

Issue Analysis Hybrid Networks

GigaPort Next Generation Network "Research on Networks"

Date: August 21, 2006

1. Introduction

Lightpaths are an important part of the hybrid networking strategy, as implemented in the SURFnet6 network through the GigaPort project. Lightpaths have the ability to provide the network's users with high speed, low latency, and transparent connections through the hybrid network and beyond.

As the use of lightpaths for high-speed networking is fairly new, there are a large number of issues related to lightpaths for which clear answers are not yet available. Many of these issues are currently under study, both within the GigaPort framework and elsewhere. However, there is no complete overview of the issues, and therefore no way to ensure that they are being addressed. Therefore, there is a need for a structured list of the issues that need to be addressed in order to make the most out of the hybrid infrastructure.

1.1. Scope

This paper provides an issue analysis highlighting relevant problems related to hybrid networking. It does not mean to resolve these issues, or analyse them in great detail. Rather, it is intended to provide a starting point for the definition of future research subjects in this area.

The results of this activity can be used to assess if the current Research Objectives in GigaPort are sufficient or if adjustments need to be made in either the objectives or the planned activities. Also, the results are scheduled to be taken to the GLIF Tech working group for further discussion on the questions raised and to start the search for their answers.

The issues shown here involve the planning, provisioning and use of hybrid networks, with a focus on lightpath services.

1.2. Common vocabulary

In order to achieve a common understanding on some of the issues in lightpath networking, a set of common definitions is essential. Even the term 'lightpath' itself is currently used to mean a variety of network connectivity options. The following definitions are used within this paper:

Lightpath: In some publications, the term lightpath is used to denote any form of wide area connection with a large bandwidth, e.g. a dedicated part of a photonic path with fixed characteristics or a VLAN across a wide-area 10GE

switched infrastructure. SURFnet and some of the other participants in GLIF follow a more strict terminology, in which a lightpath is either an entire lambda across an optical network or an end-to-end service over a next-gen SONET/SDH infrastructure supporting GFP-F and VCAT. In either case, the lightpath provides a dedicated path to the connected endpoints.

User-controlled lightpath: In a sense, any lightpath service is user-controlled as someone will have to request the service in the first place. However, we will use the term user-controlled lightpath to mean a lightpath service in which essential aspects of the service (such as topology or bandwidth) are controlled from inside the user domain without any human intervention inside the network domain.

Scheduled lightpath: a user-controlled lightpath service where requests to change essential aspects of the service are submitted well in advance (e.g. an hour or more before the change is needed), and the use of the lightpath is fairly static (e.g. using the same topology and bandwidth for several hours or longer).

Dynamic lightpath: a user-controlled lightpath where requests for changes are expected to be satisfied in a time-scale ranging from milliseconds to minutes.

Hybrid network: In essence a hybrid network is a network which provides multiple types of service. However, in this paper we will use the stricter definition as used by SURFnet: a hybrid network is a network which provides both routed IP and lightpath services.

1.3. Domains

Most of the processes involved in networking can be decomposed into a process internal to each operational domain (such as a network with its own network operations centre) and a process across domains. The process across domains links the processes within the individual domains together, usually at a higher level of abstraction. For instance, a fault found within a network has to be understood in great detail within that network, but may translate to a simple 'service not available' message across domains.

In this sense, the facilities at a user's location constitute a domain of its own, which will need to implement the same processes as any other domain. Process interfaces will have to be defined between the user's domain and one or more network domains, as illustrated in Figure 1:

