

PROMISE: Provisioning and Monitoring of Optical Services

Javier González Ordás, Pablo Arozarena Llopis, Javier García Algarra
Telefónica Investigación y Desarrollo

Gal Shabtay
Civcom

Carolina Pinart Gilberga
Centre Tecnològic de Telecomunicacions de Catalunya

Gabriel Junyent Giralt, Jaume Comellas, Eduard Escalona
Universidad Politécnica de Cataluña

Sergi Figuerola
i2CAT

Abstract

PROMISE is a pan-European project putting together the required network solutions to build and experimentally demonstrate a dynamic, self-organised, and fully service-oriented optical transport network, in the perspective to meet the needs of future optical broadband services for data traffic. The main focus of the project is to support a range of advanced high capacity services, from Bandwidth on Demand to Optical Virtual Private Networks, with assured end-to-end quality and reliability, and over an adaptable optical network infrastructure that GMPLS protocols make possible. A key aim of the work is to ensure that the GMPLS-based services, and the underlying technologies, can smoothly integrate into the operational infrastructure of network operators and service providers. The final result of the project will be a control plane enabled test-bed to showcase the provisioning and maintenance of optical services in a real world-like context, as a way to demonstrate the feasibility of providing these advanced services and to help identifying which issues need to be worked out prior to commercial deployment.

INTRODUCTION

The project PROMISE (Provisioning and monitoring of optical services) is a collaborative effort of European institutions, representing both the industry and the academic world, which intends to demonstrate the feasibility of providing dynamic and automated optical services in a realistic operational environment. The project aims at studying the conditions under which these services can be commercialised and at identifying the technological requirements that they put in the infrastructure of network operators and service providers. To meet those demands, the consortium is developing set of tools for service provisioning and supervision.

PROMISE is mainly about how to bring optical networks closer to end users. That demands, first, investigating what happens in the core of a network and finding how to best leverage the capabilities present there to provide to customers in residential or business environments appealing propositions. An optical services portfolio should encompass a range of advanced high capacity services, from Bandwidth on Demand to Optical Virtual Private Networks, with assured end-to-end

quality and reliability, over the adaptable optical network infrastructure that GMPLS protocols enable.

Secondly, there is a need to identify which applications will be the first consumers of the aforementioned services. GRID-based platforms and digital content intensive applications, due to the high capacity traffic demands they place in the networks, are the two most suitable candidates.

Another focal point of PROMISE is optical service management, as, for optical services to finally flourish, drastic changes are required in the legacy operational models and processes of operators. Therefore, the project will propose and assess innovative architectures for dealing with the automatic setup of services, and for smoothly supervising and operating them.

To achieve these goals, PROMISE adopts a holistic approach that encompasses investigation on and implementation of network elements, control software, management software and especially tailored services and applications that take advantage of the overall capabilities of future intelligent optical networks.

OPTICAL SERVICES

Introduction

Transport networks, and the optical technologies that support them, have been traditionally perceived as a fundamentally static layer whose primary role is to provide fixed capacity to higher layers, exhibiting a low degree of interaction with customers. As a result, there is a jarring disconnect between the possibilities offered by the embedded network control and network management and the way optical transport services are being assembled, which is surprisingly slow, manual, and inefficient, particularly when multiple technology layers are involved, and dynamic, automated services are expected.

The increasing trends towards higher bandwidths and adaptable networking have led to the development of distributed control plane solutions for optical transport networks. Industry push and support from the major standardisation bodies have increased the popularity of such control solutions. The next step is the creation, deployment and commercialisation of advanced high-capacity services that take advantage of the control plane capabilities, namely scalability, flexibility, resilience and overall network efficiency. But first, the relevant characteristic of the optical services (from Bandwidth on Demand to Optical Virtual Private Networks) must be identified, along with their interactions with client applications in several scenarios.

Motivation

Distributed control plane solutions for optical transport networks, sometimes referred as intelligent optical networks, are already a reality. Standardization bodies defined a set of protocols (IETF's GMPLS [1]), an architectural framework (ITU-T's ASON [2]) and implementation agreements for specific reference points (OIF'S UNI [3] and E-NNI [4]). Equipment vendors also contributed, implementing products based, to varying extents, on those standards. Network operators followed, albeit at a slow pace and with great caution, by starting the deployment of those solutions in the core of their transport networks.

However, uncertainties still exist about the most suitable architecture to use to deal with multilayer networks (different models such as overlay, peer or augmented have been proposed, see [1], and, most importantly, about how to propagate the

benefits of an intelligent optical network to customers and end users. Indeed, the dynamic capabilities of optical control plane solutions, residing at the heart of the optical network, are an enabler for the delivery of a whole range of new services directly over the traditionally static and slow to provision transport networks. Despite this potential, a number of key issues of interest to network operators, service providers, equipment vendors and GRID-based and digital content application developers remain to be solved, such as:

- Which are the most reasonable service delivery contexts?
- Which candidate services, from the most dynamic to static ones, are likely to succeed in the marketplace?
- How to add control plane awareness to applications for them to take advantage of the above services?

These aspects are explored in the following sections.

Service delivery context

In a basic control plane setup (i.e. a single administrative domain only with I-NNI interfaces) all functionality remains under the sole responsibility of a single operator. When this same control plane is to be used as a service delivery tool (via UNI or E-NNI interfaces), new stakeholders appear at different points of the chain that add their particular expectancies, requirements and constraints. The resulting context, under which optical services will be provided, is illustrated by Figure 1.

