



18-01-2011

Deliverable DS1.1.1,2: Final GÉANT Architecture



Deliverable DS1.1.1,2

Contractual Date: 30-09-2010
Actual Date: 18-01-2011
Grant Agreement No.: 238875
Activity: SA1
Task Item: T1
Nature of Deliverable: R (Report)
Dissemination Level: PU (Public)
Lead Partner: DANTE
Document Code: GN3-10-279

Authors: L. Altmanova (CESNET), T. Breach (NORDUnet), M. Campanella (GARR), M. Carboni (GARR), M. Enrico (DANTE), L. Fischer (NORDUnet), R. Pekal (PSNC), J. Radil (CESNET), R. Sabatino (DANTE), M. Scarpa (GARR), S. Sima (CESNET), T. Szewczyk (PSNC), R. Tuminauskas (LITNET), C. Tziouvaras (GRNET), J. Wojtech (CESNET)

Abstract

This deliverable is the final report on the findings of the backbone architecture study. It presents architecture options and recommendations for the future network taking into account current and future requirements (such as current architecture, services, quality, user requirements and capacity forecasts) and opportunities for improvement (such as those afforded by technology developments).

Table of Contents

Executive Summary	8
1 Introduction	11
1.1 Overview of GÉANT	11
1.2 Opportunities for Further Improvement	12
1.2.1 Technology Developments	12
1.3 NREN Input to Architecture Planning	13
1.4 Approach to Architecture Planning	15
2 Current GÉANT Architecture	16
2.1 Design	16
2.2 The Multi-Domain and Global Context	19
3 GÉANT Services and their Development	21
3.1 Existing Services	21
3.1.1 GÉANT IP	21
3.1.2 GÉANT Plus	22
3.1.3 GÉANT Lambda	22
3.2 Proposed Developments to Existing Services	23
3.2.1 Introduction	23
3.3 GÉANT Photonic Services	26
3.3.1 Outline of Photonic Services	26
3.3.2 Expected Applications	27
3.3.3 Implications for GÉANT and NREN Networks	27
3.3.4 Conclusion	28
3.4 Quality	28
3.4.1 Introduction	28
3.4.2 Quality Parameters	29
4 Capacity Forecasts	33
4.1.1 Organic Growth	33
4.1.2 Discontinuous Growth	34
4.1.3 Conclusions on Capacity Forecasts	36
5 Building Blocks	38
5.1 Introduction	38

5.2	Building Block A – Dark Fibre DWDM Infrastructure	38
5.3	Building Block B – Managed Wavelength Infrastructure	39
5.4	Building Block C – Carrier Class Transport Infrastructure	41
5.4.1	Logical Layout	42
5.5	Building Block D – Managed Carrier Class Transport Infrastructure	43
6	Technologies	45
6.1	Point-to-Point Dense Wavelength-Division Multiplexing (P2P DWDM)	45
6.1.1	Technology Overview	45
6.1.2	GÉANT Use Case	46
6.1.3	Market Analysis	48
6.2	Reconfigurable Optical Add-Drop Multiplexers (ROADMs) and Wavelength Selective Switches (WSSs)	48
6.2.1	Technology Overview	49
6.2.2	GÉANT Use Case	51
6.2.3	Market Analysis	52
6.2.4	Summary	53
6.3	Next-Generation Optical Transport Network (NG-OTN))	54
6.3.1	Introduction	54
6.3.2	Optical Transport Network (OTN)	54
6.3.3	Implications for GÉANT	56
6.3.4	Market Analysis	56
6.4	Ethernet over Multi-Protocol Label Switching (EoMPLS)	57
6.4.1	Technology Overview	57
6.4.2	GÉANT Use Case	58
6.4.3	Support for Network Services at Other Layers	60
6.4.4	Market Analysis	62
6.4.5	Summary	64
6.5	Carrier Ethernet	65
6.5.1	Introduction	65
6.5.2	Technology Overview	66
6.5.3	GÉANT Use Case	69
6.5.4	Market Analysis	70
6.5.5	Summary	72
7	GÉANT Backbone Infrastructure	73
7.1	Introduction	73
7.2	The Need to Understand Physical Infrastructure Details	73
7.3	High-Level Resilience Study	74

7.3.1	Diversity of Trunks into Geneva	75
7.3.2	Trunks into Frankfurt	78
7.3.3	Trunks into Budapest	80
7.4	Topology Enhancements and Rationalisation	82
7.4.1	Extend the Fibre Footprint	82
7.4.2	Rationalisation	83
7.4.3	Add More Meshing to the GÉANT Fibre Backbone	83
7.5	Additional NREN Access Points	84
7.5.1	Hamburg	85
7.5.2	Marseille	85
7.6	Adding Fibre Junction Flexibility Points	85
8	Architecture Options	88
8.1	Introduction	88
8.2	Fixed (P2P) transmission + carrier Ethernet + IP/MPLS	91
8.2.1	Service Delivery	91
8.2.2	Pros and Cons for GÉANT Plus Implementation	91
8.2.3	Pros and Cons for GÉANT Lambda Implementation	91
8.2.4	Overall Evaluation	92
8.3	Fixed (P2P) transmission + OTN + IP/MPLS	93
8.3.1	Service Delivery	93
8.3.2	Pros and Cons for GÉANT Plus Implementation	93
8.3.3	Pros and Cons for GÉANT Lambda Implementation	93
8.3.4	Overall Evaluation	94
8.3.5	Special Case: ODU Switching	94
8.4	Fixed (P2P) transmission + IP/MPLS	95
8.4.1	Service Delivery	95
8.4.2	Pros and Cons for GÉANT Plus Implementation	95
8.4.3	Pros and Cons for GÉANT Lambda Implementation	96
8.4.4	Overall Evaluation	96
8.5	Agile transmission + carrier Ethernet + IP/MPLS	97
8.5.1	Service Delivery	97
8.5.2	Pros and cons for GÉANT Plus Implementation	97
8.5.3	Pros and Cons for GÉANT Lambda Implementation	98
8.5.4	Overall Evaluation	98
8.6	Agile transmission + OTN + IP/MPLS	99
8.6.1	Service Delivery	99
8.6.2	Pros and Cons for GÉANT Plus Implementation	99
8.6.3	Pros and Cons for GÉANT Lambda Implementation	100

8.6.4	Overall Evaluation	100
8.6.5	Special Case: ODU Switching	100
8.7	Agile transmission + IP/MPLS	101
8.7.1	Service Delivery	101
8.7.2	Pros and Cons for GÉANT Plus Implementation	101
8.7.3	Pros and Cons for GÉANT Lambda Implementation	102
8.7.4	Overall Evaluation	102
9	Summary and Next Steps	104
9.1	An Architectural Model for the Next Generation	104
9.2	Technical Considerations and Recommendations	105
9.2.1	Optical Layer	106
9.2.2	Physical Topology	107
9.2.3	Switching Layer	108
9.2.4	IP Layer	108
9.2.5	Monitoring and Authentication, Authorisation and Accounting	109
9.3	Final Considerations and Next Steps	109
Appendix A	User Requirements	111
Appendix B	Gap Analysis	117
Appendix C	Third-Party Connectivity	130
References		132
Glossary		135

Table of Figures

Figure ES.1: Basic layers of the new GÉANT architecture	9
Figure 2.1: Current GÉANT backbone topology (“tube map” view)	17
Figure 2.2: Architectural view of the GÉANT backbone today	19
Figure 4.1: Cumulative total BW-distance product (excluding capacity for IP)	35
Figure 4.2: London–Amsterdam IP trunk during recent LHC-related data transfer activities	36
Figure 5.1: DWDM network	39
Figure 5.2: Circuit components used to deliver the wavelength service	40
Figure 5.3: Managed wavelength infrastructure	41
Figure 5.4: Carrier class transport network technologies	42
Figure 5.5: Transport network	42
Figure 5.6: Circuit component parts used to deliver transport service	43
Figure 5.7: Managed transport network	44
Figure 6.1: Wavelength Selectable Switch	49
Figure 6.2: A 2-degree ROADM	50
Figure 6.3: A “directionless” 2-degree ROADM	50
Figure 6.4: Structure of DWDM/transport architecture option	55
Figure 6.5: GÉANT network layers	61
Figure 7.1: Detail of GÉANT backbone fibre routing	75
Figure 7.2: Logical diagram of commonly routed section Geneva–Basel	76
Figure 7.3: The addition of a flexibility point in Marseille and a Marseille–Milan fibre route	78
Figure 7.4: Situation around Budapest	80
Figure 7.5: Potential additions to the GÉANT fibre mesh	84
Figure 7.6: Current configuration of parallel fibre lines	86
Figure 7.7: Placing a 3-degree ROADM at a fibre junction	86
Figure 7.8: Placing a 4-degree ROADM at a fibre junction	87
Figure 8.1: GÉANT network layered architecture	88
Figure 9.1: Representation of the next-generation infrastructure architectural model	105
Figure C.2: Capacity acquisition process overview	131

Table of Tables

Table 4.1: Trunk wavelength capacity (in Gbps) required to support forecast growth in backbone IP traffic (revisited 2009 and 2010)	37
Table 5.1: Circuit components used to deliver the wavelength service	40
Table 5.2: Circuit components used to deliver the transport service	43
Table 8.1: Pros and cons for GÉANT Plus implementation	91
Table 8.2: Pros and cons for GÉANT Lambda implementation	92
Table 8.3: Pros and cons for GÉANT Plus implementation	93
Table 8.4: Pros and cons for ODU switching	95
Table 8.5: Pros and cons for GÉANT Plus implementation	96
Table 8.6: Pros and cons for GÉANT Plus implementation	97
Table 8.7: Pros and cons for GÉANT Lambda implementation	98
Table 8.8: Pros and cons for GÉANT Plus implementation	99
Table 8.9: Pros and cons for ODU switching	101
Table 8.10: Pros and cons for GÉANT Plus implementation	102
Table A.1: Summary of user requirements	116

Executive Summary

This deliverable is the final report on the findings of the GÉANT backbone architecture study. An interim report was issued in April 2010 [DS1.1.1]. It presents architecture options and recommendations for the future network taking into account current and future requirements (such as current architecture, services, quality, user requirements and capacity forecasts) and opportunities for improvement (such as those afforded by technology developments). It also describes the architectural building blocks that can be used in different parts of the network, and describes the key technologies involved, focusing on their practical implementation, as preparation for the procurement process, and on how they can apply to the GÉANT network and service portfolio.

The approach to GÉANT architecture planning takes into account the current architecture (see Section 2 on page 16), both its design and the multi-domain and global context in which it operates, and the following aspects:

1. The contents of the GN3 white paper [GN3 white paper], which summarises the project's vision, strategic objectives and guiding principles, and outlines the rationale for the GN3 structure.
2. The services offered to and required by the GÉANT user base, how they are expected to develop, and what quality levels are associated with them, taking into account the multi-domain nature of end-to-end service provision over the extended GÉANT service area (see Section 3 on page 21).
3. An analysis of capacity demand evolution, taking into account historic growth and predictions of user demand (see Section 4 on page 33 and Appendix A *User Requirements* on page 111).
4. An analysis of the technological options that exist to fulfil those services, complemented by an analysis of the availability and maturity of technology in the market (see Section 6 on page 45).
5. An analysis of the underlying fibre infrastructure and topology to ensure optimal network resilience and performance at all levels (see Section 7 on page 73).

Three topological shortcomings have been identified: diversity of trunks into Geneva, trunks into Frankfurt, and trunks into Budapest. Three areas of enhancement have been identified: extension of the GÉANT fibre footprint, its rationalisation (where there are suitable options and circumstances that can allow this), and the addition of more meshing. The topology analysis also considered additional, building-diverse National Research and Education Network (NREN) access points, for example in Hamburg and Marseille, and adding fibre junction flexibility points.

6. Study of availability of infrastructure to augment the GÉANT dark fibre footprint (see Section 7.4.1 on page 82).

With respect to defining possible generic approaches to the GÉANT architecture (see Section 8 on page 88), three architectural components and possible implementation alternatives have been identified (these draw on the different types of network architecture building blocks described in Section 5 on page 38):

- Internet Protocol (IP) component.
- Switching component.
- Optical transmission component.

Each implementation alternative has been evaluated, for each service and overall, using a number of criteria, including reliability, user-network separation, maturity of technologies, and multi-domain deployment.

In summary (see Section 9 on page 104), the study has confirmed a pattern of constant growth in the amount of IP traffic over the GÉANT network and in the number of high-capacity circuits dedicated to projects, and a requirement for more advanced services and functionalities in the areas of authorisation and authentication, security, monitoring, and dynamic provisioning to meet user needs. The technology has evolved since the implementation of the GÉANT2 network at the end of 2005, offering new optical equipment capabilities and switching platforms, and marking the decline of Synchronous Digital Hierarchy / Synchronous Optical Networking (SDH/SONET) and a ubiquitous acceptance of the Ethernet protocol. The increasing importance of data transmission for the research and education community is placing a greater importance on the resiliency and redundancy of the services. The requirement impacts the whole infrastructure, from ensuring diverse physical routes to diverse fibres to the logical topology of the IP network.

The study has also confirmed that the hybrid infrastructure at the core of the GÉANT network represents a valid building block and provides the correct foundations for the next-generation infrastructure. This will be based on the fibre available to GÉANT, which has an enhanced role as a fundamental asset, and add the most appropriate switching layer at the packet and frame level on top of it. Figure ES.1 below shows a high-level representation of the basic layers of the new architecture; the common functions of monitoring and authentication and authorisation are part of each layer and are depicted vertically for clarity and to show the required integration. Each layer has its own control and management planes (not shown); their integration between layers is subject to technological choices and ongoing research and development.

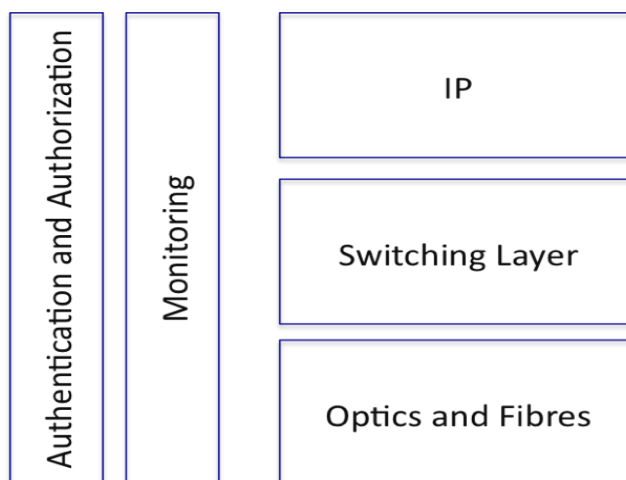


Figure ES.1: Basic layers of the new GÉANT architecture

A number of important issues are common to all layers and require, in addition to analysis at the level of each individual layer, a solution that takes into account the interaction of all the layers. These are: resiliency and robustness to failures; fast recovery from failures; ease and speed of reconfiguration. In addition, the infrastructure should be transparent to the users and allow innovation. There are additional considerations and recommendations relating to: upgrading the current optical layer; enhancing the physical topology by increasing the meshing of the GÉANT fibre footprint, ensuring that the main connections run on physically diverse trunk paths, and having more than one Point of Presence (PoP) in selected countries to open up additional access points for NRENs; switching layer; IP layer; and monitoring, authentication, authorisation and accounting.

The next generation network will be strengthened, at all layers, in the areas of resilience; agility and timely configurability; capacity; and interoperability.

From the technical studies conducted so far, there are clear preferences for the future GÉANT network, though these will be subject to further analysis:

- Availability of an agile transmission platform based on Reconfigurable Optical Add-Drop Multiplexers (ROADMs), to facilitate the resilience improvements needed, ensure the more efficient use of the topology and infrastructure, and facilitate additional access points.
- Availability of a logically separate switching layer, based on Ethernet over Multi-Protocol Label Switching (EoMPLS), carrier Ethernet (cE) or Optical Transport Network (OTN).
- Given the developments possible at the transmission and switching layer, there is now also the opportunity to review and optimise the IP layer.

The next steps are to compare the technical information and plans with vendors' contractually available solutions and reliable cost data. Further planning is required to devise an appropriate schedule for the staged approach(es) necessary to arrive at recommendations for solutions that may be implemented. This will include an assessment of the needs for further Request for Proposal work and/or commencement of some initial tendering phases. During this process the current implementations of NRENs' and international peering networks will be carefully considered to ensure that the largest number of services (including monitoring), may be seamlessly implemented. In addition, the project will monitor closely the needs of users with the most significant data-traffic demand, such as those identified in Appendix A, to ensure that the new architecture is able to meet their requirements in terms of both capacity and service provision.

The project will be cautious with regard to the possible complexities arising from novel technologies and it will ensure that the technologies selected involve low capital and operational costs, while maintaining the broadest possible compatibility and inter-operability with peering networks at all layers. Consideration will also be given to openness and interoperability. The availability of a greater number of fibres and wavelengths (either directly or provided by partners of the consortium) will help to keep complexity low, provide simpler solutions to resiliency, and enrich the services' capabilities.

1 Introduction

1.1 Overview of GÉANT

GÉANT is the high-bandwidth pan-European research and education backbone that interconnects all National Research and Education Networks (NRENs) across Europe and provides worldwide connectivity through direct links with other regional networks. It provides:

- Seamless connectivity through European NRENs to an estimated 40 million research and education users in over 8,000 campus networks in 40 countries across Europe. There are 33 directly connected NRENs.
- Interconnections with networks in other world regions, extending coverage to over 80 countries across the globe.
- Data transfer speeds of up to 40 Gbps, across 50,000 km of European network infrastructure, of which 12,000 km is based on lit fibre.
- Flexible, innovative architecture for data communications across the standard Internet Protocol (IP) backbone, and for high-capacity “private” network paths reserved for specific projects or disciplines.

GÉANT’s mission is to provide state-of-the-art services and technologies to the Research and Education (R&E) community in Europe, over and above what can be achieved from the normal commercial market. Whilst commercial 10 Gbps services are available, these are usually offered only where there are significant profits to be made, whereas for many years GÉANT has been able to deliver 10 Gbps throughout Europe, to support R&E. The setup of GÉANT also allows flexible and efficient provision of services to a wide range of user communities. The services available via the GÉANT network are as follows:

- IP (Layer 3 (L3)), with capacities up to 40 Gbps on the GÉANT backbone. 100 Gbps capacities are envisaged by 2012.
- Point-to-point transparent wavelengths, currently at 10 Gbps (10 GE or OC-192/STM-64) with 40 Gbps and 100 Gbps possible in the future. These services are commonly known as wavelengths or lambda or Layer 1 (L1) services. The official name of the GÉANT service is **GÉANT Lambda**.
- Point-to-point sub-lambda circuits. The correct technical term for these is Ethernet Private Lines (EPLs), but they are also commonly known as lightpaths, as well as sub-lambda services and L2 services. These services have 1 Gbps granularity, with GE presentation. The official name of the GÉANT service is **GÉANT Plus**.
- Extension of the three above services beyond the GÉANT European network to a global R&E context.

1.2 Opportunities for Further Improvement

The current GÉANT network was built in a cost-effective and flexible manner and has served its users for the last five years. The deployment and operation of a dark fibre infrastructure, with its own dedicated Dense Wavelength-Division Multiplexing (DWDM) transmission systems, has enabled GÉANT to develop and deliver connectivity services for a wide range of user communities connected to the NRENs.

Looking ahead to the next five years, GÉANT has the challenge and opportunity to further improve the effectiveness of its service to the user community in many ways, including:

- Increase the robustness / resilience of the underlying infrastructure. When the fibre for GÉANT was first procured in 2004/2005, there existed a limited number of options for fibre connectivity between GÉANT Points of Presence (PoPs). This has resulted in some of the fibre routes overlapping in some parts, giving rise to the possibility that a fibre cable cut will cause more than one fibre route failure. This requires a review of parts of the topology of GÉANT.
- Provide more flexibility in connecting users. At present NRENs can access the GÉANT backbone in a local PoP at the transmission layer. Only one local PoP per country is currently foreseen. NREN users utilise GÉANT via the NREN's only interconnection point. The future GÉANT backbone should facilitate alternative access points for NRENs, in order to make more efficient use of the backbone infrastructure. This may lead to additional access points in a country or, at the other extreme, shared access points between countries.
- Facilitate use of connectivity resources provided by NRENs. This is commonly known within GN3 as cross-border fibre (CBF) (although this term is not technically accurate; the correct term would be managed third-party connectivity). Administrative and operational procedures are required to facilitate the use of these resources, which could contribute to reducing the overall cost of the GÉANT network.
- Support increasing demand in terms of bandwidth requirements. From traffic forecasts and what is currently understood of demands from scientific user communities, 100 Gbps aggregate capacities will be required in parts of the GÉANT backbone. It is not clear how the bandwidth needs to be delivered to users, whether via IP or via the provision of circuit services. This implies that the architecture of GÉANT needs to be flexible, to accommodate both types of requests efficiently.
- Ensure that the technology deployed on GÉANT is future-proof and state of the art, in order to deal with emerging requirements flexibly and efficiently. This challenge is made more difficult by the challenge to contain costs.
- Explore the support of novel "photonic services" (a new type of optical service made possible in GÉANT by the availability of dark fibre and the deployment of state-of-the-art and future-proof transmission technologies, as described in Section 3.3 *GÉANT Photonic Services* on page 26).

1.2.1 Technology Developments

There is also an important practical reason for reviewing the GÉANT architecture, which is related to the technology deployed. Whilst the IP routing equipment for GÉANT was re-procured two years ago and is future-proof in terms of its ability to support 100 Gbps in a scalable manner, this is not the case for the DWDM (ALU 1626 LM) in its current configuration, or for the optoelectrical switching equipment (ALU 1678 MCC).

The DWDM equipment (ALU 1626) is in theory capable of supporting 100 Gbps transmission provided significant enhancements are implemented. However, it is not clear whether upgrading this technology to support 100 Gbps will result in the most cost-effective way of delivering such capacities, as 100 Gbps is best supported on new types of modulation schemes that would require significant re-engineering of GÉANT's optical links if it is to be cost-effectively deployed. A review of DWDM technology for GÉANT is therefore required, and options will need to be compared with upgrading the currently deployed system.

With regard to optoelectrical switching, the current platform is based on what was termed Next-Generation Synchronous Digital Hierarchy (NG-SDH) back in 2005, the 1678 MCC. ALU has not further developed this platform, and it will not support 40 Gbps let alone 100 Gbps. It is used to support the GÉANT Plus service, and whilst it has worked reasonably well, it remains a complex technology to provision and debug and is ill-suited to interface with and support bursty packet-based traffic flows. There are advances in Ethernet technologies (e.g. carrier Ethernet (cE) and Ethernet over Multi-Protocol Label Switching (EoMPLS)) that appear better suited to support the needs of GÉANT and which will scale to 100 Gbps capacities. For these reasons, a review of the optoelectrical switching technology deployed on GÉANT is required.

The IP technology, on the other hand, as mentioned above (see Section 1.1 on page 11), is state of the art and capable of supporting 100 Gbps capacities. Reviewing the photonic and switching layer provides an opportunity to optimise the configuration of the IP layer of GÉANT to ensure the most efficient use of the underlying technologies and enhance the effectiveness of the IP layer.

1.3 NREN Input to Architecture Planning

Since December 2008 NRENs have been involved in planning the GÉANT architecture via a series of workshops documented in Deliverable DS1.1.1 "Report on the Backbone Architecture Study" [DS1.1.1].

Figure 1.1 below shows the key principles that drive GÉANT architecture planning (use NREN resources, reduce costs, innovate, account for user needs) and how these principles map to the elements of architecture planning work (technology, topology, network operations and services).