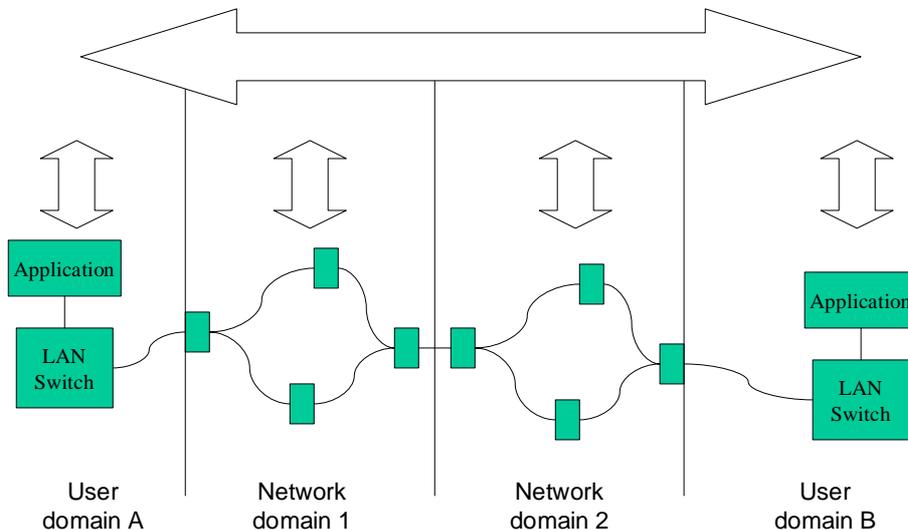


Figure 1: Operational processes within and across domains

Many of the processes implemented within and between domains in the traditional telecommunications networks also apply to lightpath networking, albeit in a more complex form. For instance, the ISDN and SS7 standards contain messages for equipment within a user's domain to tell the network what type of end-to-end connection is desired, and for domains further down the chain to report back whether this type of connection is available.

2. Business framework

2.1. End-user awareness

At this time individual researchers are not always aware of the possibilities of lightpath services, and the IT organisations at the connected institutions are not always aware of the actual needs of the researchers.

How can the provider of lightpath services improve user awareness of the advantages, limitations, and consequences of lightpath services within its network and across other networks?

How can network organization work with the IT organisations at the connected institutions to increase end-user awareness?

2.2. Costs and pricing

While the pricing of a connection to the routed network is generally based on the access port speed only, the pricing scheme for lightpaths is more complex due to the end to end and fixed performance characteristic nature of a lightpath. A form of "usage based" pricing seems to be necessary in order for the hybrid network to grow in a healthy way.

What incentives can we use to promote the introduction and use of lightpaths in a sound economically way?

How can a pricing scheme help to give access to the lightpath resources in a fair way to all users of the hybrid network, while avoiding complexity?

What is the economic viability of dynamic Lightpaths and what incentive do users have to release Lightpaths when not in use.

2.3. Roles and responsibilities

The division of roles and responsibilities between parties providing a lightpath service is less straightforward than in more traditional services. Whereas routed IP and telephony are standardised products, with very few options for the user to choose from, lightpath services may get more complicated.

As mentioned in the introduction (see 1.3), there will be a need to define process interfaces between the user's domain and one or more network domains. There are various possible models for such processes, as explored in [16]. For instance, the user domain might communicate only with a single network domain, which then communicates with the remaining network domains, as illustrated in Figure 2 (taken from [16]) and called "the parallel "master contractor" process". Other options include a cascade model, where each domain communicates with the next domain in the chain, or a user orchestrated model, where the user domain communicates with all other domains involved and called "The serial "peering relationship" process".. Obviously, the choice of model has an influence on the processes, protocols and methods that need to be developed.

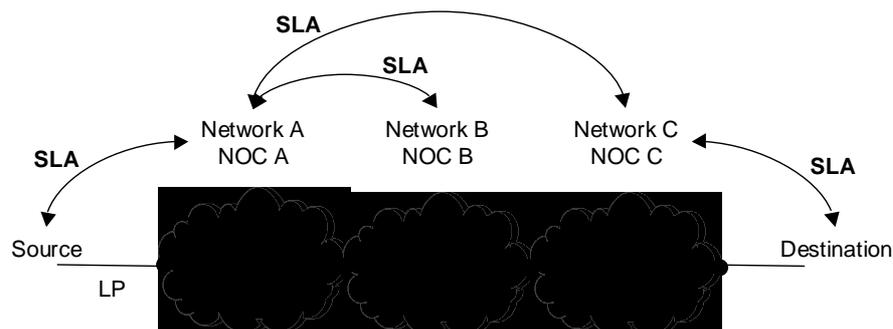


Figure 2: Network A is formally in contact with Networks B and C.

What model for intra-domain communications is best suited for the lightpath networking environment?

What are the consequences if different domains attempt to organise communications along different models?

What are the consequences of the various models and combinations of models for the processes, protocols, and methods that will have to be developed to support them?

3. Using lightpath services

3.1. User requirements

3.1.1. Static, semi-dynamic/scheduled, or dynamic

Most of the current users of lightpath services expect either a static service (order once, use for many years) or a scheduled service (order in advance, use for days or weeks). A very small number of users expect a more dynamic behaviour, as this type of use is still at an experimental stage.

