

**EARNEST Initial Workshop
Berlin, 23-24 May 2006**



Parallel Session on Technical Issues

**Discussion Paper
by Kevin Meynell and Licia Florio**

Introduction

In the technical area, SERENATE largely focused on lower-layer issues such as routing, switching and transmission speeds, but control techniques are once again attracting attention as optical switching and hybrid networks become more ubiquitous.

At the present time, there are few if any production networks in the world using 40 Gb/s. Is it likely that there will be sufficient demand to stimulate development of even higher-speed networks in the near future? Is it likely that 40 Gb/s networks will become more affordable? Which transmission technologies are best suited to such high speeds?

At the same time, it has become increasingly clear that it is not possible to use existing networks to their fullest extent due to inefficiencies within certain protocols, and unidentified problems at certain locations on networks. It is necessary to investigate these issues further, because it does not make sense to continue to throw bandwidth at problems that have more fundamental causes. Steps in this direction have already been taken with the establishment of the SA3 PERT in the GN2 project, and its experience of resolving end-to-end issues will undoubtedly provide invaluable input.

A new topic for study is middleware, which essentially provides the tools and data that help applications use network resources. As network use has increased, authentication and authorisation have become more important issues, complicated by the fact that users have become increasingly mobile and wish to use resources as they move around. Allied to this is the increasing sophistication of hackers, which makes secure methods of authentication necessary.

Finally, there are more esoteric approaches to the type of networks that may be used in future. IP has proved to be very successful, but networks have evolved tremendously since it was first introduced, and certain aspects may no longer be adequate in future. In the United States, large sums are being invested in the GENI initiative which aims to provide a *clean slate* for solving some of the current Internet's limitations. What is the potential impact of this on research networking in Europe?

Four technical areas have been provisionally selected for study in EARNEST: transmission technology, control planes, operation and performance, and middleware.



1. Transmission Technology

The past five years have seen a paradigm shift in how research and education networks are provided. NRENs have traditionally run a best-effort IP service over circuits leased from telecommunications operators (more recently based on SDH), which has limited the flexibility as to how networks can be built, and has made optimisation between the different transmission layers difficult. However, it has become increasingly possible for NRENs to lease or buy dark fibre, which allows them to implement their own transmission equipment; in particular that which takes advantage of WDM techniques. This provides them with the flexibility to configure network topologies on demand, and upgrade capacity as necessary without having to renegotiate prices.

The transmission equipment currently available commercially is able to support transmission speeds of up to 40 Gb/s per wavelength and 2 Tb/s per fibre, although there seems to be very little requirement for this at present. This is partly because 40 Gb/s interfaces are still prohibitively expensive (4 x 10 Gb/s interfaces are significantly cheaper), and because more frequent amplification and regeneration is required as speeds become higher, which obviously increases complexity and costs even further. This said, there are continual improvements in transmission speeds and distances. Will 40 Gb/s and even 100 Gb/s wavelengths become commonplace in the next few years? Will it become possible to support more wavelengths on a single fibre (at least in usable way), thus pushing up aggregate throughput into the multiple terabit range?

Access to dark fibre and/or wavelengths also means that optical switching has become increasingly important. Such switches are very flexible in terms of the transmission protocols they support, while their interfaces are less complex and thereby cheaper than equivalent interfaces on routers. Despite this, the technology employed is still not completely standardised, especially with respect to different implementations of interfaces terminating long-distance fibres. Moreover, is it better to employ optical switches that utilise electronic switching (OEO), or optical switches that utilise optical switching through the use of MEMS or other analogue techniques (OOO)? An OEO-type switch is currently cheaper and therefore more common, but necessarily requires more complex electronics as throughput increases, with all the attendant problems of power consumption and heat dissipation. Conversely, the inherent complexity of an OOO-type switch does not increase with throughput, but it cannot spot errors or undertake data format conversion so easily. Are hybrid switches likely to become more commonplace, combining the best features of both approaches?

Another issue to consider is the underlying transmission protocols. SDH has traditionally been employed on WANs, but was designed to support voice traffic and carries an expensive overhead for data networks compared to Ethernet. By contrast, Ethernet was not especially designed for resilience, although the recent WAN PHY variants of 10 GE designed to run over SDH attempt to combine the best of both worlds. However, given the fact that data networks are usually based on a mesh rather than a ring topology, perhaps this may well provide sufficient resilience by itself? Alternatively, if networks are only planning to run IP-based services anyway, might it then be even preferable to adopt the IP-over-optical protocol, particularly once the control plane integration issues have been resolved?



IP routers are mature products and are likely to remain as the backbone of production networks for the foreseeable future. Most developments are likely to centre on support for higher-speed interfaces to take advantage of increases in transmission speeds, but there is also likely to be increasing integration between the IP and optical control planes. Routers have well-developed mechanisms for supporting automatic configuration and fault management, whereas optical networks still only have crude semi-automated configuration at best. In addition, it is somewhat inefficient to maintain separate control planes that have limited knowledge about each other's topology. However, there are still physical constraints with respect to optical switching that do not permit dynamic setup to the same extent as IP networks. This situation will undoubtedly improve, but it does impose certain limitations on what can currently be achieved.