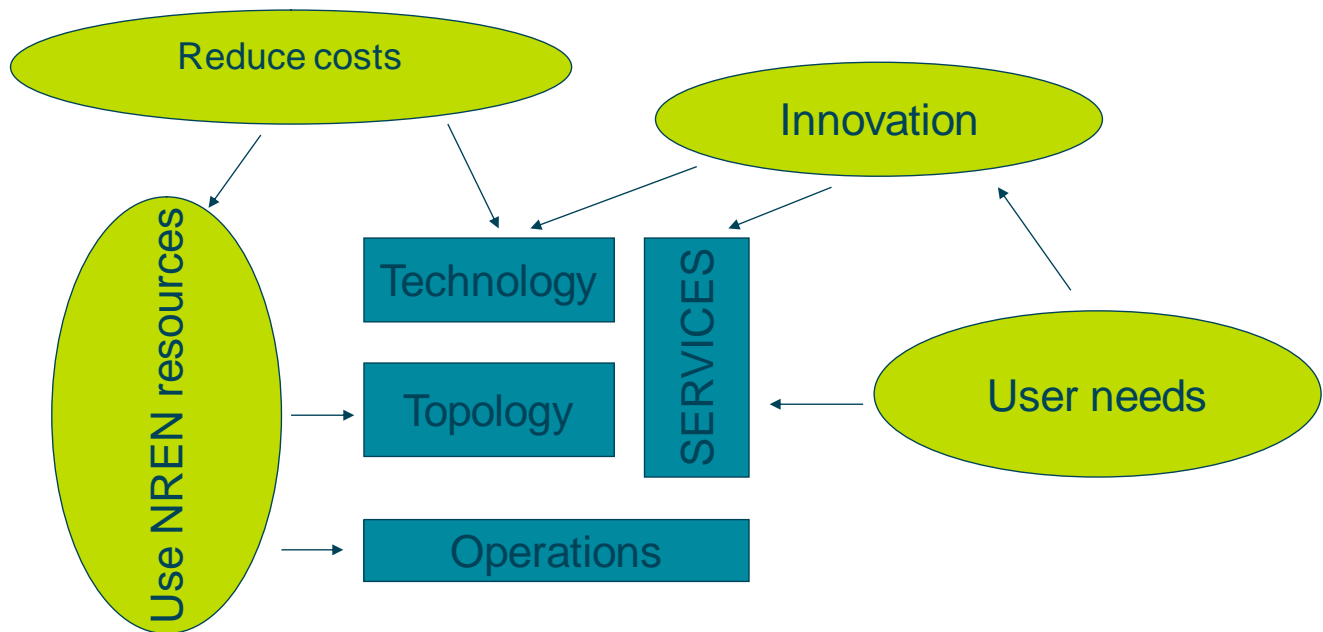


Figure 1.1: GÉANT architecture planning principles mapped to work elements

In order to ensure that the workplan for GÉANT architecture planning considered in an appropriate measure the different elements involved, members of the SA1 Supervisory Committee (SC) (appointed by the GN3 management team to oversee this work and provide strategic direction and advice) each presented their detailed view on what aspects the SA1 team needed to work on towards developing the next GÉANT architecture. This has resulted in a gap analysis (see Appendix B *Gap Analysis* on page 117), which explains how each aspect put forward by an SC member needs to be addressed by the SA1 team. The gap analysis covers the following topics: user requirements, topology issues, services, technology, virtualisation, service quality, use of NREN resources, federation, overlays, peerings, and L2 interconnections.

Considering their relative importance and whether a given aspect was adequately covered by GÉANT already, the highest priority has been placed on:

- Developing a framework for the adoption of NREN connectivity resources (CBF).
- Defining the service quality parameters for GÉANT services and how they affect technology choices and network cost.
- Examining in detail the available technology options that can fulfil GÉANT services and meet the service quality levels defined, and their cost implications.
- Examining topology enhancements to resolve resilience issues and facilitate additional or alternative access points to GÉANT.

There are some notable aspects that will either be reviewed following developments in the GN3 research activities (JRA1) or are already addressed by the GÉANT architecture planning. These include:

- Virtualisation. The important aspect in this case is that any technology choice for GÉANT should not preclude virtualisation. Meanwhile JRA1 needs to continue its work as planned, to define what services should be delivered by means of virtualisation and to outline benefits to users.

- Federation of GÉANT. There are three aspects related to this topic. The first two are about developing a framework for the adoption of NREN connectivity resources and enabling additional access points. These are addressed as high priority. The third is related to infrastructure aspects and their ownership and operation, whereby parts of the GÉANT backbone would be owned and operated by an NREN. These aspects and their benefits and consequences are part of the work plan of JRA1 T3 and will be undertaken there until practical experience over a testbed has been achieved and results that are useful, beneficial and implementable have been identified.
- L2 interconnections, “open exchanges”. The question of whether GÉANT interconnect points should be “open” is a policy matter rather than a technical matter. For the provision of point-to-point (Ethernet) circuits, the GÉANT nodes are already L2 interconnects. From a technical perspective, therefore, no further work is needed apart from normal continuous evolution. Rather, “openness” is an Acceptable Use Policy (AUP) issue and so will need to be discussed at a policy level.

1.4 Approach to Architecture Planning

As mentioned in Section 1.3 above, the approach to GÉANT architecture planning takes into account the following aspects:

1. The contents of the GN3 white paper [GN3 white paper], which summarises the project’s vision, strategic objectives and guiding principles, and outlines the rationale for the GN3 structure.
2. The services offered to and required by the GÉANT user base, how they are expected to develop, and what quality levels are associated with them.
3. An analysis of the technological options that exist to fulfil those services, complemented by an analysis of the availability and maturity of technology in the market.
4. An analysis of capacity demand evolution, taking into account historic growth and predictions of user demand.
5. An analysis of the underlying fibre infrastructure and topology to ensure optimal network reliability and performance at all levels.
6. Study of availability of infrastructure to augment the GÉANT dark fibre footprint.

Each of these aspects is considered in detail in the deliverable.

2 Current GÉANT Architecture

2.1 Design

Today the GÉANT backbone consists of connectivity based on a mixture of leased dark fibre links (lit using a carrier-class Dense Wavelength-Division Multiplexed (DWDM) system owned and operated by DANTE) and managed circuit services provided by commercial telecommunications network operators. The latter consist of wavelength services (currently up to 10 Gbps in capacity) and circuits based on Synchronous Digital Hierarchy (SDH) for smaller capacities, and are used in regions of the GÉANT footprint where it has so far not been possible to invest in dark fibre.

The overall view of the topology of the GÉANT network can be seen in Figure 2.1 below.

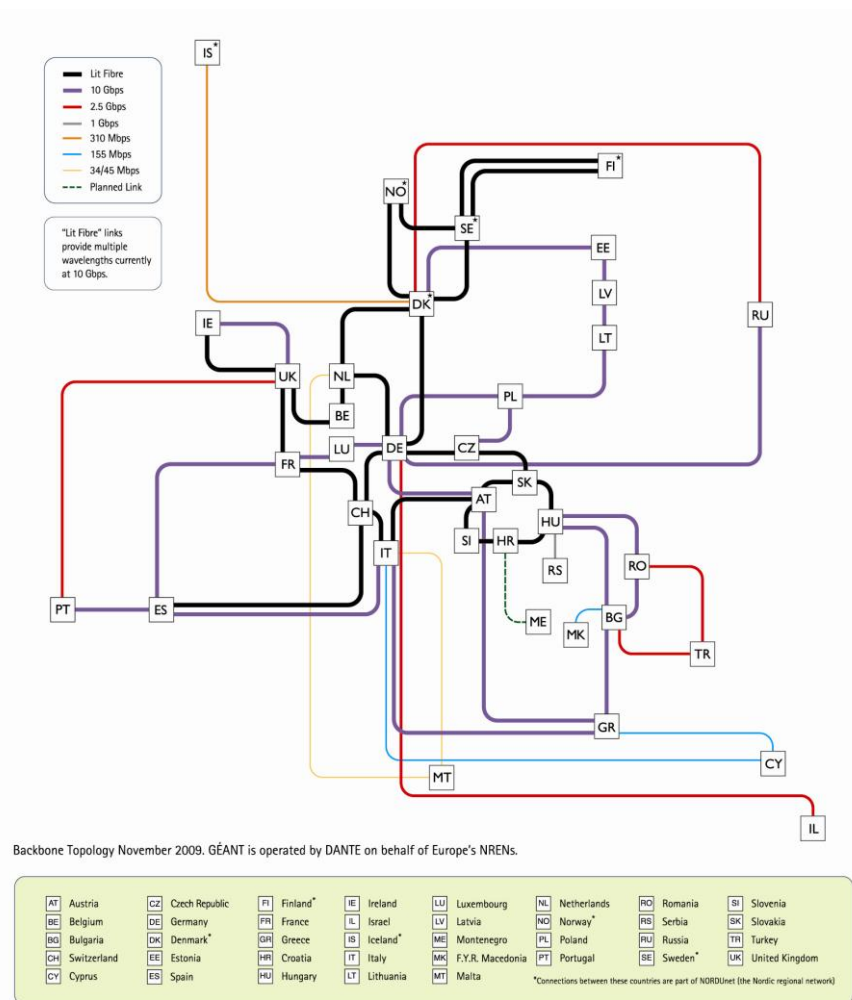


Figure 2.1: Current GÉANT backbone topology (“tube map” view)

This is, of course, a stylised (sometimes referred to as the “tube map”) view of the network and it hides a lot of physical detail. In-depth analysis of this physical detail is provided in Section 7 *GÉANT Backbone Infrastructure* on page 73.

The GÉANT Points of Presence (PoPs) on this network are variously populated with transmission, switching and routing equipment (not all nodes have all three kinds of equipment) that together provide three service-provisioning platforms. These, in turn, are used to provide the NRENs with the GÉANT backbone services (GÉANT IP, GÉANT Plus and GÉANT Lambda) which are described from a service perspective in the following section, *GÉANT Services and their Development* on page 21.

The GÉANT Lambda services (currently only at 10 Gbps and structured as either STM-64 or 10 gigabit Ethernet Local Area Network Physical Ethernet (LAN PHY)) are provided directly over the DWDM transmission platform. Hence, these services are only available to NRENs that are “on net” with respect to the fibre cloud. GÉANT Lambda service instances are usually extended across NREN infrastructures to customer sites within the relevant territories.

In addition, the transmission platform is used to provide trunk connectivity between GÉANT backbone PoPs to support two higher layer service delivery platforms.

One such platform supports the GÉANT Plus service and is based on carrier-class “next-generation” (NG) SDH switches. GÉANT Plus service instances are primarily point-to-point Ethernet private lines and Ethernet virtual private lines also known as E-Line services as defined by the Metro Ethernet Forum (MEF). As with the GÉANT Lambda service instances, GÉANT Plus service instances are usually extended across NREN infrastructures to customer sites within the relevant territories. They are used where sub-wavelength bandwidths (currently less than 10 Gbps) are required with a lead time of the order of a few working days. The mapping of the client signals to SDH Time Division Multiplexed (TDM) trails is performed using the “next-generation SDH” concepts of using Generic Framing Protocol – Framed (GFP-F) encapsulation of Ethernet into virtually concatenated groups of VC-4 trails. There is no facility for packet-oriented statistical multiplexing and limited tolerance to bursty traffic profiles.

The other platform is based on carrier-class core Internet Protocol / Multi-Protocol Label Switching (IP/MPLS) routers. This is used to provide the GÉANT IP service. This allows high-quality transit for IP traffic from European NRENs to one another and between European NRENs and associated networks globally. In addition to “best-effort” IPv4 transit, it provides native IPv6 and multicast support. Currently, GÉANT IP access is available to NRENs at capacities of up to 20 Gbps (implemented as two 10 gigabit Ethernet interfaces).

Currently the GÉANT backbone transmission architecture can best be described as a (very-)long haul, fixed DWDM system, predominantly point-to-point but including two fixed band Optical Add/Drop Multiplexers (OADMs) (one in Brussels and the other in Barcelona). Each PoP–PoP link is engineered to support up to forty optical channels at 10 Gbps, each operating with a client Bit Error Rate (BER) of 10⁻¹² or better. There is no remotely configurable space- or wavelength-selective or sub-wavelength switching function in the current platform.

The layering and heterogeneous nature of the current backbone architecture is shown in Figure 2.2. This illustrates that there are four distinct sets of PoPs:

- A “core” of PoPs that are on the GÉANT fibre cloud and are on net with respect to the NG-SDH switching and IP/MPLS platforms.
- Two PoPs (Athens and Poznan) that are off the fibre net but still have direct access to the GÉANT Plus and GÉANT IP platforms. Trunk connectivity to these PoPs relies on managed circuit (leased wavelength) services.
- A set of PoPs (those in the Baltic states, Bucharest and Sofia) that are likewise trunk-connected by leased wavelength services but have no access to the GÉANT Plus platform.
- Three sets of PoPs that are “routerless” (have no local access to the GÉANT IP/MPLS platform).
 - One set comprises those PoPs that are still on net with respect to all the other layers (including fibre), namely Dublin, Brussels, Bratislava, Zagreb and Ljubljana.
 - The second set comprises those routerless PoPs that are off the fibre net but still have direct access to the GÉANT Plus platform and hence rely on managed wavelength services for their trunk connectivity; these are Moscow, Lisbon and Luxembourg.
 - Finally there are PoPs that are not strictly speaking GÉANT PoPs at all but are NREN PoPs connected (using managed circuit services) to remote GÉANT PoPs. These are the PoPs that are part of the networks of ULAKBIM (Turkey), IUCC (Israel), the University of Malta (UoM), CYNET

(Cyprus), MARNET (the Republic of Macedonia/FYROM), MREN (Montenegro) and AMRES (Republic of Serbia).

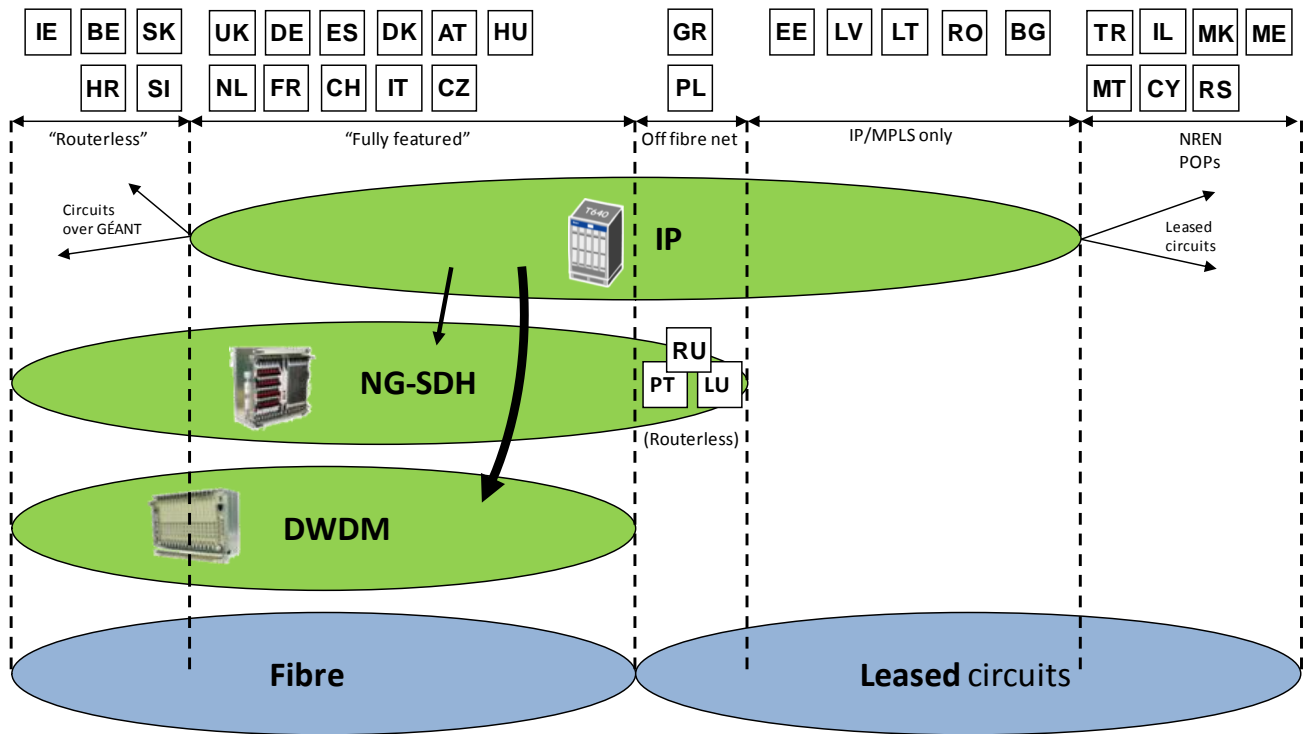


Figure 2.2: Architectural view of the GÉANT backbone today

2.2 The Multi-Domain and Global Context

GÉANT operates, of course, in a multi-domain context, supporting the interconnection between NRENs and R&E networks in other regions that are connected to GÉANT.

Figure 2.3 below shows how, in general, NRENs connect to GÉANT and GÉANT connects to other world region networks, and highlights the layered nature of the GÉANT model and of the interconnections.

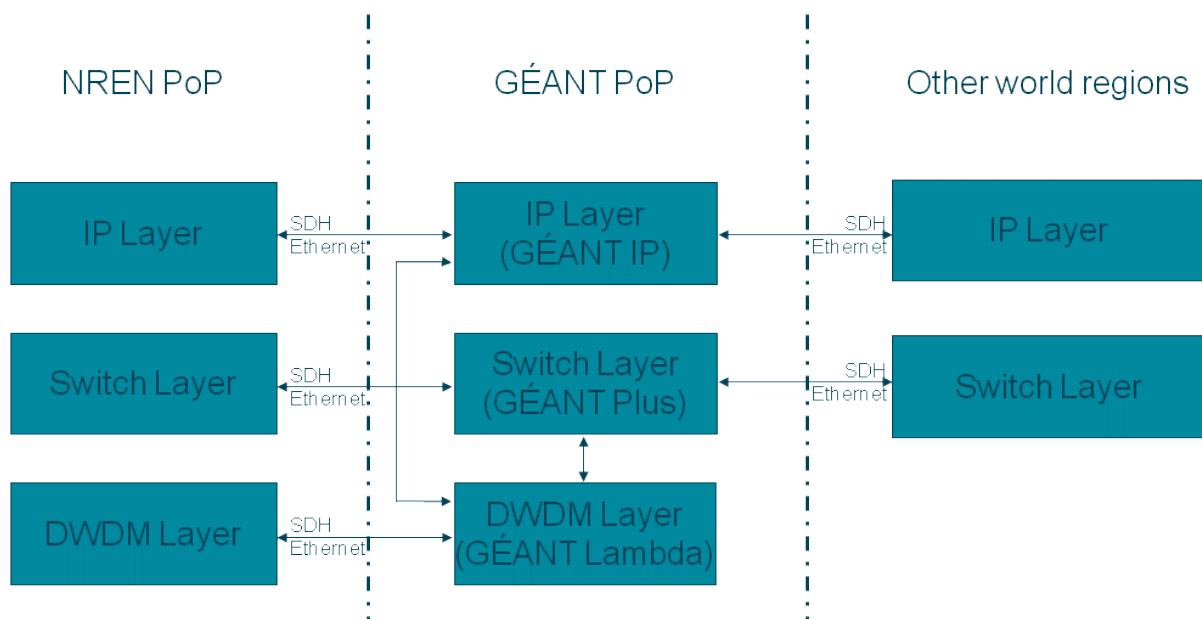


Figure 2.3: Layered GÉANT model and connections

The figure shows logical connections. For example, the connectivity between the IP layers may be via a small fibre patch in the case of co-located GÉANT and NREN PoPs, or it could be via the respective DWDM systems if the two PoPs are not co-located (this applies also to the “routerless” PoPs). The interesting fact to note is that although the figure shows that at the switching layer (for GÉANT Plus services) NRENs connect to GÉANT using either Ethernet or SDH interfaces, in reality all but one NREN connects via Ethernet. In contrast, at present it is only possible to use SDH (or SONET) as the underlying transport to connect to other world region networks at the switching and IP layers, although the interfacing may also be Ethernet. Currently there are no connections to other world region networks at the DWDM layer.

GÉANT’s switching layer supports the provision of **global lightpaths**. These are essentially point-to-point circuit services that extend into other world regions such as North America and Asia. These extensions are effected by interconnecting GÉANT Plus services with equivalent circuit services provided by other world regions’ R&E networks. At present these interconnections are provided either by direct interconnections between GÉANT and interconnection points with other world region R&E networks, or indirectly, utilising connections via a European NREN through the NREN’s GLIF Open Lightpath Exchange (GOLE), which then carries the traffic between GÉANT and another world region network. (A GOLE is essentially a switch that interconnects with GÉANT’s switch layer and supports provision of point-to-point circuit services that interoperate with GÉANT Plus services).

Figure 2.3 shows that there are three separate platforms and technology layers to support the three base GÉANT services (GÉANT IP, GÉANT Plus and GÉANT Lambda). For the NRENs and other world region networks the situation differs from NREN to NREN: some NRENs use converged routing and switching platforms, or converged transmission and switching platforms. Some NRENs use analogue transmission systems to interconnect to neighbouring NRENs or SONET connections to interconnect to the transmission layer of other world regions.

3 GÉANT Services and their Development

3.1 Existing Services

Three types of service are available via the GÉANT network:

- GÉANT IP.
- GÉANT Plus.
- GÉANT Lambda.

For all three types, but for the point-to-point services GÉANT Plus and GÉANT Lambda in particular, the multi-domain nature of end-to-end service provision over the extended GÉANT service area is a key consideration: for optimum performance, quality must be maintained from end point to end point, across each domain and including the last mile.

Each service type is described below.

3.1.1 GÉANT IP

The GÉANT IP service offers NRENs access to the shared European IP backbone. This allows transit for IP traffic from European NRENs to one another, and between European NRENs and associated networks globally. The IP service is designed to provide a robust high-bandwidth solution to the international connectivity requirements of the majority of academic users. It provides resilient service in the case of hardware failure or fibre cuts and uses advanced routing equipment to ensure fast recovery from unexpected events. GÉANT IP access is available to NRENs at capacities of up to 40 Gbps. 100 Gbps capacities are envisaged by 2012.

GÉANT IP is also referred to as a Layer 3 (L3) service, and includes:

- IPv4 (both unicast and multicast).
- IPv6 (both unicast and multicast).
- Layer 2 Virtual Private Networks (L2 VPNs).
 - Although the service provided is a Layer-2 service, it is included here as it is implemented over a platform that is predominantly used to deliver Layer-3 services.

- Backup:
 - A GÉANT IP service subscription includes a backup port on the local router which protects the NREN's access to GÉANT in the event of card failure or problems with the NREN access equipment or fibre.
 - To secure greater resiliency, NRENs may request a backup interface on a router in an adjacent GÉANT PoP. This backup access may then be back-hauled to their local GÉANT PoP via a GÉANT Plus circuit, GÉANT Lambda (where available), transmission over a CBF link or other means.

3.1.2 GÉANT Plus

The GÉANT Plus service allows user access to point-to-point circuits of between 155 Mbps and 10 Gbps across an existing pre-provisioned network.

A GÉANT Plus circuit provides guaranteed bandwidth and deterministic performance. GÉANT Plus is built on common infrastructure, but appears to its private users to be dedicated to that user's needs.

GÉANT Plus circuits are also referred to as sub-lambda circuits, Ethernet Private Lines (EPLs), lightpaths, sub-lambda services and Layer 2 (L2) services.

Features

- GÉANT Plus provides dedicated sub-wavelength point-to-point circuits configured over a network of 10 Gbps trunks and TDM (Time-Division Multiplexed) switches.
- The allocated capacity of each NREN may be used flexibly for different services to multiple locations.
- A circuit can be configured or reconfigured on the GÉANT Plus interface (where GÉANT links connect to the local NREN) at short notice.
- GÉANT Plus circuits may be requested to many European NRENs and to some non-GÉANT locations (including transatlantic destinations connected by Internet2, ESnet, CANARIE and USLHCnet) and via open lightpath exchanges (such as GOLEs).

Access to GÉANT Plus is only available to NRENs where the GÉANT backbone infrastructure supports multiple wavelengths; in most cases this is over a dark-fibre connection. Circuits can be configured at short notice, typically within five working days of receipt of request, assuming that sufficient capacity is available in both the NRENs' capacity allocations.

It is also possible for an NREN to connect to a non-GÉANT organisation/destination.

Additional GÉANT Plus subscriptions will provide a further 10 Gbps of capacity on a new interface.

3.1.3 GÉANT Lambda

The GÉANT Lambda service provides private, transparent 10 Gbps and 40 Gbps wavelengths between any two GÉANT NRENs connected to the GÉANT dark-fibre cloud and can, in principle, be extended to another

world region via an open lightpath exchange and dedicated transparent wavelength services interconnecting said open lightpath exchanges.

The GÉANT Lambda service is also referred to as wavelengths, lambda and Layer 1 (L1) services.

Features

- GÉANT Lambda provides a 10 Gbps and 40 Gbps wavelength connection between two GÉANT NRENs connected over the GÉANT dark-fibre network.
- A GÉANT Lambda is presented to the NREN as a “transparent” wavelength (with framing – currently 10 gigabit Ethernet or STM-64/256) on which they can then develop their own higher-level network layers.
- When building a network of GÉANT Lambdas, a diversely routed backup Lambda can be added to provide resiliency in the case of a fibre cut.

The implementation of dedicated GÉANT Lambdas requires additional transmission equipment to be installed at the GÉANT NRENs; this can take up to ten weeks.

A full 10 Gbps backup Lambda can be provided on an alternative, resilient route to provide protection against fibre cuts or equipment failure. This secondary Lambda will be configured over a different physical path and use different equipment to the primary Lambda.

3.2 Proposed Developments to Existing Services

3.2.1 Introduction

This section summarises the proposed enhancements to the existing service portfolio described in the previous section, which any new network architecture must be able to support. The enhancements were defined during the GÉANT Architecture Workshops, and are documented more fully in “GÉANT Backbone Service Proposals” [GÉANT_Svc_Proposals]).