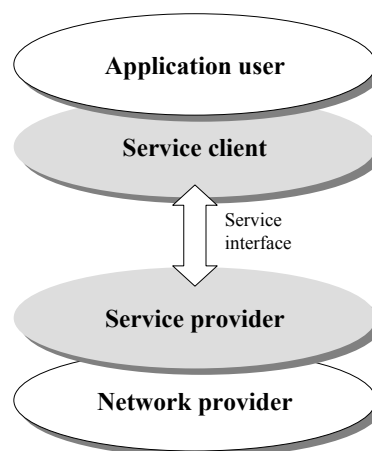


Figure 1. Service delivery context

On the provider side, it is worth differentiating between the owner of the resources, the network provider, and the organisation that actually offers the services, the so-called service provider. In most common cases the same company, often a telecom operator, will undertake both roles. However, the objectives associated to each role, and the methods and tools adopted to achieve them will be of a different nature. Indeed, a higher level of integration must be expected in this single provider case, but the fundamental separation of roles should be retained.

Similarly, on the client side, a separation of roles between the service client and the end-user of applications is proposed. The service client contracts the connectivity service from the service provider, and organises that connectivity, along with its own resources, to provide network bandwidth, typically at layer 3 (i.e. IP), to the final application users. It can be seen that the service client is more focused on the

networking aspects of the service, while the application user cares about the satisfaction and benefits that it can bring to him (from enhanced entertainment to higher productivity). Again, both roles will generally be found under the same entity, might it be a residential customer, an enterprise or an academic institution, but this is not a mandatory condition.

Specific service delivery contexts are defined using various arrangements of the above four elemental stakeholders. Several reference scenarios, from a basic setup to a more futuristic scenario, describe what PROMISE expects to be the more common cases of optical service usage. These reference scenarios are:

- Single customer-provider scenario: a single entity (typically a telecom operator) manages the service and network provider roles, and a customer acts both as the service client and application user. At this customer internal separation is kept between those two roles. It is assumed that the customer-provider interface does not carry topology or route information, resembling an overlay model where the customer essentially issues connection requests.
- Multi-provider scenario: a service provider relies on multiple network providers for the delivery of a service. In the customer side, separation between roles is assumed as in the basic scenario above, and it is also assumed that the customer-provider interface does not carry topology or route information. Two models can be adopted:
 1. The service provider has business agreements with all network providers in the chain, and it is responsible for mediation at the points where different networks interconnect. This is a hub model.
 2. The service provider has an agreement with the edge network provider. This, in turn, has a connection agreement with one or several transit providers to meet connectivity needs between non-adjacent network edge areas. This type of relationship may be repeated recursively between transit providers. This is a chain model.
- Integrated application user and service client: there is integration between the application user and service client roles. The application directly interacts with the service provider. An example of this kind of scenario is a GRID system that autonomously issues lightpath setup and tear down requests. It is not mandatory for the application to be attached to the provider optical network. There may be intermediate networking devices; the expected direct interactions happen at the control plane. An accepted option is the existence of a proxy between the applications and the service provider.
- Peering scenario: this scenario is identical to the basic one, except that the customer-provider interface can carry topology (full or partial) and path information. A high degree of trust is usually associated to this scenario. This model enables richer VPN services. The most relevant characteristic of the peering scenario is the visibility of the client over the network, and its ability to control the resources allocated to him, to an extent dependent on the particular service conditions.

Services definition

By the term optical services we refer to connectivity provided directly by an optical network (i.e. layer 1) to high capacity demanding customers. In the scope of PROMISE, the optical network only supports the connection oriented, circuit switching paradigm. Optical services will then comprise a number of optical point-to-point connections (sometimes named lightpaths or lambdas) arranged in an attractive manner for customers. The presence of an intelligent control plane is a

key feature for service differentiation, particularly in the areas of rapid provisioning and efficient failure recovery.

The main interest of PROMISE is in the dynamic services enabled by emerging control plane implementations. However, for comparison purposes and because there are circumstances when it is the only available option, more static services, governed from a centralised management systems, are considered as well.

Optical services can be characterised by assessing how they perform in each of the following service dimensions:

- **Dynamicity:** capability of rapidly changing the characteristics of a given service upon customer demand or by autonomously adapting to the service environment.
- **Security:** a secure service protects its users by implementing mechanisms that guarantee data integrity and confidentiality.
- **Capacity:** amount of bandwidth that can be transported through a single instance of the service.
- **Granularity:** minimum amount of bandwidth that may be carried by a service, as well as the increment/decrement unit for bandwidth changes.
- **Manageability:** refers to how efficiently and cost effectively a service or network can be operated.
- **Billing:** ability to bill a service by as many parameters as possible.

Attending to these dimensions, it is possible to identify the following services:

- **Provisioned bandwidth service:** defined by OIF in [5] as "enhanced leased/private line services. Provisioning is done at the customer request by the network operator.
- **Scheduled bandwidth service:** a variant of the provisioned bandwidth service where the connection activity times are pre-arranged between the provider and the customer (i.e. the customer "books" the resources), rather than provisioned on demand.
- **Bandwidth on demand service:** defined in [5] as follows: "OC-n/STM-n and other facility connections are established and reconfigured in real time. Signaling between the user NE and the optical layer control plane initiates all necessary network activities. A real-time commitment for a future connection may also be established. A standard set of "branded" service options is available..
- **Point-to-multipoint service:** allows one source to send data to a variable number of consumers at a time. Provisioning of this service could adopt a management oriented or control plane based approach.
- **Optical VPN service:** described in [5] as follows: "The customer contracts for specific network resources (capacity between OXCs, OXC ports, OXC switching resources) and is able to control these resources to establish, disconnect, and reconfigure optical connections. In effect they would have a dedicated optical sub-network under their control."
- **Network peering service:** A peer model of interconnection is used to integrate the networks of two parties with a high degree of trust. The two parties signs a peering agreement giving the client network full visibility and control over the server network. A single instance of the control plane runs in both networks.

OPTICAL SERVICE MANAGEMENT

Introduction

Once the business and technical environment which will enable the deployment and commercialisation of optical services have been identified, and the essential technological attributes of those services have been captured, the next step is to define the operation, administration and maintenance of such services, and of the underlying network that supports them. The deployment and delivery of the set of innovative optical services that control plane solutions enable, and which represent an important step forward for the exploitation of transport networks, demand significant efforts to upgrade the back-office tools and systems on which service providers rely to support these service offerings. The move towards more dynamic and self-configurable networks must be accompanied by a similar move towards more agile and automated management systems. The PROMISE management plane will exhibit such required agility and self-configuration capabilities for smooth and efficient service provisioning..

