Report on Campus Issues

January 2008

Authors: Jean-Paul Le Guigner, Martin Price, Rogelio Montañana and Michael Nowlan

ISBN 978 - 90 - 77559 - 15 - 4

Production: TERENA Secretariat
Design: Stratford Design
Printing: Noordhoek Offset b.v.

The EARNEST foresight study was funded by the European Community through the GN2 project in the Sixth Framework Programme for Research and Technological Development. This publication does not represent the opinion of the European Community; the Community is not responsible for any use that might be made of data appearing in this publication.

Contents

1. Executive Summary

The EARNEST foresight study has looked at the expected development of research and education networking in Europe over the next 5-10 years. The study was carried out between March 2006 and November 2007. EARNEST was funded by the European Union through the GN2 project, which also provides the funding for the current generation of the pan-European research and education backbone network, GÉANT2.

The aim of EARNEST was to provide input for initiatives that could help to keep the evolution of European research networking at the forefront of worldwide developments and enhance the competitiveness of the European Research Area. EARNEST has prepared the ground for the planning of the development of research and education networking infrastructure and services after the completion of the GN2 project, at the local, national, European and intercontinental level.

EARNEST can be seen as the successor of the very successful study that was carried out in the SERENATE project in the period from May 2002 until December 2003. The results of the SERENATE study, and in particular the recommendations in its Summary Report, have been very influential on the planning and development of research and education networking in Europe in subsequent years.

After an initial preparatory phase, the EARNEST work has focused on seven study areas: researchers' requirements, technical issues, campus issues, economic issues, geographic issues, organisation and governance issues, and requirements of users in schools, the healthcare sector and the arts, humanities and social sciences. Reports have been published on the results of each of these sub-studies, as well as an additional report on regulatory issues. The EARNEST study is rounded off by a Summary Report that contains recommendations for the relevant stakeholders.

The current report presents the results of a study of campus issues, looking at the status of networking within institutions. The starting point for the study was the assertion in the SERENATE Summary Report that in 2003 the campus network was often the weakest link in the chain of end-to-end services needed for research and education. The group that carried out the EARNEST study of campus issues has investigated whether campus networks are now better resourced. It has also looked at other issues, including the change of emphasis from connectivity to network services, the rollout of IPv6, training of network support staff, and collaboration between network engineers at campus level, as well as collaboration with National Research and Education Networking organisations.

The report is based on a survey of the current state of networking in research and education institutions in Europe, and looks into organisational issues as well as technologies. The study report ends with several recommendations, which, if implemented, should improve network facilities and services at campus level and the support that they can offer to the research and education communities.

2. Introduction

There is no doubt that a well-managed modern network offering up-to-date network services is a key element of the infrastructure of successful European research and education institutions in the 21st century. Staff and students are becoming more mobile as they exploit the creation of a European Higher Education Area following the Bologna Declaration of 1999, and researchers take advantage of the growing opportunities for collaboration offered by the European Research Area. Much of the success of these ventures will depend on pervasive, effective, high-quality networking – not least within campuses. Institutions exploiting the power of network services to the full will undoubtedly gain a competitive edge in their research and education activities.

The group that carried out the EARNEST study of campus issues looked at networking within research and education institutions. It took as its starting point the assertion in the Summary Report of the earlier SERENATE study (December 2003) that the campus was often the weakest link in the network chain, and the report's recommendations that:

"Therefore, universities and research institutes and their supervisory and funding authorities need to ensure that their campus networks are appropriately resourced. In general, expenditure for ongoing technical upgrade in campus networks is best treated as a budget expense on an annual basis."
and
"Research and education institutions should consider acquiring their own fibre infrastructure between their Local Area Network(s) and the point(s) of presence of key service and/or infrastructure providers, if necessary by commissioning its construction."

The group investigated whether campus networks were now resourced better, and it looked at other issues, including:

- the change of emphasis from connectivity to network services;
- rollout of IPv6;
- training network support staff;
- collaboration.

As an introduction to the current report there are some important points that should be made at the outset, concerning the categorisation of end-users and the scope of the study of campus issues.

2.1 Categories of end-users

This report refers frequently to end-users, i.e., the researchers, teachers, staff and students at research and education institutions in Europe. However, the networking requirements of those end-users differ greatly, depending on their research domains, teaching and learning disciplines etc. The report uses a widely recognised convention to divide different networking requirements into three categories - Classes A, B and C – which reflect baseline, medium and heavy usage, respectively.

- Class A broadly requires email handling and Web browsing, the basis for all professional activity nowadays. This class of usage does not require much of the network facilities, apart from reliable access, mobility and a satisfactory level of bandwidth.

Class A is estimated to consist of more than 80% of end-users in 2007.

- Class B covers two additional activities: the handling (via file transfer or database access) of significant volumes of data for the regular use of audio and video streams, and access to Grid resources. Streaming, especially when human interaction is involved, often has more demanding performance requirements than conventional file transfer. Grids also have their requirements, especially in terms of latency and jitter. These two usages may require guaranteed bandwidth and guaranteed stability for jitter and latency, but they still require only limited levels of bandwidth.
  Class B is estimated to consist of less than 15% of end-users in 2007.

- Class C covers the sustained, but often temporary, handling of extremely large volumes of data, either by transfer of very large files, or increasingly via technologies such as Grids and virtual presence.
  Class C is estimated to consist of a very small percentage of end-users in 2007.

2.2 Scope of the study of campus issues

In the strictest sense, the EARNEST study of campus issues should have been limited to networks within research and education institutions only. However, the underlying structure of most research and education networks is multi-layered, and to look at the network from the perspective of the end-user, especially when searching for potential bottlenecks, it was necessary to look also at other components - or intermediate networks - between NRENs[1] and institution networks. The infrastructure from which Europe's research and education communities obtain their network services is often quite complex; it is provided by a variety of domains, including the pan-European research backbone network[2] GÉANT2, NRENs, regional networks, Metropolitan Area Networks and campus networks.

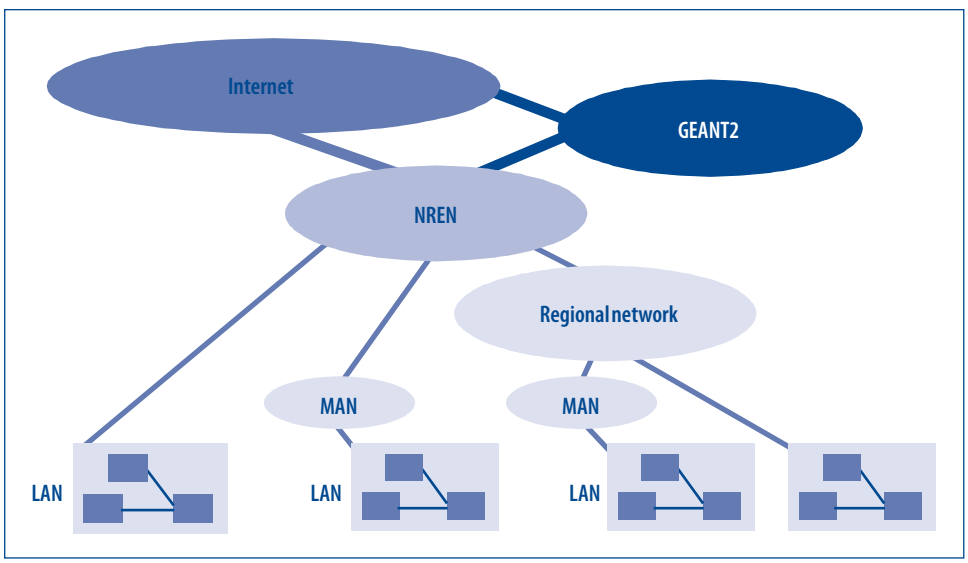

Figure 2.1: Structure of research networks in Europe

1. *In this report, the acronym "NRENs" is used to denote National Research and Education Networking organisations as well as the national networks provided by them.*

2. *Since 1993 there have been five consecutive generations of the pan-European research backbone network: EuropaNET, TEN-34, TEN-155, GÉANT and GÉANT2.*

In order to provide good network connectivity for research and education institutions, the different levels of infrastructure involved in delivering services must operate in a coherent way. In Europe, connectivity and service provision for the end-user typically depend on three, but sometimes as many as four or even five, different categories of infrastructure. These are shown in Figure 2.1. They are:

- **The Local Area Network (LAN)** infrastructure of the site where the user has his or her main workplace. This infrastructure is under the responsibility and control of the research and education institution (university, research centre etc.). Each institution may have multiple sites and in the context of this study, each site is considered to have a 'campus network'. The term 'campus network' is used for the local network infrastructure of all organisations served by NRENs and other research and education institutions, regardless of the type of institution.

- **The Metropolitan Area Network (MAN).** The MAN typically provides access to an optical-fibre infrastructure, which can be used to interconnect LANs within one city. While the LANs of individual sites are under the responsibility and control of institutions (universities, research centres etc), MANs tend to be run by commercial telecommunications operators, by a consortium where local government plays a significant role or under the responsibility of a formal collaboration of academic institutions.

- **The regional network.** In some countries, such as France, Spain and the United Kingdom, regional networks form a very important element for providing connectivity between the NREN and campuses. In those countries, NRENs usually have Points of Presence (PoPs) in large cities, often the regional capital if one exists. From the PoP, the NREN can connect to all institutions within the city, either directly or via a MAN, while other sites spread around the region will be connected via the regional network. These networks typically offer connectivity and other services, rather than just simple data transport. Usually they are built, operated and funded by a consortium of the local government and various institutions in the public and private sector.

- **The National Research and Education Network (NREN).** In each country, the NREN infrastructure typically provides national connectivity to other research and education institutions, as well as onward connection to European and intercontinental destinations. NRENs operate on a variety of different governance and funding models, depending on the national political consensus.

- **GÉANT2.** At the pan-European level, the GÉANT2 backbone network, which is planned, built and operated by DANTE on behalf of the NRENs and co-funded by NRENs and the European Union through the GN2 project, provides the primary international infrastructure for production traffic. However, it is not exceptional that for important projects NRENs also directly acquire connectivity for experimental traffic to various destinations in Europe or in other continents.

3. Methodology

The EARNEST study of campus issues consisted of two phases.

Phase 1 was a fact-finding exercise to discover the current situation of campus networking within research and education institutions, mainly universities, research institutions etc. The group of experts that carried out the study developed a questionnaire aimed at the heads of IT services in research and education institutions in Europe. The questionnaire was distributed through different channels, including the NRENs, EUNIS and TERENA.

The questionnaire was prepared in January and February 2007, and it was made available online at the beginning of March 2007, giving respondents three weeks to answer. On 2-3 April 2007, a first rough analysis of the responses was presented to a joint meeting of the groups working on the EARNEST study of researchers' requirements and the study of campus issues.

In April and May 2007, two consultation workshops were organised by EARNEST, one with NREN representatives and one with representatives of the funding bodies of NRENs. In these meetings, some of the results obtained in the study area were presented, and the workshops offered an opportunity for the two different groups to make comments and give feedback.

To clarify some responses, follow-up interviews by telephone were held with respondents in Spain, France, the United Kingdom and Ireland. These interviews took place in late April and early May, and the additional results were integrated with the initial ones.

Phase 2 looked at the strategic directions of research networking, to provide a view of the possible future of campus networking. The group took into account ongoing technological developments on the vendor side as well as various experimental projects in the research and education networking community (for example, as part of the GN2 project, in the context of Internet2 in the United States, and within NRENs). The objective was to assess the impact that these developments could have in the near future on the campus networks and on the interface between NRENs and institutions, and to provide guidance and recommendations to senior management in institutions about developing campus-network strategies.

4. Survey of campus issues

The SERENATE Summary Report said that campus networks were *"often the weakest link in the network chain"* and thus limited the performance of network services provided to end-users. This finding was the outcome of a survey of researchers, where respondents mentioned 'campus bottlenecks' as a factor inhibiting optimal use of the network. The views of network managers and heads of IT services at institutions were not sought at the time, because an investigation of campus issues was outside the original scope of the SERENATE study.

When the EARNEST study was set up, it was decided to make campus issues one of the main areas of investigation. As a consequence, the issue of campus bottlenecks (real or imagined) was one of the areas surveyed in the questionnaire sent to heads of IT services in institutions. However, the scope of the survey was made much broader and it extended to a whole range of issues affecting campus networking. The various topics are listed in the sections below, together with the main findings that have come out of the survey.

Details of the questionnaire on campus issues and of the analysis of the responses are available from the TERENA Secretariat upon request.

4.1 Bandwidth to the NREN

In the survey, the available bandwidth between institutions and their NRENs shows a very wide spectrum, ranging from 2 Mb/s to 10 Gb/s. These figures are not easy to interpret. Low bandwidth may often be adequate for small institutions with few end-users. However, a significant number of medium-sized institutions also have low bandwidth to their NREN; it appears that this is mainly due to their obligation to use an intermediate network to connect to the NREN.

In an attempt to make meaningful comparisons between institutions, the average bandwidth available per computer was calculated. Results ranged from 8 kb/s to 100 Mb/s per computer. This represents a huge difference in bandwidth, and it is not easy to explain how an institution providing 10-50 kb/s per computer can cope with this situation. It is not possible to draw any clear conclusions from the survey, but it is evident that end-users with such small bandwidth cannot exploit the full range of network services.

4.1.1 Satisfaction with low-bandwidth connections

Of the institutions taking part in the survey, 25% reported that they had relatively low bandwidth between their campus and the NREN, i.e., less than 34 Mb/s. Some of these were small institutions with few students, but some were not, with more than 5,000 students recorded. One would have expected the larger institutions to be less than content with the bandwidth available to their user community. However, surprisingly 81% of respondents considered the bandwidth to the NREN per user to be sufficient.

It is likely that many end-users, especially those in Class A, will not access external services directly but via proxy servers, caching servers, central email servers etc. In this way it is possible to cope with

many thousands of Class-A users on low-bandwidth connections. However, it is very important to manage such servers properly (see Section 4.9).

EARNEST has not been able to investigate in greater depth why these institutions appear to be satisfied with their low bandwidth. This may be a worthwhile topic for further study. One possible explanation is that the expectations of the user community in these institutions are unusually low and that those users are content with only basic network services.

4.1.2 High bandwidth to the NREN but little use of network services

The EARNEST survey of campus issues, the EARNEST survey of researchers' requirements and the most recent edition of the annual TERENA Compendium of National Research and Education Networks in Europe all indicate that the problems of low bandwidth that were identified in the SERENATE study are now largely resolved. At the present time, many institutions have adequate (or better than adequate) bandwidth to meet the current needs of their user communities. The phenomenon of network bottlenecks – whether on campus or elsewhere – is now rare.

However, there is not much evidence of innovative use of the network that takes advantage of the available bandwidth. It appears that there are several different reasons for this:

- Many end-users are not aware of the range of network services that are available to them (e.g., videoconferencing)[3].
- Some teams of network engineers in institutions are not aware of the network services that they could provide to their users.
- Some institutions are reluctant to promote new services because their network teams are too small or do not have the right skills to provide adequate support.
- Some network services are not easy to implement and to use, e.g., videoconferencing, multicast, IPv6 (see Section 4.6).
- In some cases, there are still limitations in end-to-end communications software or hardware (see Section 4.10).
- Many users do not need advanced services; a dependable, reasonably fast network providing email and Web access is sufficient for them.
- Many users are not familiar with the technical side of the network and may not even be aware that higher bandwidth has been provided for them.
- Innovative services often require a prolonged development period. This appears to be the case, for example, with some Grid projects where structuring the user community takes longer than expected.
- Often cultural changes are required. Both the end-user communities and the people and organisations that provide network services need to become more imaginative in their approach to networking. Frequently the technical infrastructure seems to change more quickly than the mindset.