2. Control Planes

Control planes are the intelligence that establishes, maintains, and tears down the different connections within a network. They should also have the ability to calculate optimal paths between end-points, distribute traffic loads, and re-route connections in the event of failure. Although the term control plane commonly refers to the switching mechanisms used in optical networking, for the purpose of this discussion it should be taken to encompass IP routing mechanisms as well.

IP routing largely works automatically with minimal manual intervention, but conventional routing protocols are becoming increasingly inadequate as the Internet grows in size. Not only are routing protocols required to handle ever greater amounts of information, but they have minimal security mechanisms to protect against abuses. Is it perhaps necessary to improve or replace the existing protocols in order to improve resilience, security, and peering and policy management? At the same time, provisioning systems for services such as Premium IP need to be developed, possibly in conjunction with some of the middleware activities (see Section 4).

By contrast, the levels of automation within optical switching are currently very limited. Customers typically have to request bandwidth at fixed rates for predefined periods of time, with (re)configuration of the network being done manually. This makes it difficult to dynamically provision bandwidth (for example, some customers may require more bandwidth at certain times), and there is a lack of resilience in the event of link failure.

These issues are being addressed by both the ITU-T (with its G.ASON recommendations) and the IETF (with GMPLS), but whilst standards exist, there are still a number of interoperability issues to resolve. In addition, research and education networks will need to consider whether to adopt an overlay or peer approach. The overlay approach keeps the optical and IP domains functionally separate, which is simpler and easier to understand. However, in the peer approach where switches and routers share topological information, there are opportunities for connection and management efficiencies. In fact, both approaches can be combined – for example, an NREN adopting a peer model for its internal network, whilst supporting overlay networks for its customers.



3. Operations and Performance

GÉANT2 and many NRENs are able to offer various guaranteed levels of service for specific types of traffic (usually known as Premium IP). While this has traditionally only been offered between the edges of networks, there are plans to offer it to users on an end-to-end basis, and as well as extend it to IPv4 multicast and IPv6. However, common pan-European approaches need to be agreed for provisioning and accessing such services, as there are currently no obligations for national and regional networks to maintain sufficient access capacity.

Beyond the actual provisioning of guaranteed service levels, it will be necessary to ensure that they are in fact being delivered. Experience has also shown that applications using high data rates or making large data transfers often do not see adequate performance despite the apparent availability of bandwidth. Even though individual networks can only be held accountable for their own edge-to-edge performance, there is a need for improved co-ordination and more integrated monitoring of all the elements in a chain in order to identify where the problems lie. In some cases, it should be possible to resolve the issues through improved management, but in other cases it may be necessary to develop and adopt new protocols.

The widespread introduction of optical networks at NREN operational level also brings new challenges. The NREN community has a lot of experience with IP routing, and it is well understood how to manage multi-domain networks. By contrast, there is much less experience of optical switching, particularly between different management domains, and this issue will become more complex as switching becomes increasingly automated. Improved tools for managing switches and monitoring light paths need to be developed, and consideration needs to be given to peering policies between optical domains.

Performance Enhancement Response Teams (PERTs) are a mechanism to co-ordinate the various NOCs at international, national, regional and campus level, in order to identify and resolve performance bottlenecks wherever they may occur on a network or networks. This includes ongoing monitoring, development and adoption of best practice, and an escalation procedure for handling problems. Should these be extended further, possibly even beyond the research and education community, similar to the way CSIRTs have developed to handle network security issues?

Finally, a number of significant user groups (the Grid community in particular) plan to establish extensive Virtual Private Networks (at both Layers 2 and 3). The manner in which such VPNs are provisioned and operated needs to be further investigated as there is yet no consensus on which VPN functions should be provided by research and education networks, nor how they should be managed.

4. Middleware

The term middleware indicates the layer of applications that connect the lower layers with directories and other multipurpose services. When talking about middleware, the emphasis falls immediately on Authentication and Authorisation for the users to access resources. The currently deployed AAs have very different capabilities, ranging from simple username/ password-based authentication to sophisticated systems for granting or denying access to resources.



Factors that have driven the deployment of middleware over the last years are the higher mobility of users and their need to access their resources from everywhere, as well as the desire to share resources (i.e. network devices, storage and applications). There is also the need to reduce the overhead of user management when users move from one place to another.

The necessity to share resources between different administrative domains (such as different departments inside the same faculty, different universities, or even different countries) has led to the creation of federations. Within federations, individual entities agree to allow access to each others' resources and adopt compatible technology in order for them to do so. In addition, different federations can make agreements to share resources (known as confederations), but the workings of the requisite trust models and how to manage the increased complexity are still open questions.

A very successful example of a federation is *eduroam*, which allows eduroam-enabled users to obtain Internet access when they move among different participating institutions in Europe and beyond (Australia for instance). This can be extended further in future, but how to integrate eduroam with other federations is still a matter for discussion.

PKI (Public Key Infrastructure) is not a new concept, but it has recently become an important issue again due to the requirement for trust relationships when building federations. The widest PKI community is represented by Grid users, where the access to the Grid resources is granted to the users upon verification of their digital certificates (X.509). It is quite clear that server certificates are becoming more and more important, and the new SCS service has made deployment of them easier. However, the use of end-user certificates outside the Grid community still remains quite cumbersome.

Will Grid applications remain bound to the use of X.509 certificates? How can the emerging (con)federations and the Grid co-exist/converge, or will they move towards other models? Will there be a real demand for end-user certificates?