Motivation

If service providers want to fully exploit the new services enabled by the control plane, then their operations support infrastructure must be accordingly adapted to migrate from a network oriented approach towards a service and customer focused one. On the one hand, upstream of the existing management systems, increased service visibility should be added to better meet the customer needs for provisioning, smooth maintenance and reparation after faults. Customers should be offered a self-service system where they can request service modifications, or set up a new connection, or monitor the status of all elements associated to the service, in an automated fashion with no involvement of the service provider personnel. On the other hand, looking downstream, the management plane must become aware of the underlying control plane capabilities and components, and manage and make use of them to be able to support the services that rely on it.

Optical transport network management systems today are not ready yet for this migration, for several reasons:

- Network models in use are too rigid to accommodate all the innovations that come with CP-enabled final service offerings. Not only a service view must be added to capture all specific parameters of each customer service instance, but the arrival of new connection types (e.g. switched connections), new conceptual entities (e.g. calls, as opposed to simple connections), and new resilience techniques (e.g. pre-planned restoration) must be addressed.
- Lack of customer focus, which complicates the addition of customer network management capabilities to the existing systems. Current generation systems were designed with the network as the main focus with little or no information stored about customers, their services and how these are impacted by network faults.
- Strict hierarchical approach to managing the network, in opposition of the demands of a service oriented optical network, in which the CP can take decisions on its own and customers should be able to implement configuration changes.

Requirements

PROMISE requirements are classified into two main groups, following the guidelines laid out in TMF's eTOM framework [6]. Requirements related to service management describe the needs of service providers with respect to creating, instantiating, activating, securing and monitoring customer services, and the capabilities that they must offer through the so called customer network management systems for customer self-management. Requirements related to resources specify the exchange of information between network resources and management systems in support of the above service management functionality and of other operational processes realized by network operators.

Delivering a service involves a wide range of activities, from service definition to Service Level Agreement (SLA) verification. A service must be first characterized by a number of parameters. How many ports or devices can it connect? What triggers the set up of connections? What are the acceptable quality objectives? Then, during the provisioning process, this definition will be instantiated into actual customer services, bound to SLA specifications. The activation phase involves configuring the control and transport planes in the network to be ready to provide the service. Connections associated to the service can then be set up via the control plane (switched connections) or the management plane (soft-permanent connections). Once a service is active, proactive supervision of it is carried out to detect service problems and measure the delivered quality, which is compared with the agreed SLA parameters. In PROMISE, all these phases will be supported and coordinated by the PROMISE service management solution.

Service management functionality is supported by actions taken on network resources, and information gathered from them. Both control plane resources (routing and signalling protocol, instances, discovery components, UNI and NNI interfaces, policy servers, etc.) and transport plane resources (ports, connections, allocations to customer services, etc.) must be properly configured. Alarms must be collected to learn about malfunctioning components and take appropriate recovery actions, as well as to map them to service problems and take them into account during the provisioning process. In the performance management area, raw measurements from both network elements and monitoring devices must be retrieved, aggregated and analyzed to assess overall network and service quality.

PROMISE service management solution

In order to meet these requirements, the PROMISE service management solution will be constructed over a generic management platform. This platform features a distributed, component-based architecture built around a CORBA bus that facilitates internal communications. System components are largely independent to enable software reuse and uncoupled evolution of components. The adopted management platform provides a number of supporting and application components, which are tuned for PROMISE purposes and complemented by specific optical service management components. Operators and customers interact with the system using its GUI, a Java applet that can be hosted by most browsers. A network database provides the persistency layer of the system, while standard SNMP interfaces support the communication towards the control plane and monitoring devices.

The overall architecture where the service management solution will be inserted can be seen in the Figure 2.

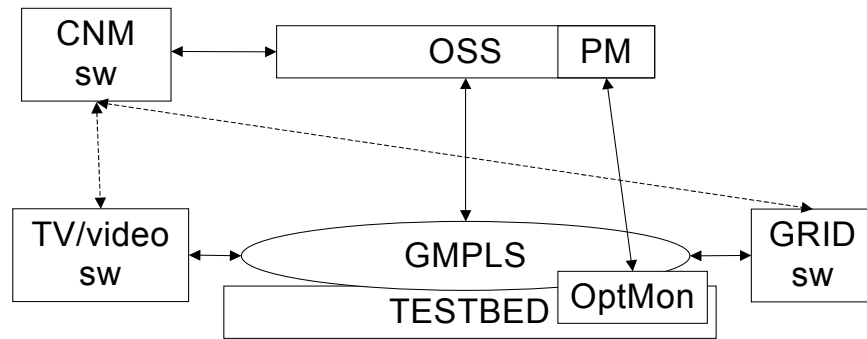


Figure 2. Overall architecture overview

The service management solution (OSS – Operational Support System plus PM – Performance Manager, in the figure), interfaces both with the GMPLS control plane agents and the optical monitoring devices in the testbed. It also provides self-management capabilities via the Customer Network Management (CNM) module, which could be further interface with the customer applications for more integrated management.

