do not definitively know what it will be at this time, or they are not talking.*

→Interconnection of regional optical networks, provided by other carriers, will be a key problem to resolve. Method of interconnection, end-to-end SLA/ performance monitoring, and billing are the key issues here.

BT thinks it will succeed in delivering value added services over their nex gen network, while managing the entire collection of network, servers, and applications. This, despite the fact that other network operators (e.g. AT&T) have failed in their attempt to combine computing and telecommunications in packaged service offerings [More in Nigel Dye's talk in part II.]

John claims that BT is already well experienced in the computer system outsourcing/ systems integration business with over one hundred customers in Europe. They won such jobs over IBM and EDS- the two leading companies offering computer system integration and maintenance services.

B. Alessandro Aiello of Telecom Italia Labs (formerly CSELT): alessandromichelle.aiello@telecomitalia.it

-OGSA has strong potential in terms of resource access and service creation

-Major issues for telcos supporting grids:

- a] Grid network research
- b] Business models and economy

-Challenge of grid interconnection to a telco network:

"Network support for optimized transfer of huge data on a wide geographical basis."

This breaks down into a network that provides: interoperability, good performance (even under heavy load), multi-domain interconnection, QOS with good granularity, dynamic bandwidth management (ASON/ GMPLS), and security.

-Virtual Organizations (VOs) are an extension of network operator's VPN technology. Network becomes another active resource in this context.

-Grid apps are another kind of service to consider for the nex Gen IP network

-Grid networks should interact with storage and computing resources

-Dynamic service aggregation amongst ASP data centers will be important

-Billing and management are key issues for internal use of grid

Question from Nigel Dye of BT: **"We've heard from several customers that IP (or IP-MPLS) based network won't support the throughput needed for many Grid applications"**

Allesandro: "Use dynamic bandwidth management (e.g. GMPLS/ASON) to set up or re-provision the end to end bandwidth needed." Note: this assumes an optical control plane for fast set-up or re-provisioning of dedicated (non-shared) point-to-point optical channels.

Alan W: That does not resolve the throughput problem, especially under heavy network load/utilization. The requested bandwidth may not be available and the SLA will not be met.

Allesandro: IP-MPLS is not likely to be used in the core network that delivers grid services to multiple users.

C. Jill Arnold: jillar@internet2.edu

There are many initiatives that Internet 2 is involved with. Here are a few:

-NSF Middleware initiative (NMI) has been a major theme for last five years, with Internet 2 taking an integrator role in that project.

NMI key goals:

- a] Scientific computing and data environment (e.g. Grids)
- b] Common campus and inter-institutional research

-Security at Line Speed is another key project. It needs to consider security over an interconnection of regional optical networks (RONs)

-To interconnect sites needing high bandwidth, Internet 2 has setup the **FiberCo** LLC- a dark fiber company.

-End to End Performance Management initiative will be looking at all aspects of SLAs/SLOs.

Note: I have received an invite to set up a conference call with the chair of this project.

-**High Bandwidth Optical Packet Infrastructure** initiative will be used by people involved in or pursuing: high energy physics research, Square Km area community, medical imaging, real time and file transfer.

4. Closing Comments of Panelists:

BT: Interconnection of RONS will be complicated and a challenge to achieving true virtualization.

Telecom Italia: Need to improve convergence (no details, but assume this is convergence of computing and communications)

Internet 2: End to End Performance initiative will be vital in monitoring SLA compliance and assessing network performance.

II. Nigel Dye of BT talk at Enterprise Grid Workshop: Sept 21, 2004

nigel.dye@bt.com

Nigel is responsible for **Transformational Business Development at BT Global Services Wholesale**. While his focus is on new business models for carriers, he is also responsible for commercial activities in BT's Grid Action Team (GAT). This encompasses managed services and solutions for European customers.

BT is evolving toward delivery of "**assured grid services**." Their current packaged solution for end to end business process outsourcing is known as "**AAI**" or "**Assured Application Infrastructure**." AAI is a standardized set of managed services that comprises five stages, throughout which BT will:

- * Audit and profile a customer's existing end-to-end infrastructure
- * Make recommendations based upon the audit to optimize infrastructure performance
- * Monitor the infrastructure to ensure that it is performing well and to identify recurring problems or trends
- * Manage an agreed program of service improvement and recommend resolutions to the root causes of poor application performance
- * Assure the performance of core enterprise applications in accordance with pre-agreed service levels

Nigel's main point is that **AAI is evolving to support multi-site grid interconnection. Hence, grid enabled AAI will provide "one stop shopping" for grid users, supplying all necessary compute, network, and management pieces of the grid.**

Information on BT's existing AAI:

http://www.btglobalservices.com/business/global/en/business_zone/business_innovations/aai.html

Key data points for BT:

a) BT's **21st Century Network+** initiative will provide the underlying network (access, switching and transport) to deliver grid services. It has already started, but will continue to evolve over the next four or five years.

b) BT has a non-exclusive partnership with HP, which will presumably supply the servers and grid middleware.

c) BT currently has network coverage in over 120 countries. Its IP network connects over 250 cities

d) BT is already providing computer system integration and operations services in Europe.

