



Technical Issues

Discussion note

Research and education networking organisations increasingly have access to dark fibre, and in some cases are now managing all aspects of transmission themselves. This means that issues such as fibre quality and the capabilities of transmission equipment have become much more relevant than before. Furthermore, although there have been advances in transmission capabilities, most research and education networks will likely be limited to existing fibre installations for the foreseeable future.

The fastest transmission equipment is currently able to support 40 Gb/s using SDH/SONET (OC-768) over a limited number of wavelengths, but this is prohibitively expensive for most research and education networks. Although prices for SDH/SONET interfaces are expected to drop significantly soon, it would seem that next-generation Ethernet will become the transmission technology of choice for those networks that do not have legacy telecommunications issues.

Most vendors appear to be focusing on 40 and 100 Gigabit Ethernet (GE) for next-generation transmission systems, and are adding carrier-class features such as OAM&P and virtual-circuit functionality. It is anticipated that the first implementations of 100 GE will arrive around 2010 or 2011, with 40 GE perhaps providing an interim solution by 2009. For both 40 and 100 GE, improved modulation techniques are expected to allow these technologies eventually to support long-haul links.

Most vendors appear to see a limited requirement for provisioning more than ~80 wavelengths over a single fibre. 50 GHz channel spacing appears to provide a good trade-off between faster line rates and longer reaches, although this is likely to be of limited concern to research and education networks that currently only utilise a fraction of the potential capacity of their fibres. However, the availability of ROADMs and wavelength-selectable interfaces on switches and routers will make WDM systems simpler and cheaper to provision and operate.

Research and education networks will continue to focus heavily on IP services, although the trend towards hybrid networks is likely to continue as networks move increasingly towards operating the underlying optical infrastructure. At the present time, the IP and optical domains largely have to be managed separately, but the introduction of GMPLS and virtualisation frameworks such as UCLP promise a more integrated approach. It is possible that production IP services may eventually become just one of many other services provisioned over a dynamic lightpath infrastructure.

With respect to IP services, there are some concerns over the future scalability of the routing system. Although these problems are not imminent, the IAB and IETF recently started to investigate whether addressing and routing could be made more efficient so that it becomes less reliant on hardware developments. In addition, there are revised predictions that IPv4 address space may be exhausted within the next five years, which is somewhat earlier than expected previously. IPv6 has already been widely adopted by NRENs, which should be well placed to support the transition from IPv4, but campuses should start developing migration strategies if they have not already done so.

It is likely that most NRENs and international backbones will continue to rely on over-provisioning to ensure reliable performance. There is currently limited demand for premium services, and there is little value in devising complex bandwidth allocation models when additional links can be established over dark fibre at marginal cost. Where customers or users have very demanding requirements, lightpaths can be used to provision dedicated private networks, and this should become a more dynamic process as better control-plane and virtualisation frameworks are developed. However, recognising that some edge networks may still have bandwidth limitations, QoS transparency should be supported in core IP networks in order to allow QoS and other traffic engineering mechanisms to be deployed over them.

Unfortunately, hybrid networks complicate network management because the IP and optical layers have evolved somewhat separately, and therefore have different management protocols, tools and operational procedures. There are a number of initiatives in the research and education community to develop tools for monitoring and managing optical networks, but these are at quite an early stage of development and offer limited integration with the IP layer. In addition, more comprehensive network monitoring is needed for understanding and managing networks, although it remains hampered by a lack of standards and is challenging to undertake at line rates above 5 Gb/s.

Another important issue for network management is that operational experience shows that around 90% of all reported problems in modern networks are due to issues at end-sites, many of which are attributable to so-called middle boxes such as firewalls, NATs, rate-shapers and intrusion detection devices. These devices, which are supposed to help manage and secure the network, can often end up making things more complicated and less secure, as well as hampering performance and delaying network upgrades. Therefore consideration needs to be given to better placement and management of such devices.

The GN2 PERT has successfully traced many of the problems that afflict users of GÉANT2 and other connected networks, in most cases demonstrating that the problems are at the end-sites. Unfortunately, it currently has a limited scope, and if the concept of the PERT is to evolve, it needs to be extended to NRENs, and possibly to regional and/or campus levels as well. This should be accompanied by the establishment of standard operating procedures and common informational and tracking systems.

In the middleware area, identity federations are becoming increasingly important for handling and supporting user access to remote services. The majority of the NRENs in Europe either already have a federation or are in the process of establishing one. Those without a federation should plan to have one in place within the next couple of years.

Because NRENs are the natural candidates for providing technical and organisational coordination for research and education communities, they should support multiple trust infrastructures in order to be able to handle different AAIs for different purposes. In addition, as identity federation implementations mature, they should consider how to support new features such as integrating Shibboleth with Grid applications, and the facilitation of inter-federation peering.

It is expected that SAML 2.0 will become the standard mechanism for exchanging identity assertions for web-based applications. Other services will probably continue to use X.509 personal certificates for the time being, at least until alternative approaches to integrate these are more fully developed. Authorisation decisions will likely be supported by schemas such as eduPerson and SCHAC, although there is still no well-established standard for communicating identity data to applications. This is therefore an area in which NRENs might be proactive.

Other AAI models such as user-centric identity management or the abstract identity framework may also yet prevail beyond the academic community, and NRENs should continue to monitor these developments. However, users might still be able to use their federated identities with services that support these models.

For discussion

How will the technical developments described above affect research and education networks at the international, national and local (campus) level? Can those networking organisations continue their current policies or are important changes of strategy necessary?

Which recommendations to various stakeholders can be formulated based on the expected technical developments?

Which of the findings and recommendations are the most important ones that should be incorporated and highlighted in the overall Summary Report of the EARNEST Foresight Study?