The internal system architecture, as can be seen in Figure 3, is based on the CORBA bus over which application and supporting components communicate with each other and, via adapters, with network elements and monitoring devices. A common Graphical User Interface hosted in standard web browsers gives both operators and customers access to the system capabilities

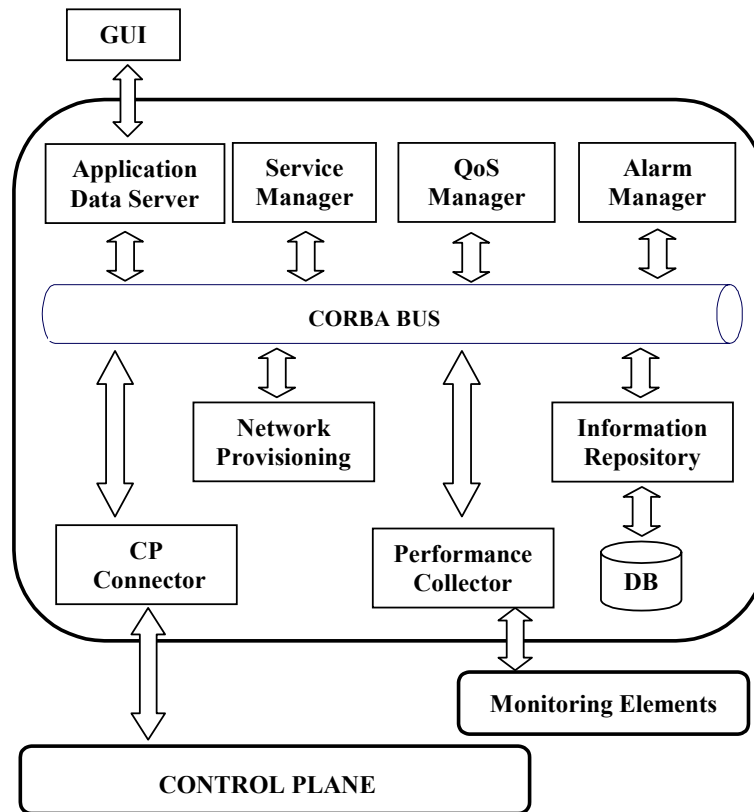


Figure 3. PROMISE MP architecture overview

OPTICAL PERFORMANCE MONITORING

The drive for reliable, high-bandwidth and multi-service networks results in stringent requirements for performance monitoring in transparent optical paths, as well as in the electrical domain. There is little doubt that performance monitoring is an essential and integral part of any optical network that provides advanced services, assures quality of service and is to be maintained in a timely fashion. In the recent years, this subject has attracted a lot of interest, encouraging significant academic research activity and numerous products from leading component vendors. This progress, along with the drive for reliable DWDM networks has paved the way for monitoring in the physical layer of optical networks, which has the potential to reduce network costs.

PROMISE has therefore devoted significant efforts to study the monitoring aspects of managing a service-oriented optical network, and how performance information can be used to guarantee service quality and reliability. This resulted in the identification of the most significant parameters that should be measured. Together with the parameter descriptions, the PROMISE consortium has presented methods - in the optical and temporal domains - to measure the various parameters that affect network performance on the one hand and assure high quality of service on the other hand.

Optical performance monitors have become common devices deployed in optical networks, as they can be used to measure a variety of optical parameters to assess network performance. While bit error rate (BER) serves as the ultimate performance metric for the health of a service, optical transmission systems can benefit from additional performance monitoring with the various DWDM channels being examined on-line in the optical domain. This will allow for:

- (a) Better network reliability
- (b) Enhanced network management and maintenance capabilities
- (c) Improved service provisioning

With this in mind, the set of key optical performance parameters targeted by PROMISE are:

- Optical power: provides information on fiber cuts as well as amplifiers' malfunction and indicates wavelength dependent loss if unexpected variance is detected in the measured optical power of the various DWDM channels
- Wavelength drifts: per channel actual wavelength variation may occur due to laser aging or wavelength locker inaccuracies. As channel spacing is being narrowed, the permitted laser drifts are increasingly limited and, as such, should be carefully measured
- Optical Signal to Noise Ratio (OSNR): a measure of the ratio of signal power to noise power in an optical channel.
- Q-factor: a numerical value that directly reflects the quality of the optical transmission signal. Q is obtained by observation of the amplitude and phase characteristics of the (analog) electrical signal.
- BER: BER measurements indicate the exact number of errors. However, it is not easily measured and cannot be used alone to indicate the type or location of a fault.
- Packet loss: this is a packet domain metric that provides the average proportion of packets lost during a measurement period.
- Chromatic dispersion: a variation in the velocity of light (group velocity) according to wavelength. This variation in velocity causes the pulses of a modulated laser source to broaden when travelling through the fiber; up to a point where pulses overlap and bit error rate increases

For each parameter, PROMISE has considered a number of ways for its measurement and has proposed how to exploit the measured information (whether by using spectral means or via direct temporal detection) for service management and resilience. It has been shown why performance monitoring should be an integral part of any optical network carrying high-bandwidth, quality services is needed. In addition means for system designers to design systems with monitored quality of service have been provided, while also making suggestions on the intrinsic parameters that constitute an optical performance monitor that may be potentially of interest to component vendors.

CONTROL PLANE

The control plane is the entity that provides intelligence to the network. A well-designed control plane entails a better network control to the service provider allowing different kind of services within an scalable and flexible architecture. The

main objective of the control plane is the establishment or release of connections upon a request made from a management system or a client (e.g. user or application). With this purpose, the functionalities can be divided into different tasks: signalling, routing, discovery and information dissemination.

Signalling

Signalling messages are used, for example, for call and connection control and management for the establishment, modification, status query and release of an end-to-end connection. The signalling communication network must terminate at each of the nodes in the transport plane.

Routing

The control plane in photonic/hybrid networks requires the explicit route of a connection to be computed at the ingress node of this connection. Spatial path computation should make use of information relative to network state in order to optimise overall resource utilization. The availability of this information, enabling traffic engineering to be performed at the ingress, is ensured by a link-state routing protocol. However, the information flooded is the result of an aggregation and is by essence not precise. The aim is to compute a spatial route taking into account both the eventual constraints requested for the connection and the information about link states.

Discovery

The auto-discovery functions performed in the control plane include neighbour, resource and service discovery. Neighbour discovery corresponds to associating two network entities within a layer network based on a specified adjacency relation. In an ASON environment, each node automatically discovers adjacent nodes both from control plane (CP) and transport plane (TP) point of view and stores the logical connectivity information in its routing database. This information can later be retrieved by the management plane (MP).