3.1.1.1. Needs for dynamic lightpaths

In the same way that a Grid can be built from computing resources, sensors, or storage facilities, one can envision a Grid of networking resources. As in other Grids, such resources could be requested and allocated on a dynamic basis.

However, at this point it is not clear what types of applications will need dynamic network resources, and specifically lightpaths, in the near future. A better understanding of these needs will help plan the development of networks and the facilities needed for dynamic lightpaths.

What applications will expect dynamic establishment of lightpaths over the next few years?

How many users, sites, connections will be involved in such applications?

Where will these sites be located?

Will other networks that serve some of these sites be able to provide dynamic lightpaths in the near future?

3.1.1.2. Expected response times

In a Grid of networked resources, lightpaths have to be switched dynamically and without manual intervention. But it is unclear what response times are expected by the applications currently in development, and whether these expectations can realistically be fulfilled with technology that will be available at that time (see also 4.4.1). It is also not clear how strong the need for such response times is, and whether it will justify the investments needed.

What response times will users expect of dynamic lightpaths? Is this in the order of minutes, seconds, milliseconds, or less?

How will applications use dynamic lightpath switching? What are the advantages to the applications if these response times can be achieved?

3.1.2. Bandwidth and topology

Lightpaths within the SURFnet6 network are currently terminated on either 1 Gbit/s or 10 Gbit/s ports. Services on 1 Gbit/s ports are provided through next-gen SONET/SDH, and can therefore be provisioned in multiples of 50 Mbit/s; services on 10 Gbit/s ports are provided as lambda service and therefore always consist of the full 10 Gbit/s. At SURFnet, an initial product portfolio has been developed in which 150 Mbit/s, 600 Mbit/s, 1 Gbit/s, and 10 Gbit/s are designated as the speeds that will be offered for lightpaths. Hardware that will allow for lightpaths to be created at speeds between 1 Gbit/s and 10 Gbit/s currently is on order for SURFnet6. There is currently very little experience with the use of lightpaths, so that the actual needs of users and applications are still uncertain.

What is the actual use of the capacity provided across the Optical Private Networks? What is the need for growth?

Is there sufficient bandwidth and routing capacity within these sites to effectively utilise the requested bandwidth?

What is the ideal topology for an Optical Private Network, in relation to the types of sites involved (e.g. users, datacentres, main routed sites) of the connected institution, such as a full mesh, star or dual star topology?

What applications can be expected in the next few years with demands for high bandwidth lightpaths? Which of these applications will require scheduled or dynamic lightpaths?

Is there a case for 'cloud bypass' models, using dynamic lightpaths to off-load large volume routed IP traffic?

3.1.3. Resource availability/transparency/predictability

Once a connected institution has invested in the infrastructure needed to connect to a scheduled or dynamic lightpath service, users will expect network resources to be available when needed. However, capacity may be large but not infinite, and there may be situations where the available capacity is not sufficient to satisfy a request.

What blocking probability would users expect or allow for dynamic or scheduled lightpath services?

If the capacity requested is not available, how will users/applications respond? Will they use a lower capacity, use a different medium (e.g. the existing IP service), delay the work, or move the work around so that the network resource is not needed?

3.1.4. Jitter, latency, etc

One of the advantages of a lightpath service, compared to the routed IP services, is a low and predictable latency, with negligible jitter. Lightpaths

enable, e.g., the use of protocols originally designed for a LAN environment, such as CIFS¹, to run over the WAN. These protocols can not cope at all with the latency introduced when running over the wide-area. While such problems cannot be attributed to the hybrid network, but to the industry that created them, the introduction of lightpaths rapidly exposes these problems, and asks for solutions.

What are the requirements of lightpath users and their applications for latency and jitter?

Which protocols are less suitable today for wide-area operations, and what are the alternatives?

What feedback can be given back to the protocol and implementation developers regarding the use of their work in the hybrid network?

3.1.5. Robustness/Availability requirements

Lightpaths, in their most basic form, are unprotected single paths through the optical network. There are various ways a more resilient connection can be provided, but in order to develop these further it is important to understand the requirements.

What are the requirements for the availability of lightpath services?