4.2 Performance of the core network on campus

Overall there have been great improvements in recent years in the infrastructure of campus networks. Nevertheless, there are still isolated pockets where there is not adequate bandwidth; this can be for a variety of reasons, including out-of-date equipment, problems with buildings, fragmented institutional management, geographic remoteness and lack of funding for networking.

3. *This is documented in some detail in the EARNEST report on researchers' requirements by Thibaut Lery and Patrick Bressler.*

A third of the survey respondents pointed at some part of their core network as an obstacle to propagating Gigabit connectivity to end-users.

4.2.1 Inter-campus connectivity within institutions

Responses to the survey indicate that some institutions have problems connecting remote campuses. Their interconnection links appear to run at unacceptably low data rates. The survey shows that only 25% of the responding institutions have internal connectivity of more than 1.2 Gb/s, while 35% of the institutions taking part in the survey have connectivity to their NREN of more than 1.2 Gb/s. There may be several different reasons for this disparity in connection speeds. It could be that there is a great difference in size of the different campuses and that smaller ones merit lower rates of connectivity. It could also be that the pace of growth of external connectivity has outstripped that of the internal network. Another reason for low bandwidth between campuses could be the obligation to use regional networks in cases where remote campuses are at some distance from the main site. It is even possible that telecommunications costs for better connectivity are prohibitively high.

4.2.2 The 'last 100 metres' problem

There are indications that there are problem areas on campuses related to outdated wiring and switches/repeaters. Of the 150 respondents who replied to the relevant question, 41% reported that it was impossible to establish a high-bandwidth connection (1 Gb/s) between any two points on the campus. Of these, 70% reported that the reason was out-of-date hardware or cabling. Possible reasons for this phenomenon include:

- the reluctance of faculties/departments to invest in upgraded equipment and cabling;
- demarcation disputes between departments and central service providers about the funding of upgrades;
- problems in making alterations to older buildings, including the presence of asbestos, for example.

Another finding of the survey is that 51% of respondents have users connected at speeds lower than 10 Mb/s. If some of these users have demanding applications, then they must be severely hindered in using the network, with poor end-to-end performance. Low-bandwidth connections such as these adversely affect network performance. For example, a person on a Gigabit connection is likely to have impaired performance when he/she is sharing an application with a colleague on a slow-speed network.

In the worst-case scenario, researchers may be unable to take part in collaborative projects requiring high-speed connectivity. The cumulative effect of such exclusions would seriously damage an institution's research reputation.

4.2.3 Progress towards improvements

The survey shows that institutions are rolling out Gigabit networks on their campuses. However, their availability is limited. Approximately 5% of institutions can provide 90% of their end-users with a Gigabit connection. This will not be a major problem for the majority of users in Class A. However, there will be serious repercussions if this distribution is uniform and only 5% of Class-B and Class-C users have access to Gigabit networking.

One of the recommendations in Table 4.1 states that a high proportion of new end-point equipment should provide Gigabit connections as a standard. There may be a slight extra cost in the purchase

of such switches today, but the labour costs in installing and managing them will be similar to that for lower-grade equipment, and the investment now will be repaid later as the demand for higher bandwidth increases. It is also strongly recommended that institutions with out-of-date cabling plan on replacement as quickly as possible.

	Recommendation
1	Aggressive replacement policies for network equipment should be put in place, with a maximum turnover period of five years.
2	An inventory of institutional wiring should be kept. Total replacement of the core and end-point wiring should take place at least every ten years. Such a period may actually prove to be too long when radical changes in technology take place.
3	A high proportion of all new network switches and end-points should be Gigabit-enabled. At this stage, purchasing lower-speed equipment is a very poor investment.
4	Even if the core network is not Gigabit-enabled, Gigabit-enabled host-network interfaces should be installed for the end-user.
5	Early adoption and rollout of equipment with higher specifications may be required for intensive services such as central servers or for high-end users in Class C.
6	There should be an internal process to define specifications that are to be followed by departments and/or faculties when installing and/or upgrading the part of the network infrastructure that is under their own responsibility.

Table 4.1: Recommendations related to the performance of the core network on campus

4.2.4 Future survey work

The current survey is a single snapshot in time and presents only the current status of campus networks. It is recommended that an annual survey is performed to show the progress in advancing end-to-end network performance. The perception of the network can be seriously degraded by a single slow connection and users may decide not to trust a network that is unreliable.

The annual survey could be used to benchmark the progress in upgrading all components of the network and it could be a valuable tool for institutions in leveraging extra resources where they identify that they are falling behind a European norm or their peers. Alternatively, it could be used by institutions as a powerful supporting argument when bidding for research projects.

A possible list of items that could be measured annually and published is as follows:

- the number and percentage of users connected by Gb/s points;
- the number and percentage of users connected via shared media;
- the bandwidth of the central core network;
- the bandwidth of the inter-campus connectivity;
- the bandwidth to the NREN.

	Recommendation
7	An annual survey should be conducted to show the progress in advancing end-to-end network performance. The annual survey could be used to benchmark the progress in upgrading all components of the network and it could be a valuable tool for institutions.

Table 4.2: Recommendation on annual survey

Some thought should be put into the campus of the future. For example, will there be much of a wired infrastructure at all? Or will many connections be done wirelessly (where appropriate), using future generations of network technology?

Users in Class B and Class C will require the greater bandwidth that cabling infrastructure (optical fibre or copper) provides, but the needs of Class-A users can probably be met with wireless access. However, there needs to be a cautious assessment of pros and cons before making a large commitment to wireless networking.

The network architecture within institutions will need to be extremely well designed to accommodate the needs of the various classes of users cost-effectively using appropriate technology.

Similarly, modern networks must operate on a device-agnostic basis. Provided there are appropriate arrangements, end-users must be facilitated to connect any device to the network without prior clearance or approval of the network operations centre. Network Access Control technology makes it possible to do this securely and efficiently (see Section 5.6).

4.3 Security issues for campus networks

4.3.1 IT security policy

There is quite a variety of attitudes to IT security policies in the institutions that took part in the survey. Some institutions do not have any formal IT security policy. This can result in management imposing ad-hoc solutions to security issues as they arise. At the other extreme, some IT security policies are too formal and rigid; they are possibly translated from companies or other organisations in the business environment and they are not appropriate for the requirements of a research or education establishment.

IT security is a vital component for any organisation using information technology. For a research and education institution, whose main asset is its intellectual capital, IT security must be seen as indispensable. European, national and institutional regulations must be adhered to, as well as laws regarding data protection, privacy, confidentiality, copyright etc. Research contracts, whether they are funded from public or from private sources, require that much of the work is performed in a secure environment and that research results are not provided to competitors or leaked to unauthorised parties.

In passing it should be noted that where IT security audits have been performed by large consulting companies, the resulting recommendations are often not appropriate for research and education environments. These companies have a good knowledge of the business sector and its requirements, but they have limited understanding of which security policies are appropriate in the research and education world.

An IT security model should be based on some recognised standard such as ISO 27001, ISO 17799 or COBIT. Some countries have national recommendations for IT security and some of the university IT organisations in these countries (for example, UCISA in the United Kingdom) have specific recommendations and tools for their implementation in the research and education community.

Educating the user community about conforming to an institutional IT security policy is vital. Users must understand why such policies are being set and what they need to do to conform. Equally, the central IT security management must do all it can to protect its users from being compromised, which is usually caused unwittingly.

It seems that standard tools are increasingly deployed, such as anti-virus, patch management, and intrusion detection and prevention. However, there are institutions that do not seem to manage these actively. In order to be effective, these tools must be actively managed, preferably from a central point in the institution.

Particular attention should be paid to identity and password management. Good implementations of these will allow users to perform their tasks and duties more easily, without interference from repeated requests for passwords. Some form of two-factor authentication is recommended, but this may prove costly in large institutions.

Too often, security is given as an excuse for why things cannot be done; local management might use that excuse if they feel that a request is difficult to meet. The vague response *"security problem"* can be a throwaway negative answer that is difficult to question. These local decisions are sometimes made despite the fact that no formal policies exist, which makes the implementation of strong IT security policies even more important.

Clear policies, coupled with Acceptable Use Statements, should enhance the service offered to users and not necessarily be considered negative. The policies should be sufficiently flexible to recognise that the institutions are at the leading edge of research and development.

In the survey, no information was gathered about the existence of formal independent security teams in the institutions.

	Recommendation
8	Formal independent security teams within institutions should be formed and they should have a wide remit, with a considerable degree of independence from the central IT service.

Table 4.3: Recommendation on security teams

As mentioned before, the questionnaire on campus issues has resulted only in a snapshot survey to gather information for the EARNEST study.

	Recommendation
9	A survey of the status of IT security in research and education institutions should be carried out on an annual basis to measure progress as institutional IT services develop. Examples of best practice in the implementation of campus security policies could then be derived from the results of these annual surveys and could be published.

Table 4.4: Recommendation on annual survey of IT security

4.3.2 Firewalls

It appears from the survey that a reasonable proportion of institutions have firewalls installed; this is both a positive and a potentially negative finding. It is positive because of the protection provided by the firewall, but potentially negative in that the firewall may introduce performance limitations that are a hindrance to end-users.

The design of a firewall and its positioning in the campus network is of critical importance. Moreover, the performance of a firewall and its ability to handle high traffic rates must be examined closely. In

a research and education environment, the installation of firewalls with the highest specifications should be considered. Their day-to-day performance should then be monitored closely to ensure that there is sufficient scope. Research traffic can peak very quickly and result in short-term bottlenecks that may not be identified by most monitoring tools.

Today, powerful firewalls are available that are capable of performing deep packet inspection at high speed, but these are more expensive than the standard commodity firewalls that many institutions are installing. It would be advisable for institutions to install high-performance firewalls to complement the high-speed networks that the NRENs are delivering.

	Recommendation
10	Direct access should be provided, under controlled conditions, to researchers who sign an agreement to ensure that privileged access is used responsibly. Such direct access could allow traffic to bypass a standard firewall.

Table 4.5: Recommendation on direct access

4.3.3 Peer-to-peer

The survey questionnaire included a range of questions around the use of peer-to-peer protocols. The results could be interpreted as indicating a rather dictatorial style displayed by network managers. In approximately 25% of the institutions covered by the survey, all peer-to-peer traffic is forbidden, with no reference to the possibility of allowing 'good' peer-to-peer traffic.

There is a perception that all peer-to-peer traffic is associated with nefarious activity on the network, i.e., things such as illegal file sharing, breaching copyright legislation etc. Of course, there is some truth in this argument and much of the peer-to-peer traffic is actually unwanted and undesirable. However, there are numerous cases where peer-to-peer traffic is valid and valuable, for example, BitTorrent for large file transfers. Nevertheless, BitTorrent would normally be banned in a regime of total prohibition. Similarly, peer-to-peer protocols are used in legitimate systems such as Skype, but these would be banned too. It is wrong to assume that, just because something is new and has its roots in the networking underworld, it cannot be useful. Skype was developed by the same people who developed Kazaa, which is widely used for illegal file sharing.

A policy that states that peer-to-peer traffic is banned generally, but that individual people may request access to the protocols in a properly documented manner might be appropriate, provided that the policy is well advertised.

A better policy for peer-to-peer would be to formally allow all peer-to-peer traffic and not impose technical limitations, but to impose policy restrictions on personal use. This would mean that illegal and antisocial use of peer-to-peer could be prohibited while allowing proper use to take place. Attempting to stop all peer-to-peer activities may be difficult from a technical point of view.

Another solution to the abuse of peer-to-peer technology may be provided by traffic shaping or metering facilities in modern networks, although this may be more relevant for inhibiting users in Class A from abusing the network than for the high-performance users in Class C.

	Recommendation
11	Do not impose blanket bans on particular types of network traffic without careful consideration of the underlying activities that are being performed.

Table 4.6: Recommendation on banning types of network traffic

4.3.4 Network Address Translation

In the survey, a significant minority of institutions (42%) reported that few (or even none) of their end-user computers were given public IP addresses.

At one time there was a shortage of IP addresses available for IPv4 users. Clearly IPv6 will eliminate this problem when it is implemented by institutions. In the meantime, one of the solutions implemented by many commercial Internet Service Providers was Network Address Translation (NAT). NAT allows many devices to be hidden behind one public IP address. NAT allows users to avail themselves of a set of IPv4 addresses that are designated 'private', and these IP addresses are never propagated out to the Internet.

NAT is useful because it conserves IP address space, enabling even a whole university to live behind one IP address. However, another use has emerged, namely using NAT to hide internal services and machines from the Internet, usually for security purposes. This practice probably poses no problems for the more traditional Internet users and in the business world, but it is a severe hindrance for many researchers and services that require their network address to be propagated across the Internet.

While it is acceptable in many circumstances that most commodity Internet users, including researchers in Class A, can be hidden behind NAT, some special treatment must be provided for any user who requires more advanced services (e.g., videoconferencing) and there must be a way for their network services to be broadcast on the Internet. It would be acceptable to have a hybrid model, where the majority of the users in an institution uses NAT and the specialists have the possibility of obtaining public IP addresses. Once again, a rigid security policy that requires NAT with no escape clause would be counter-productive to a well-provisioned research network.

Many institutions have ample IPv4 address space available to them and users who require public address space should be able to justify its provision readily. Clearly, when there is a move to IPv6, shortage of address space will no longer be a valid reason for hiding behind NAT.

4.3.5 Private network links for researchers

Many network users feel that institution networks constrain them in their research or teaching, because there are usually restrictions imposed by local policies. Sometimes they are tempted by the unrestricted network connections that can be provided by commercial Internet Service Providers, thus avoiding the restrictions of the institution's central policies. There is some limited evidence that users have purchased ADSL connections to avoid such policies.

Similarly, users at home with their own 'broadband' connection, frequently using ADSL, may feel that they have much greater freedom and access to the network than they have at their institution. Indeed, it is probably true that there are far fewer restrictions on these sorts of networks.

Most telecommunications operators have a privileged legal standing; they may not have responsibility for the activities of the end-users of their network and therefore they do not have the need to restrict users by imposing policies. In general, research and education institutions do not have such legal immunity and they have direct responsibility for the activities of all users on their network. Hence the activities of a single staff member or student could result in legal action being taken against the institution. An offended person or body may find it much more profitable to take legal action against an institution than against a financially weaker individual. As a consequence, strong policies governing network use are required by institutions and these policies have to apply to all users, whether they use the central network or lease links directly to their research unit. The legal implications cannot be evaded.

It is very important that researchers are provided with tools and solutions that meet their networking needs and that give them any special network facilities in a monitored and controlled manner. Special rules should be laid down for them and they themselves should be responsible enough to ensure that they comply and conform. Clearly, breaches of the rules should result in their specialist privileges being withdrawn.