+The **21st Century Network** trials will provide telephone and broadband services over fibre for up to 1,500 homes and businesses. The first trial participants will have their fibre lines installed and working in October 2004 and the trial will run until September 2005, with a rollout between now and 2008/9.

BT believes that Grids are a natural step in the evolution of computing and IT:

-Businesses are looking to buy services (outsourcing) that makes it easier to run their business.

-Utility computing, ASPs, Web Services will all play a role.

-Standards for interconnection, security, metering, billing, and QOS are urgently needed.

In addition to AAI/ service provider, BT is looking at also becoming a Grid Resource Broker (no details given)

Nigel could not give any time frame for when AAI or the 21st Century Network (underlying infrastructure) would be combined and made available to end users to deliver "Assured Business Services." However, he noted that both AAI and the 21st Century Network are "live" offerings at this time.

Editor's Note: While BT is fully committed, there is no (publicly disclosed) time target for offering grid services and no assurance that they will be successful at this endeavor. They are likely to meet competition from IBM and others that have combined computer and network system integration and outsourcing expertise. Nonetheless, the key takeaway is **BT's vision: telco's need to start developing virtual networks in order for distributed (multi-site) GRIDs to deliver their promise and potential. BT has starting planning for this and they are busy developing services for GRIDs.**

III. GGF12 End to End Management of Enterprise Grids: Sept 20, 2004

This panel was an interactive roundtable discussion with good audience participation. No presentations and no speeches. The panel discussion was chaired by "grid evangelist" Robert Fogel of Intel. Panel members were somewhat different from what was listed in the GGF12 plenary program. I recognized Tom Maguire of IBM and Igor Sedukhin of CA. There was a representative from HP whom I have seen at the Aug 04 OGSA interim meeting at Intel in Santa

Clara. The two end users represented were J.S. Hurley from Boeing and Mano Malayanur of FNMA.

In this report, I attempt to capture the relevant comments made by panel members. This is then followed by my own conclusions.

Tom Maguire: We need to distinguish management within the data center from management outside the data center.

It's important to use SLAs within an organization. Policy management provides the control points for setting SLA parameters. It is important to detect loss or degradation of resources, including equipment and/or the network.

SLA compliance is interesting when outsourcing your computing to a hosting environment. This implies a WAN to link your site to the hosting site. What apps are running over the WAN? Most companies do not know which apps are running on the network, in particular the mission critical apps that are on the network are not known.

What are the resources implied within a grid?

What are the standards needed for (SLA) compliance measurements? [Note: see GGF12 report of NMA RG]

J.S. Hurley of Boeing: cost is very important to us. But security is always the show stopper. It's important to be clear on what user gets from outsourcing.

For SLAs, we need standard parameters and standard methods of monitoring. We also need standards for authorization. We should treat entire protocol stack as one entity for this purpose.

Tom: De-provisioning and "cleaning" of released resources are key functions that have not been given enough attention. We need monitoring points that measure standardized attributes in a consistent way.

Policy is about how users manage risk in the resources controlled.

Igor: We talk about design, deploy and run the application, but what do we need to do in the design phase?

We should incorporate manageability into the design of the application. MSFT is providing this capability in their .NET tools.

We need to model virtualization of all resources in the grid

Tom: Only a handful of resources have been virtualized in enterprise grids that have been sold.

A "message passing architecture" is needed within domain specific apps, to permit you to get access to resource status.

J.S. Hurley: Virtualization is a means, not an end. Grid environment must be user friendly and easy to use. It should be seamless and transparent in running apps in a heterogeneous environment. Interoperability of multi-vendor equipment is very important.

FNMA: vendors must provide ease of use. Interoperability is key- we don't want a single vendor to supply the entire grid. Currently, the end user manages multiple single vendor grids. We want multiple vendors to interoperate on the same grid.

HP: We should distinguish grid middleware from management of the grid.

Tom: Open source software/ middleware levels the playing field.
Discovery of services and resources is urgently needed. (Presumably in the open source middleware)

Dealing with multiple administrative domains, each having its own policy, increases complexity by an order of magnitude.
--->Common policies are needed across administrative domains.

FNMA: We will be dependent on grid vendor management software. We will not develop our own NM software for grids. If we have to invest our own resources in such development, we will not use Grids.

Boeing: Because our computing/IT requirements as aerospace company are somewhat unique, we will have to develop some NM software, but only to augment the management software supplied to us by the vendors.

Conclusions (AW): Several items are urgently needed:

- Standardized SLA parameters and monitoring
- Common policies across administrative domains
- Multi-vendor interoperability within a single grid
- Grids should be easy to use, cost effective, and provide good security/authorization control
- Comprehensive grid network management software from vendors
- Applications should incorporate some type of manageability and resource status control/ monitoring capability
- Users need tools to determine which mission critical apps are running on the network in any point in time and for how long.

IV. Network Measurements for grid Applications RG meeting: Sept 22, 2004

Richard Hughes Jones of Manchester University chaired this well attended and provocative 1st session of the NMA RG (under the GGF ISP Area).