It is planned that the existing service portfolio will continue and that the enhancements listed below will comprise either additional features to existing services or new services. This section does not aim to provide comprehensive service definitions – most of the proposals will, anyway, require further study into their technical, operational and economic viability. Rather, it aims to set out the basic possible requirements arising from the enhancements that have implications for network architecture planning.

3.2.1.1 IP Services

Resilient IP Access

As a primary service, IP access should be provided in the most resilient, cost-effective way. Where feasible, a resilient GÉANT IP access is available today by back-hauling the service from a neighbouring GÉANT PoP. This avoids the single point of failure (SPoF) for the router equipment but not for a PoP-housing-based SPoF. There may be ways to improve the resilience by providing alternative direct access to another GÉANT node, either at an existing PoP or a new node within the network.

Implications: The new network architecture should provide fully diverse IP access.

Access to Internet Exchange Points

There has been significant interest from NRENs in accessing internet exchange points (IXPs). So far, four NRENs have arranged remote access to IXPs using the GÉANT Plus service. Whilst this approach has proved generally successful from the perspective of the NRENs concerned, it has been shown that accessing IXPs using GÉANT Plus capacity is technically, operationally and economically inefficient when applied to the whole NREN consortium. (See, for example, “Commodity IP and IXP usage in GN3”, a presentation given by M. Enrico and O. Kreiter to the Third Architecture Workshop 7–8 July 2009 [IXP_Usage], which concluded that the GÉANT Plus approach was significantly less cost-efficient than using the IP backbone, that it would not provide aggregation advantage, and that it would result in a large quantity of peerings to be managed (by the consortium as a whole).) It can also be noted that for some NRENs (and again from their perspective alone), IXP access via CBFs or GÉANT Lambda service instances is cost-effective. From the consortium-wide perspective, a more efficient, inclusive and resilient solution may be for GÉANT to peer with content providers directly and allow transit of this peering traffic over the IP backbone to/from those NRENs that want it. A proposal to connect to three IXs was approved at the September 2009 NREN PC meeting, and negotiations with potential peering partners at the IX peering point started in January 2010 for a one-year trial service.

Implications: Implications of the increased traffic loading across the GÉANT backbone resulting from the trial should be included in the capacity planning work, which in turn feeds into architecture planning.

40 Gbps and 100 Gbps IP Interfaces

For IP subscriptions greater than 10 Gbps it is possible for GÉANT to provide more than one 10 Gbps interface or 40 Gbps (STM-256) interfaces. 40 gigabit Ethernet and 100 gigabit Ethernet may be a possibility in the future, depending on the outcome of the ongoing developments in vendor support and interface costs.

Implications: The new network architecture should make available 40 Gbps interfaces for all NRENs subscribing to GÉANT IP at more than 10 Gbps. Interim upgrades of 2 and 3 x 10 Gbps should also be made available.

3.2.1.2 GÉANT Plus Service

GÉANT Plus Interfaces

Almost all GÉANT Plus NRENs now use 10 GE interfaces to receive the service. In future, 40 GE and 100 GE interfaces will also be possible, allowing the service to evolve to provide greater capacity circuits (e.g. multiple 10 Gbps) and improved flexibility.

Implications: The new network architecture should allow all GÉANT Plus services to be provided over 10 GE (and in future 100 GE) as standard and discontinue SDH interfaces. It will be possible to request GE interfaces (and potentially 40 GE and 100 GE) at additional cost. All interfaces should be able to support the relevant Ethernet OAM features.

Protected GÉANT Plus Circuits

GÉANT Plus circuits are currently offered as unprotected paths over the European backbone. It has been possible to order diverse routes to construct a more resilient Optical Private Network (OPN) structure, although there has been no take-up of this service. Although not offered so far, the current technology platform also

allows GÉANT to offer genuinely protected circuit services and allow rapid re-deployment in the case of a fibre break or other failure. The provision of any form of circuit protection requires additional investment over the standard GÉANT Plus infrastructure.

Implications: Little demand for protected circuit services has been voiced by NRENs so implementation is not felt to be an immediate priority.

3.2.1.3 GÉANT Lambda Service

40 Gbps Wavelengths

Two GÉANT NRENs have identified a requirement for 40 Gbps wavelengths. Two trial 40 Gbps links have been successfully tested on part of the GÉANT fibre footprint. 40 Gbps implementation on the whole GÉANT fibre footprint will be complex and costly. Before embarking on such an implementation it would be wise to review the long-term commitment to the technology platform supporting the service.

Implications: The new network architecture should support the implementation of a 40 Gbps wavelength service between Q3 2010 and Q3 2011 (the timing being dependent upon the outcome of the GÉANT DWDM transmission platform procurement exercise being undertaken in 2010).

100 Gbps Wavelengths

The above discussion of 40 Gbps wavelengths applies equally to 100 Gbps. For 100 Gbps, pricing is completely unknown (although there have been informal indications that the likely price will be less than twice the price of 40 Gbps transponders). The suitability of GÉANT fibre for this technology is also as yet untested.

There has been little indication so far of specific interest in 100 Gbps in wavelength band from NRENs, but this may develop during the next 12 – 24 months.

Implications: Network architecture planning should include a review of the level of interest in and the maturity of the technology for 100 Gbps in Q1 2011, with a view to supporting possible implementation from Q4 2011 onwards.

Rapidly Configurable and Restorable Wavelengths

So-called “colourless” and “directionless” ROADMs (Reconfigurable Optical Add-Drop Multiplexer) technologies offer the opportunity to effectively pre-provision wavelengths using fewer transponders than might otherwise be necessary. This pre-provisioning could be used to provide two new service features:

- Support for rapid turn-up of GÉANT Lambda services, with the aim of significantly improving service delivery times.
- Restorable GÉANT Lambda services which, in the event of a network failure, have the possibility of being automatically restored within a few seconds.

Whilst both of these services clearly have benefits, a cost-benefit analysis is needed to assess the usefulness of such services and whether they would be of interest to NRENs. It should also be pointed out that the more fibre meshing there is available, the higher the benefits achievable from the introduction of ROADMs.

Implications: Network architecture planning should include further investigation into the likely costs and benefits of these services and demand in the NREN community.

3.3 GÉANT Photonic Services

There is the potential and opportunity to develop a new type of service on GÉANT that would further distinguish the capabilities of the GÉANT network from commercial offerings; the project name for these is “GÉANT photonic services”. The potential is provided by the availability of dark fibre and the deployment of state-of-the-art and future-proof transmission technologies. The opportunity is provided by the emergence of applications that have stringent time accuracy requirements which cannot be met by the current set of services and infrastructures in GÉANT and some NRENs.

3.3.1 Outline of Photonic Services

In photonic services, an optical signal (i.e. a wavelength of a given colour) is received from the user, transported across the transmission network and handed over to the recipient in almost the same form in which it was received (“almost” because there will be optical noise and other effects on the signal; however, from a user’s perspective, nothing else changes). Interpretation of the signal depends on end users only (they can use it as an analogue or a non-standard or a standard digital signal). Nowhere along the transmission path does the signal have to be converted or processed by a method that prevents the transmitted value from being recognised (i.e. the usual devices for Optical-Electrical-Optical or OEO conversion are not applicable). Optical devices (for example, optical amplifiers) can be used to transfer the signal between end points. OEO devices that preserve “analogue” properties – for example, customised field-programmable gate array (FPGA)-based cards – could be acceptable.

The very nature of photonic systems, whereby digital signals (SDH, Ethernet) are mapped onto (or carried by) an analogue signal, enables the co-existence of photonic services with the well-established digital services in wide use today, without interference.

The European R&E community has gained relevant experience and design expertise in this area via a number of NREN initiatives (see, for example, the presentations given at the 6th Customer Empowered Fibre (CEF) Networks Workshop, Prague, 13 – 14 September 2010 [CEFWorkshop]) and the EC-funded Phosphorus project [Phosphorus].

Photonic services are not available from commercial ISPs, and therefore their potential value for the R&E community is very high. Photonic services will provide the following benefits:

- Openness (the user is enabled to participate in design and development) and transparency.
- Fixed signal delay given by fibre length only (satisfying the requirements of real-time applications).
- High speed (potentially beyond electronic transmission and processing speed).
- Low cost and low complexity.

The capability to manage analogue signals in a wide area environment can also be considered an investment in building a more flexible and future-proof network for the R&E community.

3.3.2 Expected Applications

The future backbone network should allow for providing services for new research and scientific applications that cannot be found in the portfolios of other Internet service providers. Some of these applications already exist, for example, transfer of highly accurate time and frequency, or real-time applications like processing of signals collected from earthquake sensors; some are expected as an extension of well-known real-time computer applications to real-time network applications (for example, control of unique scientific instruments or telescopes).

It can be assumed that such applications will require low jitter and transmission delays (i.e. minimise the number of OEO and buffers) and more than multi-gigabit capacity. The reason is that the degree of accuracy of time and frequency standards (like caesium clocks and hydrogen masers) is very high; for frequency stability this can be 10^{-16} . Any transmission instabilities therefore have devastating effects on stability and the accuracy of transferred time or frequency. Furthermore, new specialised modulation formats or framing structures (i.e. not Ethernet, SDH or OTN-based) should be expected. Other possibilities are open and further investigations are needed.

One of the interesting goals of metrology is to shift the output frequencies of atomic clocks from the microwave region into optical frequencies suitable for direct transmission over fibre. In this case the subject of transmission is the carrier frequency (colour) instead of a modulated signal. Therefore, the transmission path should consist of fibre and optical amplifiers only – neither OEO regeneration nor “re-colouring” is allowed. Investigations into suitable optical links and amplifier parameters are currently being carried out by several research activities.

A flexible and open photonic transmission layer is therefore very important, which can be translated into requirements for all-optical systems, supported by so-called alien wavelengths. The term alien wavelength is widely used and many vendors support such services. (For more information about alien wavelengths, see [DJ1.2.1] Section 2.5.8.)

3.3.3 Implications for GÉANT and NREN Networks

The provision of photonic services requires that the frequency spectrum of the fibre is allocated in bands (or ranges) for photonic services and bands for digital services. The band is specified by the starting and ending frequency (wavelength) and should be allocated and used in such a way as to avoid interference with other bands. From a purely technical perspective, implementation of photonic services is a simpler task than the implementation of digital transmission systems.

The requirement to avoid interference with other frequency bands must take into account the end-to-end and multi-domain scope of photonic services: an analogue optical signal is received from the user, transported across an NREN, transported across GÉANT, transported across the second NREN, and handed over to the recipient in (almost) the same form in which it was received. In most cases, the signal is transported across an

NREN; in exceptional cases, the signal may be transported by dedicated fibre from the user directly to the GÉANT PoP.

From an operations perspective, remote monitoring and setup of all active optical devices along the E2E photonic service is needed. Due to the nature of photonic signals and the lack of a digital boundary between administrative domains (as is the case for digital signals), it is important that all the parties involved in the provision of end-to-end photonic services have access to monitoring information and that agreed processes are in place for the technical design of the services and to troubleshoot operational and technical issues.

The main challenge is the all-optical reach of high-speed transmission systems. However, many vendors claim a maximal transmission distance greater than 2000 km for 40 Gbps signals without OEO regenerators. For signals with lower speeds, all-optical reach will be longer. This situation can be further improved with the help of all-optical regeneration in the future. Such methods and devices have been proposed and prototypes do exist. However, commercial and technical need for their deployment is currently very rare. Wavelength blocking can occur in any system and if all-optical E2E lambdas need to be “re-coloured”, other devices known as all-optical wavelength convertors can solve this problem. The situation is very similar for both all-optical regenerators and convertors – they are not yet commercially available as integrated elements of transmission systems. OEO devices preserving “analogue” properties – for example, customised FPGA-based cards – could solve the “reach challenge”.

10 Gbps and 40 Gbps and indeed 100 Gbps signals are carried over high-frequency signals (e.g. tens of GHz). For these signals, the reach is generally as described above. However, for real-time applications, the emphasis is on the exact timing of the event rather than speed and frequency. Frequencies for such applications are in the order of hundreds of MHz. For such low frequencies, the all-optical unregenerated reach can be expected to be in excess of 10000 km.

3.3.4 Conclusion

The opportunity and need to provision fully photonic services are maturing. At present, this type of service is not offered by GÉANT or NRENs in general. In order for GÉANT to be ready to offer this type of service, the frequency spectrum of the fibre footprint should be allocated in bands for photonic services and bands for digital services. A more thorough analysis of the implications for the technology choice for GÉANT’s digital transmission system should be conducted, and the associated costs evaluated, so that the requirements may be taken into account in equipment procurements and in the network topology. The offer of photonic services needs to be discussed further and agreed within the GÉANT and NREN community. Such discussions may lead to pilot projects of photonic services in GN3.

3.4 Quality

3.4.1 Introduction

GÉANT’s mission is to offer high-quality network services, over and above what can be obtained from the commercial market. This is a very broad statement, and it really depends on the definition of quality. In general,

there will be some quality parameters where GÉANT undoubtedly exceeds corresponding quality parameters that can be achieved from the commercial market, others where they are comparable, and others where the commercial market exceeds the quality of GÉANT. For example, with GÉANT essentially being a closed network, the number of directly connected hosts is not a relevant quality parameter, yet this might be a distinguishing quality factor for a commercial network operator. On the other hand, the ability to perform data transfers of ~10 Gbps between any two locations on GÉANT's fibre footprint is of primary concern for GÉANT. It is clear that for the parameters that are relevant to GÉANT, the quality GÉANT offers must be over and above what can be achieved from the commercial market.

GÉANT has traditionally succeeded in maintaining very high quality standards for the parameters that are relevant. The parameters vary per service and depend also on the maturity of and experience with delivering such services. The challenge for GÉANT is to continue to achieve high quality levels for its core service parameters and to improve further on the quality of the developing services and their associated parameters.

This section analyses the parameters that are relevant to GÉANT on a per service basis, elaborates on their values (and the rationale for such values) and the implications for the GÉANT architecture.

Two key principles are to be noted:

- Where values are set, these are intended as target values rather than Service Level Agreements (SLAs).
- It must be possible to measure the values. In most cases, GÉANT is able to measure the value of quality parameters. In other cases, it is not possible for GÉANT to do so, only the end user of the service. In such cases, measurement methodologies for that particular parameter must be agreed between GÉANT and the service end user.

3.4.2 Quality Parameters

3.4.2.1 IP Service

The quality parameters for the GÉANT IP service are as follows:

Availability of NREN Access to GÉANT

Many NRENs have both a primary and a backup access link to GÉANT. In cases where a backup link exists, the target for this parameter is 99.9%. This is a challenging value, as the demarcation point of NREN access to GÉANT is on the NREN access router; it therefore also includes the physical interconnection between GÉANT and the NREN.

Experience in GÉANT has shown that it is not always possible to achieve this value. Indeed, in a very limited number of cases (3) the availability has been <99.5%. These cases should be examined and dealt with separately.

The implications for GÉANT architecture are:

- Use equipment with high availability components.

- Simplify interconnection between GÉANT and NRENs.

Round-Trip Time (RTT)

This parameter refers to the Round-Trip Time (RTT) between any two NREN access PoPs in normal conditions and in error conditions. The rationale for the parameter is that Transmission Control Protocol (TCP) performance is inversely proportional to RTT. The higher the RTT, the lower the TCP throughput that can be achieved. GÉANT supports many bandwidth-hungry applications, therefore high TCP throughput is mandatory.

The current design rule for GÉANT is maximum 50 ms RTT between any two NREN access PoPs. This allows TCP flows of 1 Gbps using a standard TCP stack on a workstation. This applies both in normal conditions and conditions of outage of one circuit, where a re-route of traffic will take place. This in turn affects the configuration of the IP layer, the strategic placement of routers in the backbone and the amount of meshing in the network (topology).

For the next-generation GÉANT architecture, taking into account the availability of 100 Gbps and the realistic possibility of TCP transfers of 10 Gbps using GÉANT IP, the RTT should be such as to permit 10 Gbps between any two NREN access points, albeit using larger, non-standard TCP window sizes. This again implies a maximum of 50 ms RTT between any two NREN access PoPs (where there is infrastructure to support such capacities). It should be noted that each 1000 km of distance measured by fibre route adds 10 ms of signal transmission delay to RTT *without* any further delay due to signal processing (e.g. due to FEC). However, even the most complex FEC processing should not add more than a few ms in total. Even if it were as high as 10 ms (in terms of RTT) then this would allow distances of 4000 km to be spanned within the aforementioned 50 ms limit, which is largely commensurate with the longest PoP-PoP distances on the GÉANT backbone, provided that the current level of meshing is maintained or improved upon.

RTT is also affected by congestion on routers, therefore this should be as low as possible. See *Packet Loss and Re-ordering* below for further analysis.

Packet Loss and Re-ordering

This parameter refers to packet loss between any two NREN access PoPs in normal conditions and in error conditions, and the re-ordering of packets. The rationale for the parameter is that loss of a packet triggers retransmissions of parts of a TCP window, which ultimately reduces TCP throughput. Re-ordered packets have the same effect, to some extent, of lost packets. There is also increasing use of loss-sensitive applications such as video-conferencing, the usability of which is severely affected by hence packet loss.

Packet loss and re-ordering should be 0 in normal and in error conditions. A small loss can be tolerated only during the brief period during which a circuit fault is detected and packets are re-routed by the standard routing protocols (Border Gateway Protocol (BGP) convergence).

The implication for GÉANT architecture planning is that the upgrade process of a GÉANT backbone link should be initiated when its average utilisation surpasses 30% of the link capacity. This is to accommodate, as far as possible, bursty traffic and additional traffic in case of outages in the network that require traffic to be re-routed over other links and avoid congestion on links. When an upgrade is actually implemented, utilisation levels are likely to be > 40%.

For re-ordering, the implication is to use router architectures that do not give rise to packet re-ordering.

Jitter or Delay Variation

Jitter, or variation in delay, affects real-time applications such as video-conferencing. Jitter must be minimised, and ideally should be 0. It is not, however, possible to guarantee 0 jitter in statistical multiplexing technologies where buffering of packets is involved.

The implication for architecture planning of reducing jitter is to deploy hardware that is optimised to keep jitter at negligible levels.

3.4.2.2 GÉANT Lambda and GÉANT Plus – Fixed Circuit Services

The quality parameters for fixed circuit services delivered over the GÉANT backbone, i.e. GÉANT Lambda and GÉANT Plus, are as follows:

Response Time to User Request (Yes/No Decision)

This parameter relates to the business process of requesting an instance of GÉANT Lambda or GÉANT Plus services. A yes/no answer must be given within a defined timeframe. The process must be simple from the user's point of view. Ideally the response should be delivered within a matter of days.

This quality parameter does not affect the GÉANT architecture.

Implementation Time

This parameter relates to the lead time for implementing a GÉANT Plus or GÉANT Lambda service instance after an affirmative answer to a request has been received.

At present, the implementation time for GÉANT Lambda is approximately three months, as implementing a new GÉANT Lambda service requires the purchase and installation of new hardware. New technologies (such as Wavelength Selective Switching (WSS)) combined with new business process models (involving keeping a stock of unused hardware, for example) may lead to reduced implementation times. This parameter may therefore have an impact on the technology choice for GÉANT. Appropriate values for this parameter require further discussion.

For GÉANT Plus services, the current implementation time is in the order of approximately five days, mainly for business process reasons. It may be possible to improve this. However, in order to do so, the impact on the Operational Support System (OSS) for GÉANT needs to be examined.

Availability

The availability of GÉANT Plus and GÉANT Lambda services is distance dependent. Other dependencies include the combination of support and maintenance contracts with the equipment vendors, Mean Time Between Failures (MTBF) of the equipment, location of fibre (i.e. how likely is a fibre cut), Mean Time to Recovery (MTTR) of fibre cuts, and the protection embedded in the network design. No value for this parameter has been set for now.

The implication for the architecture of maximising the availability of these services is to reduce the number of network components being used. In particular, this means using ROADMs where possible at fibre junctions and possibly also in GÉANT PoPs in order to minimise the length of the circuits and the number of components used.

Throughput, Loss, Jitter

These are typical “packet” quality parameters. The rationale for these is that circuit services are mostly used by their users to transport IP traffic. However, GÉANT cannot monitor these parameters in this case, because the service delivered by GÉANT does not have such counters. There are, however, underlying parameters that have an effect on these packet-based parameters and which have implications for the GÉANT architecture. These are:

- Underlying data link Bit Error Rate (BER) of 10^{-15} .
- User input traffic should be shaped to end-to-end (E2E) link capacity.
- Backbone link capacity should not be statistically multiplexed. Statistical multiplexing technologies can be used, however, so long as the allocated capacity per backbone link does not exceed link capacity (this is automatic with TDM technology). This will also help meet “low” one-way delay (OWD) and jitter targets.

As mentioned above, it is not possible for GÉANT to monitor these packet-based parameters. However, it is reasonable to set out what the performance experienced by the user should be, in defined conditions. For example, packet loss should always be 0 provided the user does not overload the link capacity. Likewise the TCP throughput should be 95% of link capacity, provided a defined set of conditions are met.

These pre-conditions for the performance experienced by the user will need to be mutually agreed with the user prior to handover of the service.

With regard to throughput, this is intended as TCP throughput. It should be 95% of link capacity. Loss should be 0.

3.4.2.3 GÉANT Lambda and GÉANT Plus – Dynamic Circuit Services

GN3 is in the process of developing dynamic circuits. These are a complement to the GÉANT Plus and GÉANT Lambda services that are currently available and provisioned manually by the Network Operations Centres (NOCs). They will be provisioned automatically via software tools after requests initiated by users.

As dynamic provisioning as a service is currently still under development, it is not yet clear what parameters are associated with it and what the implications for the architecture might be. The following are the current expectations for the parameters:

- Provisioning time. As this is a differentiator from fixed circuits, the value should be in the order of 10 minutes to start with.
- Availability, once the circuit has been set up. The value should be the same as for fixed circuits, and will be measured over the duration of the reservation.
- Call setup success ratio statistics.
- Delay, jitter, loss, burst size.

4 Capacity Forecasts

Periodically DANTE performs surveys in which it solicits input on future backbone network service requirements from all of the GÉANT NRENs. The last of these was a questionnaire distributed to the NRENs at the end of 2008 seeking input on their expectations for GÉANT capacity usage throughout the period of the GN3 project. This was reported in some detail in Deliverable DS1.1.1 “Report on the Backbone Architecture Study” [DS1.1.1].

The conclusions drawn **at the time** were that when taking into account the need to support the ongoing organic growth of the IP traffic carried by the GÉANT backbone (which had been consistently growing at a rate of 40%–50% per annum) and the expected demands for GÉANT Plus and GÉANT Lambda services, then it was clear that the GÉANT transmission layer needs to be upgraded in terms of the totality of the “end of life” capacity we should expect it to support and its capabilities. A key part of the latter is the requirement to be able to cost-effectively supply 100 Gbit/s wavelength services.

4.1.1 Organic Growth

At the time that DS1.1.1 was written, a year had passed since the original capacity forecast had been made so this gave the opportunity to check whether the original backbone transmission capacity forecast had been borne out by real growth figures. The results of this for the IP component of the transmission capacity forecast are shown in Table 4.1. As with previous versions of this table, the link utilisation levels (actual and forecast) are shown in the columns labeled “Gbps”. These figures are broadly based on 95th percentiles of 5-minute samples taken from the traffic level monitoring systems. As such they present a good measure of the “peak” usage of the links rather than averages calculated over a diurnal cycle. The columns in Table 4.1 labeled “cap” show the capacity of that particular link in Gbps and the resulting percentage utilisation is shown in the column labeled “%”.

Note that very often in monitoring IP networks, values of traffic levels based on 5-minute samples are referred to as “peak values” since such a sampling period is commonly the shortest used. Clearly there can be higher peaks on shorter time scales (of the order of seconds) but these are not usually recorded by the monitoring systems used on production IP networks. Hence, if it is the intention of a network operator to accommodate such peaks whilst imposing as little queuing delay – resulting in jitter – (even under certain failure conditions) as possible, then capacity planning processes need to be designed with this goal in mind.

Now, at the time of writing this deliverable (October 2010), the IP capacity forecast has been checked again. These figures (again reflecting 95th percentiles of 5-minute samples) are also shown in Table 4.1. This shows that some links have undergone significantly more growth in traffic levels than others. An analysis of the growth

factors on a link by link basis reveals that they range from 0.2 to 10. The values that are less than one indicate that traffic levels on that particular link have dropped off significantly during the preceding year. Usually this is the result of shifts in traffic flows on the backbone due to changes made to the routing setup. Typically there will be a corresponding trunk that has seen a big increase in utilisation (e.g. one of the trunks for which the growth factor is above 2. The average of these values for 2009 comes out at 1.7 and for 2010 (or rather the 10 month period to date) the average growth factor is 1.5.