4.4 Obstacles and bottlenecks

Few respondents to the survey questionnaire said that their campus networks were the root cause of bottlenecks. Bottlenecks arise for many different reasons: obsolescent cabling and equipment in buildings was often quoted, as well as obsolescent core network equipment (routers, switches, etc.). Other reasons include a shortage of networking expertise, links between buildings and between campuses, and security measures. However, no dominant reason emerged from the survey.

Many respondents thought that raising the awareness of senior management in institutions would significantly help resolve some of their difficulties. It would also help if guidelines and recommendations were published at European or national level. This would assist institution managers to understand better the benefits of up-to-date network services, and how to acquire them.

The idea of benchmarking by national or European organisations was also welcomed by respondents. Finally, the same holds for setting up formal procedures for identifying and incorporating the needs of end-users into networking policy and funding plans.

	Recommendation
12	Strengthen the collaboration between National Research and Education Networking organisations and institutions to improve the deployment of key services: share strategic information, raise awareness of innovative services at senior levels, co-ordinate working groups, and obtain feedback from end-users and especially from those with demanding requirements.

Table 4.7: Recommendation on collaboration between NRENs and institutions

4.5 Network services provided to end-users

When establishing policy for the provision of network services at an institution, one should take into consideration the different categories of usage of the network. Therefore, categorising end-user requirements[4] is an important element of the strategic process to gather pertinent inputs for

4. *Baseline usage, medium usage, heavy usage: see Section 2.1 for more details*

policy making. Only 10% of the institutions that responded to the survey questionnaire include in their policy-making processes actions for identifying demanding or very demanding projects and appropriate procedures for requesting particular services. In many cases there are mechanisms for taking informal soundings, but these are not totally satisfactory.

The survey questionnaire mentioned 26 services, which for the sake of clarity were divided into three categories: 'upper-layer services' close to the applications, 'lower-layer services' that depend on the basic network infrastructure (links, switching and routing equipment etc.) and 'middleware and security services'. IPv6, multicast and Quality of Service were dealt with separately (see Section 4.6).

For each service, it was asked if the service was provided to students and/or staff. In case a service was not provided, it was asked why that was so. Possible reasons could be internal (related to the organisation of the institution, policy, security considerations, lack of expertise, lack of resources etc.) or external (i.e., reasons beyond the control of the institution, e.g., the NREN or the intermediate network not supporting the service).

4.5.1 Upper-layer services

Almost all respondents to the survey questionnaire run personal email services as well as mailing-list management services and Web services (Web storage, WebDAV, wikis, collaborative environments, etc.) for staff. For students, things are a bit different; they usually have personal email services, but only half of the respondents provide mailing-list management and Web services for students.

Between 50% and 75% of the institutions participating in the survey provide videoconferencing and voice-over-IP (VoIP) for staff. Audio/video streaming, Grid-dedicated services, services dedicated to virtual organisations and HDTV (High-Definition TV) are provided by fewer institutions, the proportions decreasing from 41% to 5%. For the same services being provided to students, the figures are much lower.

4.5.2 Lower-layer services

Wireless connectivity is almost ubiquitous. If the survey were to be repeated one year later, coverage would probably reach 100%, for staff as well as for students. Wireless roaming services (e.g., eduroam) are progressing, but so far they are provided by only 40% of the institutions, which means that there is still some way to go to reach a global service. IT managers possibly consider the provision of VPNs to be a good alternative to eduroam, because VPNs are provided in the majority of institutions.

Bandwidth on demand and provisioning of lightpaths are still very rarely provided (8% to 15% of institutions).

4.5.3 Middleware and security services

Authentication and authorisation for network access, environments for securing email (antivirus, anti-spam etc.), protection against denial-of-service and other attacks and backup mechanisms to ensure reliability are among the services run by most institutions, for students as well as for staff. They are 'traditional' or 'classical' services, and IT managers and chief security officers increasingly consider them to be essential if not mandatory.

It is not satisfactory that only about 60% of respondents provide confidentiality of data communications and disaster recovery. These services should be considered essential for protecting the assets of the institution and ensuring continuity of crucial activities.

Identity federation services are one area on which much emphasis has to be put. Only one third of the institutions in the survey are running such a service. It is a 'young' area, and it needs to be strongly supported.

4.6 Rollout of videoconferencing, multicast, IPv6 and QoS

4.6.1 Videoconferencing services

Videoconferencing services have been available in research and education institutions in Europe since the early 1990s. Early versions were expensive, but in the last ten years the arrival of IP connectivity, reduced hardware costs and better video-compression algorithms have changed the scenario enormously.

In spite of these improvements, the current use of videoconferencing is below expectations. There are several reasons for this, the main one being that videoconferencing is not as easy to use as the telephone, which is often a perfectly acceptable alternative. There is also evidence that many end-users are not aware that their institutions offer a videoconferencing service.

The use of videoconferencing is likely to grow, but at an unpredictable rate. To increase its use, institutions should take steps to promote awareness of their videoconferencing services, provide suitable training for users and offer good technical support, especially in the early stages of implementation. This should include not only the network-related aspects of the service but also the audio-visual set-up of the room: recommendations about lighting, projectors, locations of cameras and microphones etc. Once users become acquainted with the service, the level of support can be reduced.

Embedded systems should be made available for group videoconferences in two versions: specially arranged ready-to-use rooms for virtual meetings conveniently distributed through the institution, and portable units ready to be installed on request when necessary, e.g., for remote participation in conferences, symposiums etc. Those facilities should be managed by the IT service department and leased to the users on request.

The use of Multipoint Control Units (MCUs), either owned by the institution or leased from an external company, should be considered part of the service. Sharing MCUs between institutions should be encouraged, either at regional or at national level (in the latter case the NREN itself should help set up the service).

Most potential users are not aware of the possibilities and benefits of videoconferencing. In order to increase its use, 'demo' dissemination sessions could be organised as part of user meetings and other public events. Training courses oriented towards end-users should be organised frequently, and documentation for non-specialists about the use of videoconference services should be made readily available to everybody, for example, in the form of webpages.

	Recommendation
13	Institutions should actively promote videoconferencing and provide full technical support in the early stages to users of videoconference systems, especially to untrained occasional users.

Table 4.8: Recommendation on videoconferencing

4.6.2 Multicast

The history of multicast is similar to that of videoconferencing. The technology was first specified in 1989. There were expectations that it would quickly become widely used, especially for multipoint videoconferencing services and video streaming. However, multicast did not become popular with commercial Internet Service Providers, and many of those providers still do not support multicast. At the present time, multicast is not widely used in the research and education community. The main application for this technology would be video streaming, but there appears to be little demand from potential users. According to the 2006 edition of the TERENA Compendium of National Research and Education Networks in Europe, only 60% of NRENs offer multicast. Those that do not offer multicast tend to be the smaller NRENs.

The major benefit of multicast is reduced load on networks and servers. However, there are several drawbacks, namely :

- complex and unreliable routing protocols;
- complex and incipient address allocation schemes;
- low levels of security.

These issues make commercial Internet Service Providers reluctant to adopt multicast. The future growth of multicast is very difficult to predict, but multicast is likely to become more popular when the current drawbacks are resolved. As in the case of videoconferencing, the issue of raising user awareness needs to be addressed.

	Recommendation
14	NRENs should encourage institutions in their country to implement multicast in their networks and organise awareness campaigns to promote its use and increase the multicast content in their networks. For example, some universities already have TV channels on the air with a scope that is in most cases limited to the internal network or the metropolitan area. Those TV channels could be delivered by multicast with good-quality video using 4-5 Mb/s and be made available worldwide in a very simple way and at very limited additional cost.
15	In those cases where the NREN does not provide multicast, a plan should be devised to introduce support of multicast gradually in the network, with the goal of making multicast available to all the connected institutions.

Table 4.9: Recommendations on multicast

4.6.3 IPv6

The story of IPv6 is similar to the stories of videoconferencing and multicast. In this case, protocols have been around since 1995, but they have not fulfilled expectations.

The pan-European research backbone network GÉANT2 provides complete support of IPv6, but it is not clear whether all NRENs do the same. In 2004, the TERENA Compendium of National Research and Education Networks in Europe showed that 70% of NRENs supported IPv6, but there is no later information available. Similarly it is not known whether intermediate networks support the protocol. It is possible that an institution and its NREN are able to offer IPv6, while the intermediate network does not offer it.

According to the EARNEST survey of campus issues, only a small proportion of academic institutions in Europe have deployed IPv6 in their core networks. The main reason for not deploying IPv6 is again the lack of end-user demand. Other reasons are lack of compatible hardware or software and lack of training.

It must be acknowledged that IPv6 has essentially only one benefit: a much bigger address space. The advantage of a bigger address space is not especially attractive for research and education institutions in Europe, because IPv4 currently meets their needs. Moreover, a significant number of those institutions are using NAT in their networks today, presumably as a basic firewalling procedure (blocking incoming connections; see Section 4.3.4) rather than as a mechanism to save address space. But a side effect of using NAT is to reduce the need for IP addresses, thus reducing the pressure to migrate to IPv6.

Migration to IPv6 has to take place[5], because the Internet continues to grow quickly, and with it the need for more address space. Moreover, more and more host implementations support IPv6. The use of NAT is not a satisfactory solution since it is not completely transparent to the applications and requires extra effort from software developers.

It will not be possible to put off implementing IPv6 indefinitely. Institutions should begin making preparations to migrate soon.

	Recommendation
16	All NRENs and intermediate networks (MANs and regional networks) should provide IPv6 services natively to allow institutions to prepare for a smooth transition to IPv6.
17	Institutions should provide their users with IPv6 services at all layers (network, transport and application) at least in their core network, and ideally in the entire network. They should request address allocation from the responsible entity (either their NREN or the RIPE NCC). Their IPv6-enabled network should be connected to the rest of the IPv6 Internet.

Table 4.10: Recommendations on IPv6

4.6.4 Quality of Service

Quality of Service (QoS) is implemented by a set of protocols (e.g., Diffserv) that are at least 10-15 years old but do not come anywhere close to meeting initial expectations. However, there are some peculiarities in this case.

According to the survey, a significant proportion of academic institutions in Europe have deployed QoS inside their campus networks, but in only a few cases is QoS available in the connection to the Internet as well. In cases where there is no QoS, the main reasons are no need (i.e., there is plenty of bandwidth available) or no demand from users.

The most important application using QoS is IP telephony, followed by videoconferencing and multimedia streaming. In most cases, IP telephony has probably been the main reason for deploying QoS. This would explain why so many institutions have QoS support inside their network only: IP telephony normally generates IP traffic only inside the institution network and uses gateways to connect with the traditional telephone network for external calls.

In most institutions, IP telephony is a relatively new service offered by IT service departments. Therefore one can predict increasing deployment of QoS in campus networks in the near future as these services become more popular, but the deployment is likely to be restricted mainly to the institution network. In a later phase, QoS could also be implemented at the NREN level if institutions start to exchange IP telephony traffic natively.

5. *See the recent statement by the RIPE community at www.ripe.net/news/community-statement.html*

In order to facilitate the implementation of QoS services, institutions should ensure that their network equipment supports the relevant standards and protocols. They should also carefully evaluate the benefits (and possible drawbacks) of having equipment from one single vendor in the core, especially from the point of view of simplicity of configuration, management and maintenance.

4.7 Operational structures and control of services

About 80% of the institutions participating in the survey have central Network Operations Centres (NOCs). Most of them provide services like incident handling, metering, performance monitoring and testing. However, only 65% of them have formal links with the NOC of their NREN, and less than 50% have control of the complete institution infrastructure. 'Peripheral' parts of the campus network (cabling, local servers etc.) are often the responsibility of departments, faculties or laboratories, and not of the central NOC. In such situations it is extremely difficult to take a global vision of end-to-end services and to ensure appropriate levels of security.

	Recommendation
18	Central NOCs in institutions should have formal links with the NOC of their NREN, and a formal responsibility to manage the campus network.
19	The special needs of end-users should be taken into account when defining networking policy. There should be formal procedures for end-users to request special services.

Table 4.11: Recommendations on operational structures and control of services

4.8 Obstacles to seamless end-to-end connectivity

Inevitably in this complex multi-domain topology there are sometimes obstacles to seamless end-to-end connectivity. Examples are:

- **Shortage of bandwidth.** In the survey it was found that some institutions of moderate size have low bandwidth to the NREN. The majority of these are connected to their NREN through intermediate networks, most often a regional network. Regional networks have evolved significantly in recent years, and current initiatives (in the United Kingdom, France and Spain) appear to show that these infrastructures are being upgraded to Gigabit speeds. This will be sufficient for users in Class A and possibly for users in Class B or C. However, there may still be a limitation in the near future for institutions with Class-C users who require 10 Gb/s or even more. The NRENs will soon be able to deliver such throughput, but regional networks risk lagging behind unless they make the necessary investment.

- **Lower level of reliability and performance.** Each network domain has its own Network Operations Centre. Multiple layers of management inevitably increase the risk of failures of communication, making co-ordination sometimes difficult. As a consequence, service disruption may take more time to be fixed. There is no common set of tools or procedures to identify the cause of poor performance or to monitor end-to-end connectivity. If the different levels of infrastructure are under the same management responsibility, the problem diminishes.

- **Discontinuity of services.** Harmonisation of services at layers 2 and 3 of the OSI model is crucial to facilitate and enhance the collaboration of researchers at the national and international level. Any unavailability of service may undermine the participation of researchers in important projects. The following services are offered by NRENs, and most of them should also be provided by intermediate networks:
  - IPv6, multicast and QoS should be available throughout the research and education networking infrastructure, but their deployment is not always guaranteed in every component between the NREN and the campus of the institution.
  - Virtual Private Networks. In some cases, only the NREN is providing Layer-2 and/or Layer-3 MPLS services as well as Traffic Engineering facilities. These are the basis for virtual networks for specific projects or communities. Layer-2 VPNs could also be used as a platform by the NREN to directly connect institution sites across a regional network, for example. In that way they would be able to provide all the services to end-sites that they normally can offer only to sites with a direct physical connection.
  - Bandwidth on demand is requested mainly by users in Class B or C. It is crucial to support projects that request end-to-end guaranteed bandwidth. Researchers who are not able to gain access to this facility may be barred from participating in national or international strategic scientific projects. Bandwidth on demand is provided by the majority of NRENs, but it is not very often provided by intermediate networks. This holds in particular for regional networks that employ transmission technology and link capacity that are not capable of delivering this service.
  - Lightpath allocation could be thought of as a particular type of bandwidth on demand. However, lightpaths offer very powerful additional features: the ability to multiplex several transmission channels per lightpath, each one offering guaranteed speeds up to 10 Gb/s or even 40 Gb/s; a neutral transport in a transparent way of any type of Layer-2 protocol; bypassing costly routing infrastructure etc. Undoubtedly, lightpaths open large windows of opportunity for the scientific community. The number of NRENs offering lightpath allocation is growing steadily. Intermediate networks must address this development more seriously.

  The problem of intermediate networks not offering the same level of services as those provided by NRENs could be addressed by actions such as the following:
  - The senior management of local and regional networks should organise serious evaluations of the specific needs of researchers, identify what the scientific stakes are, consider the strategic importance of the services mentioned above, and include the proper requirements in the procurement documents for future calls for tender.
  - NRENs should co-ordinate strategic initiatives better, with the aim of raising awareness of the importance of innovative network technologies among managers of local and regional networks.