Local resource discovery is the procedure that a node takes to find out what resource it has for service provisioning. Service Discovery describes the auto-discovery feature that is used for verifying and exchanging service capabilities of a network. As client devices are connected to the network, they get via CP information on the transport network available service (e.g. supported call rates, available call protection schemes and so on). The MP must be able to define the services and their capabilities on each port of the network, enabling or disabling selected functionality.

Information Dissemination

Together with the Routing Controller the control plane uses OSPF-TE to disseminate traffic-engineering information. This information is advertised throughout the network as Opaque Link Advertisements (OLSA) using the flooding capabilities of the protocol. Once all OSPF daemons have the global information, a new database is constructed, the Traffic Engineering Database (TED) with the OLSAs of the whole network. When a change occurs (i.e. a data-link has been allocated) the new state is advertised and updated in the TEDs all over the network. This updated information will be used to find free routes with updated information about the resources. This is especially interesting in the case that no wavelength converters are available. With traditional GMPLS and specifically using the LABEL_SET object, just spatial routes are computed, so the use of global resource information dissemination improves the wavelength assignment.

Control Plane architecture

The ASON G.8080 recommendation states that the control plane architecture can be logically divided in some functional blocks:

- Routing controller (RC): The RC calculates routes from a source to a destination node given some specific parameters as wavelength to use, bandwidth needed, lambda continuity (no wavelength converters), etc. The protocol implemented in PROMISE for the information dissemination in order to calculate updated routes is OSPF-TE.
- Link Resource Manager (LRM): It implements the interface between the OCC and the optical node (Connection Controller Interface, CCI) and maintains the state of the logical resources defined. It has a real-time updated database of the link state of the Forwarding Adjacencies (FA) of a node. It also configures the optical equipment to reserve the resources for each connection by means of SNMP.
- Connection Controller (CC): The CC is the responsible of receiving the connection requests and to process them upon some parameters given. It accepts connection requests from a network management system (NMS) through the NMI (SNMP) and interacts with the other processes to trigger a soft-permanent connection.
- Signalling Controller (SC): It implements the NNI interface using the RSVP protocol. When the request is processed, it creates a PATH message, forwards it through a pre-calculated route (built in the Explicit Route Object – ERO), and sends a RESV message back by the inverse route. On each node, it transfers to the CC some information needed to update the MIBS and reserve the corresponding resources.




Figure 4. Control Plane Architecture

The relations between the different blocks are shown in Figure 4.

AGILE OPTICS

Chromatic dispersion (CD) is considered as a major limiting factor in 10Gb/s networks. Therefore, CD compensation techniques are being explored intensively, involving different technologies ranging from dispersion compensating fibers (DCF), free-space etalons, fiber bragg gratings and planar lightwave circuits.

Being an optical phenomenon, CD, which is regarded as linear filtering in the optical domain, can be best compensated for optically, i.e., via a tunable optical dispersion

compensator (TODC). Thus far, TODCs have been considered as bulky, costly, dissipate high power, and often did not provide the necessary required tunability.

PROMISE will use a TODC based on free-space architecture comprising of cascaded Gires-Tournois etalons that offers the following uncompetitive features:

- Ultra small TODC packaged in a standard and compact butterfly package
- Widely tunable – introduces any dispersion value ranging from -1,700ps/nm to +1,700ps/nm
- Less than 2dB insertion loss
- Low power dissipation
- Bit-rate and OSNR agnostic
- Enables both single and multi-channel compensation
- Based on reliable and mature free-space optical components
- Less than 2dB OSNR penalty at 200km of 10Gb/s NRZ transmission in SSMF
- Interoperable and complementary to dispersion tolerant line codes as well as electronic dispersion compensation (EDC) for further extending systems' dispersion tolerance.

Having a low-loss TODC simplifies the dispersion map of a network and can be used in various applications:

- Integrated in 300-pin MSA transponders or transceivers – compensates single channel residual dispersion, allows for up to 200km routes with no external dispersion compensation
- Multi-channel compensation for in-line dispersion, e.g., every 80km – saves EDFA's, simplifies amplifiers' architectures and allows for smooth upgrade of 2.5Gb/s to 10Gb/s networks
- Integrated in ROADMs or amplifiers – enables low-loss cost effective solutions

CONTROL PLANE AWARE APPLICATIONS

Some emerging applications are fuelling the growth of data traffic which in turn asks for high capacity and flexible services. Examples of such applications are: multi-media applications (both real-time and non-real-time), shared access to remote resources, network-wide computation and data services (grid-computing, data-grid), data storage networking, disaster recovery, etc. The characteristics and the network requirements raised by these applications are highlighted in the following sub-sections.

Video applications

INTRODUCTION

Taking into consideration that HDTV applications are real-time (or near real-time) they will impose strict conditions to the optical network. These constraints will be different when considering two-way (videoconference), one-way point-to-point (Video on Demand - VoD) or one-way point-to-multipoint (distribution) communications. In this way, High Definition (HD) videoconference is the most restrictive application as only very low delay is allowed. On the other hand, HD VoD

or HD distribution have also to be provisioned in the range of seconds (users can not be waiting for minutes to see a sports match or a movie). From the capacity point of view all of the applications can potentially occupy a whole optical channel.

VIDEO APPLICATIONS AND NETWORKING

Different network aspects will be taken into account and the requirements of every application is estimated.