Note that the cells highlighted in blue indicate links that are currently running at utilisation levels that exceed 30%. In the past this has been used as a level at which the process of considering and executing a capacity upgrade has been triggered. Although this sounds like a low threshold the fact is that the whole upgrade process is quite lengthy and typically by the time that the upgrade is performed the utilisation level has reached 40% or even more. This is illustrated in Table 4.1 where link capacity upgrades that were starting to be considered at the time that DS1.1.1 was written are still not yet in place (e.g. most of those highlighted in blue). With the exception of the Vienna–Athens trunk (Vie–Ath), all of those links highlighted in blue should have been upgraded by the end of the year: those at 10 Gbps to 2x10 Gbps and Amsterdam–Frankfurt and Frankfurt–Geneva to 40 Gbps. (Part of the delay here is in getting the latter two upgraded since the 10 Gbps resources freed up by so doing will be used to implement some of the upgrades to 2x10 Gbps on the rest of the western ring.)

Table 4.1 also shows how the IP trunk utilisations are projected to grow over the next two years (based on a simple growth rate of 50% pa). This table assumes that no significant changes to the present day routing setup are made and that trunk capacities do not change from those planned to be implemented in the coming months (shown in the columns for 2011 and 2012). Clearly the capacities as shown in the table will be insufficient for some of the trunks by factors of about two, so they are very likely to change by more than is indicated in the table – e.g. 40 G to 100 G (especially when taking into account the expectations of discontinuous growth in IP demand as discussed below).

4.1.2 Discontinuous Growth

In addition to considering the growth of the GÉANT IP service, the original forecast addressed capacity expected to be needed to support growth in GÉANT Plus and GÉANT Lambda services. 2009 and 2010 growth figures for backbone trunk capacity to support GÉANT Plus and GÉANT Lambda services are considerably down on what was originally forecast at the end of 2008 (which were broadly based on the same annual growth factors of 1.5 as applied to the IP capacity forecast). Despite this, significant growth in demand for all the GÉANT backbone services is still expected to take place during the GN3 project as it did during the first three years of the GN2 project. This is illustrated in Figure 4.1 which shows the cumulative total of BW-distance product for capacity provisioned on the GÉANT (fibre) backbone that supports the provision of GÉANT Plus and GÉANT Lambda services (in other words, excluding capacity provisioned for the support of GÉANT IP services).




Figure 4.1: Cumulative total BW-distance product (excluding capacity for IP)

As new rounds of projects that need substantive network connectivity service are approved for funding and start to come on line then it is expected that the capacity growth profile in time just exhibits delayed yet sizeable upward steps rather than the continuous steady growth profile that was used in this simple modeling. This was certainly seen to be the case throughout the GN2 project.

Discontinuous growth is not confined to the bandwidth provisioning services (GÉANT Plus and GÉANT Lambda). There is now a general expectation amongst many in the R&E networking community that there will be a wave of new demands hitting the IP backbone over and above the organic growth described above. Some of this is expected to come from the HEP physics community – especially related to traffic that will be attributed to the LHC community. A case in point is illustrated in Figure 4.2, which shows the impact of some sustained data transfers on one of the GÉANT IP backbone trunks on the western ring – namely London–Amsterdam. Although the activity giving rise to this particular traffic profile (which was actually the result of some data transfers between the US and the UK) had ceased by the time that the September 2010 figures in Table 4.1 were recorded, it should be noted that the resulting flows were continuous and sustained for a few days at a time. Had a significant reroute occurred during that time (say due to a trunk outage elsewhere on the network) then there would have been a high likelihood that the London–Amsterdam trunk would have been running close to its full capacity or even saturated.




Figure 4.2: London–Amsterdam IP trunk during recent LHC-related data transfer activities

Events like this illustrate the necessity to be prepared for the IP backbone to support “unexpected” step changes in traffic levels (aside from those associated with temporary shifts of traffic due to rerouting during trunk outage scenarios). Sometimes these events can be relatively shortlived (such as the one shown above) but we might also expect to see step changes followed by much longer lived flows (e.g. during a prolonged period of data gathering and distribution by a large scale scientific experiment like those associated with the LHC).

4.1.3 Conclusions on Capacity Forecasts

Taking into account the historical data, while also projecting the future needs of new high-data demands that could materialise from users such as those identified in Appendix A (see page 111) and from the European Strategy Forum on Research Infrastructures (ESFRI) projects, where data is observed in some scientific fields to double every year, the main conclusions that can be drawn from this capacity forecasting work are as follows:

- The current GÉANT DWDM transmission platform will most likely need to undergo upgrade or replacement before the end of the GN3 project.
- 40 Gbps transmission capability on the GÉANT backbone is needed now on some specific routes. Further work is needed to evaluate whether this capability will be required across the network before the availability of 100 Gbps transmission capability is commercially realistic.
- Based on the initial demand forecasts across the current topology and architecture (which need further evaluation), 100 Gbps transmission capability on the GÉANT backbone would be predicted to be needed from the end of 2011/start of 2012.

The project will maintain close contact with the user communities that have the most significant data-traffic demand, to ensure the capacity forecasts are as up to date and accurate as possible.

Link	Actual 95th percentile usage									Forecast 95th percentile usage (made at end of 2009)								
	end of 2008			end of 2009			Sep 2010			end of 2010			end of 2011			end of 2012		
	cap	Gbps	%	cap	Gbps	%	cap	Gbps	%	cap	Gbps	%	cap	Gbps	%	cap	Gbps	%
Lon-Ams	10	3.2	32%	10	3.5	35%	10	3.5	35%	20	5.3	26%	30	7.9	26%	40	11.8	30%
Ams-Cop	10	2.6	26%	10	1.5	15%	10	1.0	10%	10	2.3	23%	20	3.4	17%	20	5.1	25%
Cop-Tal	10	1.5	15%	10	1.2	12%	10	2.0	20%	10	1.8	18%	10	2.7	27%	20	4.1	20%
Tal-Rig	10	0.1	1%	10	0.5	5%	10	1.0	10%	10	0.8	8%	10	1.1	11%	10	1.7	17%
Rig-Kau	10	0.1	1%	10	0.5	5%	10	1.0	10%	10	0.8	8%	10	1.1	11%	10	1.7	17%
Kau-Poz	10	1.2	12%	10	0.7	7%	10	1.4	14%	10	1.1	11%	10	1.6	16%	10	2.4	24%
Ams-Fra	10	3.4	34%	10	3.6	36%	10	4.5	45%	20	5.4	27%	30	8.1	27%	40	12.2	30%
Cop-Fra	10	2.0	20%	10	2.5	25%	10	4.5	45%	20	3.8	19%	20	5.6	28%	30	8.4	28%
Lon-Par	10	2.4	24%	10	3.4	34%	10	3.0	30%	20	5.1	26%	30	7.7	26%	40	11.5	29%
Par-Mad	10	0.7	7%	10	1.0	10%	10	1.0	10%	10	1.5	15%	10	2.3	23%	20	3.4	17%
Par-Gen	10	3.3	33%	10	2.5	25%	10	4.0	40%	20	3.8	19%	20	5.6	28%	30	8.4	28%
Mad-Gen	10	3.8	38%	10	4.0	40%	10	5.0	50%	20	6.0	30%	30	9.0	30%	50	13.5	27%
Mad-Mil	10	0.0	0%	10	0.0	0%	10	0.0	0%	10	0.0	0%	20	0.0	0%	20	0.0	0%
Mil-Gen	20	3.1	16%	20	4.5	23%	20	4.5	23%	30	6.8	23%	40	10.1	25%	50	15.2	30%
Fra-Gen	20	5.6	28%	20	6.5	33%	20	8.0	40%	40	9.8	24%	50	14.6	29%	80	21.9	27%
Fra-Poz	10	1.9	19%	10	1.5	15%	10	2.0	20%	10	2.3	23%	20	3.4	17%	20	5.1	25%
Fra-Pra	10	1.5	15%	10	2.5	25%	10	5.5	55%	20	3.8	19%	20	5.6	28%	30	8.4	28%
Poz-Pra	10	0.4	4%	10	0.2	2%	10	0.2	2%	10	0.3	3%	10	0.5	5%	10	0.7	7%
Fra-Vie	10	0.1	1%	10	0.1	1%	10	0.1	1%	20	0.2	1%	20	0.2	1%	20	0.3	2%
Mil-Vie	10	1.7	17%	10	1.0	10%	10	1.5	15%	10	1.5	15%	10	2.3	23%	20	3.4	17%
Vie-Pra	10	0.6	6%	10	2.0	20%	10	5.0	50%	10	3.0	30%	20	4.5	23%	20	6.8	34%
Vie-Bud	10	2.6	26%	10	3.0	30%	10	1.8	18%	20	4.5	23%	30	6.8	23%	40	10.1	25%
Bud-Pra	10	1.3	13%	10	1.2	12%	10	1.0	10%	10	1.8	18%	10	2.7	27%	20	4.1	20%
Vie-Ath	10	2.3	23%	10	3.7	37%	10	3.5	35%	20	5.6	28%	30	8.3	28%	50	12.5	25%
Ath-Sof	10	2.8	28%	10	0.5	5%	10	2.0	20%	10	0.8	8%	10	1.1	11%	10	1.7	17%
Sof-Bud	10	2.9	29%	10	0.6	6%	10	2.5	25%	10	0.9	9%	10	1.4	14%	10	2.0	20%
Sof-Buc	10	0.1	1%	10	0.8	8%	10	0.0	0%	10	1.2	12%	10	1.8	18%	10	2.7	27%
Bud-Buc	10	2.8	28%	10	0.7	7%	10	1.3	13%	10	1.1	11%	10	1.6	16%	10	2.4	24%

Table 4.1: Trunk wavelength capacity (in Gbps) required to support forecast growth in backbone IP traffic (revisited 2009 and 2010)

5 Building Blocks

5.1 Introduction

This section outlines the different types of network architecture building blocks that can be used in different parts of the GÉANT network. These building blocks can be categorised as follows:

- Dark fibre DWDM infrastructure.
- Managed wavelength infrastructure.
- Carrier class transport infrastructure.
- Managed carrier class transport infrastructure.

These blocks will be used later on in the deliverable (Section 8 on page 88) to build the options for re-engineering the GÉANT architecture.

5.2 Building Block A – Dark Fibre DWDM Infrastructure

This refers to “self-provided” wavelengths provisioned over “owned” (actually, leased) dedicated fibre that is lit using a carrier grade DWDM transmission system that is itself fully owned and operated by DANTE (on behalf of the GÉANT consortium). Clearly the fibre must conform to a minimum quality specification.

The DWDM transmission system can be based on:

- Fixed add/drop or simple point-to-point capability. (As has already been stated, this is the situation for the GÉANT transmission system today but with the addition of tuneable transponders.)
- Basic Reconfigurable (Optical) Add/Drop (ROADM) capability.
 - Typically in conjunction with tuneable transponders.
- Advanced ROADM capability, supporting so-called “directionless”, “colourless” and multi-degree flexibility.
 - Details of implementation vary from vendor to vendor, but typically this requires tuneable transponders, flexible filters of one kind or another, various splitters and combiners, and wavelength selective switches or similar.

The architectural layout is based on a DWDM network that has at least one flexible Add/Drop (T-ROADM) in each GÉANT PoP where an NREN domain connects to GÉANT. In order to achieve resiliency and to support service restoration, the T-ROADM must have as a minimum 2 branches into the network. The overall network could require strategically placed Reamplification, Reshaping and Retiming (3R) generation in order to cope with resiliency and restoration and deal with the long transmission distances inherent in the GÉANT backbone.

In principle, a DWDM network based on ROADM is more appealing than the situation today (where Reamplification, Reshaping and Retiming (3R) regeneration cannot be avoided), for a number of reasons as outlined in Section 6.2 on page 48, but full flexibility at all locations may prove to be impractical for reasons of upfront cost and the amount of colocation space that may potentially be required.




Figure 5.1: DWDM network

The schematic outline of the DWDM network in Figure 5.1 shows a ring/meshed structure.

5.3 Building Block B – Managed Wavelength Infrastructure

Managed wavelength infrastructure (also known as grey fibre DWDM infrastructure) is provided by a third party (i.e. an NREN or a commercial operator) offering full wavelength services. It is a dedicated unprotected point-to-point wavelength, and complements the wavelengths that are delivered by GÉANT's DWDM system (building block A) in order to build a complete network. It consists of three different circuit component parts:

- **Wavelength:** an optical wavelength specified according to the Optical Transport Network (OTN) architecture, which has been defined in the ITU-T G.872 standard and the interfaces defined in G.709 [ITU G.709].
- **Client Interface:** the compliant interconnection interface between the NREN/operator (provider) and the GÉANT network (user).
- **Local Access:** a dark fibre connection between the GÉANT network equipment interface and the provider equipment.

These components are shown in Figure 5.2 and described in Table 5.1 below.

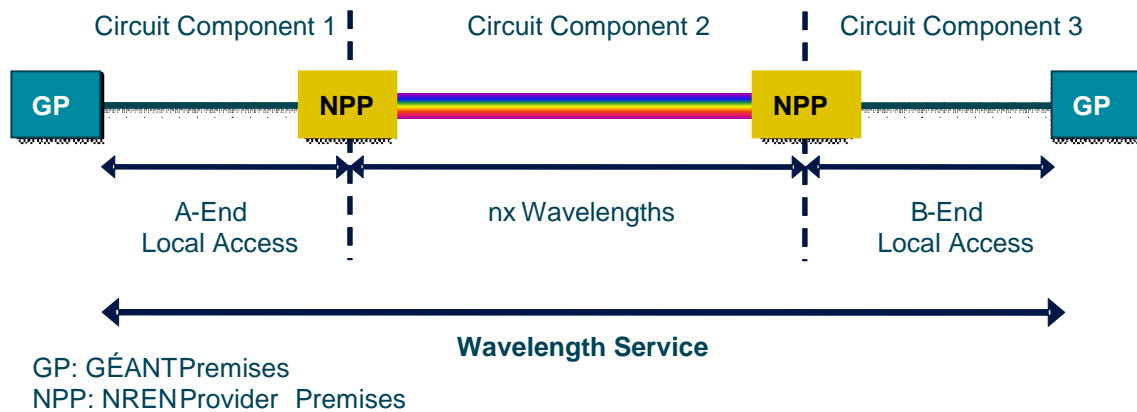


Figure 5.2: Circuit components used to deliver the wavelength service

Circuit Component	Description
1	The originating local access (A-End), which is provided by the NREN provider using NREN provider dark fibre or leased dark fibre from a third-party operator.
2	The circuit using the NREN provider's DWDM network.
3	The destination local access (B-End) provided on behalf of the NREN provider (same procedure as for A-End is applicable).

Table 5.1: Circuit components used to deliver the wavelength service

Delivery of the wavelength service should be in close proximity to GÉANT network equipment, preferably at colocated NREN provider and GÉANT premises. The reason for this is that traditionally the local access circuit components (1 and 3 in Table 5.1) are based on black and white optics (meaning at 1310 nm or 1550 nm depending on the reach that is required) running over dedicated fibre pairs with an OEO function at the NPP points indicated in Figure 5.2). In this case, when multiple wavelength services are required, then multiple fibre pairs will be required for circuit components 1 and 2. This is clearly much less of a problem when these components are simply fibre patches within a colocation facility as opposed to fibre pairs that go across town and may be many kilometres long and costly.

Of course, where there is support for so-called “alien wave” access to the provider's DWDM network (circuit component 2 in Table 5.1) then there is no longer the expectation/restriction that access must be via black and white optics and OEO conversion. Currently the expectation is that, although there may be some NREN providers who are willing to support alien wave access in production, there are still NRENs who will not and the majority of commercial network operators will not do so either.

The wavelength service can be provided by several NRENs interconnected via cross-border fibre (CBF) and is therefore seen as two or more independent DWDM networks interconnected at demarcation points via back-to-back client interfaces (based on black and white optics). Again, there may be providers that do this based on an alien wave approach, or even interworking at the DWDM layer where such OEO conversion at the demarcation points is avoided, but it remains to be seen how much this will occur in practice (especially where commercial operators are involved).




Figure 5.3: Managed wavelength infrastructure

The logical layout is identical to Block A, DWDM infrastructure. In this case, however, the wavelengths between the nodes are provided by the third parties considered above. Note that in this case the nodes in Figure 5.3 are less likely to be ROADMs unless they have integrated OEO switching functionality.

5.4 Building Block C – Carrier Class Transport Infrastructure

A carrier class transport infrastructure (CCTI) is an infrastructure based on packet, Time-Division Multiplexing (TDM) or Internet Protocol/Multi-Protocol Label Switching (IP/MPLS), carried by a DWDM or a managed wavelength infrastructure. Carrier class transport will be used for the GÉANT Plus service and sub-circuit interconnects to a centralised IP architecture if the IP architecture is changed from distributed to centralised.

JRA1 T1 has defined Carrier Class Transport Network Technologies (CCTNTs) as technologies designed to provide transport for network services and protocols. “Carrier class” denotes that the technologies are extremely reliable, support a wide range of speeds up to the current industry maximum, and are well-tested and proven in their capabilities.

A transport network technology must meet the following criteria to qualify as a CCTNT:

- Effectively support diverse types of traffic such as the elastic traffic of data applications and time-sensitive multimedia traffic.
- Effectively support all popular customer services such as Internet access, Virtual Private Networks (VPNs), Voice over IP (VoIP), IPTV and others.
- Be manageable by providing diverse and feature-rich Operations, Administration and Maintenance (OAM) functionality.
- Be reliable by providing resilience and fast restoration for transport connections.
- Be scalable to support numerous customer connections through a carrier network.
- Be able to provide Quality of Service (QoS) and bandwidth guarantees when necessary.
- Provide separation of customer and provider networks in terms of operation and configuration parameters such as address spaces, connection IDs and others.
- Be cost-effective. This is a major requirement for service providers. Network costs are extremely high and a good argument for service providers to change their legacy technologies is the ability to provide better and more flexible services at the same or, if possible, lower cost.

- Deliver high bandwidth and performance up to the current industry limit (i.e. up to 40 G today and up to 100 G in the near future).
- Conform to the appropriate standards.
- Be multi-protocol. A CCTNT should be capable of transporting any kind of customer traffic and so should support different if not all existing protocols.

Taking into account these criteria, the technologies considered are the following.

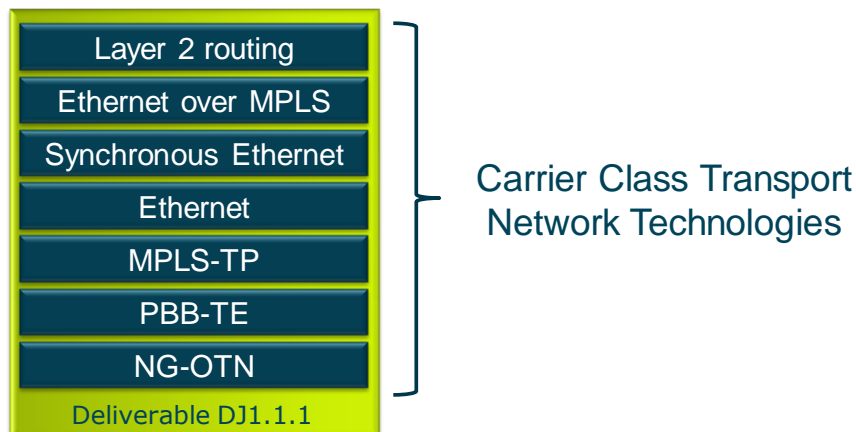


Figure 5.4: Carrier class transport network technologies

These are described in detail in “Deliverable DJ1.1.1: Transport Network Technologies Study” [DJ1.1.1] and are also discussed in the next section, Section 6 *Technologies*, where the emphasis is on how they apply to GÉANT.

5.4.1 Logical Layout

The logical layout foresees a transport network with adequate meshing in order to achieve resiliency and to support the restoration of services. The transport network is composed of transport nodes owned and operated by DANTE (on behalf of the GÉANT consortium). It is based on an underlying wavelength infrastructure which is composed of GÉANT’s DWDM infrastructure and a managed wavelength infrastructure (i.e. some combination of building blocks A and B).




Figure 5.5: Transport network

The schematic outline of the transport network in Figure 5.5 shows a meshed structure.

5.5 Building Block D – Managed Carrier Class Transport Infrastructure

Managed carrier class transport infrastructure is provided by a third party (i.e. an NREN or a commercial operator) offering services based on packet, TDM or IP. It is a dedicated unprotected point-to-point transport link, which complements the GÉANT-owned transport infrastructure as described in Section 5.4. It consists of three different circuit components parts:

- **Transport:** a service specified according to Section 5.4.
- **Client Interface:** the compliant interconnection interface between the NREN/operator (provider) and the GÉANT network (user).
- **Local Access:** a dark fibre connection between the GÉANT network equipment interface and the provider equipment.

These components are shown in Figure 5.6 and described in Table 5.2 below.

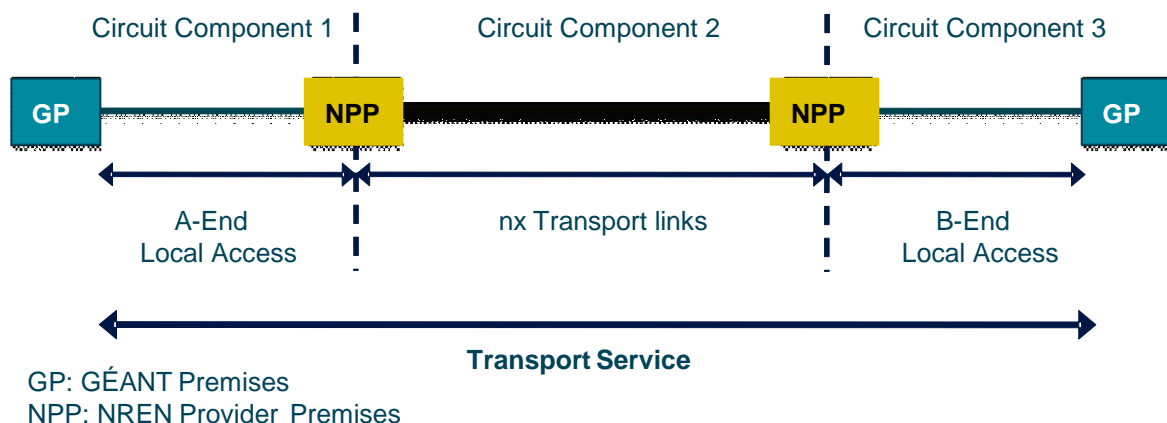


Figure 5.6: Circuit component parts used to deliver transport service

Circuit Component	Description
1	This is identical to circuit component 1 in Section 5.3.
2	The circuit using NREN provider transport network.
3	This is identical to circuit component 3 in Section 5.3.

Table 5.2: Circuit components used to deliver the transport service

Delivery of the transport service should be in close proximity to GÉANT network equipment, preferably at colocated NREN provider and GÉANT premises.

The transport service can be provided by several NRENs interconnected via CBF and is therefore seen as two or more independent transport networks interconnected at demarcation points. (CBF is discussed in the JRA1 T3 deliverable DJ.1.3.1 “Architecture Considerations for Federated Backbone Networks – Study” [DJ1.3.1]. An overview of GÉANT’s approach to evaluating the CBF resources offered by NRENs is provided in Appendix C *Third-Party Connectivity* on page 130.)

The architectural layout is based on a transport network with adequate flexible interconnection representation in order to achieve resiliency and to support the overall restoration topology of the network. The transport network is owned and operated by DANTE (on behalf of the GÉANT consortium), with DANTE-owned and -operated nodes (again on behalf of the GÉANT consortium) facilitating the GÉANT network services except for the full wavelength service, which can be provided either by GÉANT or by the third-party DWDM network.




Figure 5.7: Managed transport network

The schematic outline of the managed transport network in Figure 5.7 shows a meshed structure.

6 Technologies

This section provides a brief overview of each technology in the network architecture, describes how it supports the GÉANT backbone, and presents an analysis of the current market situation. The technologies covered are:

- Point-to-Point Dense Wavelength-Division Multiplexing (P2P DWDM).
- Reconfigurable Optical Add-Drop Multiplexers (ROADMs) and Wavelength Selective Switches (WSSs).
- Next-Generation Optical Transport Network (NG-OTN).
- Ethernet over Multi-Protocol Label Switching (EoMPLS).
- Carrier Ethernet (cE).

The section does not describe in detail the technologies or the theory behind them; rather it focuses on the devices' practical implementation, as preparation for the procurement process, and on how they can apply to the GÉANT network and service portfolio (GÉANT IP, GÉANT Plus, GÉANT Lambda) with respect to 10 Gbps, 40 Gbps and 100 Gbps capabilities and dynamic provisioning/protection/restoration.