In future research and education networking policy it must be a priority to guarantee end-to-end functionality, including bandwidth, services, supervision etc. To reach this goal in a complex networking infrastructure of many domains, there should be strong co-ordination of the activities of NRENs, regional networks, MANs and institutions.

Closer collaboration between domains is needed:
- to identify the requirements of researchers, teachers, students etc. before finalising the main strategic policy orientations, architecture, topology, services etc. of each project;
- to plan greater homogeneity and continuity of technologies and services at each level of the

global infrastructure; this may mean having common elements in procurement documents;
- to set up common operational procedures with formal interfaces between NOCs, and common or interoperable tools for performance metering and monitoring, incident handling etc.

	Recommendation
20	The specific requirements of end-user communities should be recognised and incorporated into the operational specifications in the different domains serving those communities.
21	Continuity of the services provided by NRENs should be guaranteed by all the networking domains involved in order to enable seamless end-to-end connectivity when a particular quality of service is requested.
22	The interface between different management domains should be formalised, e.g., by common Service Level Agreements between domains and especially between regional networks, MANs and the NREN.
23	High priority should be given to the following services: IPv6, multicast, QoS, MPLS (Layer 2, Layer 3) as well as Traffic Engineering.
24	Infrastructures for optical continuity (links as well as basic optical equipment) should be planned in areas where communities of users are identified as potentially heavy users of network services.
25	There should be strong and formalised structures for collaboration between the different management domains.

Table 4.12: Recommendations related to obstacles to seamless end-to-end connectivity

4.9 Connectivity to sites served by commercial Internet Service Providers

In recent years, research and education networks in Europe have improved their capacity by several orders of magnitude. There has been significant investment in infrastructure and technology by many organisations and at different levels: campus networks, intermediate networks (MANs or regional networks) and NRENs as well as at the European level. At that level, DANTE set up various generations of the pan-European research backbone network: EuropaNET, TEN-34, TEN-155, GÉANT and GÉANT2. DANTE also provides connectivity to (and between) several national research and education networks in North and South America, the Mediterranean region and Asia.

As a consequence of these ongoing improvements, today the vast majority of the research and education community in Europe enjoys excellent connectivity with most of their peers in the rest of the world, with short delays, high bandwidth and high-quality services operating on a 24x7 basis.

However, the situation is not so good when the remote destination is not within the research and education community. Sometimes there is a noticeable increase in response time when the user tries to reach a server located in the commercial part of the Internet. The fact that the user is accustomed to short delays within the research and education networking environment makes the experience even more frustrating.

Long response times may be caused by bottlenecks somewhere in the network path or by an overloaded server. In many cases, solutions to the problem are outside the scope of the research and education networking domain. But in spite of the complexity of the problem, it is important to resolve it, because from the user's point of view it does not matter where the delay is. He/she is suffering poor network performance and therefore a low quality of service.

To overcome this problem, caching proxy servers are sometimes deployed. If the hit ratio of the proxy server is high (i.e., if there is high repetition in the files downloaded by the users), this is a good

solution. However, proxy servers have several important drawbacks that outweigh the benefits in most cases. These drawbacks include the following:

- Nowadays, many pages are dynamic and non-cacheable, which reduces the hit ratio significantly. A list of known HTTP proxy/caching problems is compiled in RFC 3143.
- The hard-disk performance of the proxy server becomes a limitation in high-load conditions.
- If Internet access is configured through a proxy server, the server becomes a critical part of the service.

	Recommendation
26	Do not use a caching proxy server, except in very specific conditions when the benefits clearly outweigh the drawbacks.

Table 4.13: Recommendation on caching proxy servers

Another way to improve the response time in low-bandwidth situations is to use Content Delivery Networks (CDNs). While caching proxy servers are reactive, i.e., they bring the content when the user requests it, CDNs are proactive, in the sense that they try to bring the content in advance to have it located near the user before he/she requests it. Unlike proxy servers, which are configured and maintained exclusively by the hosting institutions, CDNs require close co-operation and co-ordination with the content provider in order to be effective.

There are some obvious similarities between proxy servers and CDNs, and some of the drawbacks of proxy servers also apply to CDNs. However, it is generally accepted that the CDN approach is more efficient because it is not completely automatic. CDN is managed by the content provider who knows which content is cacheable and thus can take an informed decision to make the process more efficient. Moreover, the proactive strategy of CDN brings the benefit of a fast response even for the first user.

Some European NRENs (e.g., RedIRIS in Spain, SWITCH in Switzerland and RENATER in France) have agreements for hosting Akamai servers in their network nodes. Akamai decides which content is replicated in the servers, and the NREN decides which users can access the servers. Having a CDN infrastructure owned by the NREN would allow the NREN to decide the content according to the interests of the users.

	Recommendation
27	Each NREN should consider establishing a pilot CDN infrastructure to test the scope for improved response times in accessing popular content in the commercial part of the Internet.

Table 4.14: Recommendation on Content Delivery Networks

4.10 Software restrictions

Many research and education institutions in Europe have network infrastructures providing connections at 100 Mb/s or even Gigabit Ethernet to the desktop. However, users are sometimes not able to obtain the full benefit of the bandwidth available, due to limitations in end-to-end communications.

The limitations happen for a variety of reasons, including problems in the communications software (typically at the transport or application layer), poor system performance produced by some

operating system component, or even hardware issues. Fine-tuning of the software is needed to improve performance, but institutions often do not have sufficiently skilled engineers to do the work.

	Recommendation
28	Careful selection of protocols and implementations is recommended to avoid bandwidth limitations. Skilled tuning of software and systems can improve performance considerably. However, this requires expertise beyond the skills of the average user. Because the user gets his/her work done anyway, he/she may not even be aware that there is a problem and does not seek specialist support.
29	Institutions and NRENs should provide support, training and documentation about performance optimisation issues, especially for the needs of users in Class C.

Table 4.15: Recommendations related to performance optimisation

4.11 Governance, funding and operations of network infrastructures and services

4.11.1 Connections to the NREN

The majority of respondents to the questionnaire for the EARNEST survey of campus issues were connected to the Internet via an NREN, either directly (50%) or through an intermediate network (40%). A small proportion use connections obtained from a commercial Internet Service Provider. The funding of NREN connections varied considerably; only in a minority of cases was the connection to the NREN fully subsidised by sources outside the institution.

Respondents were asked to identify the benefits of NREN services. The most often cited advantage was the ability to get more capacity and higher bandwidth at reasonable prices and in a flexible manner, without having to undergo time-consuming negotiations with Internet Service Providers.

Most NRENs provide more services than Internet Service Providers do, and the services provided are of better quality. The services are normally state-of-the-art, which is often not the case for services offered by Internet Service Providers. Moreover, NRENs often give some flexibility to run extra services or pilot services that are not yet considered to be operational. The services are tailored better to their connected institutions. Internet Service Providers are often said to be more business-oriented. NRENs generally know what an institution's business is, and are more aware of the needs of the academic community.

4.11.2 Network management, policy design and funding at institution level

The majority of institutions in the EARNEST survey of campus issues have an Acceptable Use Policy (AUP) that is supported by the governing body of the institution. However, a significant minority (26%) either have no AUP or, if it exists, the AUP has not been approved at a senior level in the institution. This should be a cause for concern, because it suggests that the institution's senior management does not recognise the strategic importance of networking. It is also significant that only about 30% of respondents formally consult their end-user community about networking policy. In some institutions even the head of IT services does not appear to know how to influence the institution's networking strategy.

Interestingly, in response to a question about the effectiveness of recommendations and guidelines from NRENs, the majority of respondents said that they believed such guidance had real influence on their institution's networking policy.

	Recommendation
30	Each institution should publish an Acceptable Use Policy document, which should be communicated to end-users in order to make them aware of the services available to them and which should state the end-users' rights and obligations regarding the way in which they use those services. The governing body of the institution should endorse this document.
31	End-users should be represented in the process of developing policy documents that directly impact on services offered to them. National recommendations and guidelines for best practice would help to create such procedures inside institutions.
32	National recommendations are perceived to be influential in developing institutional networking policies, and the creation of such recommendations should be encouraged where they do not already exist.

Table 4.16: Recommendations on the development of policies

4.12 Organisation of networking expertise, staffing and training

The level of quality and performance of network services at institution level, as well as their range of functionality, depend crucially on the size and the level of expertise of the network team that supports these services. The EARNEST survey of campus issues found that there is a wide diversity of situations in this area. Therefore it is impossible to draw general or generic conclusions. In this section, only a few observations are highlighted, which identify issues for serious consideration.

About 66% of respondents to the questionnaire felt that their networking team was not adequately resourced. In 50% of the institutions in the survey the size of the networking staff was less than 5 fulltime equivalents. There were several large institutions in this group.

A majority of the IT managers who responded to the questionnaire said that they experienced difficulties in recruiting staff with a good level of expertise. This makes the network service vulnerable in key areas such as security, installing new services, performance monitoring, informing users of new functionalities, and keeping pace with developments in technology.

A significant minority of the institutions in the survey provide very little training for support staff. The majority do not have systematic training plans for staff in IT service departments.

	Recommendation
33	There should be a clear link between an institution's networking policy and the resources available to deliver its objectives. The policy should be drawn up by senior management, working closely with end-users and the network support team. The policy should specify the services to be provided to the end-users and the quality of these services (level of reliability, coverage in time etc.), providing a 'charter' between the institution and the end-users. The network support team should be adequately staffed and have appropriate expertise.
34	IT keeps evolving at a fast pace. Technicians and engineers frequently have to adapt to new technologies (equipment, software, protocols, architectures, services etc.). It is vitally important that they are properly trained to carry out the tasks expected of them.

Table 4.17: Recommendations related to network support staff

4.13 Raising awareness, and training and support for network users

The EARNEST survey of researchers' requirements has provided strong evidence that many network-related services are underutilised[6]. Researchers, teachers and students are either not aware of the services currently available or they make only limited use of them.

For example, the survey of researchers' requirements showed that only a minority of respondents uses IP telephony, videoconferencing, bandwidth reservation and encryption of data, and then infrequently (i.e., less than once a week). When asked about their use of specific network-related tools, a significant number of respondents replied that they did not understand some of the terms used, implying they were not aware of the availability of some commonly available network services. The survey also asked respondents to assess aspects of the network infrastructure within which they worked. Only 41% agreed they received adequate training in network use to improve the quality of their research or teaching.

There is supporting evidence from the EARNEST survey of campus issues that, although it is technically feasible to provide many network services, they are not available as widely as might be expected. Such services include eduroam, multicast, bandwidth on demand, QoS and PKI (Public Key Infrastructure).

4.13.1 Possible reasons for low level of use

There are a variety of reasons for the under-utilisation of network tools and services, depending on the local situation. They include the following reasons provided by participants in the survey:

- A network service is available, but it is not publicised and many users do not know of its existence.
- A service is available, but it is badly supported because the network support team is too small to support the service or there are not enough network staff members with the right level of skills.
- A network service is not available, even though it is technically feasible to offer it. Some service providers at institution level are reluctant to offer a service, such as videoconferencing, because they know they will not be able to cope with the demand for support.
- The network infrastructure is incorrectly configured and thereby prevents the service being offered, sometimes for questionable reasons of 'security'.
- There is a culture of low expectations in the user community which inhibits users from demanding new services.
- The service provider has not adapted to its new role of offering services rather than basic network connectivity.
- There is a lack of customer focus in the organisation providing services.

4.13.2 How to raise awareness and how to train and support end-users

To improve the take-up of network services there is clearly a need for better dissemination of information about developments in this area. There should be more publicity, better training and higher levels of support for many of the network-related tools and services that are currently available. A relatively modest investment in promotional videos, online training material and brief explanatory documents would yield enormous benefits throughout the European higher-education sector.

6. *See the EARNEST report on researchers' requirements by Thibaut Lery and Patrick Bressler.*

It is also evident that some institutions have not increased the size of their network teams to enable them to support network-related tools and services adequately. Some institutional network teams are too small and cannot hope to provide more than basic support for network connectivity. They spend much of their time 'fire fighting'. Even if they wish to introduce new services, many wisely hesitate to do so because they are already overstretched. In recent years, institutions have invested significantly to upgrade their campus network infrastructure. However, there is evidence that many do not have enough human resources to exploit fully the benefits of their increased investment in the technology.

It is also important that end-users and network staff are properly trained to use and support the new services. Some training materials will be specific to the situation in individual institutions; others are more generic and could be developed for national or even international audiences.

It should be self-evident that 'human networking' is an important element in the research and education networking environment. At the strategic and management levels there is much co-operation and collaboration. However, there should be closer collaboration at grassroots level too, in order to share skills and keep up to date with new developments in networking technology.

Finally and most importantly, there is a need for cultural change. There are signs that some campus network teams have not extended the scope of their service beyond providing connectivity. A paradigm shift in the nature of network support is taking place, from providing connectivity to providing network-related services, and some network support teams and their user communities do not appear to have recognised this.

	Recommendation
35	Promote awareness of network-related tools and services in both the end-user community and the organisations providing services to them.
36	Provide well resourced skilled network teams in institutions to provide network-related tools and services and to support the users.
37	Encourage all network support staff to collaborate with colleagues at other institutions, regional networks and NRENs.
38	Encourage a change of focus from connectivity to network-related services.

Table 4.18: Recommendations on awareness raising and the support of end-users

4.14 Collaboration and dissemination of information

4.14.1 Collaboration between NRENs and institutions

In most countries there is a formal arrangement (or there are several formal arrangements) at national level to bring together the network experts of NRENs and institutions. The purpose of these arrangements is to foster collaboration, ensuring a better process of information dissemination and the exchange of experience and expertise.

Sixty percent of the institutions in the EARNEST survey of campus issues reported that their network experts meet their NREN colleagues once per year or even less often, which gives little opportunity for collaborative work. There seems to be a pretty good level of dissemination of information about announcements and events of all kinds. However, it seems that the results of workshops are not very well publicised, or perhaps workshops are rare events in some countries. Information dissemination

and workshops are crucial for the coherent development of network services in every domain (NREN, MANs, campus networks), especially as research networking is a fast changing branch of information technology.

The survey shows that there is little feedback from end-users to the NREN. If such feedback is obtained at all, it may certainly be improved. This would help the NREN as well as the intermediate network to leverage the quality, reliability and security of their services. According to the survey, incident handling, security, troubleshooting, performance, videoconferencing and authentication are among the highest on the list of topics that are submitted as problems to the NREN. It is not clear how many NRENs have formal channels to collect feedback from the user community, but such formal channels would be welcomed by end-users. Better channels of communication would be beneficial all round.

4.14.2 Collaboration between network engineers from different institutions

The most surprising result of this part of the survey came from answers to a question about the level of collaboration between engineers from different institutions. Only one third of respondents considered there was enough collaboration, and half of the others said that levels of collaboration had not improved in recent years. However, many respondents recognised the potential benefits of closer collaboration with their counterparts in other institutions. In particular, they said that their own expertise would be increased by sharing experiences with their peers.

	Recommendation
39	Organise more regular workshops between NREN engineers and institution engineers on strategic issues for the future of networking in the research and education community. In case no such workshops exist, the NREN should take steps to create regular cycles of workshops on critical issues.
40	Set up formal procedures (processes) for gathering feedback from end-users in order to improve the quality and reliability of services. The creation of such procedures relies on networking management decisions in each institution. End-users will not spontaneously guess how to contact their NREN and transmit their complaints or proposals for improvement.
41	NRENs, funding bodies and national organisations that have administrative supervision of institutions for research and higher education should encourage institutions to organise better the collaboration between engineers and networking managers from institutions.