1. Introduction:

Richard initiated the meeting by reviewing the RG charter. The key point is the relationship between network measurements (both for faults and SLA/SLO compliance) and grid middleware.

Abstract of the Charter: The Network Measurements for Applications Research Group (NMA-RG) will focus on the relationship between network performance characteristics and Grid middleware. The objective of NM-RG is to explore what aspects of the network are most important to Grid middleware, and explore how the knowledge of network behaviour such as bandwidth, latency, or jitter, might be used to build network-aware middleware.

Scope- Specific topics of interest include, but are not limited to:

- What network characteristics are most useful to Grid Applications?
- What network characteristics are current network-aware Grid applications using/planning to use?
- Are there "derived characteristics" (combinations of characteristics) that are even more useful?

- When can passive monitoring be used, and when must active monitoring be used?
- When and how should Grid services should collect and publish passively collected network monitoring data?
- How to monitor "Next Generation Grid Networks".-----

Richard suggested that, perhaps, some "lightweight measurements" might be used to predict "heavy measurements."

Some key open issues for the NMA RG:

- Passive vs Active Monitoring
- How to monitor Nex Gen Grid networks?
- SLA performance monitoring- how and by whom?

2. Presentation on Grid Labs and Ibis: What do application tool kits need to know about the network, by Thilo Kielmann kielmann@cs.vu.nl

a] Grid Labs application tool kit (GAT): consists of simple, high level application oriented API which is independent of grid middleware (e.g. Globus, Avaki, Unicore, etc).

GAT uses a set of services.

Use cases covered by GAT: Remote file transfer, replica selection, resource selection (CPU close to data), remote data visualization

Estimated completion time of a job, e.g. file transfer, along with optimization parameters are most important to grid users

From the GFD.23 NM-WG Hierarchy document, several parameters need to be measured/ known: **bandwidth, r-trip delay, jitter, max TCP buffer size (usually 64K or 128K bytes).**

b] Ibis: a Java environment for support of Grid Applications

In this environment, small sized messages (a few K bytes) are typically sent.

Bounded Jitter is important due to the fine grained parallelism of jobs submitted.

Also important to measure 1-way and r-trip delay of the small and medium sized messages sent.

Ibis issues for grids: low performance (throughput/ goodput realized with TCP is inadequate), security (firewalls and encryption needed), use of CIM?

Important characteristics:

- actual bandwidth used, spare capacity, achievable bandwidth
- Delay, 1-way and r-trip + jitter

-Configuring TCP buffers, firewalls and NATs (to translate private to public IP addresses in order to connect to outside world/ public IP networks).

Open issues that need to be solved:

-Performance monitoring is not scalable or reliable

-Application monitoring is too slow, takes too much overhead (especially if the application does self monitoring)

3. On the Deployment and Execution of Component Apps on the Grid, by C. Perez of IRISA/INRIA, Rennes, France

Examples of compute intensive apps running on computational grids include: Tera Grid, UK eScience, European Data Grid, Grid 5K (France)

Grid consists of interconnected LANs, SANs, and WANs

Network oriented information: message size, frequency, security

Network oriented constraints: latency, bandwidth

Deficiency of the MDS (Metadata Directory Service) of Globus: Inadequate network description- no network related information is provided. Need to know the network topology for best mapping of grid components to the network. Network information provided should include: asymmetric links, firewalls, ability to find and bypass/transit non-IP networks.

Need to build a non hierarchical network to support Grid computing

Open issues:

-Accuracy of measurements is needed to predict future network performance

-Information access is the API suitable for finding adequate resources?

-Scalability? Is a p2p approach a viable solution?

User level constraints must be considered: Want results either: fastest possible, or want it by a specific date/time, or want low cost (< some amount of money)

4. What Do Grid Apps Need from Network Measurements, by Martin Swarry of Univ of Delaware

Long term goals of using network performance measurements in grid applications:

-Adaptive applications, based on current network resources available (ascertained from network measurements)

-Autonomous systems

-Ease of debugging

Predictive network performance - could be used to predict application performance

Key issues:

- How current is the performance information obtained?
- How would apps use a density/ distribution function for network performance?
- What does topology mean (how can apps make use topology information?)- both physical and logical
- What performance parameters should be measured?
- What do we need from the information service? It should be fast and not introduce unnecessary overhead
- Some apps (e.g. medical imaging) require very low latency and jitter. How can we measure those parameters end to end?
- Performance is determined by application, protocol stack, network. These all need to be measured separately and distinctly.

We need two documents to be produced by the NMA RG:

- a] **Use cases**- what does your app need and do with performance measurements now?
- b] **Goals**- what do end applications want to do with performance measurements?

5. Proposed NMA RGMilestones from Richard Hughes Jones:

Feb 05 - draft of document describing current use of network measurement/ monitoring information

Mar 05 - GGF 13 meeting in Seoul, Korea: presentations and check point our progress, refine workplan

Oct 05 - Use of derived performance characteristics

Feb 06 - draft of document describing issues for monitoring Nex Gen networks