6.1 Point-to-Point Dense Wavelength-Division Multiplexing (P2P DWDM)

6.1.1 Technology Overview

Point-to-point Dense Wavelength-Division Multiplexing (P2P DWDM) technology provides a way of implementing wavelength paths with the following characteristics:

- The route of each wavelength path can only be changed with manual intervention and on-site visits including actions such as cards re-cabling.
- Between two consecutive regeneration nodes, each wavelength path uses a specific frequency that can only be changed with manual intervention and on-site visits including actions such as re-cabling cards.
- The capacity of each wavelength path can be any of 2.5 Gbps, 10 Gbps, 40 Gbps or 100 Gbps. The modulation schemes that can potentially be used vary from the traditional amplitude-based Non Return to Zero (NRZ) to the state-of-the-art coherent Polarization Modulation Quadrature Phase Shift Keying (PM-QPSK).

6.1.2 GÉANT Use Case

6.1.2.1 IP Services

P2P DWDM technology has been used in the GÉANT network and many other European NRENs since about 2005 to simply provide fixed trunks and access links (based on both packet-over-SONET/SDH and, at 10 Gbps, Ethernet framing) between the routers and switches that provide the IP services.

6.1.2.2 Fixed Circuit Service: Full Wavelength Capacity

P2P DWDM technology is currently used in the GÉANT network for providing the GÉANT Lambda service (again based on a mix of SDH and Ethernet framing, although increasingly Ethernet framing is proving more popular).

6.1.2.3 Fixed Circuit Service: Less than Full Wavelength Capacity

P2P DWDM technology can be used for providing guaranteed-capacity circuits of less than the full wavelength capacity. This can be accomplished via the use of muxponders, which are able to perform sub-wavelength multiplexing (usually via SDH or OTN) of several input signals and transmit the aggregate at a wavelength of the DWDM grid. Note that transmission cards with Ethernet multiplexing capabilities also appear on vendors' roadmaps. The main disadvantage of implementing the fixed circuit service via muxponders is that all input signals are obliged to have the same end points. This restriction could be overcome by introducing sub-wavelength switching capabilities at selected DWDM nodes; in this case each circuit can be routed independently as long as a set of wavelength paths are pre-installed and there is available bandwidth. It should be noted that there is at least one vendor offering sub-wavelength switching capabilities on its DWDM platform based on ODU-1; other vendors are planning to include this feature by 2011 or 2012 and extend the granularity down to ODU-0 (~1Gbps).

6.1.2.4 Dynamic Circuit Service

The dynamic circuit service has the same characteristics as the fixed circuit service, with the addition that a shorter lead time is expected (i.e. a few minutes) and demand is less predictable. Assuming that spare transponder/muxponder cards are already installed on the network, the following cases are possible:

- Cards are installed at the appropriate end points and a wavelength path has already been configured between the end points. In this case, wavelength path instantiation is straightforward and involves little or no lead time.
- Cards are installed at the appropriate end points but a wavelength path has not been configured between the end points. This is the typical case, since the cost of maintaining a set of any-to-any spare lit wavelength paths is not affordable. In this situation, in order to set up a wavelength path, card re-cabling and equipment configuration are required. This is usually implemented by a vendor's specialist engineers. Response time depends on the SLA defining the relations between contractors; however, it cannot be less than a few days.

Given the points outlined above, it may be concluded that using P2P DWDM technology for implementing the dynamic circuit service entails significant risk of failing to provide a service with an acceptable lead time. For this reason it is not recommended.

The lead time could potentially be minimised (e.g. reduced to the order of seconds depending on the size of the network) by including ROADMs with directionless and colourless features as well as an advanced control plane (e.g. GMPLS) in the network (as described in Section 6.2). In order to provide a dynamic circuit service of less than full wavelength capacity, the points raised in Section 6.1.2.3 would also apply in this case.

6.1.2.5 GÉANT Services Capacity

The current GÉANT network is implemented with point-to-point 10 Gbps wavelengths using NRZ modulation. The capacity of individual wavelengths can be further upgraded to 40 Gbps / 100 Gbps using advanced modulation schemes, such as coherent PM-QPSK. For most equipment vendors, 40 Gbps and 100 Gbps wavelength transmission using a flavour of PM-QPSK appears either in their current portfolio or on roadmaps for availability within 2011. It should be noted here that some vendors have clearly expressed their intention to bypass the coherent PM-QPSK variety of 40 Gbps and go directly to 100 Gbps wavelength capacity. However, transmission of 40 Gbps / 100Gbps has some requirement implications for the existing DWDM network – at least some of the following:

- Addition of new In-Line Amplification (ILA) sites that were previously skipped.
- Use of Raman amplification (not an absolute requirement per se but it helps improve OSNR which in turn helps to extend transponder reach).
- At some ILA sites, introduction of ROADM filters that include per wavelength Variable Optical Attenuators (VOAs) for gain equalisation purposes.
- Avoidance of G.655 fibre, which has degraded performance compared to G.652, when it comes to high-capacity phase-modulated signals.
- Removal of Dispersion Compensation Modules (DCMs), which are used for managing the accumulation of chromatic dispersion in a fixed way.

With reference to the avoidance of DCMs, it should be noted that although DWDM networks with 10 Gbps wavelengths require DCMs for managing chromatic dispersion in the *optical* domain, state-of-the-art coherent receivers (40 Gbps / 100 Gbps) are managing chromatic dispersion and other optical impairments (e.g. polarisation mode dispersion) in the *electronic* domain. Hence DCMs are not necessary for pure 40 Gbps / 100 Gbps DWDM transmission. Moreover, deployment of DCMs in 40 Gbps / 100 Gbps DWDM networks brings an additional burden to the optical design because:

- Span attenuation is artificially increased.
- The strategy of keeping the chromatic dispersion low over the network, which is extensively used in 10 Gbps DWDM networking via the deployment of DCMs, increases the impact of harmful non-linear effects, such as cross-phase modulation.

Given the points outlined above, it may be concluded that a DWDM network based purely on coherent transmission (40 Gbps / 100 Gbps) can be expected to have improved regeneration performance compared to a DWDM network used for coherent transmission *and* 10 Gbps wavelength transmission, both having the same

amplification technology and dimensioning. This finding leads to a trade-off: supporting a mixture of 10 Gbps, 40 Gbps and 100 Gbps wavelengths brings an additional regeneration cost to the DWDM network (for the 40 Gbps and 100 Gbps wavelengths). On the other hand, demand for 10 Gbps wavelengths is not expected to disappear in the next few years and it is not clear whether the prices of 100 Gbps / 40 Gbps transponders will be less than 10 / 4 times the price of a 10 Gbps transponder. Use of muxponders for carrying 10 Gbps services (10 Gigabit Ethernet or STM-64) over coherent wavelengths could be a way of managing this trade-off, but has the disadvantages already stated in Section 6.1.2.3.

6.1.2.6 *Interaction with Other Layers*

Data Plane

With regard to the data plane, the P2P DWDM network is a transparent layer that can carry traffic produced by any of the well-known network protocols.

Control Plane

Typically, transport equipment that implements P2P DWDM technology does not include an advanced control plane (e.g. GMPLS). This is because the technology's inherent inflexibility in dynamically creating/modifying paths reduces the value of using an advanced control plane.

6.1.3 **Market Analysis**

See Section 6.2.3.

6.2 **Reconfigurable Optical Add-Drop Multiplexers (ROADMs) and Wavelength Selective Switches (WSSs)**

This section describes all-optical devices – Reconfigurable Optical Add-Drop Multiplexers (ROADMs) and Wavelength Selective Switches (WSSs) – and their potential benefits for existing and future services offered by the GÉANT network for the NREN community.

Deliverable DJ1.2.1, “State of the Art Photonic Switching Technologies” [DJ1.2.1], produced by JRA1 T2, describes ROADMs, WSSs and many other photonic (optical) components and devices. It is used as the primary reference for the information in this section.

As stated above, this section does not describe in detail the technologies or the theory behind photonic equipment; rather it focuses on the practical implementation of relevant devices as preparation for the procurement process and on how they can apply to the GÉANT network and service portfolio (GÉANT IP, GÉANT Plus, GÉANT Lambda) with respect to 10 Gbps, 40 Gbps and 100 Gbps, and dynamic provisioning/protection/restoration.

However, underlying technologies do influence important characteristics like insertion loss, bandwidth (i.e. whether 40 Gbps / 100 Gbps signals can pass), channel grid (100 GHz, 50 GHz, 25 GHz, 12.5 GHz), switching speed, number of ports.

ROADMs are mentioned and briefly described in Section 6.1 *Point-to-Point Dense Wavelength-Division Multiplexing (P2P DWDM)* on page 45 – clear evidence these devices are fundamental because of their all optical nature.

6.2.1 Technology Overview

ROADMs and WSSs (investigated by JRA1 T2) differ from other technologies such as NG-OTN, EoMPLS and carrier Ethernet (investigated by JRA1 T1). ROADMs and WSSs are the basic all-optical building blocks of modern optical DWDM transmission systems and can transport and work with any of these data transport protocols. On the other hand, control plane and integration with other network components are essential to provide all advantages of all-optical equipment.

ROADMs and WSSs can be characterised as Optical-Optical-Optical (OOO) circuit switches or crossconnects (not packet switches in the L2/Ethernet sense, as there is no packet-header recognition). WSSs can be considered as a building block of ROADMs, so this section will discuss the features and benefits of ROADMs only.

Both ROADMs and WSSs are photonic, wavelength selective devices. They are analogue, like other photonic devices – optical amplifiers, dispersion compensators, wavelength convertors and so on. As far as the authors know, only one vendor produces “digital ROADMs”, as part of their Digital Optical Network offering.

WSSs switch DWDM channels between ports called composite or line ports. The switching is called lambda routing in some references. WSSs (or Photonic Cross Connects (PXC)) are used to build ROADMs.

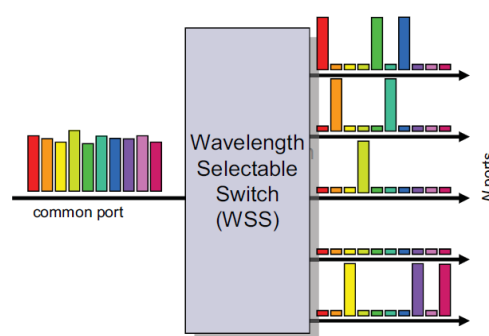


Figure 6.1: Wavelength Selectable Switch

ROADMs switch DWDM channels, like WSSs, but also provide adding and dropping of individual channels (one channel at a time) from DWDM transceivers like XFPs or XENPAKs. Usually there is one group of ADD inputs and DROP outputs together with “network” or “line” interfaces. The number of line interfaces determines the so-called “degree” of a ROADM. A 2-degree ROADM is illustrated in Figure 6.2.

Modern ROADMs can be truly multi-degree (up to 9 line interfaces is common; one vendor has announced 23; theoretical architectures for 256 have been published), “directionless” (i.e. ADD/DROP ports are not dedicated to one direction) and “colourless” (i.e. any wavelength can pass through any tributary port). The term “contentionless” is also often used. It refers to a situation when “same wavelengths carrying different information can be received/sent simultaneously with the help of multiple tunable transponders from/to different input/output fibre ports” [ROADM_Sandesha]. A directionless 2-degree ROADM is illustrated in Figure 6.3.

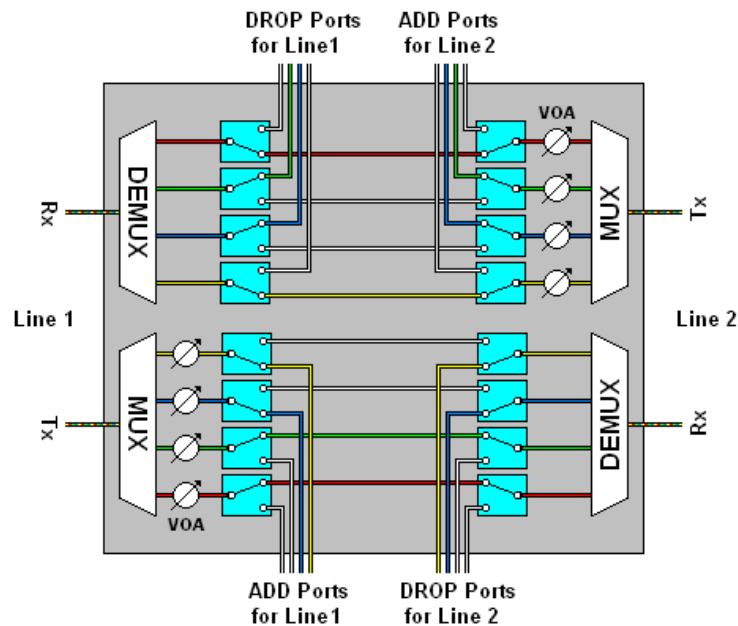


Figure 6.2: A 2-degree ROADM

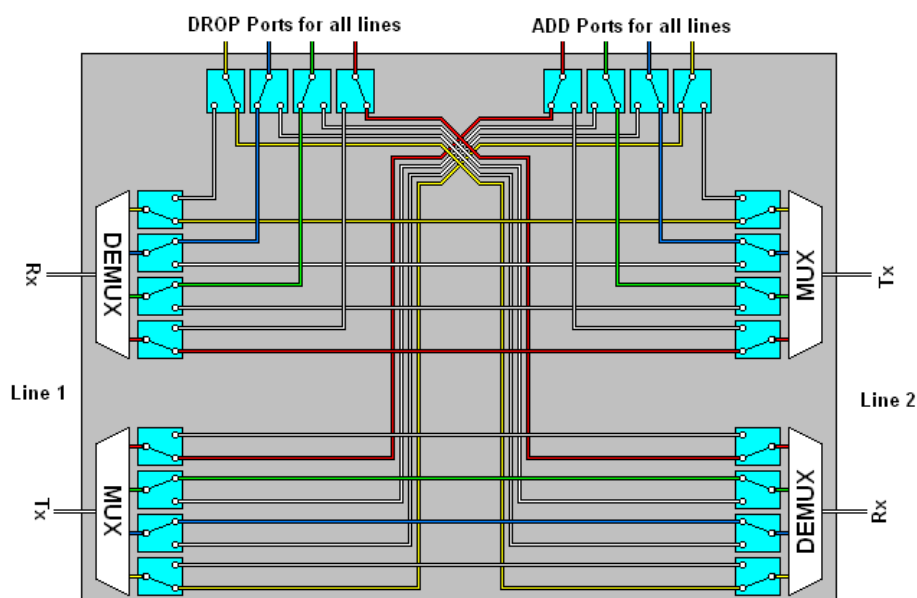


Figure 6.3: A “directionless” 2-degree ROADM

ROADMs (or WSSs) are deployed in modern optical networks to:

- Provide better scalability.
- Provide better flexibility when deploying new services.
- Provide automated connectivity (Bandwidth on Demand (BoD)) and multipoint connectivity.
- Reduce OPEX and CAPEX (can be true or not).

Some vendors claim their ROADMs provide network/protocol/speed transparency, which is true but is also an inevitable consequence of their all-optical nature, which removes the need for Optical-Electrical-Optical (OEO) conversions. All of these features could be considered important for enabling the GÉANT backbone to provide new services (for example, those considered in Section 3.3).

6.2.2 GÉANT Use Case

As mentioned above, ROADMs (and WSSs) are relevant to the GÉANT network(s) and services – i.e. GÉANT IP, GÉANT Plus and GÉANT Lambda services – both fixed and dynamic. Protection and restoration are important for all of these services, while quick and dynamic E2E provisioning is important for GÉANT Plus and GÉANT Lambda.

Only a few ROADMs have been deployed in the GÉANT backbone up to now and their limitations are clearly evident in certain areas, especially when provisioning new GÉANT Plus and GÉANT Lambda services (the GÉANT IP services are not affected to the same extent because they are more fixed).

Some of the benefits of using ROADMs within GÉANT apply to the full service portfolio; others are more important for GÉANT Plus and GÉANT Lambda. Those that apply to all GÉANT services are:

- Improved scalability, flexibility.
- Reduced network design complexity.
- Dynamic protection/restoration/disaster recovery.
- New optical paths for planned maintenance (no outages) – proactive possibilities.
- Load balancing.
- Support of mesh, multi-ring topologies with multi-degree devices – at least 4, but 9 desirable.
- Control plane (global) must be ready, integrated into operational support systems (OSSs).
- Lower power consumption.
- Lower cost (per bit).
- Higher degree of integration (improved port density).

The benefits that are more important for GÉANT Plus and GÉANT Lambda services are:

- No manual intervention needed when provisioning new channels or lambdas (e.g. no re-patching, re-attenuation for single wavelengths).
- True lambda/wavelength/circuit on demand, automated, quick and flexible (A2Z provisioning with only a few clicks).

- Multipoint connectivity (a.k.a drop & continue, familiar from SDH, true optical multicast) for HD/4K/8K video.
- Elimination of OEO.
- Protocol/service/bit rate transparency.
- Optical monitoring (e.g. Chromatic Dispersion (CD), Polarisation Mode Dispersion (PMD), Optical Signal-to-Noise Ratio (OSNR), pass band), not Bit Error Rate (BER) and other electrical parameters. Essential for lambda and alien lambda services.
- Possibility of integration with other all-optical devices in one linecard.

It is the authors' opinion that really useful and functional GÉANT Plus and GÉANT Lambda services (i.e. quick, on-demand, dynamic) can be implemented only with ROADMs with the advanced features discussed in the following paragraphs (to be colourless, directionless, etc.) and with a control plane implemented. If ROADMs do not have certain of these features, some benefits cannot be achieved; it is therefore very important to select ROADMs with the right features.

ROADMs can support only optical quality parameters (e.g. noise, distortion, power levels), not electrical parameters such as BER, errored seconds, link utilisation, or even parameters related to Quality of Service (QoS). This is no surprise and monitoring of optical parameters is very important in modern optical networks; reliance on BER (and similar parameters) is not enough.

ROADMs can be used with any other technology (e.g. CCTNTs such as OTN, cE, EoMPLS, SE) because of their all-optical nature and in this sense they are not dependent on any of them. ROADMs are the lowest layer in an optical transmission system and must be integrated; standalone ROADMs without a control plane can offer only limited advantages. As the lowest layer of a network, ROADMs interact with optical wavelengths only, totally agnostic to what data is being transmitted.

ROADMs can be used to build and deploy many network topologies, including mesh and multi-ring topologies. This is dependent on the degree of ROADMs. With degree 2 ROADMs, no multi-ring or meshing is possible. With degree 9, however, the possibilities are broad and sufficient to meet most requirements; with higher numbers (like 23 from one vendor), the possibilities are almost infinite. Alternative NREN access points (or adding additional PoPs) are enabled by the use of ROADMs' all-optical features and can be implemented very easily; even CBF connections can be implemented in a similar way.

One thing that may need to be taken into account (although this looks a bit further into the future) is to ensure that any ROADM-based solutions that are deployed will support variable/flexible wavelength grids (sometimes referred to as "FlexiGrid").

6.2.3 Market Analysis

This section presents information gathered from the RFI regarding ultra-high-capacity transmission (40 Gbps /100 Gbps), ROADM availability and control plane availability on optical transmission platforms.

These technologies are all-optical – in other words, analogue. Only one vendor manufactures and promotes so-called digital optical networking and equipment. Some vendors offer a variety of optical amplifiers (Erbium

Doped Fibre Amplifier (EDFA), Raman) and other components like compensators of chromatic (CD) and polarisation (PMD) dispersion.

All vendors support P2P DWDM and ROADMs technologies, although there are some important differences that have to be taken into account, for example, the degree of ROADMs, maximum optical reach, control plane developments, and availability of 100 G interfaces.

The information obtained from the RFI responses may be summarised as follows:

- Point-to-point DWDM transmission is implemented by all vendors.
- Industry converges on the implementation of coherent polarisation multiplexed transmission for 40 Gbps / 100 Gbps wavelengths using PM-QPSK. Relevant 50 GHz spaced transponders are either already available or will be available during 2011. It should be noted that some vendors state that they have changed their plans for 40 Gbps wavelengths implementation and will skip directly to 100 Gbps implementations.
- 40 Gbps /100 Gbps transmission places more stringent requirements on the optical domain, especially when coupled with 10 Gbps wavelength transmission, leading to degrading of regeneration performance and Raman amplifiers usage. Some vendors state that they plan to release new amplifier families in order to improve 40 Gbps / 100 Gbps transmission performance.
- ROADMs are currently available from most of the vendors with degree at least 8. Moreover, ROADMs with integrated colourless and directionless features are either already available or appear on vendors' roadmaps for release during 2011.
- An advanced control plane (e.g. GMPLS) for the optical transmission platform is either already available or appears on vendors' roadmaps for 2011, in most cases. It should be noted that the RFI responses did not provide very specific information on the control plane implementation.

6.2.4 Summary

ROADMs are beneficial and perhaps essential to all GÉANT services, especially for the dynamic services because of their all-optical nature and truly reconfigurable capabilities. ROADMs can offer the following benefits, which cannot be achieved with other components:

- Very good scalability and flexibility compared to old solutions with fixed OADM.
- Dynamic protection/restoration/disaster recovery.
- Proactive possibilities for planned maintenance – no outages.
- Automated bandwidth-on-demand services.
- Reduced network design complexity.
- Protocol/service/bit rate transparency: 10 G / 40 G or even 100 G Ethernet or OTN signals.
- Optical monitoring, which is important for new all-optical services.
- Support of mesh, multi-ring topologies.
- Higher degree of integration (improved port density).
- Lower power consumption.
- Lower cost.

There are many ROADMs available on the market, integrated in modern optical DWDM transmission systems, with many diverse features. ROADMs are not dependent on other technologies; rather they serve as the basic building blocks at the lowest layer.

6.3 Next-Generation Optical Transport Network (NG-OTN))

6.3.1 Introduction

This section presents the key features of Optical Transport Network (OTN) and Next-Generation Optical Transport Network (NG-OTN), which are seen as possible technologies in the evolution of digital and analogue backbone transmission for both the commercial and NREN community, providing support for Time-Division Multiplexing (TDM), packet and IP services.

It then outlines how these technologies can support the GÉANT network architecture and services.

6.3.2 Optical Transport Network (OTN)

The OTN architecture concept was initially developed by the ITU-T a decade ago, to build upon the SDH and DWDM experience and provide bit rate efficiency, resiliency and management at high capacity. OTN therefore looks a lot like SONET/SDH in structure, with less overhead and more management features. It does, however, bring many developments and advantages over SDH:

- **Transparent Client Signals.**
This allows the end user to view exactly what was transmitted at the far end and decreases the complexity of troubleshooting.
- **Better Forward Error Correction.**
OTN has increased the number of bytes reserved for Forward Error Correction (FEC), allowing a theoretical improvement of the Signal-to-Noise Ratio (SNR) by 6,2 dB. This improvement can be used to enhance the optical systems in the following areas:
 - Increase the reach of optical systems.
 - Increase the number of channels in the optical systems.
 - Ease the introduction of transparent optical network elements, such as OADMs, Photonic Cross Connects (PXCs), splitters, etc.
- **Better scalability.**
The old transport technologies like SONET/SDH were created to carry voice circuits, which is why the granularity was very dense – down to 1,5 Mb/s. OTN is designed to carry a payload of greater bulk, which is why the granularity is coarser and the multiplexing structure less complicated.
- **Tandem Connection Monitoring.**

The introduction of additional Tandem Connection Monitoring (TCM) combined with the decoupling of transport and payload protocols allow a significant improvement in monitoring signals that are transported through several administrative domains, e.g. a meshed NREN topology where the signals are transported through several other NRENs before reaching the end users.

In a multi-domain scenario – “a classic carrier’s carrier scenario” where the originating domain can’t ensure performance or even monitor the signal when it passes to another domain – TCM introduces a performance monitoring layer between line and path monitoring allowing each involved network to be monitored, thus reducing the complexity of troubleshooting as performance data is accessible for each individual part of the route.

Finally, a major drawback with regards to SDH is that a lot of capacity during packet transport is wasted in overhead and stuffing, which can also create delays in the transmission, leading to problems for the end application, especially if it is designed for asynchronous, bursty communications behavior. This over-complexity is probably one of the reasons why the evolution of SDH has stopped at STM 256 (40 Gbps).

OTN has all the capabilities required to monitor, manage, and control each client signal transported on a per wavelength basis in the network. In this way, OTN adds operations, administration and maintenance (OAM), and provisioning and troubleshooting functionality to optical carriers.

6.3.2.1 ODUflex

The current trend is that DWDM systems are enhanced – or have an equipment add-on – to be capable of handling digital OTN switching or grooming in direct integration with their traditional analogue handling of OTN. In addition, packet and TDM, as edge technologies, are groomed into OTN using ODUflex, as shown in Figure 6.4 below. In this way, DWDM and transport systems converge into a single DWDM/transport unit.