Table 4.19: Recommendations on collaboration and information of dissemination

5. Survey of emerging network technologies

This chapter must start with an extensive note of caution. It is extremely difficult to identify network technologies that are likely to revolutionise campus networking in the next 5-10 years. Many technologies are evolving constantly in an incremental way and it is not clear which of them will endure and which will bring much additional functionality to the end-users. Experience elsewhere in the field of information technology indicates that there could suddenly be a dramatic change in the networking landscape that no one is predicting today.

A few examples may help to illustrate the complexity of the issue:

- Vint Cerf, known as 'the father of the Internet', acknowledged that he never imagined what a revolution the Web would bring to the Internet.
- Wireless networking has been a major innovation, but six years ago not many people foresaw that it would be a key technology for communicating with mobile phones and PDAs, for instance.
- Audio-video technologies have been around for a while, but videoconferencing has not yet inundated the Internet. However, YouTube, launched just three years ago, is now the largest video transmitter over the Internet, using technologies that cannot be considered to be innovative.

The observations above imply that institutions need to be ready to react when a new technology (or even a novel use of an existing technology) comes along. Robust replacement policies and constant monitoring and evaluation will help follow the path of technological development. Institutions that are not ready to respond risk losing ground in an increasingly competitive world.

It is important to observe that innovations in the use of networking stem more from the application layers and from the way of packaging and delivering contents than from evolutions in basic transmission technologies or even, at least until now, middleware. These innovations are often triggered by new business concepts using the network, rather than new networking technologies as such, but they do require the transmission technologies used to be upgraded to support increases in traffic.

End-users usually have the ability to assimilate innovations automatically, with occasional technical assistance or training from their institution's network team. This can be achieved because such a new service often has consequences only for the end-user's desktop (new application software), or at most for the local server of the department. There may be no immediate impact on the wider campus network.

In the long run, the cumulative effect of the introduction of such new applications may cause a substantial increase in network traffic and put some pressure on the institution to upgrade its core network infrastructure as well as the bandwidth of the connection to the NREN. This situation often leads the IT management (or the network team) to react by implementing tougher security measures (e.g., changing the firewall rules) to 'contain' the explosion of traffic.

From what is said above, it is easier to understand that campus networking expertise should focus primarily on introducing and/or upgrading technologies (software, equipment, links) that cannot be manipulated directly by end-users. The functionalities expected from these technologies include

ubiquitous connectivity with mobile facilities, availability of higher bandwidth, seamless end-to-end services with QoS for specific applications, identity management to support collaborative platforms etc.

The main aim of the second phase of the EARNEST study of campus issues was to look at evolving network technologies that are likely to have a significant impact on networking in institutions, and to predict the future of campus networking as a consequence of those developments. An additional aim of this phase was to identify the requirements that need to be met so that institution networks and NRENs can benefit from these new technological developments, and to draw from this analysis a list of recommendations for best practice for the near future.

All areas of network technologies were considered for investigation. However, emphasis was put on technologies that support services at the interface between an NREN and a campus network, because this is the place where innovative services provided by NRENs might have difficulty 'crossing the border'.

The Internet was originally designed and built by and for the research community, and it has rapidly become a vital tool for sharing resources, knowledge and information, and for strengthening collaboration in research and higher education. At the European level, GÉANT2 and the NRENs are vital components of the global Internet for research and education. The institution networks that they interconnect are themselves crucial parts of the overall infrastructure, and the quality of service offered by institution networks to end-users relies greatly on the profile and performance of technologies deployed at campus level.

Distributed applications, communication services[7] or network services available to end-users depend on functionality from several different categories of technologies. The technologies needed to run an application end-to-end will depend on its degree of sophistication, its level of security, its capacity to be distributed, the performance and quality of service offered etc. The variety and number of fields of technology required may be very large.

For the purpose of the study it was necessary to distinguish between what can be identified as 'application-layer technologies' and what can be considered to be 'network technologies'. This distinction was made for two main reasons: to limit the size of the investigation and to focus primarily on technologies that come within the scope of institution networking.

Technologies for institution networking are shared across the campus. As such, they cannot be provided and run by individual end-users. The institution network is the responsibility of a dedicated central network team and requires strong involvement from the institution's management to set up and implement various policies for delivering network services to end-users.

The current chapter deals with a selected subset of technologies (either shared technologies or those under the management of network teams) that are considered to be innovative and/or under evolution.

Because there is no clear evidence of innovative technologies on the horizon that might profoundly affect campus networking, it has been necessary to devise another strategy to select fields of technology for investigation. In addition to the selection on grounds of the criteria mentioned above, the areas of technology investigated in this chapter have been selected because they are considered to have strong end-user requirements:

7. *For the majority of end-users it is more appropriate to use the terms 'communication services' or 'collaborative tools' rather than 'network services', because they are not aware (or should not be aware) of the details of the network environment but are in contact only with the upper-layer service. Only end-users in Class C are likely to have a closer relationship with the environment of network services and usually they work together closely with their colleagues in the central networking team.*

- **Lower-layer technologies for the transport infrastructure:**
  GÉANT2 and NRENs currently have the ability to provide high-performance dedicated end-to-end services for specific projects or communities. Examples are high-bandwidth capacity, guaranteed bandwidth, bandwidth on demand, performance monitoring, incident handling etc. Campus networks have to identify end-users with demanding applications and start to plan for the deployment of the following technologies:
  - continuity of optical technologies (links and equipment);
  - technologies for bandwidth management over the optical layer:
    - Ethernet technologies, GMPLS;
    - control-plane technologies;
  - technologies for end-to-end service control: performance management technologies.

- **End-user connectivity:**
  - Providing efficient access to all end-users and enabling mobility inside the institution. This implies widespread coverage with current and new wireless technologies. The increasing popularity of the laptop computer is the major driver behind this requirement.
  - Researchers, teachers and students have become more and more mobile; they want to have ubiquitous, guaranteed access and connectivity on demand. Secure roaming technologies like eduroam must be employed for this purpose. Mobile IP technologies will be very helpful for roaming inside campuses when using telephone-over-IP, video-over-IP or any other application that needs continuous connectivity.
  - Institutions need to control access to their internal network from laptop computers, PDAs or any other connecting mobile device. Technologies like Network Access Control (NAC) are well suited for automatically controlling access and providing access to the end-user.

- **Middleware technologies:**
  - identity management systems;
  - identity management and mobility;
  - identity management and identity federations.

5.1 Lower-layer technologies

This section covers optical transport infrastructures, Ethernet transmission technologies, control-plane technologies, control and performance. It presents the results following the same categorisation as in the EARNEST study of technical issues; more technical details are provided in the report on that study[8].

NRENs are moving increasingly towards a portfolio of services that includes end-to-end connectivity (wavelengths, Layer-2 and Layer-3 VPNs, Ethernet) and guaranteed bandwidth on demand at very high speeds reaching 40 Gb/s and soon the 100-Gb/s rate. Control-plane tools as well as performance tools are being implemented to control the allocation and management of circuits and bandwidth and to control and monitor network performance.

5.1.1 Policy-related issues

Often, the services and functionalities mentioned above are provided as far as the gateway to the institution network, and specific action is then required to ensure they extend to the final host. Thus only a very small minority of research projects truly can use end-to-end services as described above.

8. *See the EARNEST report on technical issues by Kevin Meynell et al.*

So far, most campus networks are providing 'best-effort' IP services to the vast majority of their end-users, based on traditional switching and routing equipment, which is interconnected, most of the time, by means of optical fibres and/or copper cabling. IP is the ubiquitous protocol for end-to-end connectivity for the type of service that is characteristic for Class-A users.

There are several reasons for this situation, including:

- End-users on campus only request access to services that are characterised as Class A in Section 2.1. Either they are not aware of additional potential services or they are not allowed to use them.
- There are no expressed needs or there is no widespread culture to share costly resources between institutions (e.g., high-performance computing, clusters, large databases, information systems for institution management, backup systems for security purposes etc.).

Clearly, this is not a satisfactory situation and institution managers should not be content with current practices. They should pay more attention to

- increasing the awareness among end-users of the potential of network services to support better collaboration and resource sharing in research and higher education; globalisation of the research community is already taking place with the virtualisation of resources;
- making a business case and developing a model for sharing costly resources between institutions; this would save funds and expertise that is better used elsewhere, and at the same time it would increase the quality of services.

Heading towards these goals implies a strong evolution in the range of network services offered to selected communities within the institution. Transparent, high-speed end-to-end connectivity is one of the most strategic services that should be available; it would allow the building of dedicated virtualised networks for specific purposes. These types of services will not rely solely on IP but also (perhaps even mostly) on Ethernet.

As mentioned in Chapter 4, the basic network infrastructure (optical fibres, Ethernet equipment etc.) is mostly in place, and even in cases where the infrastructure needs to be upgraded this should not be an insurmountable problem - especially if the institution's senior management recognises the need.

Apart from upgrading the infrastructure, the most important actions to be undertaken will be obtaining the expertise for new services, creating strong collaborative ties with NREN engineers and changing the culture from providing 'best-effort IP' to end-to-end connectivity and virtualised networks.

Existing and upcoming technologies will allow institutions to provide seamless end-to-end connectivity, higher bandwidth, guaranteed bandwidth, bandwidth on demand, performance control etc. Campus networks already have the main parts of the underlying infrastructure (optical fibres, Ethernet core switches, etc.) on which they can build the provision of advanced services.

Different technologies permit this and these are identified in the next sections.

5.1.2 Transport infrastructure: optical fibre

From the early 1980s onwards, institutions tended to install multi-mode fibres on their campuses for interconnecting their buildings because that was the best and least expensive solution. This is no longer true today[9]. Multi-mode fibre has two advantages: it is easier to install (especially for

9. *See the EARNEST report on technical issues by Kevin Meynell et al.*

wiring in buildings all the way to the desktop) and it requires cheaper end-point equipment than other solutions. However, this category of fibre brings a very severe limitation as regards higher bandwidth. In principle, it is able to transport 10 Gb/s, but only for less than 300 metres. On large campuses, this is a real handicap when end-to-end circuits (wavelengths or Ethernet) are requested for interconnecting distant services/applications with very high bandwidth.

Single-mode fibre has the ability to carry higher bandwidth than multi-mode fibre and for much longer distances. Transport bit rates of up to 40 Gb/s are possible, and when the specifications are completed, 100 Gb/s will become available.

This clearly shows that institutions should, whenever possible, interconnect buildings and campuses with the most recent fibre technology in single-mode.

	Recommendation
42	Institutions should identify their user communities that need high bandwidth, and they should start planning to upgrade the fibre technology on routes where high bandwidth is likely to be required[10].

Table 5.1: Recommendation on upgrading fibre technology

5.1.3 Transmission technology: Ethernet

Ethernet has become the basic transmission technology for campus LANs. Research and education networks were early adopters of Ethernet in the mid-1980s and nowadays Ethernet has not left any space for other legacy technologies. Ethernet has the potential to carry any type of application or service content (data, voice, video etc.), thus demonstrating that no limitation is imposed by Ethernet as the only lower-layer transmission technology at campus level in case no SONET/SDH transmission technology is deployed.

Most NRENs offer direct access to wavelengths and therefore can establish optical paths between campuses at national and international level with guaranteed Quality of Service. However, often these optical paths with QoS do not extend beyond the connection to the institution network or even sometimes only as far as to the entrance of a MAN interconnecting institution networks, and therefore the Quality of Service then terminates at that point. There are two main ways of extending QoS to the end-user's desktop or to the final server. One way is to continue the optical path up to end-point, but this may prove to be expensive and not always very flexible to implement. The other is to implement Layer-2 circuit-switched Ethernet with QoS, which may require upgrading the Ethernet infrastructure, but often that can be done at limited cost.

Several evolutions have brought Ethernet to the forefront recently, and Ethernet has an opportunity to become the transmission technology of the future on medium to long distances. Possibly Ethernet will replace SONET/SDH, at least in the research and education networks (LANs, MANs and NRENs).

Some of the latest improvements are:

- The ability to transport Ethernet on much longer distances over optical fibres without having to apply any transformation to Ethernet frames. This allows building large-scale Ethernet networks, as well as creating seamless Ethernet end-to-end paths. This brings the advantage of shorter transmission delays, and it reduces the cost of intermediary equipment for end-to-end links.
- The introduction of Layer-2 circuit-switched Ethernet with QoS and reliability.

10. *It should be noted that many institutions upgrading their fibre infrastructure are already selecting single-mode fibres.*

- The drafting and publication of OAMP specifications that will make the management of long-distance (multi-domain) Ethernet networks easier - certainly easier than the quite complex management of SONET/SDH networks. These specifications also have a strong influence on resilience, possibly pushing Ethernet networks to the same level of resilience as that of SONET/SDH networks, and with significantly less investment.
- The continuous increase in bit-rate capacity. Starting from 10 Gb/s in 2002, there are now possibilities for 40 Gb/s, and soon 100 Gb/s.

These new developments could have a strong impact on campus networking for a variety of reasons:
- They bring more flexibility for establishing seamless host-to-host connections over long distances, as well as easier implementation of bandwidth on demand with traffic engineering (QoS).
- In some cases they lead to simplification of the routing architecture, thus lowering latency and diminishing the overall cost of networking.
- They facilitate the establishment of Layer-2 VPNs with QoS and reliability.

More details about these different developments can be found in the EARNEST report on technical issues.

	Recommendation
43	Institutions hosting research and education projects that require high bandwidth, end-to-end seamless connectivity and guaranteed QoS should start planning for the early introduction of new Ethernet equipment offering the latest functionality.

Table 5.2: Recommendation on introducing new Ethernet equipment

5.1.4 Control planes

The Global Grid Forum[11] defined the concept of control planes as follows: the control plane is the infrastructure and distributed intelligence that controls the establishment and maintenance of connections in the network, including protocols and mechanisms to disseminate this information as well as algorithms for engineering an optimal path between end-points. From this definition it is clear that control planes include IP routing (IPv4, IPv6), MPLS, IP multicast, PBB-TE, T-MPLS, GMPLS, ASON, UCLP and probably more.

Some of these technologies have been deployed on campus LANs for some time, e.g., IPv4 routing, which is currently the main protocol for Layer-3 networking and has proved to be simple to operate. Some are implemented, but have less coverage, e.g., IP multicast and IPv6, which has even lower coverage. These last two technologies have been around for many years, without significant increase in coverage. Even when they are deployed on campuses, they are limited to few groups of users. The main reasons for this situation come from the complexity of deploying the technology and the lack of demonstrable value-added functionality to improve the quality and diversity of services for end-users. It has been quite difficult to convince network managers to give these technologies high priority when there is a shortage of staff and a prolific list of other, more urgent tasks.

MPLS is implemented mainly by NRENs and commercial Internet Service Providers with the aim of providing Layer-2 or Layer-3 VPNs, and may be associated with Traffic Engineering and Quality of Service. The deployment of MPLS on campus networks is rather limited; the operation of MPLS on campuses normally takes place in close co-ordination with an NREN to deliver end-to-end services for specific projects.