- Resilience: two different parameters have been analyzed: availability and recovery time. We consider that an availability of 99.5% (which means outage periods of hours during a year) is enough for the different HDTV applications. With respect to the recovery time in case of failures, 100 ms is considered an assumable value.
- Provisioning: A provisioning time of some seconds is considered appropriate for all the HDTV related applications. When looking at the lease time it is expected that values can range from minutes (in case of HD videoconference) to hours (a movie usually takes around two hours).
- Interfaces: HD TV distribution is point to multipoint while the rest of the applications are point-to-point. The interfaces are supposed to be layer 3 (IP) for all the applications.
- Physical characteristics

Table 1 shows the main physical characteristics of the different applications that will be considered in the PROMISE test-bed. It is important to highlight that they match perfectly with the optical network below so it seems feasible to adapt these applications to the optical network

Application	Degree of symmetry	Data rate	Key performance parameters and target values		
			One-way Delay	Information loss	Other
HD Videoconference	Two-way	270-1500 Mbps	< 150 msec preferred < 400 msec limit	< 1% FER	Lip-synch : < 100 msec
HD video on demand	One-way	20-1500 Mbps	< 10 sec	< 1% FER	
HD distribution	One-way	20-1500 Mbps	< 10 sec	< 1% FER	

Table 1 Performance demands from applications

- QoS issues: No particular QoS issues are devised when comparing to other applications. Targeted provisioning and recovery times are compliant with the typical values considered in GMPLS oriented networks. On the other hand, packet-loss ratio has to be dealt at the IP level then it has no implications on the optical network.

GMPLS SUITABILITY TO TV APPLICATIONS

Many universities and research institutions are now deploying applications of a next-generation IP and optical infrastructure for distributed and dynamic services such as grid computing, distributed supercomputing, high-definition (HDTV) imaging, streaming video and video-on-demand. These innovators are driving new

high-bandwidth applications and spurring the development of "bandwidth broker" management tools. So GMPLS is a good way to these applications and innovators.

GMPLS set of protocols define a control plane between multiple optical layer network elements, such as SONET/SDH, wavelength, and fibre switched devices. GMPLS uses IP-based routing and signalling protocols to establish connections between access points in the OTN with sufficient information and control to support carrier-class inventory, provisioning, and traffic engineering operations.

In PROMISE, the final objective is to run these applications over GMPLS controlled optical channels in a real test-bed, improving the interactions between network, services and applications. Optical networks are in principle suitable for these applications as they are able to provide large amounts of bandwidth on demand, assuming a GMPLS model.

Point-to-point HDTV applications can be mapped directly over optical channels in a GMPLS/ASON scenario. Therefore, no new extensions are required from current protocols and only an interface between the applications and the optical network would be necessary. Then, the PROMISE work will focus on the development of this interface, which will be responsible for generating a request to the optical network control plane for setting-up an optical channel.

The study of L1 VPN [7] provides interesting insights in relation with the concept of HDTV distribution. It is possible to envisage that for providing a HD Video broadcast (IP-HDTV) application we need of a point of presence that "radiate" the video to the customers, that usually is a replica located in the metro area where customers are; the link between the replica and the customer site may be provided by VPN-L2 or Public IP and it develops inside a metro area, but the data transferred from the source of video signal to the replicas must be carried on a more available and fast media, such as VPN-L1 to minimize excessive delays and too low availability.

GRID applications

INTRODUCTION

A Grid is a collection of distributed computing resources available over a local or wide area network, which appears to an end user or application as one large virtual computing system. The idea is to create virtual dynamic organizations through secure, coordinated resource-sharing among individuals, institutions, and resources.

Grid computing is an approach to distributed computing that spans not only locations but also organizations, machine architectures and software boundaries, to provide unlimited power, collaboration and information access to everyone connected to a Grid.

It is widely accepted that GRID applications will create new business models and will force network providers to evolve their networks and support large-scale grid traffic and other novel requirements posed by Grids. On the other hand, network providers will probably use Grid technologies to improve their own internal operations (e.g., billing, events' analysis, and modeling). In summary, GRID applications will imply major changes in the networking world.

GRIDS AND NETWORKING

Applications over grid architectures try to get advantage of the potential resources available in a network. This results in the possibility of sharing them in order to accomplish a task in an optimal way, this is, dividing tasks and assigning them to

network components obtaining much faster results. However, the bottleneck in such architectures has usually been the network speed. For complex tasks which deal with large amounts of data, transmission speed has been the main difficulty. Optical grid networks solve the problem, allowing data to be shared regardless the size and bringing the possibility of an effective resource-sharing between heterogeneous grid components.

There is another advantage from these fast networks: once the transmission time is not a problem, new virtual organizations can emerge in this scenario. Distance is a problem no more so information sharing in real time of large amounts of data becomes possible, allowing interaction between agents where it was not possible before.

Finally, not only applications which need real time data take advantage of this, but also batch processes and backups between organizations.

Given these mentioned premises, the Grid High-Performance Networking Research Group (GHPNRG) (see [8]) has studied the relationship between network research and Grid application and infrastructure development. The objective of its work is to cover the space between the networking and grid research communities. It serves as a forum for information exchange on advances and requirements in both fields, as well as by providing a focal point for liaison activities between the Global Grid Forum (GGF) and the different networking standards bodies. One of the PROMISE objectives is to provision optical channels for GRID applications being then tightly related to the GHPNRG activities. The GRID applications characteristics have to be taken into account as well as the network capabilities (high bandwidth traffic transport, dynamic resource allocation, availability and QoS provisioning). In summary, in the context of the PROMISE project GMPLS based optical network, the whole network can be considered as a GRID service and then be used as another GRID resource.

Grid applications using optical network infrastructure are currently only utilized by a small number of organizations which need to process extremely large jobs (e.g. large data file transfers between known users or destinations). As this kind of GRIDs is static, long-lived wavelength circuits between clients with centralized management are usually deployed. This type of Grid networking relies on carrier provision of optical network resources while the Grid users have no visibility of the optical network infrastructure. In other words, the Grid user is not able to setup paths over the optical Grid network.

As Grid applications evolve, the need for user controlled network infrastructure increases, in order to support emerging dynamic and interactive services. Examples of such applications may be high resolution home video editing, real-time rendering, high definition interactive TV, e-health and immersive interactive learning environments. These applications need infrastructures that makes vast amount of storage and computation resources potentially available to a large number of users. Key for the future evolution of such networks is to determine early on the technologies, protocols, and network architecture that would enable solutions to these requirements.