Figure 6.4: Structure of DWDM/transport architecture option

This can be described more precisely as sub-port level grooming, which offers a high level of grooming flexibility and optimises transport efficiency for networks with a significant volume of point-to-point traffic in the service mix. The sub-port grooming option allows maximum grooming flexibility by enabling VLANs or pseudowires within a port to be logically mapped to optical links using ODUflex.

Sub-port level grooming also enables finer granularity by supporting sub-port interfaces or virtual interfaces such as VLANs, which means that ports do not have to consume the full capacity of a physical port or wavelength. Different VLANs from the same transport node are mapped to different virtual containers through the transport network to different destinations using ODUflex and VLAN shaping.

ODUflex is a new technology that allows grooming of traffic between optical transport equipment and routers in a manner that efficiently addresses incremental bandwidth growth in steps as granular as 1 Gbps. Carriers no longer have to allocate a full ODU container to each connection, but rather can increase capacity in increments for connections that require them. Through control plane integration, dynamic bandwidth adjustments can be made to optimise the network as needed over time. Further opportunities for enhancing overall network resiliency and fault isolation also emerge from control plane integration.

NG-OTN will provide improvements to service-layer networking efficiency, protection and restoration functions, scalability and flexibility.

6.3.3 Implications for GÉANT

6.3.3.1 Network Service Support

GÉANT Lambda Service

The GÉANT Lambda service is already based on OTN on the DWDM line side and is therefore supported in the current and future GÉANT network. In the context of considering OTN technology, the GÉANT Lambda service in a multi-domain environment would be supported by using transponders with OTN support *on the client side*.

GÉANT Plus Service

The GÉANT Plus Ethernet service is supported by the NG-OTN service. Ethernet is mapped into ODUx and is switched and carried across domains. For further information please see [DJ1.1.1].

GÉANT IP Service

The GÉANT IP service is supported by the NG-OTN service. For further information please see [DJ1.1.1].

6.3.4 Market Analysis

This section outlines equipment functionality capable of switching and multiplexing ODUx including ODUflex. Point-to-point ODUx transport is treated in sections 6.1 and 6.2.

From the RFI study it has been identified that several vendors have recently launched digital OTN switches ranging from Gigabit to Terabit switching capacity.

The most common available functionality is:

- TDM support.
 - Full non-blocking OTH switching and multiplexing.

- Full non-blocking SDH/SONET switching and multiplexing.
- Line rates from 155 Mbps to 100 Gbps.
- Packet support.
 - 1 GbE, 10 GbE, 100 GbE.
 - Note that this support is more mature for point-to-point and “switching on a blade” rather than in terms of device-wide packet switching.
- WDM Support.
- Close integration with WDM systems either as an independent switch or as a subsystem in a DWDM system.
- Control plane for TDM and packet.

OTN switches support the following building blocks as defined in Section 5 *Building Blocks* on page 38:

- Block B – as the glue between domains when interconnecting on the OTN level. Especially suited to larger interconnects where several domains interconnect, e.g. in Hamburg where DFN, NORDUnet, PSNC and SURFnet are present.
- Block C – matches the Carrier Class Transport infrastructure requirements with excellent integration into a DWDM or grey DWDM infrastructure. In addition, supports packet, TDM or IP/MPLS traffic flows.
- Block D – same as block C. However the best match is when the grey transport infrastructure is OTN-based. If that’s not the case, it will be more relevant to find a device matching the grey transport network.

The OTN switches support the following GÉANT services:

- GÉANT Lambda. The full wavelength service is supported, but would typically be provided by a native DWDM point-to-point link.
- GÉANT Plus.
- GÉANT IP as a transport for trunks between IP routers and switches. OTN devices do not have any kind of IP functionality.

6.4 Ethernet over Multi-Protocol Label Switching (EoMPLS)

The main purpose of this section is to explain how EoMPLS can be used in the GÉANT network to implement transmission services.

6.4.1 Technology Overview

Ethernet over MPLS (EoMPLS) is one of the technologies that can be used to transport Ethernet frames over a provider’s backbone network. Details of how EoMPLS works, key concepts and definition of terms can be found in Section 3.5 of the GN3 deliverable DJ1.1.1 “Transport Network Technologies – Study” [DJ1.1.1], prepared by JRA1 T1.

6.4.1.1 *Advantages and Disadvantages*

The main strengths and advantages of EoMPLS technology are:

- Flexibility – The service can be adapted to user needs (P2P, any-to-any or P2MP transmission).
- Security – End-user traffic isolation (based on inner label).
- Network security – For P2P connections there is no need to enable MAC learning, which protects the core network against loops on the customer side.
- High resilience – Based on MPLS mechanisms (Standby Path, Fast Reroute).
- QoS assurance – Based on TE and queueing technologies.
- Interoperability – Solution based on RFC standards.
- Technology independence – Ethernet interfaces are currently most popular, but any other technology can be used in the MPLS core to carry EoMPLS traffic.
- Scalability – Support for hierarchy based on labels (not on frame-in-frame encapsulation).
- Multi-service network – Most MPLS switches that support L2 VPNs also support IP/IPv6 routing.
- The same QoS scheme for IP/IPv6 and MPLS-based services.

The main disadvantages or weak points of EoMPLS are:

- Maintenance overhead – Some more maintenance tasks are required in order to deploy and maintain the MPLS network.
- Higher cost of network devices compared to simple L2 Ethernet switches.

6.4.2 **GÉANT Use Case**

6.4.2.1 *Potential Topology Developments*

As a general requirement, the network topology should provide at least two independent paths to each node in the network. Moreover it is recommended that the topology allows at least two paths between each pair of network nodes to be established using different resources (from the physical point of view).

The most efficient way to implement EoMPLS services in the GÉANT network is to use IP/IPv6/MPLS routers. This allows the building of a single network infrastructure providing multiple services.

It is also possible to use separate devices for EoMPLS services and for IP/IPv6 routing. This solution might be more attractive from a financial perspective in the GÉANT environment, even though an additional MPLS switch would be required in each Point of Presence (PoP). This option requires further financial analysis. However, maintaining two separate infrastructures for IP/IPv6 and MPLS-based services introduces maintenance overhead.

6.4.2.2 Service Mapping

For NRENs connected to the GÉANT network with Gigabit- or 10 Gigabit-Ethernet interfaces, EoMPLS can be used to provide the GÉANT Plus service and services based on GÉANT Plus that might be offered in the future (such as Bandwidth on Demand or Fixed Dedicated Capacity).

Transmission Services

EoMPLS in general allows two kinds of service to be provided:

- Point-to-point transmission of Ethernet frames (Virtual Private Wire Service (VPWS)). This maps to GÉANT Plus services.
- LAN emulation (Virtual Private LAN Service (VPLS)). GÉANT does not at present offer such services. However, the development plans of SA2 and JRA2 may make use of this functionality.

Depending on the functionality implemented on the switching device, additional options are available regarding the type of transferred Ethernet frames:

- Untagged frames (with or without adding VLAN Identifier (VID) at one end).
- Single tagged frames (with or without VID rewrite).
- Double tagged frames (with or without outer VID rewrite).

Moreover, for distributing multicast traffic the VPLS service can provide point-to-multipoint (P2MP) transmission over dedicated P2MP LSP in order to optimise traffic flow through the MPLS backbone (Virtual Private Multicast Service (VPMS)).

Service Quality Parameters

Service quality parameters are the same as in packet switched networks with additional traffic engineering functionality offered by Resource Reservation Protocol – Traffic Engineering (RSVP-TE).

In general, QoS for Layer 2 MPLS-based VPNs (EoMPLS) can be deployed using two kinds of mechanisms:

- Packet queuing on physical interfaces.
- Traffic flow path selection.

Packet queuing on physical interfaces covers all technologies relating to the differentiation of traffic transmitted over the same physical interface (like queuing, Weighted Random Early Detection (WRED), Hierarchical QoS (H-QoS), etc.).

Traffic flow path selection implements traffic engineering (TE) and covers:

- Path selection based on different constraints (available bandwidth, hop limit, etc.).
- Resource reservation and priority for established path.
- Traffic to path mapping based on QoS policies (and/or packet marking).

Usually TE policies are implemented by the RSVP-TE protocol and QoS-related parameters are mapped to the **EXP** field in the MPLS packet.

Fixed and Dynamic Services

Currently there is no widely used UNI that could be recommended for establishing circuits dynamically upon end-user request. In most cases, all EoMPLS circuits are established as fixed entities. The path between end points can, however, be established dynamically.

Dynamic EoMPLS services (provisioned on demand) can be provided using third-party automatic provisioning tools (like AutoBahn).

Moreover, multi-domain connectivity with EoMPLS can be implemented by stitching EoMPLS instances provisioned in multiple domains.

6.4.2.3 Capacity Forecast

Currently EoMPLS can be provided using GE or 10 GE as the access interface. This means that in some cases 100 Gbps core interfaces may be required to provide appropriate capacity for the core network.

Link aggregation in the core is not recommended because of the different hashing technologies used by each vendor and the possibility of exceeding the capacity of the participating interface by a single data stream.

6.4.2.4 Functionality Required in the GÉANT Network to Make Use of EoMPLS

EoMPLS implemented with TE mechanisms requires an underlying inter-domain routing protocol with TE extension. This functionality is widely implemented by all leading vendors.

EoMPLS is usually implemented on high-end routers. The routers (switching devices) require an underlying transmission system in order to provide long-distance transmission.

Moreover, an underlying transmission system is also required in order to deliver the GÉANT network services portfolio with GÉANT Lambda services and GÉANT Plus services for NRENs connected to GÉANT with SDH interfaces.

6.4.3 Support for Network Services at Other Layers

EoMPLS can be used as a transmission medium for other network devices equipped with Ethernet interfaces. This means it can connect routers, switches and servers using Ethernet interfaces.

In most cases, the devices implementing the EoMPLS service allow IP/IPv6 routing as well. This allows a multi-service network to be built on a single hardware platform. The EoMPLS service can be used to provide connectivity between NRENs (for example, direct IP/IPv6 exchange).

The convergent architecture makes it possible to use the same redundancy scheme for IP/IPv6-based services and for EoMPLS circuits.

The layered model of the GÉANT network is presented in Figure 6.5. The network requires three main elements:

- Switching devices (IP/IPv6 and MPLS).
- Optical Transport Network (OTN) (transport).
- Physical medium (fibre).

The service-to-resource mapping is also shown in Figure 6.5. Moreover the diagram shows the relations between the GÉANT services and the network infrastructure element used to implement those services.

As noted before, the switching devices (Switching layer) require an underlying transmission system (OTN layer). The OTN layer requires a physical medium (Physical layer).

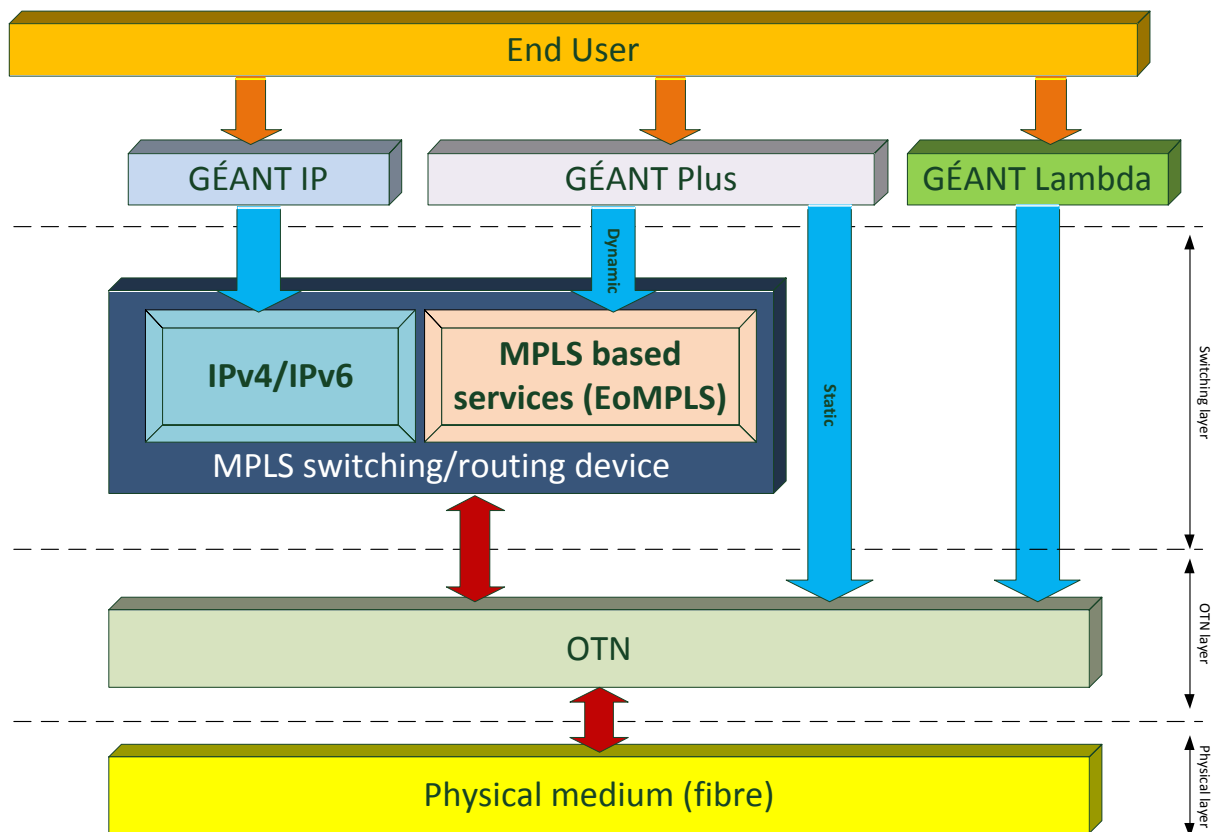


Figure 6.5: GÉANT network layers

6.4.4 Market Analysis

This section summarises RFI response(s) related to IP and MPLS switching technologies. Based on the RFI, the support for GÉANT services is shown.

6.4.4.1 IP Services

Three kinds of hardware platforms were proposed in order to build the network providing IP service:

- High-performance IP core router.
- Multi-service router.
- Core MPLS switch (Label Switch Router (LSR)).

The high-performance IP core router is a hardware platform designed for “carrier of carriers” networks. The devices are 100 Gbps ready with slot capacity up to 200 Gbps. They are able to support many kinds of network interfaces (Packet over SDH (PoS), Ethernet, etc.), but as a consequence are more expensive.

Multi-service routers are designed to provide not only IP services but also MPLS-based Layer 2 VPNS (VLL/VPLS). In most cases the devices are equipped with Ethernet interfaces (1 GE, 10 GE, 100 GE). The devices allow transiting the IP or MPLS traffic and service termination and aggregation. Because of unified architecture (focused on Ethernet interfaces/switching) they are less expensive than IP core routers.

The core MPLS switch (LSR) is an MPLS-optimised platform which has been designed specifically to address the scale at which next generation networks are growing. The platform has been optimised for speed whilst maintaining network economics. This has been achieved by placing the emphasis on MPLS forwarding while still providing the full MPLS/IP control plane. In this manner the platform does not need to perform costly routing table IP lookups but can rely upon more efficient MPLS label lookups.

As a consequence of the proposed platforms, the GÉANT IP service can be provided in different ways:

- Without change to the model currently used in the GÉANT network).
- In a hierarchical network topology.

The existing GÉANT network can be upgraded with new devices to continue providing IP services at higher capacity levels. In this case, high-performance IP core routers or multi-service routers can be used. However, using multi-service routers is more economically efficient, because it allows other services (such as GÉANT Plus) to be provided on the same platform.

The proposed platforms allow two types of hierarchical network topology to be built:

- 2-level hierarchy.
- 3-level hierarchy.

Using a 2-level hierarchy, the services are terminated at multi-service routers. The multi-service routers are then connected to IP core routers. The IP core routers are the core of the network, which will be optimised for efficient switching and redundancy.

The 3-level hierarchy introduces dedicated MPLS switching levels in the core network. The IP core routers (and possibly multi-service routers) will mark packets with MPLS labels. Based on those labels, the MPLS core switches will perform high-efficiency and low-cost packet switching. However, this solution introduces more maintenance and design complexity, which makes it not economically justified.

6.4.4.2 *Fixed Circuit Service: Full Wavelength Capacity*

No change compared to current provisioning methodology is proposed.

6.4.4.3 *Fixed Circuit Service: Less than Full Wavelength Capacity*

All the principles mentioned for IP services (see Section 6.4.4.1) apply to fixed circuit services requiring less than full wavelength capacity, with the following additional remarks:

- Access interface (for the end user) should use Ethernet technology.
- The circuits must be provisioned over core links that are not over-subscribed, with bandwidth usage of less than 70% (in order to avoid statistical multiplexing).

6.4.4.4 *Dynamic Circuit Service*

All principles mentioned for fixed circuit services requiring less than full wavelength capacity (see Section 6.4.4.3) apply to dynamic circuit services.

In order to provide more interactivity and robustness in dynamic service provisioning, the dynamic provisioning tool could be used.

6.4.4.5 *GÉANT Services' Capacity*

A wide range of core and access interfaces is available:

- 1 GE (optical Small Form Factor Pluggable (SFP)-based and/or copper).
- 10 GE (optical 10 Gigabit Small Form Factor Pluggable (XFP)/SFP+ based).
- 100 GE (optical 100 Gigabit Small Form Factor Pluggable (CFP)-based).

It is not currently clear whether services can be terminated at 100 GE interfaces or not. An initial assumption should be made that early 100 GE interfaces will have some service termination limitations.

6.4.4.6 Control Plane Issues

For IP services, fixed circuit services requiring less than full wavelength and dynamic services, the well-known and standard protocols are used (IS-IS, OSPF, Multiprotocol Border Gateway Protocol (MBGP), RSVP, Traffic Engineering extensions to IGPs). No issues are expected within single autonomous system signalling.

Currently some limitations exist with multi-domain signalling for fixed circuit services requiring less than full wavelength and dynamic services; an external provisioning tool can be used as a workaround.

6.4.4.7 Protection/Restoration

The GÉANT IP Service uses protection and restoration basing on IGP protocols (e.g. IS-IS, OSPF) and supporting mechanisms (such as Bi-directional Forwarding Detection (BFD)). This allows IP traffic to be rerouted to the backup link in case of failure. The switchover time depends on the network topology and on the detailed parameters of the protocols (such as the values for **hello** or **hold** timers).

Fixed circuit services requiring less than full wavelength and dynamic services use protection and restoration mechanisms available in MPLS networks:

- Global protection (standby LSP).
- Local protection (Fast Reroute).

These MPLS-based protection and restoration mechanisms can also be used for IP services when IP over MPLS switching is used.

6.4.5 Summary

EoMPLS technology provides transparent transmission of Ethernet frames between service end points.

The EoMPLS service can be enabled on the same devices that are used for IP/IPv6 routing. For this reason, the GÉANT IP and GÉANT Plus services can be implemented on the GÉANT network using the same infrastructure. It is important to note that EoMPLS can be used to implement only the GÉANT Plus service. However, in the case of GÉANT, a separate device to run EoMPLS technology to deliver GÉANT Plus services may be more attractive financially.

QoS for EoMPLS can be implemented not only in terms of defining traffic priorities, but also in terms of resource reservation provided by Traffic Engineering mechanisms.

The packet switched based services can be deployed in GÉANT network in few ways. Proposed platforms can be easily adapted to GÉANT network requirements. High capacity platforms and high speed interfaces are available right now. 100 Gbps interfaces became available soon (1H11).

6.5 Carrier Ethernet

6.5.1 Introduction

This section introduces “carrier Ethernet” (cE) and considers the options for its use in GÉANT.

In the context of the discussion below, carrier Ethernet is a general term that encompasses several technologies. Generally it refers to the extension of the functionality of Ethernet service support to carrier grade operation and includes traditional Ethernet with carrier grade extensions such as OAM, routing and QoS. Similarly, Provider Backbone Transport (PBT, also known as Provider Backbone Bridge Traffic Engineering (PBB-TE)) and MPLS-TP are included in this discussion. EoMPLS, however, is not included here; it is discussed in a separate section (the previous section, 6.4 on page 57).

To create a definition of wide-area Ethernet services the MEF has identified the following service categories:

- **E-Line:** a service connecting two customer Ethernet ports over a wide area network (WAN). Two E-Line services are defined: Ethernet Private Line (EPL) and Ethernet Virtual Private Line (EVPL).
- **E-LAN:** a multipoint service connecting a set of customer end points, giving the appearance to the customer of a bridged Ethernet network connecting the sites.
- **E-Tree:** a multipoint service connecting one or more roots and a set of leaves, but preventing inter-leaf communication.

All these services provide standard definitions of characteristics such as bandwidth, resilience and service multiplexing, allowing customers to compare service offerings and facilitate Service Level Agreements (SLAs). MEF does not specify how Ethernet services are to be carried; this may be done using various technologies.

Traditional Ethernet has a number of limitations when operating in the WAN. For example, it has no ability to provide QoS guarantees; bridge and spanning tree do not scale in large networks; it lacks OAM functionality such as protection and performance monitoring. A range of extensions to Ethernet have evolved to fill all or parts of these gaps. These include:

- **Ethernet OAM, QoS & >10G.** OAM functions simply make the technology more appealing from an operational perspective, especially where many service instances are provisioned, QoS extends the range of (data plane) service levels that can be supported (more than a single “best effort” service) and support for >10 G interfaces (link-bonded $n \times 10$ GE, 40 GE and 100 GE) clearly makes the technology more appealing for deployment in high-capacity backbones. (Further details can be found in Section 3.2 of the GN3 deliverable DJ1.1.1 “Transport Network Technologies – Study” [DJ1.1.1].)
- **Ethernet Layer-2 routing, SPB, etc.** These extensions mainly allow a meshed and ring-based Ethernet network to be more efficiently utilised from the perspective of bandwidth utilisation (than would otherwise be the case in the presence of Spanning Tree Protocol / Rapid Spanning Tree Protocol (STP/RSTP)) and are also arguably more suited to supporting multicast traffic profiles. (For further details, see DJ1.1.1 Section 3.3 [DJ1.1.1].)

- **Ethernet congestion notification and priority-based flow control.** These protocols, respectively IEEE 802.1Qau and IEEE 802.1Qbb, are useful for providing storage-over-Ethernet services. They are currently still in the standardisation process but many vendors of Fibre Channel over Ethernet switches rely on these protocols. Some within the GÉANT community expect that there will be user demand for these types of applications during the lifetime of GN3 project.

In addition to these Ethernet enhancements, there are two new protocols that standardise a packet-oriented transport technology: **Multi-Protocol Label Switching Transport Profile (MPLS-TP)** (see DJ1.1.1 Section 3.6) and **Provider Backbone Bridge Traffic Engineering (PBB-TE)** (see DJ1.1.1 Section 3.7). These are intended to appeal to large-scale carriers and are more suited to carrying data-centric services than traditional telecom transport technologies like Plesiochronous Digital Hierarchy / Synchronous Digital Hierarchy (PDH/SDH).

In order to support the full suite of services, it will be necessary to support multiple protocols since different services have different requirements that cannot be addressed by a single, universal approach to network design. The added flexibility provided by the wide range of services offered by carrier Ethernet allows NRENs to design the network according to their own design guidelines and supports network virtualisation.

6.5.2 Technology Overview

6.5.2.1 Ethernet Extensions

This section summarises the new protocols that have been added to the IEEE standards 802.1 and 802.3 to make Ethernet suitable for use as a carrier grade transport protocol, namely, OAM, OSS, QoS and Layer-2 routing.

OAM

Ethernet OAM provides end-to-end service management. This is important in the multi-operator environment of GÉANT. Ethernet OAM will allow NRENs to create, monitor, and troubleshoot Ethernet links and services in a standardised fashion. The following protocols are the required building blocks for Ethernet OAM:

- IEEE 802.1ag: Connectivity Fault Management (CFM) [IEEE 802.1ag].
- ITU-T Y.1731: OAM functions and mechanisms for Ethernet-based networks [ITU-T Y.1731].
- IEEE 802.3ah: Ethernet Link OAM (EFM OAM) [IEEE 802.3ah].
- MEF 16: Ethernet Local Management Interface [MEF 16].

Compatibility between NREN OAM and GÉANT OAM is desirable but not necessary. If we consider GÉANT Plus services to be conventional “carrier’s carrier” services, then the NREN OAM would normally be carried transparently over GÉANT and GÉANT would not view the carrier’s OAM. The quality of GÉANT maintenance may be improved if GÉANT Plus services and NRENs agree on a common OAM method. For example, if GÉANT uses OTN OAM, to achieve OAM compatibility the NRENs would need to implement OTN OAM also.

Operational Support System (OSS)

An additional requirement for carrier class operation is a fully featured Operational Support System (OSS). This GUI-based interface will allow rapid service deployment and maintenance, and is needed for:

- Path creation/deletion/configuration.
- Alarm diagnosis.
- Accessing performance monitoring data.