Is there a need for protected paths within the Optical Private Network, or can the higher layers provide sufficient resilience?

3.2. Consequences of payload protocols

3.2.1. Ethernet

Ethernet and related control protocols, such rapid spanning tree and EAPS, were not originally designed for dynamic topologies or high latency connections.

What happens when protocols designed for rapid convergence times such as 50ms are deployed across links with latency above this time, and with a dynamically changing topology?

What 'race conditions' can occur in a single Ethernet bridged network with a high latency?

3.2.2. IP

In contrast to Ethernet, IP is designed as an end-to-end protocol across high latency networks. However, the addressing and routing mechanisms used at the IP layer are designed to deal with fairly static networks with occasional topology changes, and a fixed hierarchy (see also [7] and [10]).

¹ CIFS is a Common Internet File System, developed by Microsoft. More information can be found at: <http://www.microsoft.com/mind/1196/cifs.asp>

3.2.2.1. IP addressing and address resolution

IP addressing in general is a commonly understood mechanism, but the use of lightpaths may require new ways of dealing with the IP addressing.

What are the consequences of dynamically connecting LANs while using public IP address space?

How can the use of link-local addressing and zero configuration methods simplify the use of IP in dynamically changing networks?

3.2.2.2. IP routing

Most connected institutions currently use lightpaths to connect multiple sites to a single main site, which provides outside connectivity. However, in the future they may want to get external connectivity to multiple sites, which makes routing slightly more complex. If they use dynamic lightpaths on top of that, sometimes linking to one institution and sometimes to another, the BGP setup may become fairly complex.

Do connected institutions have the capabilities to deal with the complexities of using BGP in a complex network?

What are the effects of using BGP in a dynamically changing network topology? How do the various timers used in BGP interact when the network topology changes rapidly?

How can traffic flows be controlled when using BGP in a changing topology, e.g. for 'cloud bypass'? How can a router set up a bypass path without creating an undesired traffic sink?

3.2.3. TCP

It is well known that standard TCP does not perform well across connections with a very high bandwidth-delay product. This is often resolved by employing many parallel TCP sessions (e.g. gridFTP), but there is a need for a more practical solution. This is especially true in distributed applications which may change their topology, thereby changing both delay and bandwidth, using scheduled or dynamic lightpaths.

How can TCP performance over 'long fat pipes' be improved in a way that is transparent to the application?

How can TCP parameters best be tuned in a changing topology (e.g. Microsoft recommends to adapt the TCP window size to the bandwidth-delay product, but there is no method to change this dynamically)?

What happens when packets arrive out of order due to 'cloud bypass' kicking in, as a large number of packets is in transit on the previous route while later packets take the shorter route? Will this result in TCP decreasing its window, just when capacity is improved?

3.2.4. Problems arising from protocols designed for low latency LANs

Many of the protocols in use within LAN environments expect a very low latency (e.g. lower than 3 milliseconds). This is especially true for 'chatty' client/server protocols (such as the SAP client, and Windows/CIFS).

Although lightpath services tend to have a fairly low and constant latency, this may still not be good enough for some applications.

This issue is not specific to lightpaths, as it can occur with any type of wide area network. However, lightpath service can be positioned to replace existing connections (e.g. leased lines) and, depending on the topology of the underlying optical network, actually increase the latency. Lightpaths will also encourage layouts where users and applications are in separate sites. Therefore, there is a need to explore the consequences of a longer latency for the application.

What specific protocols are currently in use within local area networks which may well run over lightpaths in the future?

What are the consequences of a higher latency for the application, based on the properties of these protocols?

What alternatives do users have to compensate for a higher latency (e.g. local caching, application accelerators, protocol changes, Windows tuning, server based computing)?

3.2.5. Application

Running applications over lightpaths ensures that the application uses a path with predictable characteristics. However, these characteristics might change over time, e.g. if rerouting sets in due to a fiber cut or when the topology of the underlying photonic network changes due to changes being applied to the network.

What needs are there for the application to be aware of actual available bandwidth, latency etc.? How can the application respond to changes?

What classes of applications are most sensitive to changes in available bandwidth, latency, and other network characteristics? What strategies should these applications implement to respond to changes?

4. Providing lightpath services

4.1. Discovery and description

In order to manage any type of network, there is a need to discover the topology of the network at various layers, and to describe the results in a human readable or machine readable format. Hybrid networks present new challenges in this area, mostly because the communication between user

domain and network domain needs a common representation of the topology at some level.