There are then two main common requirements generated by a large number of Grid applications:

- Large amounts of bandwidth provisioned and scheduled on-demand
- User or application management and control of the network resources (i.e. set-up self-organized distributed computing resources and facilitate bulk data transfers)

A number of other requirements concerning throughput, priority, latency, QoS and storage capacity will also influence the Grid network design but they are more specific to the type of application. Grid applications are also likely to differ in the number and type of participants, and also in the degree of trust between the participants.

A new network concept is now emerging to satisfy Grid application requirements. This is a network where resources such as ports, whole equipment, even bandwidth are controlled and maybe owned by the user. Furthermore, in contrast to traditional (telecommunications) networks where applications are allocated resources and routed over fixed network topologies, in Grid networks, resources under user/application control are organized in an automated way to provide connectivity without getting the permission from a carrier or a central authority. In other words, the user will drive its own virtual network topology.

G-UNI service

The signalling and data transport to be developed between Service provider and the underlying optical transport network constitute the GUNI, (i.e. the interface between the Grid user and the optical transport network. This same interface can also support other applications, such as VoD). The GUNI functionalities are grouped in the following categories [8]:

- Signalling
 - Flexible bandwidth allocation
 - Support for claiming existing agreements including
 - Scheduled services, i.e. advance reservations
 - Incorporation of AAA-information
 - Automatic and timely provisioning
 - Light-path setup
 - Automatic neighbour hood discovery
 - Automatic service discovery
 - Fault detection, protection and restoration
 - Propagation of service and agreement related events
- Transport
 - Traffic classification, grooming, shaping and transmission entity construction
 - Data plane security

The signalling mechanism will be responsible for requesting, establishing and maintaining connectivity between Grid users and Grid resources while the data transport mechanism will provide a traffic/bandwidth mapping between the Grid service and the optical transport network. Most of these functionalities are already given by the optical network control plane so a simplified interface will be developed which only generates a request to the optical network.

PROMISE G-UNI DESIGN OPTIONS

In this section the design architectures proposed for the grid enabled interface between the application layer and the GMPLS layer are discussed.

First approach: framework layer

The first solution is based in an intermediate layer, called Grid Framework (see Figure 5, which is responsible for deciding which type of net resource is being used for each petition of the application.

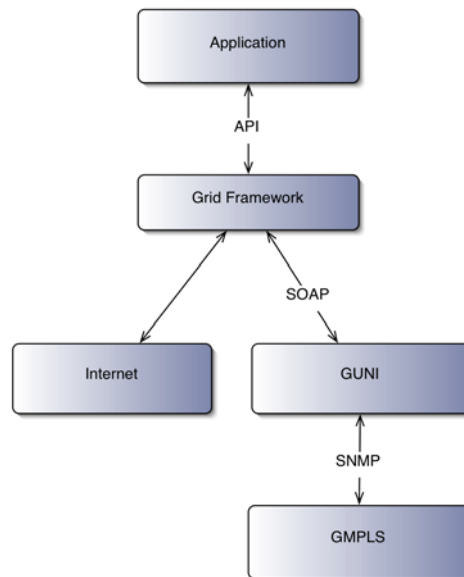


Figure 5. Framework layer

The framework would consist, basically, of a user-friendly API interface used by the application layer. The application would use its functions to demand a resource, specifying the parameters. Then the framework, according to this parameters and its real-time information about the nodes of the system, will redirect the job to the best node possible. If the framework decides that the optical net is the best solution, it will contact the GUNI interface in order to obtain the best optical net route. If not, it will redirect the job to the internet accessible grid node.

Second approach: direct access

In this scenario, the application has to decide in each moment which type of connection is to be used to obtain the best performance. The application has direct access to the internet accessible grid nodes and to the GUNI interface, as illustrated in Figure 6.

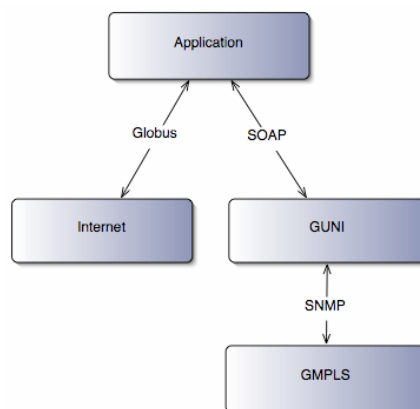


Figure 6. Direct access

The solution preferred is the latter. Although the framework layer provides abstraction and reusability, it is to be importantly considered that in the

applications to be developed, performance is a critical factor. A framework layer is usually too much generic and less controllable or predictable by the final developer, as it should be AI-enabled and should require an intensive supervised learning. This project is focused in specific types of applications in which the type of connection to use in each case can be decided correctly in the developing process, so the automatism provided by a framework layer is not so influential.

HIGH LEVEL GUNI SERVICES

The system use cases will be responsible for optical network connection management and monitoring functionalities. The former will be called manageConnection, while the latter will be called monitorConnection. Among the functionality supported we can mention the following use cases (see Figure 7):

- getConnection: will open a connection with the optical net, given a few parameters
- endConnection: will close a given open connection
- isReachable: will verify that a given node is reachable within the optical network
- isAvailable: will verify that a given node is able to open a connection, given a few parameters