QoS

All of the technologies are inherently based on statistical multiplexing of packets. If used to provide GÉANT Plus services, this represents a change from the TDM-based technology currently used and would result in a so-called “statistical multiplexing gain”, i.e. more efficient use of trunk capacity. The down side of statistical multiplexing is the risk of packet loss due to congestion. To ensure that services have guaranteed throughput, Ethernet needs quality of service protocols.

The IEEE 802.1D-2004 specification [IEEE 802.1D-2004] Annex G “User priorities and traffic classes” describes traffic classes and gives a foundation for “soft” QoS, where some traffic classes receive better treatment than others.

IEEE is preparing some standards to define the implementation of hard QoS in Ethernet:

- IEEE 802.1Qat: Stream Reservation Protocol (SRP) (Draft) [IEEE 802.1Qat].
- IEEE 802.1Qav Forwarding and Queuing Enhancements for Time-Sensitive Streams (also known as Queuing and Forwarding Protocol (QFP), approved in December 2009) [IEEE 802.1Qav].

The SRP protocol is similar to RSVP: it is able to book enough resources along the traffic path to guarantee the required QoS metrics. QFP defines algorithms for QoS mechanisms, including ingress metering, priority regeneration, shaping and time-aware queue draining.

Routing

Layer-2 routing is designed specifically for Layer 2 devices such as Ethernet. These build on the well-known L3 routing protocols. There are currently several initiatives that aim to develop the Layer 2 routing protocol:

- Provider Link State Bridging (PLSB): proprietary routing protocol from Nortel, proposed to the IEEE 802.1aq Shortest Path Bridging Working Group [NortelPLSB].
- Shortest Path Bridging (SPB): from the IEEE (802.1aq). This project group is working on two different flavours of SPB: SPB VID (SPBV) and SPB MAC (SPBM) [IEEE 802.1aq, SPB IS-IS].
- Transparent Interconnection of Lots of Links (TRILL): from the IETF [IETF TRILL].

In summary, several new protocols that enhance Ethernet to provide carrier grade performance have either been recently released or are still in standardisation. Vendors may offer some combination of these protocols to offer an Ethernet-based carrier grade solution. However, not many such solutions have been offered in the RFI process.

6.5.2.2 PBT

PBT, which extends Ethernet to achieve carrier class transport, is standardised in IEEE. PBT is based on the layered VLAN tags and MAC-in-MAC encapsulation defined in the PBB standard IEEE 802.1ah [IEEE 802.1ah], but it differs from PBB in eliminating flooding, dynamically created forwarding tables, and spanning tree protocols. PBB-TE OAM is based on IEEE 802.1ag [IEEE 802.1ag].

PBT is a technology that supports connection-oriented deterministic paths with PBB framing and fast-path protection switching technique.

To achieve these objectives, PBT makes the following changes in PBB functionality:

- Switches off learning of backbone MAC addresses.
- Switches off STP functionality of backbone switches.
- Uses a combination of B-VID and B-DA as a path label.
- Assumes a manual or Network Management System (NMS)-based building of forwarding tables.

Momentum on PBT seems to have stalled with the demise of Nortel Networks.

6.5.2.3 MPLS-TP

Unlike Ethernet and PBT, which are standardised in IEEE, MPLS-TP is standardised in IETF operating as a joint working team with ITU-T. MPLS-TP is intended to be a low-cost technology. This is to be achieved by modifying existing MPLS RFCs; features that are not relevant are removed and new extensions are added to support carrier transport functionality are added. As with the protocol additions to Ethernet and PBT, MPLS-TP is also designed to support E-Line, E-Tree and E-LAN services.

MPLS-TP is a profile of IP/MPLS designed to meet transport network operational requirements. It takes key elements from IP/MPLS such as MPLS / Pseudowire Emulation Edge to Edge (PWE3) architecture and forwarding mechanisms, and, optionally, GMPLS control plane, and provides additional functionality such as performance monitoring, OAM, Tandem Connection Monitoring (TCM), protection switching and ring protection.

The main characteristics of MPLS-TP are:

- Connection oriented: it uses Label Switched Paths (LSPs) and Pseudowire (PW) to deliver point-to-point, point-to-multipoint and multipoint-to-multipoint services.
- Client agnostic: it is able to carry any type of client traffic such as Asynchronous Transfert Mode (ATM), SDH.
- Transport layer agnostic: it can run over Ethernet, OTN, etc.
- Provides OAM and protection mechanisms.
- Network provisioning via centralised Network Management System (NMS).

MPLS-TP is being pursued by a group of telecommunications equipment vendors, including Alcatel-Lucent and Huawei.

6.5.3 GÉANT Use Case

6.5.3.1 GÉANT Lambda

The GÉANT Lambda service is currently delivered using the DWDM equipment and provides 10 Gbps wavelength services. It is expected that this service will continue to be provided with DWDM equipment.

6.5.3.2 GÉANT Trunks

All the technologies under this heading will be suitable to provide trunks (in the form of “virtual circuits” of one form or another) to the IP network as an overlay. The term “virtual circuits” means VLANs, LSPs and pseudowires (PWs).

6.5.3.3 GÉANT Plus

As described in Section 3.1.2 on page 22, the GÉANT Plus service provides point-to-point circuits with capacity 155 Mbps to 10 Gbps. The interconnect protocols offered are both SDH and Ethernet. In the case of Ethernet, the service types supported are Ethernet Private Line (EPL) and Ethernet Virtual Private Line (EVPL) in accordance with ITU-T G.8011.1/2 [ITU-T G.8011.1, ITU-T G.8011.2]. GÉANT Plus is currently delivered using Alcatel’s 1678MCC SDH platform.

Given the ubiquity of Ethernet in NRENs’ networks, users are increasingly interested in connection to other networks over Ethernet line services. In recent years, GÉANT Plus service requests have been almost exclusively for Ethernet rather than SDH services. Given our experience in GÉANT2, TDM service (e.g. SDH or OTN) are not considered a high-priority service for GÉANT Plus (though these may be required for GÉANT Lambda service). In fact, there has been a discussion about categorically **not** supporting non-Ethernet based GÉANT Plus service instances in the future (e.g. SAN-related services, conventional low-capacity private line services like those in the PDH and SDH hierarchies). Broadly speaking, pseudowires need to be used and there has been little or no experience in DANTE of providing non-Ethernet services over pseudowires.

While in GN2 and to date in GN3 EPL and EVPL services have been considered sufficient, we should aim to upgrade the GÉANT Plus platform to a technology that is capable of delivering a wider range of services. A flexible platform that can also deliver the E-LAN, E-Tree and IPTV/VoD services is desirable.

The use of the frame-based carrier class transport technologies described in this section makes carrier Ethernet a suitable technology for the GÉANT Plus service. Indeed, it offers some clear benefits over the existing SDH platform. For example, the existing SDH-based GÉANT Plus service has exhibited throughput problems handing VLANs between operators. This is typically due to poor buffer management at network bottleneck points. Also, the existing SDH-based GÉANT Plus service has the limitation of not having machine-readable Ethernet frame counters (and performance management) which are necessary on a per-VLAN (or QinQ service tag) and per-port basis for all technology types. Well-implemented carrier Ethernet should be able to overcome both of these issues.

Carrier Ethernet can be used to reduce network costs through the use of statistical multiplexing gains to reduce the number of trunk wavelengths. Furthermore, GÉANT Plus instances can (in principle) be overbooked at ingress and egress if required.

6.5.3.4 GÉANT IP

None of the technologies can be used to provide the GÉANT IP service itself (except logical trunk connections – see Section 6.5.3.2 above). The arguable exception to this is the largest of the core Ethernet switches that can double up as capable IP routers.

6.5.4 Market Analysis

This section reviews the carrier Ethernet component of the responses to the RFI and considers how this carrier Ethernet equipment could be used in GÉANT.

6.5.4.1 Overview of Carrier Ethernet in RFI Responses

No stand-alone carrier Ethernet solutions were described in the RFI responses. Four vendors (F, G, H and J) did not include plans for carrier Ethernet products, either because:

1. Their switch offering is EoMPLS-based.
2. Their switch offering is OTN-based with L2 functionality either not present or far in the future.
3. They did not respond to the switching part of the RFI.

Five vendors (A, B, C, D and E) all have either ROADMs containing L2 switching cards or separate OTN switches containing L2 switching cards. For all of these vendors, all or part of the carrier Ethernet functionality is currently unavailable; a fully featured carrier Ethernet offering in these products will become available in the next year or two.

6.5.4.2 Usage of Carrier Ethernet in GÉANT

A review of the available transport products' technology reveals that many of the vendors are moving towards a Packet Optical Transport Service (POTS)-type architecture in which access and transport networks are integrated in a single solution. cE and TDM are combined to deliver a flexible and reliable access grooming and transport layer.

Control plane implementations take advantage of mature IP control plane protocols. In fact, many of the implementations rely on an IP control plane where some useful extensions have been added to open the CP to the requirements of the next-generation network (NGN). At the moment, MPLS-TP seems to be the de facto standard switching technology available on almost all cE equipment in the RFI responses. The Traffic Engineering (TE) extensions to MPLS make it possible to implement a CO-PS (Connection Oriented – Packet Switched) network where cost, reliability and throughput are the key transport network requirements.

None of the POTS offerings includes a fully mature carrier Ethernet component. This may be ascribed to the POTS concept being relatively new; these products therefore require an integration phase in which carrier Ethernet functions are incorporated into the ROADMs/OTN equipment. For this reason a more cautious approach may be to opt for a stand-alone carrier Ethernet or EoMPLS solution with a more mature implementation.

The value of the POTS concept of integrating cE into ROADMs is not yet clear. A possible benefit is flexibility. ROADMs can be installed without any L2 functionality, and then L2 cards can be rapidly added as and when needed.

Typically, the ROADMs that are POTS enabled offer either OTN or L2 switching – but not both. We therefore need to be careful how we make use of this function. For example, OTN may be useful for managing sub-lambda services (such as 10 G point-to-point services over 100 G wavelengths) but, if used in this way, the L2 switching functionality may not be available for GÉANT Plus.

Closer analysis is required to fully understand any constraint imposed by the architecture of POTS products. A poor integration of L2 into a legacy ROADM could result in blocking features of the switching function.

The OSS software available with current cE platforms needs to be sufficiently mature and reliable to allow a real no-touch switched network. The information provided in the RFI responses is not sufficient to confirm that this is the case; deeper investigation is required in the next stage of procurement. It should be noted that the OAM protocols available on the equipment are key to efficient network supervision and management and fast recovery in case of major faults. Complete statistical information needs to be available to keep track of bandwidth usage and service quality in order to support sophisticated network planning. Alarms and recovery schemas need to be sufficiently mature and reliable to support traffic rerouting for planned upgrade or repair events.

The integration of carrier Ethernet into ROADMs and/or OTN switches makes it possible to deliver the GÉANT Lambda and GÉANT Plus services with a single Packet Optical Transport Service solution. The GÉANT network architecture that best leverages the value of carrier Ethernet is the model in which IP is a centralised service with L2 + TDM transport layer providing access. The IP layer can be centralised and the services can take full advantage of having the cE layer as the transport layer. In order to perform efficient and flexible L2 VPN service deployment, MPLS becomes a key technology in the access layer as in the transport and service layer. cE is completely integrated in this kind of NGN view.

In general, carrier Ethernet supports flexible and efficient service delivery for GÉANT IP, GÉANT Plus and GÉANT Lambda. The GÉANT IP service can be carried over cE – L2 frame switching can be performed before the IP packet payload is processed. For GÉANT Plus, cE may be sufficient on its own for delivery of L2 E-LAN and E-Line services.

For GÉANT Lambda cE could be used to aggregate circuits with bandwidth requirements lower than 10 Gbps i.e. the typical interface speed for DWDM equipment.

6.5.5 Summary

In conclusion, the carrier Ethernet implementations presented as part of POTS solutions in the RFI responses are not fully mature. However, in general, cE is becoming sufficiently mature and flexible to be used in a future-proof NGN, especially if a good and complete OSS system is available.

7 GÉANT Backbone Infrastructure

7.1 Introduction

Part of the process of reviewing the architecture of the GÉANT backbone has been to study the physical infrastructure (the detailed routing of physical connectivity and location of PoPs) and how this affects the reliability of the higher layer services that make use of this infrastructure. The term physical connectivity is used to cover the fact that such analysis should not be confined to studying the GÉANT fibre footprint alone but also needs to consider the routing of underlying fibre that carries any managed bandwidth services (e.g. leased wavelengths) used to construct parts of the GÉANT backbone. Included in the category of managed bandwidth services are what are known within the GÉANT community as cross-border fibre (CBF) services.

The GÉANT backbone topology was studied as viewed from the highest level (the “pan-European” view). This was done to consider macroscopic shortcomings and the analysis was broadly divided between two categories:

- Identification and mitigation against “gross shared risks with respect to connectivity infrastructure” (taking into account how various failure scenarios can be expected to affect the performance of the GÉANT services).
- Identification of hitherto missed opportunities to enhance and/or rationalise the topology of the GÉANT backbone.

The process of reviewing the detailed routing of GÉANT connectivity infrastructure reveals opportunities to open up additional or alternative access points to the GÉANT backbone. The rationale for doing this is described below.

7.2 The Need to Understand Physical Infrastructure Details

The stylised GÉANT “tube map” (shown in Figure 2.1 on page 17) is very familiar but it hides a lot of physical detail that cannot be ignored when it comes to undertaking serious infrastructure-based resilience studies. In order to perform such studies, then, it is essential to understand the detail of the physical routing of underlying physical fibre and duct infrastructure. Ducts are the tubes that carry optical fibre cables and are usually buried in a “trench” with a number of other ducts in parallel. Often, if the ducts are large enough, they are subdivided into a number of smaller “sub-ducts”. The cables carry many optical fibres. (These are usually used in pairs but not exclusively so. However, they are typically traded in pairs.) Fibre providers often trade fibre pairs, cables,

sub-ducts or ducts amongst themselves (in the form of swaps or long-term leases or outright purchases). Often, groups of fibre providers will have originally undertaken shared proprietary builds (sharing the costs of a common network build and then dividing up the resulting fibre infrastructure at the level of ducts or sub-ducts or cables or maybe even groups of pairs within a cable). All of this means that there is a great deal of scope for situations where fibre pairs (or unprotected managed bandwidth services) that are sourced through different operators are commonly routed over some distance. This means that they exhibit “shared risks” of varying degrees of severity depending on how much common routing there is and the precise nature of the common routing (i.e. whether the fibre pairs in question are in the same cable or duct or trench, or whether they are in ducts that are separated by a few metres but running down the same street). The spectrum of circumstances that can be considered to present a shared risk of one kind or another is very much a continuum, and usually each situation has to be examined in some detail on a case-by-case basis.

As stated above, all of this means that it is essential to understand the physical routing of the underlying fibre that supports the GÉANT backbone connectivity. Actually, it is not really sufficient to consider the GÉANT backbone in isolation (just looking for shared risks between links that are, on the face of it, physically diverse, as indicated in Figure 2.1); it is also desirable to extend such analysis to include infrastructure belonging to the NRENs. This is especially important when it comes to considering the resilience of end-to-end services that extend over the larger GÉANT service area (traversing NRENs and the GÉANT backbone). This having been stated, the main focus of this deliverable, and indeed of SA1 Task 1 itself, is to address the architecture of the GÉANT backbone, so it has not yet been possible to examine all the overlaps with NREN infrastructure that may have a bearing on the end-to-end services that are the consideration of SA2.

7.3 High-Level Resilience Study

The detail of the routing of the GÉANT fibre- and managed bandwidth-based trunks is shown in Figure 7.1. This is useful on a number of counts. It shows a topology that relates much more closely to real distances (and hence transmission delays) than can be appreciated by looking at the stylised tube map alone. It also illustrates the macroscopic topological shortcomings referred to above. Three have been identified and are expanded below.

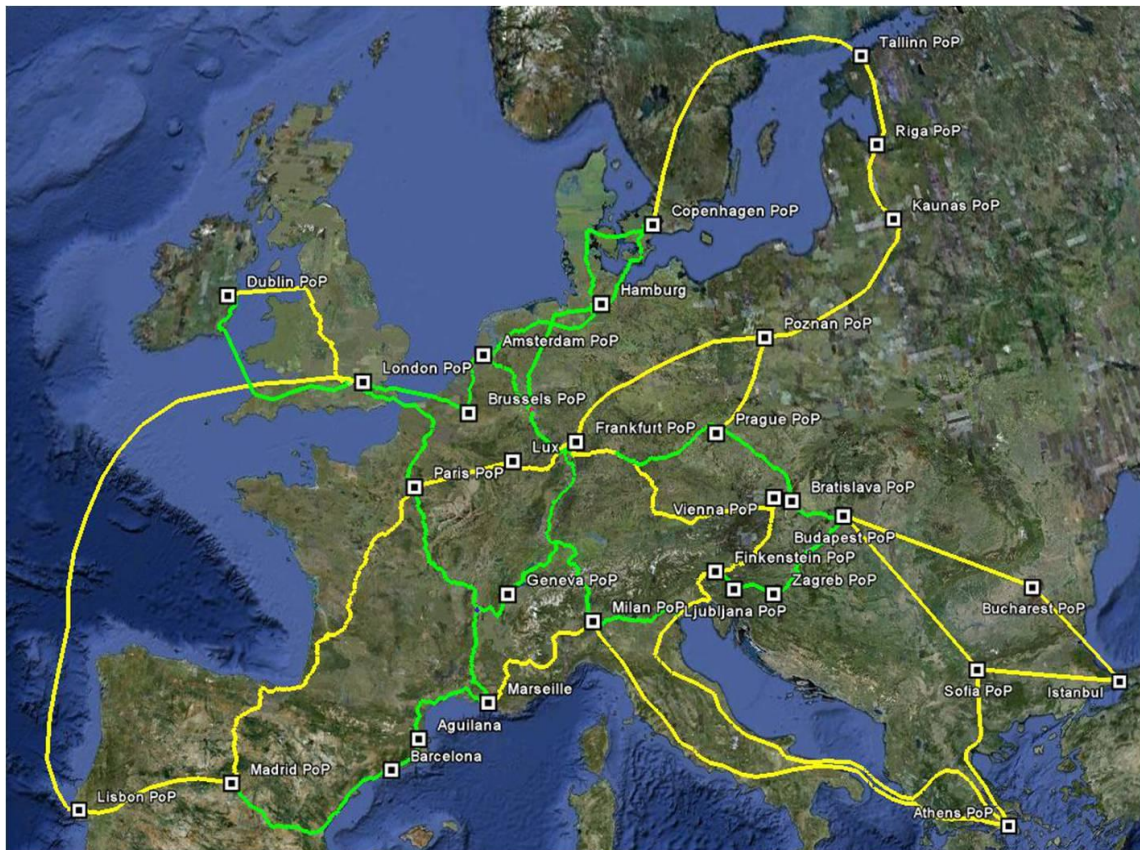


Figure 7.1: Detail of GÉANT backbone fibre routing

7.3.1 Diversity of Trunks into Geneva

There are four fibre trunks in and out of the GÉANT PoP at CERN in Geneva. These are Geneva–Frankfurt, Geneva–Milan, Geneva–Paris and Geneva–Madrid, the first three of which are provided by one supplier and the last by another. Of these, Geneva–Frankfurt and Geneva–Milan are commonly routed (along the highways) between Geneva and Basel, where the two routes diverge. Likewise, Geneva–Paris and Geneva–Madrid are commonly routed between Geneva and Lyon. This is illustrated (for Geneva–Basel) in Figure 7.2.

This situation is not a surprise. It was consciously entered into at the beginning of GN2. At the time it was mainly down to limitations in what was offered to the GÉANT consortium during the initial fibre procurement (undertaken during the main connectivity procurement at the beginning of the GN2 project during 2004/5). This, in turn, is largely down to there being limited options in terms of real fibre infrastructure in the ground. Although many operators can and have offered connectivity (fibre and leased wavelength services) into the Geneva PoP, very often these are based on the same underlying fibre. There are, however, some alternatives and these are discussed below.

It should be noted that the situation for the Geneva PoP is still reasonably robust. The two pairs of routes in and out of Geneva are physically diverse right up to the PoP itself (housed in building 513 – the CERN data centre) and this includes diverse routing across the CERN campus.

In addition, for both commonly routed sections, the two fibre pairs comprising the two routes use separate amplifier shelves. This is illustrated for Geneva–Basel (actually Giebenach) in Figure 7.2. Clearly both shelves rely on the same (dual) power feeds and, residing in the same hut, are exposed to the same environmental conditions (dust, cooling, etc.) but nevertheless there is an element of diversity between the inline amplification of dual routes.

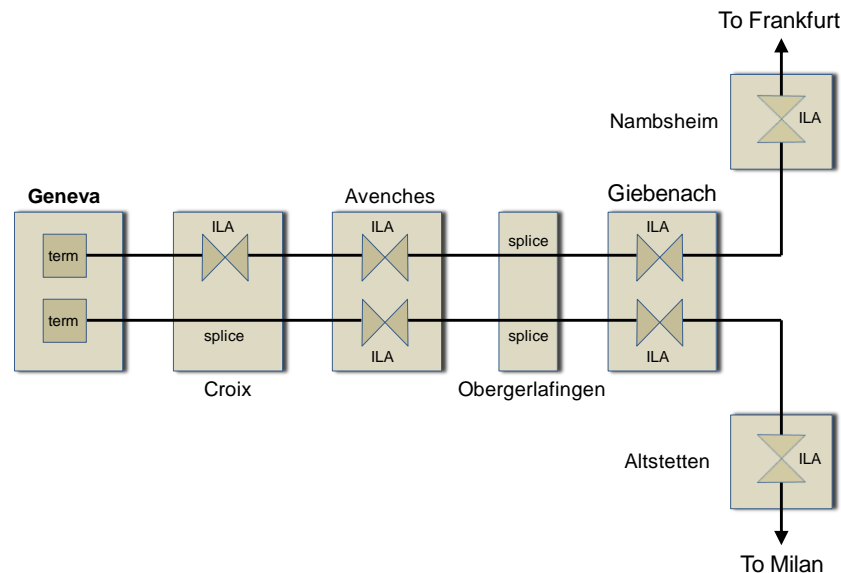


Figure 7.2: Logical diagram of commonly routed section Geneva–Basel

Various failure scenarios that can arise (some of which have done so during the last few years) have been studied. These include complete fibre cuts (in which both parallel transmission pairs have been cut), complete amplifier hut outages, partial fibre cuts (in which one of the parallel transmission pairs remains intact), and individual amplifier shelf outages.

The failure scenario that needs to be addressed is the more serious complete fibre cut or amplifier hut outage that leads to a double trunk failure (e.g. Geneva–Frankfurt and Geneva–Milan at the same time). The obvious way to do this is to acquire a third physically diverse fibre or managed bandwidth route out of Geneva. Given the fact that Geneva is the PoP at which the western ring is connected to the eastern ring (via Milan), it is preferable that this third route is essentially eastbound in as much as it adds diversity to the Geneva–Milan trunk. In addition to the acquisition of a simple diverse third route between Geneva and Milan, another option has been identified.

Various options have been considered, including the addition of a new diverse fibre trunk and the use of CBF solutions. It is worth noting the following fact about CBF solutions (and indeed any managed wavelength solutions). In essence their use to augment the GÉANT fibre backbone mesh will detract from the deployment and operation of an agile optical (photonic) network. This is because such trunks cannot be connected into a wavelength selective switching node (e.g. ROADM) and take part as a bearer link in a Wavelength Switched Optical Network (WSO) domain. This point is expanded in the Agile transmission sections 8.5, 8.6 and 8.7 beginning on p. 97.

There are possibilities to lease another pair of fibres between Geneva and Zurich that is diverse from the current Geneva–Milan route passing through Basel and Zurich. In Zurich, these fibres could be connected to the existing GÉANT fibre pair going to Milan. This would also enhance the possibility of opening up an additional access in Zurich should the local NREN (SWITCH) want to take up any opportunities that this would present (see Section 7.5 *Additional NREN Access Points* on page 84). Possible fibre providers have been identified. This would mean that the routes Geneva–Frankfurt and Geneva–Milan would be physically separated. Being a fibre solution, it would also fit well into a solution based on photonic wavelength switching (using ROADM/WSS elements – see discussions in Section 6.2 *Reconfigurable Optical Add-Drop Multiplexers (ROADMs) and Wavelength Selective Switches (WSSs)* on page 48). Note that care would need to be taken to ensure that the routing through downtown Geneva and onto the CERN site is diverse from the routing taken by the current Geneva–Frankfurt and Geneva–Milan routes.

Another option is to use a CBF solution via SWITCH and GARR infrastructures routed Geneva–Lausanne–Brig–Domodossola–Lugano–Como–Milan. In order to support a larger number of lambdas, 100 G and non-Ethernet transport mechanisms, SWITCH have indicated that this option is likely to require equipment upgrades and various enhancements to the fibre infrastructure. For example, SWITCH would most likely need to lease additional fibres on parts of the link in order to support highest speed long-distance connections and the increasing number of small sites along the line. As a result, they could offer such a "large scale CBF service" in 2012 at the earliest. It is worth noting that sharing the "express fibres" with the European backbone network should lower the cost for both SWITCH and DANTE. This route would also be more direct (to Milan) than the current route passing through Basel and Zurich.