4.1.1. Network discovery within the domain

Network discovery is the process of discovering which elements, connections and endpoints exist within a network, and discovering their main characteristics.

In general, an effective hybrid network will need some form of discovery mechanism. The alternative, using documentation from the network build process as a representation of the network, can not be guaranteed to be sufficiently reliable. However, not all relevant data can be discovered automatically, so that some form of manual data insertion will always be necessary.

4.1.1.1. Network elements / endpoints

The network elements that are implemented at the edges of the network determine the "face" of the service offered to the end user or his/her application. Network Elements that are inside the network also play an important role. The abstraction of all elements that matter to lightpath provisioning in the hybrid network needs to be discovered and represented at the middleware layer that is responsible for offering the end to end services to the user.

Which types of network elements within the hybrid network can be discovered through an automatic discovery process?

What mechanisms have to be implemented for each type?

What output format does each mechanism provide?

What elements can not be discovered through an automated process (e.g. purely passive network elements)?

Which network elements need to be discovered in order to provide effective services?

Should network design take into account network elements that can not be discovered (eg by including active instead of passive elements at the edge of the network)?

4.1.1.2. Network connections / network services

Besides discovering the network elements, the abstraction layer also needs to be aware of the connections that run between the elements.

Which types of connections, between which network elements, can be discovered through an automated discovery process? (eg physical fibre, lambda, SDH section and path)?

What mechanisms have to be implemented for each type of connection?

What output format does each mechanism provide?

4.1.1.3. Shared risks

Network connectivity is usually provided with a form of redundancy at some level (eg, path protection, or redundant paths with spanning tree across them). Whatever mechanism is used to provide redundancy must take into account any shared risk between components or connections. Therefore, an awareness of these shared risks is needed, at the very least within the domain and possibly across domains.

Which types of common risks can be identified through the automated discovery process, and which types can only be identified from network build information or other sources (e.g. common equipment, common SDH container, common lambda, common fibre, cable or duct)?

What if a domain shares risks with another domain (e.g. two networks offering a service actually using the same duct)? Can that shared risk be discovered? If not, how can this risk be mitigated?

4.1.1.4. Description

The output of the network discovery process will result in a representation of the network in some format, and possibly several different formats for different parts of the network.

In order to plan services across the network, and to provide abstract service information to other domains, there will be a need to translate this representation into a common description format. Several formats exist (see [17] for an example), but it is not clear that these will cover the functionality required.

At what level of abstraction should the output of the discovery process be presented to other processes within the domain?

Can existing formats capture all the data required at this level of abstraction?

What are the pros and cons of existing formats and methods? Is there a need for additional or improved formats?

4.1.2. Service discovery across domains

4.1.2.1. Description language

A network domain can be described at various levels of abstraction, with high detail at the low abstraction level and low detail at the high abstraction layer. Tools will be needed for generating the description in the right level of detail.

At what level of abstraction should a service description language represent the underlying network? Is there a need for the requesting process to have an understanding of the network topology, or should the network present itself as a 'black box'?

What tools need to be developed to efficiently describe the network infrastructure in an automated way?

4.1.2.2. Conversion between representations

Different network domains may well come up with different representations to describe the available services. In order to define end-to-end services, there will be a need for common formats or conversion between formats. Even when using the same description language, the level of abstraction or the definitions used may still vary.

What service description languages are currently being developed? What are the pros and cons of each, and how do they map unto each other?

4.1.2.3. Service comparability

Different networks provide different services. It may not always be clear to the requesting user or application which services can be combined into an end-to-end service. For instance a lightpath delivered over a next-gen SONET/SDH section across one network may be linked to an Ethernet Private Line Service or a routed IP service on another network.

If services can be linked together into an end-to-end service, it may not be immediately clear what the constraints of the resulting service are. For instance, if next-gen SONET/SDH-based lightpaths in separate domains are linked through an Ethernet Private Line Service through a third domain, the resulting service may not have the properties the requestor expects. Such a constraint will have to be clear to the requesting process, which can then decide how to deal with it.

What are the possibilities and constraints for interworking between different types of lightpaths or other transport services?

How can a formal description of various services help compare services, and establish interworking scenarios?