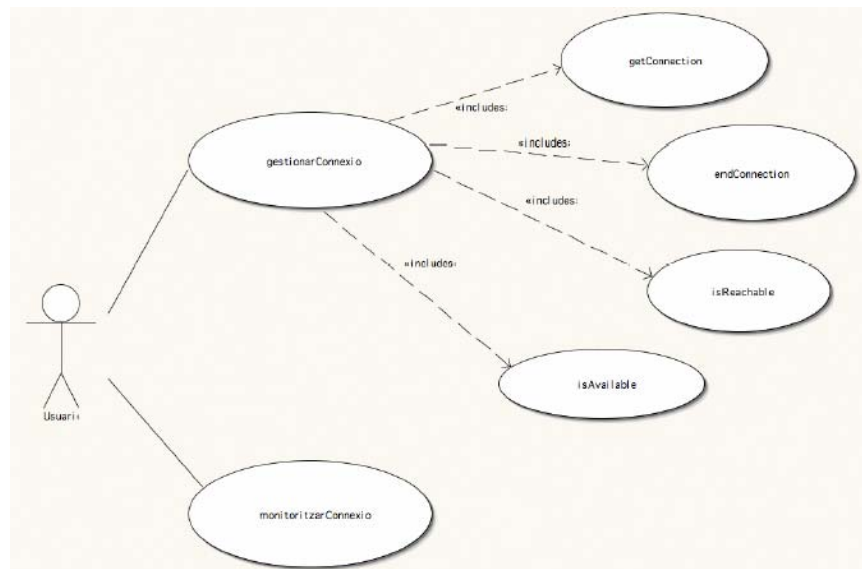


Figure 7. Application use cases

There will be a high level service for each of this use cases. They will be implemented using the software design pattern of use case controllers.

PROMISE TESTBED

PROMISE is building a service oriented test-bed where the most relevant aspects of an optical service offering will be covered. The test-bed will be constructed through the integration into an existing network of different building blocks, at the transport, control and management planes. In addition, GMPLS-aware applications

implemented by the project will be an essential part of the test-bed as the means to highlight the potential value of the technologies involved.

The testbed will integrate, in an optical metropolitan network, the following elements:

- Photonic cross-connects, enabled with GMPLS routing and signalling protocols
- Optical network and service management systems, with support for provisioning and monitoring and SLA visibility
- Optical measurement devices
- GMPLS-aware multimedia applications and GRID middleware

Figure 8 illustrates the architecture of the optical nodes that will comprise the testbed




Figure 8. sROADM Architecture

Experimentation at the testbed will help producing recommendations and best practices for migrating to and operating service-capable dynamic optical networks, along with a processes and systems architecture for optical service fulfilment and assurance.

Finally, the project plans to carry out validation tests looking at the carrier-grade readiness of control plane protocols, focusing in particular in their resilience aspects

CONCLUSION

PROMISE will address, from a service centred perspective, the topic of the evolution of optical networking to help giving shape to a path from today's architectures and technologies towards all optical, service driven, intelligent optical networks.

PROMISE intends to act as a driver for the commercialisation of optical, GMPLS-based, services. For operators and service providers reluctant to take the step that separates technological availability from full service deployment, validation and demonstration in the project test-bed should help to obtain a better understanding of the potentiality and challenges that these services present.

The project will also contribute to adding GMPLS awareness to end user services and applications, such as Video on Demand or GRID middleware, and will showcase how those applications take advantage from GMPLS. Support of these capabilities at client devices will serve as an additional incentive to deploy optical services.

Another problem that optical services face is the severe architectural limitations of many legacy OSSs. PROMISE will build service fulfilment and assurance tools to facilitate a smooth transition and cover identified functional gaps in current architectures. Service visibility will be a common attribute of these fulfilment and assurance systems; both will use agreed SLA parameters to govern their provisioning and supervision strategies.

In conclusion, the advent of innovative optical services must follow the deployment of optical control plane solutions in order to realize the full potential of intelligent optical networking. All participants in the delivery chain, i.e. network and services providers, customers and end users, can benefit from such deployment, either as a result of reduced costs, increased revenues, higher efficiency or more satisfaction.

For this scenario to come true, it is essential that both sides move towards the common goal. On the providing side, services must be created and tailored to suit the needs of the targeted user applications. On the client side, control plane aware applications must be built to take full advantage of the service capabilities. PROMISE will work on both sides and demonstrate the benefits to each of them in order to serve as a catalyst to trigger the widespread availability of optical services.

Acknowledgement

PROMISE is a Celtic initiative project partly funded by the Spanish Ministry of Industry, Tourism and Trade as part of the PROFIT programme

References

1. E. Mannie (Ed.), "Generalized Multi-Protocol Label Switching (GMPLS) Architecture", IETF RFC 3945, October 2004
2. ITU-T Recommendation G.8080/Y.1304, "Architecture for the Automatically Switched Optical Network (ASON)", November 2001
3. J. Jones (Ed.), "User Network Interface (UNI) 1.0 Signaling Specification, Release 2: Common part", OIF Implementation Agreement, February 2004
4. L. Ong (Ed.), "Intra-Carrier E-NNI Signaling Specification", OIF Implementation Agreement, February 2004
5. Yong Xue, Monica Lazer (Editors), "Carrier Optical Services Framework and Associated Requirements for UNI", OIF contribution oif2000.155.1, September 2000
6. Telemanagement Forum Guidebook 921, "Enhanced Telecom Operations Map (eTOM). The Business Process Framework for the Information and

Communications Services Industry. Addendum D: Process Decompositions and Descriptions", Release 4.5, November 2004

7. ITU-T Recommendation Y.1312, "Layer 1 virtual private network generic requirements and architectures", September 2003
8. Dimitra Simeonidou and Reza Nejabati (Editors), "Optical Network Infrastructure for Grid", GGF Informational document, GFD-I.036, August 2004

Acronyms

ASON	Automatically Switched Optical Network
BER	Bit Error Rate
CD	Chromatic Dispersion
CP	Control Plane
GMPLS	Generalized Multiprotocol Label Switching
GUNI	GRID User Network Interface
MP	Management Plane
NNI	Network to Network Interface
OSNR	Optical Signal to Noise Ratio
PROMISE	Provisioning and monitoring of optical services
SLA	Service Level Agreement
TODC	Tunable Optical Dispersion Compensator
TP	Transport Plane
UNI	User Network Interface
VoD	Video on Demand
VPN	Virtual Private Network