11. *The Global Grid Forum was the community of users, developers and vendors leading the global standardisation effort for Grid computing. In 2006, the Global Grid Forum merged with the Enterprise Grid Alliance and became the Open Grid Forum.*

MPLS no longer presents interoperability issues on multi-domain infrastructures. Nevertheless, before deploying MPLS on campuses, careful attention must be given to alternative technologies, like providing dedicated VLANs with QoS to reach the end-point of the MPLS circuit provided by the NREN. On a well-provisioned campus network, this solution will offer the same Quality of Service, will not require new expertise at campus level and will not introduce the extra complexity of multicast services.

GMPLS is a generalisation of MPLS for heterogeneous networks. It has the ability to work with fibres, lightpaths, SDH/SONET and packets. It is intended to provide signalling, routing and discovery of network resources. Its first objective is to have a unique control plane for different kinds of transmission technology. As for MPLS, the goal in terms of service is to deliver end-to-end circuits and VPNs with guaranteed Quality of Service (bandwidth, delay etc.). One can see GMPLS as a bandwidth-broker protocol, especially for optical networks. A great deal of effort is put into the development of standardised tools for the optical control plane. However, GMPLS and ASON still have interoperability issues that need to be fixed before GMPLS can be broadly deployed. Running GMPLS is challenging in practical terms, and that is why it has been used only on a limited scale (National LambdaRail, test beds, very few carriers). NRENs should pioneer putting GMPLS into operation before campus networks deploy it.

5.2 Wireless Networks

The second phase of the EARNEST study of campus issues looked at wireless technologies based on three IEEE standards: 802.11, 802.16 and, a more recent one, 802.20.

5.2.1 802.11 (or Wi-Fi)

The term Wi-Fi has been promoted by the Wi-Fi Alliance as the brand name of any 802.11-based wireless LAN. The Wi-Fi Alliance is an organisation of providers of wireless equipment and software; its mission is to promote and certify interoperability of 802.11 products.

802.11 networks operate in two frequency bands: 2.4 GHz (S-band, 802.11b/g) and 5.7 GHz (C-band, 802.11a). The S-band is more popular, but is more prone to interference because it uses the same frequency as microwave ovens, cordless phones, video cameras and Bluetooth devices. An amendment, 802.11n, planned to be approved in September 2008, will operate in both bands at higher bit rates. There are already commercial products based on the 802.11n draft document.

Several channels are defined for transmission in each band, with different channel numbers used in each regulatory domain. Luckily, the EMEA region (Europe, Middle East and Africa) has a common regulation, so the same channels are used all around Europe and there is complete equipment compatibility across national borders (something especially important for mobile equipment). Some differences exist with the rest of the world, especially in the 5.7 GHz band, but interoperability is possible in most cases.

The maximum bit rates are 11 Mb/s for 802.11b, 54 Mb/s for 802.11a and g, and 600 Mb/s for 802.11n (although commercial products today have a maximum of 300 Mb/s). The bit rate can be reduced under hostile environmental conditions, such as a weak signal or high interference.

Typical throughput of wireless networks is significantly lower than the nominal bit rate. While in wired networks one can normally expect throughputs of over 90% of the nominal bit rate, in wireless networks typical figures are in the order of 35-40%. For example, a 54 Mb/s 802.11g system has a maximum throughput of 19 Mb/s (35%). An 802.11n system with a nominal bit rate of 300 Mb/s has a maximum throughput of 100 Mb/s (33%). This is due to the overhead and peculiarities of the radio transmission protocol at the physical and link layer.

The effective throughput of wireless networks is reduced even further by the following factors:

- the transmission is always half duplex (only one channel is used);
- the channel (and therefore the bandwidth) is shared among all simultaneous users covered by the same cell;
- design mistakes (for example, incorrect assignment of radio channels in adjacent cells) can produce interference that reduces the bit rate.

Thus, in spite of the impressive bit rates achieved in the latest wireless networks, their throughput will remain well behind that of wired networks. Because of this, wireless networks are appropriate only for Class-A users or, in optimum conditions (strong and high-quality signal and few simultaneous users per cell), also for Class-B users.

As far as security is concerned, since 2001 it has been well-known that WEP is insecure[12] and its use is discouraged. The Wi-Fi Alliance specified WPA as an interim solution, which was later replaced by 802.11i (also known as WPA2). Those protocols are considered secure enough for home and office use.

Wireless access points for campus networks must support advanced features, particularly the ability to provide multiple SSIDs simultaneously. This is necessary for providing different access profiles to end-users.

Another interesting characteristic of wireless campus networks is the ability to configure channel and transmit power of the access point automatically and in co-ordination with the rest of the network. This enables a simple and effective provision of the service. Sometimes it is also interesting to configure access points as monitoring stations, with the objective of supervising the correct operation of the wireless network.

Some vendors of wireless equipment propose architectures based on centralised systems, where the network intelligence and management is concentrated as much as possible in specific appliances called wireless-LAN switches or controllers. This reduces the function and configuration of the access point to the minimum, to the extent that the access point is no longer able to operate standalone. These access points are known as 'thin', in contrast to the more traditional 'fat' access points that can perform all the functions related to the wireless network themselves. Centralised systems are easier to set up and manage when large networks are involved. Today, only a few vendors offer thin AP (Access Point) systems and all of them use proprietary protocols, which means that the whole system must be provided by the same vendor. However, the IETF has set up the CAPWAP (Control and Provisioning of Wireless Access Points) Working Group, which is specifying a set of standard protocols for such systems and has already produced several RFCs. Hopefully, with standards specifications, more vendors will offer interoperable thin AP systems, so that prices will decrease and consumers will be able to choose their vendors with more freedom than today.

12. *See, for example, www.isaac.cs.berkeley.edu/isaac/mobicom.pdf.*

5.2.2 802.16

802.16 is the standard proposed by the IEEE for wireless MANs, also known as WiMAX. The WiMAX Forum is an organisation of providers of wireless equipment and software; its mission is to promote and certify interoperability of 802.16 products. It is developing for the 802.16 standard a function that is equivalent to that of the Wi-Fi Alliance. One of the tasks developed by the WiMAX Forum is to define profiles restricting the wide possibilities currently offered by the 802.16 standard.

The first 802.16 standard was approved in 2001. Three amendments (a, b and c) were approved later. In 2004, a fourth amendment (802.16d) was approved and all the previous documents were withdrawn, with the 'd' amendment then being renamed 802.16-2004. In 2005, a new amendment, 802.16e-2005, was approved. Hence in its present form, 802.16 comprises only two parts: 802.16-2004 and 802.16e-2005. Although strictly speaking the names 802.16d and 802.16e are incorrect, they are frequently used to refer to those documents.

802.16d (802.16-2004), also sometimes called 'Fixed WiMAX', is designed for fixed and nomadic (not mobile) users. No roaming functionality is included in the service. The main goal of 802.16d is to provide 'last mile' access, i.e., high-speed Internet access to residential users, competing with ADSL and CATV in areas with a low-density population where it can be cost-effective, or in cases where fast or temporary provision of the service is required. LOS (Line Of Sight) or near-LOS is required. Frequency bands are between 2 and 11 GHz. Bit rates up to 70 Mb/s and ranges up to 30 kilometres are possible, but they cannot both be achieved simultaneously.

802-16e (802.16e-2005), also called 'Mobile WiMAX', is designed for mobile users. Roaming mechanisms at 'vehicular speeds' are incorporated. The aim is to provide a service somewhere between UMTS and 802.11 networks. Thanks to advanced radio transmission techniques, LOS is not required. Frequency bands are in the range of 2 to 6 GHz. Bit rates up to 35 Mb/s at ranges up to 10 kilometres are possible, although - as with 802.16d - they cannot both be achieved simultaneously.

Regulatory conditions for 802.16 have not yet been completely specified. However, it seems likely that Europe will not be as lucky as in the case of 802.11, because different radio frequencies are being used in different countries. This will probably result in variants of the same equipment adapted to each particular case, even if the only difference is the underlying radio component. This is unfortunate, because interoperability, and thus the portability of equipment, could be affected. However, manufacturers can provide multi-spectrum devices that could operate across different regulatory domains. This is similar to the situation today for mobile telephony services, with dual-band, tri-band and even quad-band terminals.

The technical complexity and sophistication of 802.16 (and specially 802.16e) is higher than that of 802.11, both at the physical and at the link layer. As a consequence, it has a spectral efficiency, robustness and throughput superior to 802.11. Interestingly, the latest additions to the 802.11 standard, like 802.11n, incorporate at the physical layer some of the features that are present in 802.16.

In principle, 802.16 is designed for carrier and service provider networks. However, the technical characteristics, the range and the inclusion of unlicensed frequency bands make it an interesting alternative for wireless networks on campus. Because 802.16 standards are more recent than 802.11, there are fewer products on the market and their deployment is still very limited.

Although 802.11 and 802.16 may be seen as competing technologies, they can also complement each other, for example, connecting Wi-Fi access points in remote locations by 802.16d bridges.

802.16 uses a cell architecture like the one of 802.11, where all users in the cell share the same channel. The longer range of 802.16 permits large cells to be configured, where potentially there would be a large number of simultaneous users. Care must be taken in those situations, because this would probably give low performance to the end-user in spite of the hypothetical superior throughput of 802.16.

5.2.3 802.20

IEEE 802.20, also called Mobile-Fi or MBWA, is another wireless 802 standard. Its working group was approved by the IEEE on 11 December 2002, with the goal to prepare a formal specification for a packet-based air interface designed for IP-based services.

So far, no standards have been approved by this working group. The proposed draft includes the following characteristics:

- IP roaming and hand-off at more than 1 Mb/s;
- optimised for full mobility up to speeds of 250 km/h;
- operation in licensed bands below 3.5 GHz.

As can be seen, there is some overlap in functionalities between 802.20 and 802.16e. The main difference is that 802.20 specifically addresses roaming at higher speeds, presumably offering a better service while moving in cars, trains etc.

It remains to be seen how successful 802.20 will become on the market. Even then, it is doubtful if it will have any relevance for campus networks, especially taking into account the fact that 802.20 will only operate in licensed bands and hence it will not be easy for academic institutions to provide the service themselves.

5.2.4 Mobile IP

Mobile IP is a standard protocol developed by the IETF to allow mobile devices to move from one network to another while maintaining the same IP address. This allows the device to keep transparent connections at the transport and upper layers while moving. A VPN tunnel or a VoIP connection, for example, will not survive a change of IP address on the fly. Mobile IP uses an efficient and scalable encapsulation mechanism, which does not propagate host-specific routes through the Internet.

Mobile IP is used mainly (but not exclusively) in wireless networks like 802.11, 802.16e and GPRS/UMTS. All those networks include roaming functions at the link layer. In principle, Mobile IP is not needed, but there are situations when a change in the IP subnet is necessary and then roaming at the network layer is needed. For example:

- Usually all the APs in an 802.11 network share the same IP subnet. However, this is not feasible in large networks because an unacceptable amount of broadcast traffic would be generated. The solution is to divide the wireless network into zones that are served by different IP subnets.
- When two 802.11 networks belonging to different institutions overlap, roaming is in principle feasible. However, since the two networks have different SSIDs and belong to different IP subnets, roaming has to be done at the network layer.

- A user moves away from an 802.11 network to an area served by an 802.16e or GPRS/UMTS network. Roaming at link layer is not possible between the different technologies, but roaming at network layer would permit keeping all sessions transparently.

In order to implement Mobile-IP functions in a network, three elements are needed:

- Mobile Node (MN): the mobile device;
- Home Agent (HA): the device in the home network (normally a router) that takes care of the traffic directed to the MN and re-routes it encapsulated towards the Foreign Agent;
- Foreign Agent (FA): the device in the foreign network (normally a router) that receives the traffic re-sent by the HA and forwards it to the MN. The FA is optional; if it is not present, the HA re-directs the traffic directly to the MN.

Most routers can be upgraded to serve as HA or FA. If the existing router cannot be upgraded, FA or HA software can be run on a Linux or Unix server. On the host side, there are several third-party MN implementations available on the market. Currently, Windows XP has no built-in support for Mobile IP. It is likely that in the future Mobile IP will be supported by default in all main operating systems.

Mobile IP has been developed by the IETF also for IPv6, providing the same functionalities as in IPv4, although the mechanisms are not exactly the same. One of the differences, for example, is that while in IPv4 packets always travel to the home network and from there go encapsulated to the mobile node, in IPv6 the packets go directly from the source to the mobile node using the routing header. This facility is not available in IPv4. Therefore, Mobile IP is one of the few cases where there is a real benefit in using IPv6.

5.3 Identity management challenges on campus

Providing centralised identity management facilities in an institution for research or higher education is a tremendous challenge. However, there are enormous benefits to be gained from such systems, and real problems to be encountered when they are not in place.

The absence of dependable identity management systems in institutions may lead to situations where

- multiple authentication systems exist, driven by local system administrators; different applications (email, Web, proxies, e-learning, management information systems etc.) acquire their own authentication and authorisation schemes, forcing users to operate multiple logins and passwords;
- the management of access control rights becomes extremely complicated, potentially causing breaches in the overall security architecture of the institution's systems;
- there is no automatic way to suppress access rights from all applications at once when someone leaves the institution;
- individuals are barred from accessing useful scientific, educational or administrative resources, because there is no guarantee they will be granted their rights by local managers;
- some users are obtaining access to applications that they should not be permitted to use;
- there might be a duplication of identity management tools in use: Active Directory, Oracle Identity Manager, Novell etc.;
- there is no authoritative central source for personal information.

Linked to the institutional challenges of managing user authentication, authorisation and accounting are the growth of user mobility across Europe and the globalisation of research, which pose similar problems. The sharing of electronic resources for e-learning and the deployment of Grid platforms are also greatly dependent on identity management environments.

An identity management system provides authoritative identity records for all members of an institution – teachers, researchers, support staff, students, visitors etc. Each entry should contain the appropriate identity information needed to grant automatic access to all permitted services, resources etc., depending on the person's role in the institution.

Good practice requires that all users of a campus network or information system should have some form of secure access to services. This would normally mean that they have a unique identity, usually augmented by a password, to gain access to particular services. Different people will have different access rights to various services and databases, e.g., some can read only part of a database while others may have wider read-access rights and may even be permitted to write to the database.

In some cases, user authentication may go beyond the username and password, and require another form of authentication, such as a hardware token or key that provides an extra level of security for accounts with greater privileges and responsibilities.

The starting point for sound identity management is good knowledge of potential users. Experience has shown that there are often problems in maintaining accurate records centrally. It is strongly recommended that personal information should be derived from central core databases in the institution. Typically this would be the staff/payroll database or the central student management database. There might also be a special category for visitors to the institution.

All rights and privileges should be derived from these databases and appropriate accounts should be created when users are added to the databases as they join the institution. Similarly, and very importantly, user accounts should be removed when users leave the institution and are deleted from the central databases.

Many institutions have a number of autonomous units on campus running different identity management systems. It is good practice to move to a single identity management system as quickly as practicable. The implementation of a campus-based federated identity management system would help to merge the various systems on campus and to present a unified view to the outside world. Even if a federated approach is adopted, the whole issue of identity management needs to be tackled by the institution in order to protect its intellectual property.