How can a requesting process combine a service representation from different domains into a service representation of the resulting end-to-end service?

4.1.3. Service discovery combining network services and other services

Scheduled or dynamic lightpaths can be used in combination with other scheduled Grid resources such as computing, storage, sensor networks or data repositories. In order to schedule such combinations effectively, the requesting process will need to have a representation of all the relevant available resources. The representation of these resources will have to link the end-points of the network service to other resources, so that a process can, for instance, find both an available CPU and a lightpath service to connect to that CPU.

There have been a number of developments in the area of Grid resource discovery; however it is not clear that the methods developed in these initiatives can be used without change for network connectivity services.

How suitable are existing Grid resource discovery methods and protocols to represent network resources?

What improvements are needed in order for a requesting process to gain a complete view of the relevant resources, both network based and others?

4.2. Resource reservation and allocation

4.2.1. Capacity management

In order to plan for the deployment of static lightpaths, and to fulfil requests for scheduled or dynamic lightpaths, an operational domain will need knowledge of its available resources. While it is possible to set aside an amount of capacity for specific services, such as scheduled lightpaths or routed IP services, this may not be the most effective use of resources. Rather, different services should be able to draw from a 'pool' of resources, including bandwidth on any given segment.

How can a management system get an overview of available network resources, taking into account existing lightpaths, scheduled reservations, and forecasted use of all available services (including routed IP)?

How can a representation of available and planned resources be used to feed the resource allocation process?

How can a representation of capacity allocated to current use, reservations, or forecasted use be fed into the network build process?

4.2.2. Resource request and response

The handshaking between the entity that requests the lightpath and the layer that controls the allocation and use of lightpaths needs to ensure that there is a transparent way of signalling the requests and their responses, as the requestor needs to be informed about what happened to the request issued.

What information is needed in a request for a scheduled or dynamic lightpath (e.g. desired bandwidth, latency constraints, starting time, duration, priority, other)?

What information is needed in a response to a requestor (e.g. actual starting time, actual bandwidth, expected latency, service identifier, other)?

4.2.3. Best path determination

A lightpath between two end locations on a network will in most cases be constructed along one of two or more possible paths. What determines which

paths is best will most likely vary from network to network and from the actual end locations in the topology of the network.

What is the most effective method to determine the 'best' path through a network to satisfy a given request?

What boundary conditions should best path method take into account?

What are the costing parameters involved, and on what basis can they be established?

How should the path determination method take into account existing and forecasted fill grade of the individual resources?

How can the path determination method take into account any shared risks (e.g. same SDH path, same cable or duct, etc)

How can different domains work together to determine an overall best path across multiple networks? Can this be realistically achieved?

How can the requesting domain work with the network domains to combine costing functions for network resources and application resources? For instance, if there are paths available to multiple servers each providing the same service, how can the use of servers and network resources be optimised overall?

4.3. Network planning and build-out

4.3.1. Feeding capacity planning data into network build process

In order to plan future network build-out, an overview of current and forecasted use of capacity will be needed. The design process is usually partly manual, but could be improved by tools that use the capacity planning data and the current network configuration to establish strategies for expansion.

How can capacity planning data be used to feed the network design process?

4.3.2. Optimising topology across layers

Network planning includes optimising the design based on the cost of different options and on future capacity needs. For instance, a network planner may have to choose between adding lambdas on a WDM system or acquiring additional fibre routes.

This optimisation problem gets more complex if application level options are included. For instance, providing a protected path towards a server may turn out to be more expensive than setting up the application on multiple servers in different locations, and using unprotected paths. However, this requires the optimisation process to be aware of the possibilities of the application, the layout of the network, and shared risks in the network infrastructure.

How can cost data for network expansion options be used to optimise network design choices through automated tools?

How can application level options be included in the optimisation process? How can user domain and network domain work together to achieve a global optimisation rather than local?

4.4. Provisioning / deprovisioning / changes

4.4.1. Set-up time within customer requirements

Although it is not yet clear what set-up times users of scheduled and dynamic lightpaths will require, there have been requests for very short set-up times, and there should be a clear view on what is achievable.

What are the delays in currently available systems (e.g. DRAC) from user request to set-up of a lightpath, for SDH control or lambda switching?

What are the physical limitations on set-up time within the network equipment currently used for the next-gen SONET/SDH or for lambdas?

Can part of these limitations be overcome, for instance by storing and re-using tuning parameters for each configuration rather than waiting for control loops to converge? What are the disadvantages?

What are the limitations if other available equipment is considered?

4.5. Incident management

4.5.1. Alarm monitoring

On the routed IP network a well established way exists to monitor the health of the network and to act upon the information gathered from the network management systems. For lightpaths this fault information will be received at the NOC in great, and perhaps too great, level of detail to be used in fault communication processes, as correlating the fault information along the actual path of the lightpaths affected may be needed first.

How can a Network Operations Centre correlate alarms from network elements at various levels to establish what connections are affected?

How can fault information be (automatically) translated into consequences for services, including lightpath services?

How can service fault information best be provided from one domain to others, including the user domain? What representation forms exist that can be processed automatically? How can this information be correlated to information exchanged in service discovery and set-up?

How can service fault information be fed back to the requesting application? What are the options for the application to respond?

4.5.2. Service assurance

Existing tools for service assurance concentrate on the IP layer. For instance, “ping” and “traceroute” are used by both network operations centres and by end-users to localise networking problems. As a lightpath service is a layer 1 service, any IP based tool will see it as a single hop, making it more difficult to localise a problem.

There are several tools in development to assist in fault location within a layer 2 network. For instance, the IEEE 802.1 working group has specified a set of tools and mechanisms for Connectivity Fault Management under 802.1ag. While useful for lightpath users, these mechanisms do not solve the problems associated with fault location in a layer 1 service.

When an end-to-end lightpath fails, or provides a lower quality of service than expected, how can a user pinpoint (at least) the domain within which the fault has occurred?

Can layer-1 loopback features, which are usually geared towards the operator, be used to provide fault location capabilities to the user?

4.6. Performance management

4.6.1. Measuring the quality of the service

As a lightpath in principle is an end to end pipe, in which you cannot examine what is in the type, most equipment used will have ways to gather data for measuring the quality of the service.

How can error rates (bit error rates, errored frames) be monitored and translated into measurements representing the quality of the service provided?

How can end-to-end latency and jitter across a lightpath be determined? Is it possible to measure the actual values, or is there a need for some type of proxy measurements?

What other measurements are relevant to define the quality of a lightpath service, once established?

4.6.2. Measuring the utilisation of services

Monthly and live reporting about the use of the network resources is what all network organizations do. Many tools (like MRTG) and best current practices (like the 5-minute reporting interval) have been developed over the past years. Measuring the use of lightpath services and reporting on this is not understood yet.

How can the actual use of lightpath services be measured? Is it useful to measure frames or octets transported, or is there a need for more sophisticated measurements?

How can data on the utilisation of lightpath services be used to forecast future growth as an input into the capacity management process?

4.7. Authentication and authorisation

Besides access to the lightpath resources that should be protected, the access to the control plane should be protected as well, in a multi-domain environment in which one network has to accept credentials of another network to make the multi-domain lightpath happen.

What existing security mechanisms can be used to control access to control plane functions, within and across domains?

How should each domain define trust boundaries, and how can trust best be extended across domains?

What tools are needed to define the rights each user or domain has to interact with the control plane of another domain? How can different users be assigned different rights to specific resources?

5. Developing lightpath services

5.1. Testing

For the sound operations of a routed IP network, the test bed is a necessary tool to test new software or hardware. For the deployment of lightpath services, this is also very much true, moreover since many of the operating paradigms are still unknown.

What facilities are needed to test and improve lightpath services and tooling?

What degree of separation is needed between production traffic and the testing environment?

What tools are needed to test lightpath services, both in test beds and in the production environment? How can 'real life' use be simulated as accurately as possible?

5.2. Lightpaths across multi-technology domains

The current practice of creating multi-domain "lightpaths" shows that a "lightpath" is capable of traversing domains that use different technologies. Experience shows that not all technologies are equally suited or applicable. A clear example of "lightpath" that crosses multi-technology domains is a "lightpath" that runs between a next-gen SONET/SDH domain and a domain that uses VLANs on a overprovisioned switched 10GE infrastructure or even "premium IP" on a layer 3 infrastructure.

What technologies are in use for creating lightpath services in networks today, and what are their characteristics?

What interworking problems may arise between two domains each using a different technology?

What are the consequences of creating lightpaths using more than one technology to the end to end service, e.g. on a per technology pair basis?