In research environments and in particular in the area of Grids, the historical use of X.509 certificates for user authentication is now being replaced by identity federation technologies. There are ongoing projects in Switzerland and in the United Kingdom, for example, exploring the use of Short Lived Credential certificates as proxy for federated identities for Grid applications. However, the integration of these research federations back into campus systems and onwards into national federations is rarely taking place.

In several countries (for example, Switzerland, the United Kingdom, Spain, the Netherlands, France, Norway, Finland, Croatia, Austria) there have been concerted efforts to build national federation

services. Major benefits will accrue to the users of such systems, because access to local and remote services can be facilitated by those systems. The sharing of national resources can also be facilitated by such distributed identity services.

In countries where there are no national initiatives, NRENs should take an active part to establish initiatives among their customers and to sponsor the national element of an international federation.

5.3.1 Mobility and identity management

One of the biggest trends in the computing and networking industries over the last few years has been the implementation of mobile solutions for communications. The use of personal communications devices has become universal among knowledge workers, with almost everyone using a mobile phone for voice calls and messaging. This trend in personal communicators will continue in the coming years. Personal devices are becoming more powerful and usable, and the difference between personal computers and mobile phones is rapidly diminishing.

Mobile devices pose special challenges to network service providers in research and higher-education institutions. Users want to connect seamlessly to their institutional campus networks wherever they are, for example, at home, on a journey or at another institution. It is possible that users will understand that there are different capacities on the various networks they engage with, but they will not expect to have to reconfigure their mobile devices as they move about between networks.

New wireless technologies are emerging all the time and capacities are increasing as new services are implemented, so current bandwidth restrictions may be gradually eroded in the future. It is likely that a wired network will always have higher capacities than a wireless one. However, the greater convenience of a wireless network may mean that users opt for this when they do not have enormous data-transfer requirements. Even where users have high capacity requirements, they may use a more convenient wireless network to control their large experiments on a 'remote control' basis.

As well as catering for mobile devices, network service providers must cope with mobile users. It is becoming increasingly common for staff and students to roam among a range of institutions in the normal course of their work. Increasingly, researchers are working on projects on a collaborative basis with other institutions and wish to take their 'identity' with them as they move about. It is increasingly important that users can be authenticated at another site and have their authorisation permissions set at appropriate levels. This could include access to research databases or physical access to laboratories and equipment in the remote locations, all authorised and authenticated via their home institution.

Similarly, the Bologna Declaration encourages students to spend time at different European institutions during their studies and they require access to their home university when they are away. A federated identity management system would facilitate this.

Most institutions have identity management processes in place, but they are often fragmented. Current systems of federated identity management seem to encompass small groups of collaborators with special requirements such as joint access to databases. It is now vital that planning is put in place for the integration of individual campus identity management systems into one large federated European (or even global) system, so that staff and students can migrate seamlessly through institutions across the continent.

5.3.2 eduroam

eduroam® is a system, set up and co-ordinated under the auspices of TERENA, that enables co-operating institutions to allow staff and students to roam wirelessly among the institutions taking part. All users are authenticated by their home site and, depending on their remote status, are allowed onto the local network to access the Internet. Accounting of the connections is collected at various points along the network.

While eduroam provides a tremendous service to the users, its facilities are somewhat limited: large numbers of European universities have not yet implemented the system, and therefore the roaming service is still of limited use. Where it is implemented, it is an invaluable service for the itinerant user.

One limitation of eduroam is its relatively poor authorisation facility. This is restricted to either allowing the end-user to access the local network or not. The levels of service that can be obtained locally are set by the local network and may not provide all the facilities that the user requires. Therefore, functions such as VPN or SSL may be available on some networks and not on others. A truly federated network access system would enable users to access the services they require, after authorisation by their home institution. Another, perhaps even bigger, limitation of eduroam is the difficulty of deploying it. The user's laptop must be reconfigured, and therefore the technology is not easy for students to use unless appropriate user support services are available.

	Recommendation
44	In order to promote mobility across Europe, the eduroam service should be provided by all institutions. Several benefits accrue from this. One is the increase in mobile services, another is the creation and management of an identity management database of users.
45	Institutions for higher education should provide support to deploy eduroam services within the student community.

Table 5.3: Recommendations on eduroam

5.3.3 Identity federation

Until recently, an institution's user-identity repository was used only by applications within the perimeter of that particular institution. The situation has changed now that there is greater mobility and collaboration between universities; for example, collaboration between universities allows students to access remote electronic learning environments at institutions other than their own. To meet wider needs, universities are adopting new identity federation technologies that make it possible to share user information in a controlled way. These new architectures make it possible for users at one institution to access Web resources provided by other institutions using their usual means of authentication. This is convenient for both the user and the service provider, because identity federation reduces resource managers' administrative tasks of registering external users.

Identity federation technology benefits from an infrastructure provided by an NREN (to build the trusted links), but institutions still need to set up the appropriate technical service. Each participating member of a circle of trust acts as a Service Provider, or an Identity Provider, or both.

An Identity Provider should be considered to be an identity data producer, so this service needs to be connected to the institution's user-identity repository. Within a national circle of trust and even at a European level, all participating institutions should share common syntax and semantics for user attributes. This may be addressed by common schemas, such as eduPerson (and its national variants) or SCHAC[13]. Setting up an Identity Provider for an institution often leads to standardisation and extension of the data gathered in the identity repository.

® *eduroam is a registered trademark of TERENA.*

13. *The Schema Harmonisation Committee (SCHAC) is a subgroup of TERENA's task force TF-EMC2.*

At the other end, a Service Provider is the identity data consumer and, as such, intercepts Web applications requiring access control provided by the identity federation infrastructure. Currently only a limited number of applications are compatible with identity federation solutions; for others, significant work is required to make software 'federation aware'.

If identity federation infrastructures allow new forms of use, they also bring additional constraints for the institution. From the user's perspective, the Identity Provider provides access to remote services, and this access provision should be widely available through the appropriate technologies.

	Recommendation
46	If they have not already done so, institutions should plan to introduce a unique identity management system, replacing any existing identity management services run in different areas of the institution.
47	The creation of a campus-wide identity management system should be considered a prerequisite for a campus-wide security management system.
48	Institutions should participate, in co-ordination with their NREN, in initiatives to deploy national and international identity federations, for the benefit of greater mobility, roaming, resource sharing, virtual organisations for research collaboration etc.
49	In countries where there are no national initiatives for identity federations, NRENs should play an active part in establishing initiatives among their connected institutions and they should sponsor the national element of the international federation.

Table 5.4: Recommendations on identity management systems and identity federations

5.4 Monitoring the performance of networks

Today, information and communication technologies are crucial in any organisation, including institutions for research and higher education. In a world of distributed resources, network services are the cornerstone of an organisation's operation. Any dysfunction or failure of the network will lead to serious loss of time, money etc. Hence, high levels of reliability and availability are expected from network services and whenever problems occur, they have to be fixed very quickly. Usually this can be done when there is only one network management entity providing the service. However, when a problem occurs (e.g., loss of connectivity, a drop in performance etc.) on an end-to-end connection crossing multiple network management domains, it becomes difficult - if not impossible - to identify precisely where and why the problem has occurred. Fixing it then becomes time consuming.

An important feature of the research and education networking world is that it is truly multi-domain. Even when each domain on the network path monitors its own infrastructure, there is no easy way to resolve problems. All domain administrators must be contacted to get appropriate information to point to the source of the trouble. To improve the quality of service to end-users there is clearly a need to minimise the time between the occurrence and detection of a problem and the moment when it is fixed. One facility that will assist reaching this goal is to have a constant clear picture of the behaviour of the different networks on an end-to-end path, with the ability to identify immediately the link, equipment etc. causing degradation. This is very crucial in the context of Grids, where multiple sites are involved and the probability of trouble increases.

PerfSONAR is a collaboration between the GN2 project and Internet2 and ESnet, in which open-source tools are developed to meet the objectives above. PerfSONAR allows end-users to learn the reasons for degradation much more quickly. It gives the operators of the different domains the necessary data to identify and fix a problem, providing standard diagnosis tools that are ready to use.

One important feature of PerfSONAR is its three-tier architecture:

- The lower layer runs in the network equipment and collects the data. It may be a NetFlow collector, SNMP collector or a proprietary collector.
- The intermediate layer formats the data collected and implements a standard interface to publish them to the upper layer.
- The upper layer aggregates data from the different domains (by accessing the intermediate layer) and offers a visual interface accessible to authorised end-users. Strict rules are enforced for accessing the data, depending on their level of confidentiality.

Web services, clients and servers that communicate using XML messages that follow the SOAP standard, are integrated in PerfSONAR to ensure the interoperability of the different components. As an open project, PerfSONAR allows for constant improvement of the tools by any contributor.

	Recommendation
50	In the context of research and education networks, NRENs have a special responsibility to bring PerfSONAR to operational status. They need to co-ordinate its overall deployment and organise the necessary training as well as the dissemination of information to institutions, MANs and regional networks in order to promote the active contribution of those institutions and networks.
51	In order to provide PerfSONAR with pertinent, usable information, institutions should install and run the appropriate tools at the lower and intermediate layers.

Table 5.5: Recommendations related to PerfSONAR

5.5 Peer-to-peer applications

Peer-to-peer applications are probably the most influential development in the Internet at large since the introduction of the World Wide Web in 1991. The decentralised architecture of peer-to-peer applications provides robustness and scalability at a very low cost for the service provider, because there are no central servers or critical links that could become traffic bottlenecks.

The peer-to-peer model has been present in the Internet since the very beginning. At the network layer, the IP protocol operation is peer-to-peer, and the same is true at the application layer with the server-to-server communication of DNS, SMTP and NNTP protocols. However, in all these cases, the user is not aware of the peer-to-peer nature, either because it is at lower layers (IP) or because he/she is accessing the service through a client program (DNS, SMTP and NNTP).

Today, the term peer-to-peer is used mainly to refer to protocols, applications and networks intended for file sharing between private users. The first well known peer-to-peer application was Napster, created in 1999. The Napster network was shut down in 2001 due to legal problems related to copyright infringement, but other, more sophisticated programs followed that are virtually impossible to control. Today the most popular peer-to-peer networks are BitTorrent, eDonkey, FastTrack and Gnutella. The use of these networks has increased enormously with the wide availability of high-speed Internet access following the deployment of ADSL and CATV networks in residential areas.

Trying to estimate the amount of peer-to-peer traffic on the Internet is not an easy task, because the results depend a lot on where the measurement is taken, and on the use of traffic shaping devices by some Internet Service Providers to control peer-to-peer traffic. Depending on the source, the figures go from 35-40% up to 50-60% of all Internet traffic[14].

14. *See www.slyck.com/story1502.html .*

Although file sharing by peer-to-peer networks is in itself obviously a legal technology, many use it to download and upload copyrighted material without the owners' permission. This has led to attacks against peer-to-peer networks from many copyright owners. As a consequence, many institutions restrict, or even completely ban, access to peer-to-peer networks by their users. However, there are many uses of peer-to-peer where file sharing is legal, for example, for distribution of open-source software, works in the public domain or works for which the distributor owns the copyright.

Apart from the debate about the legal situation of some peer-to-peer file sharing, there is no doubt that peer-to-peer applications are very useful for large file transfer and much more convenient than traditional protocols like FTP, thanks to their high reliability, resilience and user-friendly interface. The benefits of peer-to-peer are especially interesting when the user has to work in 'hostile' environments with intermittent connections and relatively low bandwidth. For example, BitTorrent, the most popular peer-to-peer protocol, is often used to distribute free/open software and for the download of Linux image distributions.

Another, more recent use of peer-to-peer applications is the distribution of multimedia streams (audio and video). The most popular of these applications is Skype, an IP-telephony network developed by the founders of Kazaa, one of the first file-sharing peer-to-peer applications. The terms 'peercasting' and 'P2PTV' are commonly used to refer generically to those networks that attempt to get the benefits of multicast IP when the Internet Service Provider does not provide such a service. Peer-to-peer TV applications are being used increasingly by broadcasters, because they can reach a much larger number of users with higher quality and lower cost than conventional unicast distribution. Popular applications for peer-to-peer TV are: Alluvium, CoolStreaming (based on BitTorrent), Cibersky-TV, Joost, Octoshape, PeerCast, Vuze and Zattoo. Peer-to-peer TV is also used by private users with a broadband connection to share TV channels that they are able to receive. In this case, the copyright infringement issue arises again, especially when pay-TV content is involved.

It is interesting to note that all the peer-to-peer networks that have appeared since 1999 have developed new protocols and services only at the application layer. By comparison, the protocols at the network and transport layers (IP, TCP and UDP) have remained almost unchanged for the last ten years.

5.6 Network Access Control

One of the challenges facing service providers in academic and research institutions is the dilemma between providing easy access to the network and protecting the information on and around that network.

The classification of users in Section 2.1 of this report can help resolve this dilemma. High-intensity users (Class C) tend to be few, but very demanding. On the other hand, Class-A users are many, but with much lower demands on an individual basis. Clearly, on a cumulative basis, they present large, and increasing, demands on the network. Because Class-A users tend to be numerous and tend to have larger numbers of small problems with their software, hardware and network components, some special mechanisms must be provided to ensure that they can work efficiently without introducing vulnerabilities into the network.

Mobility and the growth of commodity computing and network devices also lead to new challenges, with users moving around and connecting to a number of networks in the course of a day or week. This may mean that they sometimes connect to networks that are not as well protected as their usual research and education networks. Thus it is feasible that vulnerabilities are picked up on these other networks and introduced into the users' home networks with damaging effects.

An emerging solution to satisfy the need to connect to a range of networks is called Network Access Control (NAC) or some variant of that name. NAC is usually implemented in an appliance that sits on the network and intercepts new machines connecting to the network. When a new connection is detected, the device determines the hardware and software characteristics of the machine and decides whether to grant the machine access or not.

The NAC box is configured with rules about what types of machines can connect, what types of operating systems are allowed, what patch levels are required, what antivirus software is required etc. If the machine meets these requirements, then it will be allowed to connect to the network. Otherwise it can be put into a quarantine network where it can be upgraded, either automatically or manually by the user.

The big advantage of NAC boxes is that they can be configured to check the status of machines each time they connect to the network, to make sure that no viruses or spyware are being introduced into the network. If a new 'zero day' event is detected, the NAC can force each connected machine to check and remedy any vulnerability on an on-demand basis.

In one institution it used to take up to four hours to connect a student's PC to the network, including full virus check and the installation of patches. This meant that at busy times students would have to wait up to four weeks to connect their machines to the network. It also meant that students had to physically bring their personal computers to the helpdesk area during normal working hours. Following the installation of a NAC system, the students were able to connect their machines whenever they wanted to, and this almost totally eliminated the need for support staff to look at the machines. In just a small percentage of cases, some special action had to be taken to resolve esoteric problems. As a result of the NAC system, the majority of students connected outside normal office hours and more students connected in the first four weeks of the university year than had been connected in the first three months of the previous year. Linking the NAC to the central identity management system provided extra control and authentication. The requirement that only particular operating systems could be used by students was also removed. The switch to NAC also freed up staff time, so that extra facilities could be provided to the students. This would not have been possible in the pre-NAC environment.

The information and technology research and advisory firm Gartner, Inc. has been forecasting that NAC systems will be widely used to connect machines in the future. Today, many network companies as well as machine-security companies sell NAC equipment. These include Cisco Systems, Extreme Networks, Juniper Networks, Checkpoint, Symantec, McAfee, Sophos and Trend. Microsoft has announced a NAC solution, but few details are currently available and there is a suspicion that it will only apply to their Vista operating system.