What are the consequences for troubleshooting lightpaths using more than one technology?

5.3. Future network architecture

Many of the issues discussed earlier are related to the architecture of current advanced optical networks. Some of these issues may disappear, or change considerably, if the next generation of these networks turns out to have a fundamentally different architecture. Understanding future developments is therefore important to all the other subjects of study.

Will future generations of optical networks use substantially different fibre topologies?

Will wavelength plans (filters, lasers, etc.) become more flexible (e.g. tunable lasers and filters)? What are the consequences for network design?

Will dynamic equipment such as Wavelength Selective Switches replace more static equipment?

Will the economics change to the point where radically different design choices start making sense, for instance if adding wavelengths becomes so cheap that it is easier to add an wavelength for a new connection than to multiplex streams across SDH containers?

Will 40 Gbit/s connections be provided as a lambda service or by wrapping multiple 10 Gbit/s lambdas?

6. Literature

[1] Performance Optimization of Dynamic All-Optical Networks, Wolff, Repasky Mumey, Green, and Lin, in <http://www.coe.montana.edu/ee/rwolff/performance-paper-final.pdf>

[2] OCSN 2005, Transient Effects in MEMS-switched All-Optical Networks with EDFAs, Wolff and Achar, in <http://www.coe.montana.edu/ee/rwolff/transients-csav1-hdr.pdf>

[3] Analysis of Physical Constraints in an Optical Burst Switching Network, in http://edocs.tu-berlin.de/diss/2005/buchta_hao.pdf

[4] UCLP Roadmap for creating User Controlled and Architected Networks using Service Oriented Architecture, Bill.St.Arnaud, in http://www.canarie.ca/canet4/uclp/UCLP_Roadmap.doc

- [5] Circuit vs Packet switching - Facts from the network, Campanella, in http://www.terena.nl/events/tnc2006/core/getfile.php?file_id=955
- [6] Developments in Transmission Technologies, Graham, in <http://www.terena.nl/activities/earnest/ws1/presentations/graham-trends-in-transmission-technologies.ppt>
- [7] Routing integrity in a world of Bandwidth on Demand, Wilson, in http://www.terena.nl/events/tnc2006/core/getfile.php?file_id=1078
- [8] The network infrastructure at iGrid2005: lambda networking in action, Grosso, de Boer and Winkler, in Future Generation Computer Systems, Vol.22 Issue 8, pp. 915-919
- [9] G-lambda: Coordination of a Grid Scheduler and Lambda Path Service over GMPLS, Takefusa e.a., in Future Generation Computer Systems, Vol.22 Issue 8, pp. 849-1054
- [10] Using Zero Configuration Technology for IP addressing in Optical Networks, Dijkstra, in Future Generation Computer Systems, Vol.22 Issue 8, pp. 908-914 and in http://www.science.uva.nl/~fdijkstr/publications/Link_Local_Addressng.pdf
- [11] Token based networking: Experiment NL-101, Gommans, , in Future Generation Computer Systems, Vol.22 Issue 8, pp. 1025-1031
- [12] A G.709 Optical Transport Network Tutorial, Barlow, in http://www.innocor.com/pdf_files/g709_tutorial.pdf
- [13] Dynamic Provisioning of LightPath Services for Radio Astronomy Applications, Sobieski ea., in <http://dragon.maxgigapop.net/twiki/pub/DRAGON/DynamicProvisioningOfLightPathServicesForRadioAstronomyApplications/iGrid2005Paper-v8r2-final.pdf>
- [14] Common Service Definition, Sobiesky and Lehman, in <http://dragon.maxgigapop.net/twiki/bin/view/DRAGON/CommonServiceDefiniton>
- [15] The Ongoing Evolution from Packet Based Networks to Hybrid Networks in Research & Education Networks, Martin, in <http://www.jnr.ru/NEC/NEC-2005/proceeding2005/Martin.doc>
- [16] A fault resolution process for multi-domain Lightpaths networks (draft), Hatem, Giesberts and Bos, in <http://www.glif.is/list-archives/tech/docA8099W8jh9.doc>
- [17] Semantics for Hybrid Networks Using the Network Description Language, van der Ham, Grosso and de Laat, in <http://staff.science.uva.nl/~vdham/research/publications/0603-NetworkDescriptionLanguage.pdf>