Using a NAC box should eliminate spurious security requirements restricting users' choice of hardware or software and it has the potential to enhance software while liberating the users to select systems best suited to their work.

Network Access Control offers more functionality than only providing access control and security checking when connecting a host to a network. The Policy Enforcement Point[15] (PEP), which processes the connection of a host to the network, can behave as a switch, a firewall, a router etc. If the PEP is an Ethernet switch and 802.1X is in use, the host will be directed to a specific VLAN (or remediation VLAN) after host checking / authentication / authorisation takes place. If the PEP is a firewall, it can dynamically apply complex policies up to Layer 7 (e.g., this host is allowed to connect to a limited set of remote servers, or it can only use protocols HTTP and VoIP etc.). The PEP could even add interesting capabilities such as a dynamic encrypted tunnel (IPsec) between the host and the PEP (firewall), to provide protection from an unsecured access network (e.g., wireless). Some of these functionalities are already available on the market.

It is feasible to think about

- using NAC technology to allow any end-user to access specific services, either by using Layer 3-7 PEP, or with 802.1X by connecting the end-user to a specific VLAN connected to a lightpath service, thus connecting the campus to remote locations in the context of a particular research project;
- exploring future development or extensions of the NAC technology; for example, if the PEP is a router (or a Layer-3 switch), it could enforce new policies such as QoS, rate limiting, logical router (in line with virtualisation initiatives) etc.

This shows that NAC technologies should be considered first as a solution to the current security risks faced by campuses, and second as a framework to help deploy end-to-end services automatically, bridging the campus with NREN services.

To deploy NAC environments in an interoperable way, attention should be paid to two issues:

- enabling NAC environments to work properly requires the deployment of 802.1X technology as the basic standard protocol for authentication when accessing a network resource;
- before selecting any NAC product or environment, one should carefully study reports on the latest developments published by the TCG relating to this area of technology.

	Recommendation
52	Install appropriate NAC systems so that users can connect a diverse range of equipment and software without introducing security risks into the institution. Deploy technologies that are provided to be interoperable, so as to facilitate the mobility of end-users. Deploying Authentication and Authorisation Infrastructure (AAI) environments along with NAC technology should be considered fundamental for better conditions of work for mobile end-users.

Table 5.6: Recommendation on Network Access Control

5.7 Network requirements of Grid applications

Since the beginning of this century, NRENs have looked at Grid applications as very demanding applications, especially as regards their requirement for high bandwidth. Therefore, the fast development and growing use of Grid infrastructures have served as a rationale for the deployment of multi-Gigabit links and hybrid networks in several NRENs, so that they could cope with the specific requirements of Grid communities.

These specific adaptations at NREN level are satisfactory only for a subset of Grid applications, namely those that need to transfer huge amounts of data in a timely manner, and for which

15. *PEP is Trusted Network Connect terminology; Trusted Network Connect (TNC) is an open architecture for Network Access Control, promulgated by the Trusted Network Connect Work Group of the Trusted Computing Group (TCG).*

permanent high-speed circuits (end-to-end) have been established between different end-points of the Grid platform, across the NREN and up to the hosts on the different campuses. However, permanent high-bandwidth connections are not the only requirement of Grid applications. Some applications need end-to-end bandwidth in the network for shorter periods of time (a few hours, or even less), using dedicated and specific software interfaces. These applications need to be guaranteed that their data is available at the remote destination at a precise moment in time. Unfortunately, users with such requirements usually have to manage with the 'best-effort' networks that they have been using for several years.

As a consequence, developers of Grid applications have had to constrain their work to cope with 'best-effort' networks, where there is always the risk of losing routing packets. The drawback of this approach is that advanced network services (like AMPS, which is currently being developed in the GN2 project) are not integrated as crucial components in the design of an application, whereas such services could help to enhance the performance of applications and even change the way in which they are operating.

Network experts or people in charge of network services are often asked to act and give advice after applications are already developed and are in use, and when users are facing unexpected bad experiences. In most cases this is too late, because integrating complementary network services at this stage is then too complicated to take advantage of better functionality.

Fostering the use of advanced network services by Grid applications means:

- promoting their implementation in the early stage of software specification; it implies providing good software interfaces to network services that can be addressed directly from applications, with configuration lead time as short as possible, under the control of the application, and on an end-to-end basis;
- deploying, in NRENs and on campuses, the appropriate network infrastructures and services to allocate bandwidth dynamically on demand;
- making campus networks more transparent for Grid-specific use, by designing campus network architectures that will permit, for instance, Grid flows to 'bypass' middle boxes that may hinder the provision of high bandwidth; however, these campus network architectures must be safe to guarantee that no potential new security breaches are introduced.

6. Conclusions

6.1 Findings

The EARNEST study of campus issues undertook a survey of the current state of networking in research and education institutions in Europe. The questionnaire also asked respondents to give some indication of their future plans and expectations in this area.

Evidence from the survey and from other sources, such as the parallel survey conducted as part of the EARNEST study of researchers' requirements, indicates there has been a significant investment in network infrastructure in institutions since the SERENATE study in 2002-2003. There is little evidence of persistent bottlenecks at the campus level, although there are still pockets of low bandwidth and obsolescent equipment and cabling that need to be upgraded. It is also likely that some bottlenecks arise from unnecessarily stringent security arrangements.

Results from the survey show that the network infrastructure in institutions is generally satisfactory, although in many cases this state of affairs appears to have happened accidentally rather than by good planning. Network strategy needs to be addressed systematically at the highest level within institutions. There needs to be a coherent network policy, which is developed in conjunction with end-users and preferably with formal mechanisms to consult them. This is especially important where there are demanding applications and projects, for example, in particle-physics research. Institutional networking policies should include topics such as:

- the network services that are made available;
- replacement policy for cabling and equipment;
- security;
- Acceptable Use Policy;
- relationships with other Internet domains, e.g., NRENs;
- raising awareness of network services;
- training;
- arrangements for consulting users.

There are clear indications from the EARNEST surveys that some of the potential beneficiaries of networking services are not aware of their availability, or in some cases not even of their existence. Institutions should take steps to actively promote greater awareness of the scope and availability of network services (e.g., videoconferencing), and they should offer training to help users make better use of networking facilities.

In many institutions the network support team is too small to meet current-day demands. As a consequence, team members spend much of their time 'fire fighting' and cannot risk offering new services to their user community, even though it is technically feasible to do so, because they would not be able to cope with the extra workload. In particular, they would not have sufficient staff or expertise to provide adequate training and support. As senior managers plan future strategy, it is vitally important that they recognise the increasing dependence of their institutions on network services, and that they allocate sufficient funds and resources to exploit the network effectively.

Responses to the survey of campus issues indicated that there is good collaboration between NRENs and institutions in some countries, but room for improvement in others. Effective networks depend as much on good human collaboration and co-operation as on the technical infrastructure.

Finally, it is important to emphasise the paradigm shift that is taking place in networking, from only providing connectivity to offering and supporting network-related services. Institutions should carefully assess whether they are sufficiently equipped to support the rapidly changing nature of networking and they should take appropriate steps if necessary.

6.2 Technologies

The survey of technologies that was carried out as part of the EARNEST study of campus issues focused essentially on network layers that require greater involvement from network teams in institutions, because of their specific characteristic of being shared and needing substantial funding and often strong central co-ordination of implementation in order to fulfil specific demands from the research and teaching community. Application layers are not dealt with in this report. Nevertheless, institutions should be aware that several significant innovations in the use of the Internet have appeared at the application layer, often generating increased traffic, and thus demanding a constant upgrade of basic infrastructures.

One important observation of the survey is that there is no convincing evidence of emerging technologies that are likely to have a major impact on campus networks in the near future. However, it would be a great mistake for IT management and network teams to be complacent, because many areas of technology are incrementally evolving at a relatively fast pace. This trend calls for a culture of constant monitoring of these developments by network teams in order to be prepared for appropriate upgrading, especially where institutions are hosting end-users with demanding applications.

Services to end-users will be better if strong co-operation and an exchange of expertise takes place between institutions and NRENs, for planning and organising the deployment of appropriate technologies for higher speeds, dedicated bandwidth (circuits or wavelengths), end-to-end guaranteed performance, secure roaming, identity federations etc.

It is a prerequisite that IT management of institutions should initiate serious action in different directions to

- upgrade the core infrastructures of their networks (optical fibres, WDM equipment when needed, Ethernet's new functionalities for increasing end-to-end performance etc.);
- upgrade wired and wireless infrastructures for total coverage of the institution in order to provide ubiquitous, high-speed connectivity;
- set up and run crucial services like centralised identity management systems, identity federation, secure roaming infrastructures, Network Access Control (NAC);
- collaborate with NRENs to introduce monitoring tools, such as PerfSONAR, in order to guarantee end-to-end level of services, especially in institutions hosting end-users with demanding applications.

6.3 Main recommendations

6.3.1 Recommendations for institutions

Networking policy

- Define networking policy at the highest level within institutions, covering the following topics:
  - strategic plans to meet the aims and objectives of end-users;
  - annual budgets to deliver the planned objectives and keep infrastructure and services up to date;
  - provisions for a well resourced network support team;
  - rules for network security;
  - arrangements for end-user participation in policy making;
  - adoption of national guidelines;
  - arrangements for collaboration with NRENs, other institutions etc.

Infrastructure and services

- Set aggressive replacement policies for equipment with a maximum life expectancy of five years.
- Provide users with IPv6 services at all layers, and undertake the migration to native IPv6.
- Adopt institution-wide specifications for network infrastructure, including elements controlled by departments or faculties.
- Ensure seamless end-to-end connectivity where a particular quality of service is required.
- Provide support and training for performance optimisation, especially to Class-C users.

Security

- Adopt security measures that are appropriate for the purpose and do not hinder the effective use of the network.
- Establish an institution-wide security team with a high degree of independence.

End-user representation and awareness

- Ensure that end-users are involved in defining networking policy.
- Establish formal procedures to identify end-user requirements.
- Circulate an Acceptable Use Policy (AUP) to all end-users that sets out clearly their rights and obligations when they use the network.

Collaboration

- Establish strong, formalised arrangements for collaboration with other management domains, e.g., NRENs, intermediate networks, other institutions.
- Encourage network teams to share their expertise with colleagues in other management domains.

Raising awareness and training

- A cultural change in networking is taking place with the emphasis moving from providing connectivity to providing network-related services. To speed up this change of focus, institutions should
  - provide training courses and documentation of good quality for end-users to raise awareness of the network services available and promote their use, including videoconferencing, multicast, video broadcasting, video on demand, voice-over-IP telephony;
  - make arrangements for network support teams to retrain in order to keep up to date with fast changing technologies.

6.3.2 Recommendations for NRENs

- Collaborate more closely with institutions in the following areas:
  - deploying key services;
  - sharing strategic information;
  - organising training for innovative services;
  - co-ordinating working groups of network staff to share expertise;
  - understanding the demands of high-end users.
- Assist institutions to provide support and training to end-users about performance optimisation, especially Class-C users.
- Provide guidelines for institutional network policies.

7. References

EARNEST	http://www.terena.org/activities/earnest/
GN2	http://www.geant2.net/server/show/nav.749
SERENATE	http://www.serenate.org/
TERENA Compendium	http://www.terena.org/activities/compendium/

8. Acronyms

AAI	Authentication and Authorisation Infrastructure
ADSL	Asymmetric Digital Subscriber Line
AMPS	Advanced Multi-domain Provisioning System
AP	Access Point
ASON	Automatically Switched Optical Network
AUP	Acceptable Use Policy
CAPWAP	Control and Provisioning of Wireless Access Points
CATV	Cable TV
CDN	Content Delivery Network
COBIT	Control Objectives for Information and related Technology
DANTE	Delivery of Advanced Network Technology to Europe
DiffServ	Differentiated Services
DNS	Domain Name System
EARNEST	Education And Research Networking Evolution Study
eduroam	Education Roaming
EMEA	Europe, Middle East and Africa
EUNIS	European University Information Systems
FA	Foreign Agent
FTP	File Transfer Protocol
Gb/s	Gigabits per second
GÉANT	Gigabit European Academic Network Technology
GHz	Gigahertz
GMPLS	Generalised MPLS
GN2	Multi-Gigabit European Academic Network
GPRS	General Packet Radio Service
HA	Home Agent
HDTV	High-Definition Television
HTTP	Hypertext Transfer Protocol
IEEE	Institute of Electrical and Electronics Engineers
IETF	Internet Engineering Task Force
IP	Internet Protocol
IPsec	IP security
IPv4	Internet Protocol version 4
IPv6	Internet Protocol version 6
ISO	International Organization for Standardization
IT	Information Technology
kb/s	kilobits per second
km/h	kilometres per hour
LAN	Local Area Network
LOS	Line Of Sight
MAN	Metropolitan Area Network
Mb/s	Megabits per second
MBWA	Mobile Broadband Wireless Access

MCU	Multipoint Control Unit
MN	Mobile Node
MPLS	Multi Protocol Label Switching
NAC	Network Access Control
NAT	Network Address Translation
NNTP	Network News Transfer Protocol
NOC	Network Operations Centre
NREN	National Research and Education Network
NREN	National Research and Education Networking organisation
OAMP	Operations, Administration, Maintenance and Provisioning
OSI	Open Systems Interconnection
PBB-TE	Provider Backbone Bridges – Traffic Engineering
PC	Personal Computer
PDA	Personal Digital Assistant
PEP	Policy Enforcement Point
PerfSONAR	Performance Service Oriented Network Monitoring Architecture
PKI	Pubic Key Infrastructure
PoP	Point of Presence
QoS	Quality of Service
RFC	Request for Comments
RIPE	Réseaux IP Européens
RIPE NCC	RIPE Network Coordination Centre
SCHAC	Schema Harmonisation Committee
SDH	Synchronous Digital Hierarchy
SERENATE	Study into European Research and Education Networking As Targeted by *e*Europe
SMTP	Simple Mail Transfer Protocol
SNMP	Simple Network Management Protocol
SOAP	Simple Object Access Protocol
SONET	Synchronous Optical Networking
SSID	Service Set Identifier
SSL	Secure Sockets Layer
T-MPLS	Transport MPLS
TCG	Trusted Computing Group
TCP	Transmission Control Protocol
TEN-155	Trans-European Network Interconnect at 155 Mb/s
TEN-34	Trans-European Network Interconnect at 34 Mb/s
TERENA	Trans-European Research and Education Networking Association
TF-EMC2	Task Force on European Middleware Co-ordination and Collaboration
TNC	Trusted Network Connect
TV	Television
UCISA	Universities and Colleges Information Systems Association
UCLP	User Controlled Light Path
UDP	User Datagram Protocol
UMTS	Universal Mobile Telecommunications System
VLAN	Virtual LAN
VoIP	Voice-over-IP
VPN	Virtual Private Network

WDM	Wavelength Division Multiplexing
WebDAV	Web-based Distributed Authoring and Versioning
WEP	Wired Equivalent Privacy
Wi-Fi	Wireless Fidelity
WiMAX	Worldwide Interoperability for Microwave Access
WPA	Wi-Fi Protected Access
XML	Extensible Markup Language