On a much shorter time scale (within 4–6 months), SWITCH have indicated they could provide three 10 GE links from Geneva to Manno, where GARR could pick them up and extend them to the GÉANT PoP in Milan. In order to route the lambdas diversely from COLT's links, SWITCH would lease an additional fibre pair between Lausanne and Geneva and take care to minimise parallel routing with GÉANT links in the last mile to CERN.

Finally, there is another scenario which does not add a third diverse route out of Geneva towards Milan but rather takes advantage of the fact that the existing **Geneva–Madrid** fibre route passes through a location in Marseille in which it is known that a major fibre infrastructure route comes in from Milan via Turin. GÉANT currently makes use of this fibre route as there is a managed wavelength service between Milan and Madrid (used as a resilient path for the IP layer) that runs over it. The idea is that by placing a ROADM/WSS node in Marseille (where there is currently an inline optical amplifier) it would be possible to establish alternative trunks between Geneva and Milan that are diverse from those that use the Geneva–Basel section and are not too lengthy. At the same time, adding the necessary flexibility point in Marseille would help to open up additional GÉANT access in Marseille (e.g. to access MOLEN – see Section 7.5 *Additional NREN Access Points* on page 84). This option is illustrated in Figure 7.3.

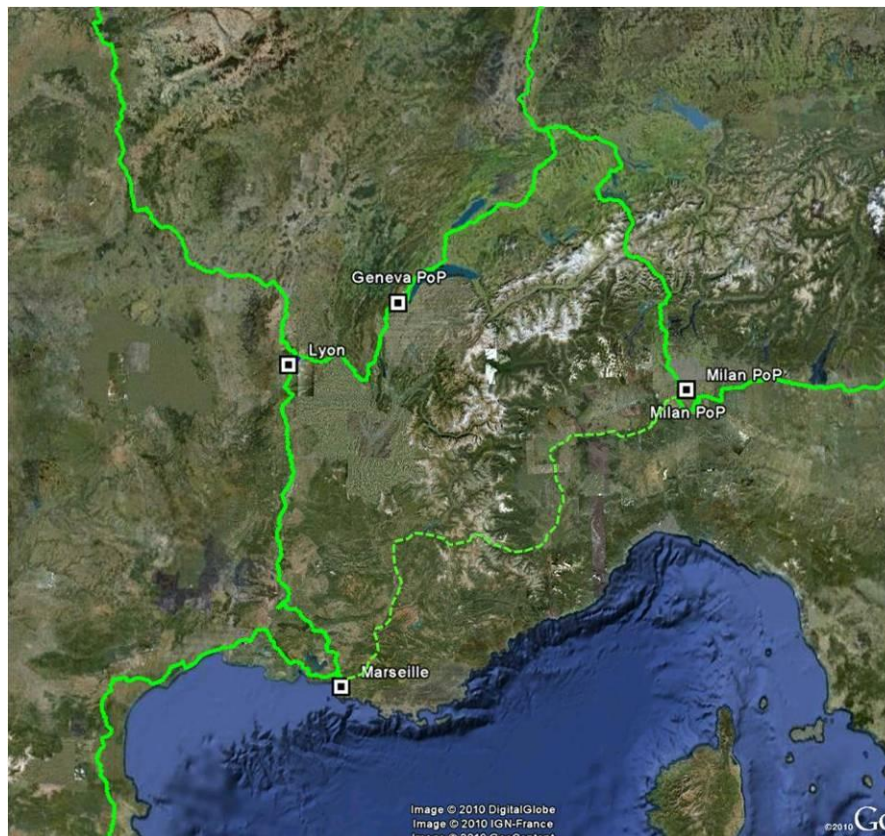


Figure 7.3: The addition of a flexibility point in Marseille and a Marseille–Milan fibre route

7.3.2 Trunks into Frankfurt

Frankfurt has always been a major nexus on the GÉANT backbone (as indeed it is on many pan-European backbones). There are four fibre trunks in and out of the GÉANT PoP in Frankfurt. These are Frankfurt–Amsterdam, Frankfurt–Geneva, Frankfurt–Copenhagen and Frankfurt–Prague. In addition, there are a number of important trunks based on managed wavelength services. These are Frankfurt–Poznan, Frankfurt–Moscow and a link between Frankfurt and Vienna that exists to provide resilience at the IP layer against the possibility of a double fibre disconnecting the eastern and western GÉANT fibre rings. All of these trunks are provided by different operators with the exception of Frankfurt–Copenhagen (fibre) and Frankfurt–Moscow (managed wavelength), which are provided by the same operator. There are even more links that come into Frankfurt including one from Israel (architecturally an IP access link), a trunk from Paris via Luxembourg, and one of the transatlantic links (currently a 10 Gbps link coming from Washington).

As with Geneva, there are some common routing problems. First, in a similar manner to the Geneva–Basel situation described above, Frankfurt–Amsterdam and Frankfurt–Copenhagen (although provided to DANTE by different suppliers) are commonly routed between Frankfurt and Dusseldorf, where they diverge. Again, this situation is not a surprise. It was consciously entered into at the beginning of GN2 and for the same reasons as with Geneva (i.e. limitations in what was offered to the GÉANT consortium during the initial fibre procurement, which is itself largely due to there being limited options in terms of real fibre infrastructure in the ground along the Rhine valley).

However, since then at least two other shared risks have been identified or have arisen due to the addition of new links and/or the changing of providers. Most notably these are:

- A shared risk between the Frankfurt–Moscow and the Frankfurt–Copenhagen trunks. This is unsurprising given that both trunks (one a fibre pair and the other a managed wavelength) are provided by the same operator and that they have, also unsurprisingly, routed the Frankfurt–Moscow wavelength via Dusseldorf, Hamburg, Copenhagen, Stockholm, Helsinki. With this shared risk alone, there would not be too much of a problem (see below) but we must bear in mind that the Frankfurt–Dusseldorf section of this particular shared risk is also shared with the Frankfurt–Amsterdam trunk.
- A shared risk between the Frankfurt–Prague trunk and the Frankfurt–Vienna “IP resilience trunk”. This has only recently been identified and arises because the two suppliers that provide these trunks use a common underlying fibre provider between Nuremberg and a point to the south of Frankfurt near Darmstadt. Fortunately, there is little or no shared risk between these routes and that taken by the Frankfurt–Geneva fibre trunk.

The various failure scenarios that can arise here have been examined, with the result that a number of options to improve on resilience can be considered. As in the case of diversity around Geneva, the consequence of this is that a requirement for an additional Frankfurt-bound trunk that is physically diverse from the shared Frankfurt-Dusseldorf section has been identified.

Such a trunk already exists in the form of the trunks passing through Luxembourg. This is based on managed wavelength services at present but potential fibre providers are known about. There are a number of reasons why a fibre upgrade could be considered desirable here. First, it would enable the establishment of a diverse trunk to either Paris or Brussels (preferably both) that would improve the latencies experienced by rerouted traffic under the conditions of the most serious outage scenario (a “double trunk fibre cut” between Frankfurt and Dusseldorf). Second, by adding to the optical mesh of the backbone, it significantly improves the economics of being able to operate an affordable dynamic/automatically restorable wavelength-provisioning platform in the backbone. This is because many more optical restoration paths become feasible without the need for 3R regeneration (which, in turn, requires the geographically strategic up-front deployment of not-insignificant numbers of DWDM transponders). This is explained more in Section 6 *Technologies* on page 45. Third, it allows for the provision of lower latency “cross-Europe” services. Unfortunately, fibre routes through Luxembourg are unlikely to be cheap as they attract a significant price premium due to the popularity of reduced latency telecommunications services used by the financial sector to support algorithmic trading.

Another option would involve adding a diverse trunk to Dusseldorf to remove the shared risk. As with Geneva, this could be based on a CBF solution or on acquiring fibre and optionally adding a 4-degree ROADM in Dusseldorf enabling it to act as a “fibre junction flexibility point.” The route could also be used to provide a second access to DFN, but Dusseldorf is not of sufficient interest to DFN for them to do this.

Finally there is the option of a CBF solution between Frankfurt and Amsterdam, although it would be essential to ensure physical diversity from the existing Frankfurt–Amsterdam fibre trunk.

There is also the problem of the shared risk between Nuremberg and Darmstadt. This is quite simple to solve. The current supplier has the option to move the Nuremberg–Frankfurt section of the Frankfurt–Vienna trunk onto their own fibre infrastructure rather than using the underlying provider in use today. This is relatively easily done.

7.3.3 Trunks into Budapest

The Budapest GÉANT PoP is connected by fibre trunks to the north to Bratislava (and from there to Vienna and Prague) and to the south to Zagreb. This is shown in Figure 7.4. The IP layer topology is incongruent to this fibre topology, in part because both Bratislava and Zagreb are “routerless PoPs”. The two westbound IP trunks out of the Budapest PoP are towards Vienna (the solid red line in Figure 7.4) and Prague (the solid blue line in Figure 7.4).

It would have been possible, given the reality of the fibre footprint, to implement the Budapest–Vienna IP trunk as an unprotected wavelength routed via Zagreb, Ljubljana, Finkenstein (a location near Klagenfurt in western Austria where there is a three-way regenerator) but this is a particularly long way round, as illustrated by the dashed red line in Figure 7.4. Instead, this IP trunk was implemented as a 1+1 protected service (using the MCC NG-SDH cross-connects in Budapest and Vienna) where the short working path was routed via Bratislava (and hence has a shared risk with the Budapest–Prague IP trunk between Budapest and Bratislava) and the long protection path routed as shown by the dashed red line.

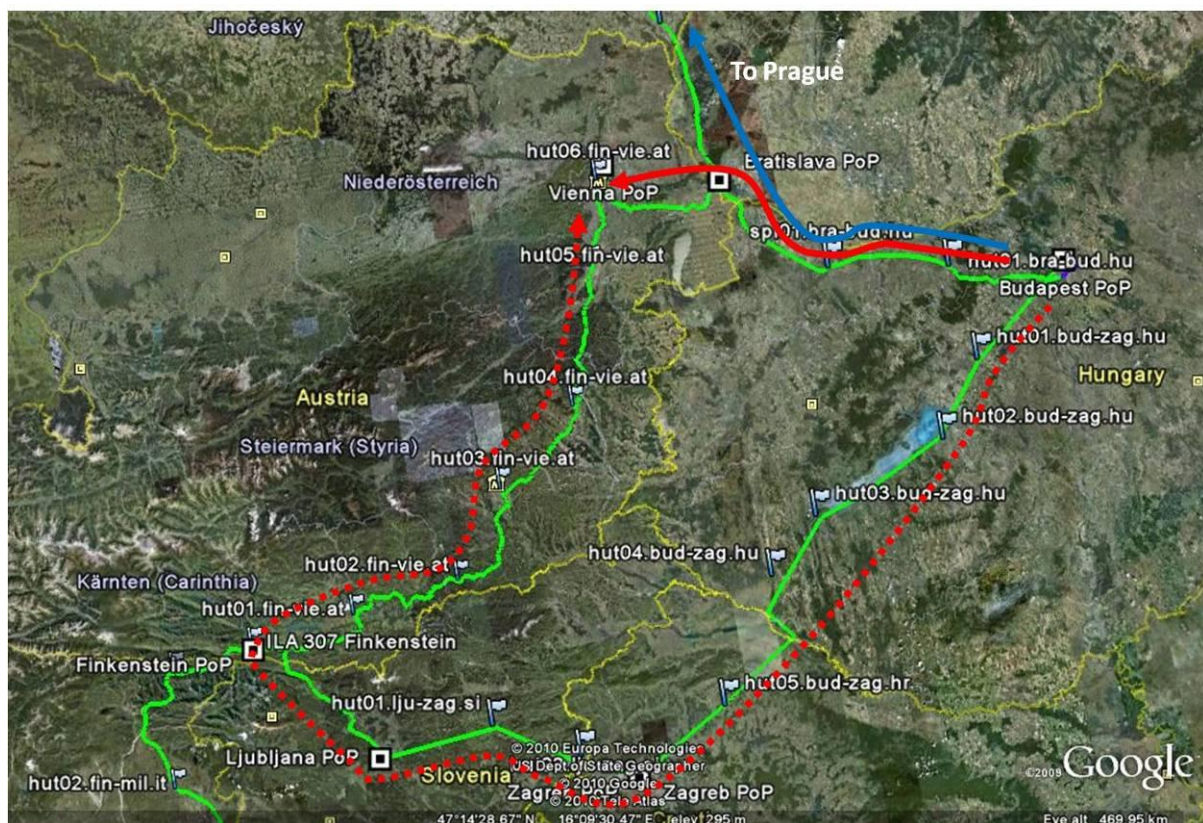


Figure 7.4: Situation around Budapest

Unlike Geneva–Basel, Geneva–Lyon and Frankfurt–Dusseldorf, there are not two parallel pairs with separate transmission systems along the Budapest–Bratislava section. There is just a single pair with a three-way terminal system in the Bratislava PoP. Hence there is only one failure scenario to consider here and this is where there is a fibre cut or simple amplifier failure between Budapest and Bratislava. This causes an outage

on both the Budapest–Prague IP trunk and the working path of the Budapest–Vienna IP trunk. Of course the latter is protected (based on SNCP) and switches to the much longer protection path within the usual ~50ms.

This failure scenario is not particularly serious in terms of trunk capacity required in and out of Budapest. After all, the Budapest–Vienna IP trunk should not actually fail at the same time as the Budapest–Prague trunk unless there is an unfortunate double fibre cut; there are also IP trunks from Budapest towards Bucharest and Sofia that would provide some resilience to the GÉANT IP service for Budapest connectors in this event. What is of more concern at this stage is the scalability of the current solution.

Currently, the GÉANT DWDM platform is fundamentally point-to-point in nature, as described in Section 2 *Current GÉANT Architecture* on page 16. This means that most GÉANT PoPs house n -way terminals and all transiting wavelengths need to be regenerated (i.e. undergo OEO conversion). This is because n -degree (R)OADMs were not available at the time the transmission platform was rolled out – only 2-degree with local add/drop. The upshot for the protected Budapest–Vienna IP trunk is that it takes no less than 12 transponders to implement it (4 on the working path and 8 on the protection path). When it is necessary to upgrade the capacity of this trunk, to say 20 Gbps, that would entail deploying another twelve 10 Gbps transponders (assuming the transmission platform were not upgraded in any way) and six 10 Gbps interfaces in the two MCC switches that facilitate the protection switching. This would obviously be unnecessarily costly. If it were necessary to upgrade the trunk capacities to 40 G or 100 G, then the regeneration-related costs would clearly be very high and indeed the MCC platform would need to be replaced by some other element that could cope with protection switching trunks at these higher capacities.

Options to improve on this situation are now presented. An obvious option is to seek a CBF solution to interconnect Budapest and Vienna directly. Ideally the route will avoid Bratislava but if that is not possible then a tripartite CBF provisioning setup transiting Slovakia is possible, but great care would need to be taken to ensure (or at least maximise) physical diversity from the existing GÉANT Budapest–Bratislava–Vienna route as shown in Figure 7.4. As with all CBF solutions (which are effectively managed wavelength solutions), this would hinder the possibility of adding to the GÉANT optical mesh for the purpose of adding wavelength agility into the transmission layer.

Another approach is to retain the same or a similar logical setup to that in place today. By implementing a suitable upgrade to the photonic layer (e.g. adding ROADMs/WSS capability and perhaps optimising for coherent transmission – see Section 6.2 *Reconfigurable Optical Add-Drop Multiplexers (ROADMs) and Wavelength Selective Switches (WSSs)* on page 48 for more details), it should be possible to retain the 1+1 protected nature of a higher capacity trunk ($n \times 10$ G or 40 G or even 100 G) and to do so at lower cost than would otherwise be the case were the point-to-point nature of the current transmission platform retained. Ideally, fast protection switching is needed (as is the case today) so it would be necessary to ensure that whatever equipment performs this protection switching does so as cost-effectively as possible. Many optical transmission platforms can do this themselves (using an approach that looks like the SNCP used by the MCCs in today's solution) but they typically tie up four DWDM transponders (or more if either of the working or protection paths is unfeasible without intermediate regeneration).

A variation on this theme would be to seek to make the trunk automatically restorable using optical agility capability in the transmission layer (ROADM switching in conjunction with control plane capability). It should be noted that restoration times can be of the order of minutes (which in itself should not be a problem as the routing re-convergence time of the GÉANT IP platform should be no more than a second). However, westbound IP traffic would have to reroute to the east and south via Sofia and Athens during the restoration

process. This is arguably a severe shortcoming of this solution unless it proves possible to provision sufficient spare trunk capacity on this route. A benefit of this approach is that it should tie up fewer DWDM transponders than for the 1+1 protection case above (i.e. two or more depending on how much intermediate regeneration is required).

In both cases, it should be noted that this approach becomes more attractive as the reach of transponders over the GÉANT fibre footprint improves. This point is discussed in some depth in Section 6 *Technologies* on page 45.

7.4 Topology Enhancements and Rationalisation

Three areas can be considered: extension of the GÉANT fibre footprint, its rationalisation (where there are suitable options and circumstances that can allow this) and the addition of more meshing. These are discussed below.

7.4.1 Extend the Fibre Footprint

Where possible (and cost-effective) the project intends to extend the fibre footprint to locations that are today served by leased wavelengths. In principle, all the trunks in yellow in Figure 7.1 on page 75 are candidates to be “upgraded” to fibre. Clearly it may not be possible to upgrade them all. See Deliverable DS1.1.2 “Procurement Plan” [DS1.1.2] and Deliverable DS1.2.2,1 “Follow-Up on Short-Term Procurement Plans” [DS1.2.2,1] for more details on the priorities when it comes to procurement, but it is worth noting that the most pressing need is to identify and hopefully procure fibre infrastructure passing through SEE and on to Greece (Athens) and Turkey (likely to be Istanbul).

Other links that are candidates for upgrade to fibre routes are:

- Paris–Madrid (currently two leased 10G wavelengths).
- Two diverse links to Lisbon.
- Two diverse links to Poznan.

When Paris–Madrid has been considered in the past, it was concluded that it was not cost-effective to light up a fibre route. Clearly this can be re-examined but unless there is a clear and pressing need to exceed the 2x10G capacity that is currently in place along this route, it is not likely that the economics will be much more compelling today than they were five years ago. The case for making this upgrade may be improved somewhat when considering options for the upgrade of the GÉANT trunks to Lisbon (e.g. opening up a flexibility point somewhere along this route to facilitate the connection of an overland link to Lisbon as an alternative to the current Lisbon-London 2.5 G managed “wavelength” service that runs over the Tyco WE sea cable today).

At the time of writing, the current expectation is that links to Lisbon and Poznan are prime candidates for CBF-based solutions. However, as has been noted previously in this document, CBF solutions (and indeed any managed wavelength solutions) will detract from the deployment and operation of an agile optical (photonic)

network. This is because such trunks cannot be connected into a ROADM/WSS and take part as a bearer link in a WSON domain.

7.4.2 Rationalisation

Given that both the Amsterdam–Copenhagen and Frankfurt–Copenhagen trunks pass through Hamburg (close to each other but not coinciding) and then diverge before coming together again in Copenhagen, and the fact that four NRENs have PoPs in Hamburg, then there is a clear rationalisation that can be made centering on Hamburg.

DFN are at three locations and in recent years SURFnet, NORDUnet and PSNC have established PoPs in a common colocation facility. NORDUnet are dually connected into Hamburg via their “Southern Cross” fibre infrastructure (one leg of which is almost fully commonly routed with the Hamburg–Copenhagen section of the Frankfurt–Copenhagen GÉANT trunk. Hence there may be some merit in opening a new GÉANT PoP in Hamburg which would replace the Copenhagen PoP (in terms of providing a primary GÉANT access to NORDUnet) and, at the same time, provide additional accesses to the other three NRENs present.

One difficulty that would need to be addressed in this scenario is that a number of trunks and access circuits that currently land in the GÉANT PoP in Copenhagen would have to be B-end shifted to Hamburg (or potentially extended there by NORDUnet).

A potentially useful outcome of this scenario would be that the then “freed-up” fibre pair running from Hamburg up to Copenhagen and back could be used for experimental purposes – possibly in conjunction with the transport and transmission equipment test lab that is being constructed at the NORDUnet premises in Copenhagen.

At the time of writing, the details of how this scenario might be achieved have been partially considered by the relevant parties but it is not yet appropriate to include the details here.

7.4.3 Add More Meshing to the GÉANT Fibre Backbone

Analysis has shown that in order to be able to make cost-effective use of an agile optical layer (based on the deployment of ROADM/WSSs in all GÉANT PoPs and some other strategic locations such as fibre junctions – see Section 7.6 *Adding Fibre Junction Flexibility Points* on page 85), it is preferable to add to the meshing of the fibre network. A number of potential new fibre links (in addition to those that can be considered to be part of the footprint extension) have been identified during the course of this investigation into topology and these are shown in blue in Figure 7.5 below.

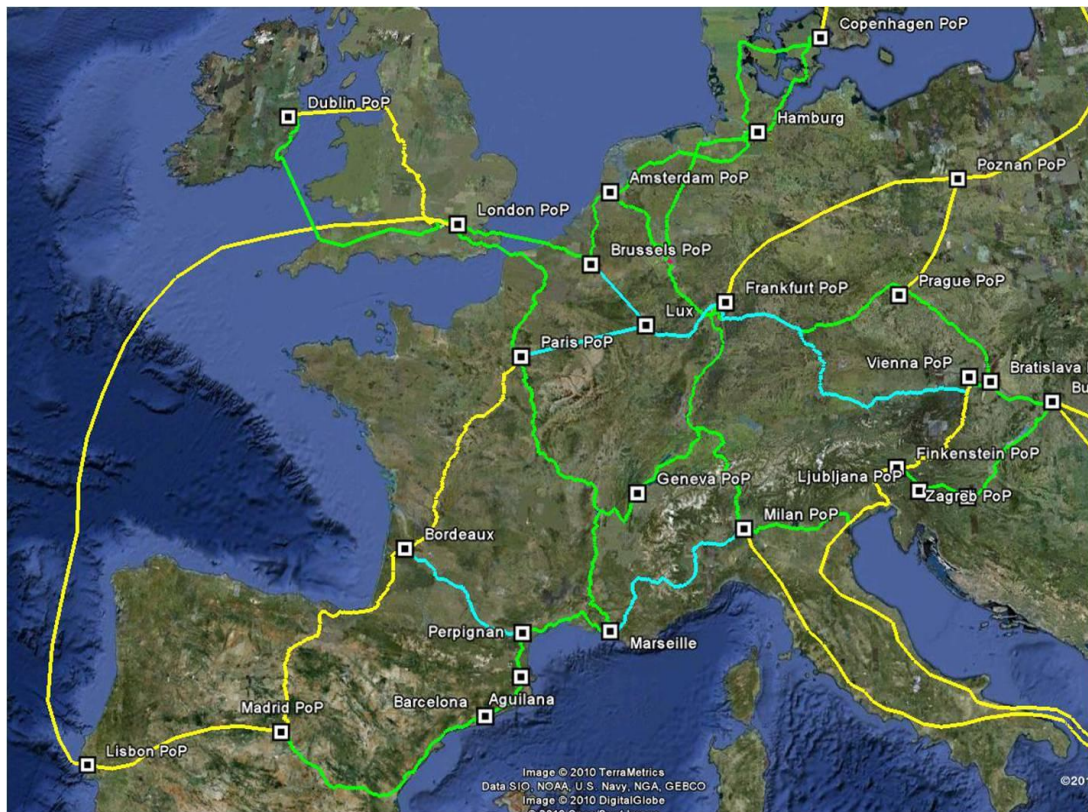


Figure 7.5: Potential additions to the GÉANT fibre mesh

The links are: Frankfurt–Luxembourg, Paris–Luxembourg, Brussels–Luxembourg, Milan–Marseille and Frankfurt–Vienna (provided the shared risk with Frankfurt–Prague between Frankfurt and Nuremberg is removed).

As described elsewhere in this section, there are other benefits to adding these links into the fibre mesh. Chief amongst these is the fact that they help mitigate against the various shared risks that have been identified and described above (see Section 7.3 *High-Level Resilience Study* on page 74).

One piece of meshing that is clearly missing is a fibre link interconnecting the two long lines from Madrid (to Paris and Geneva) around the area of the Franco-Spanish border. One possibility is a Bordeaux–Perpignan fibre route (shown in blue in Figure 7.5). Without a link here, it may mean that there is not much scope to derive the benefits of an agile optical layer (e.g. restoration capability without the need to deploy standby regeneration) for services that source/sink in Madrid. However, in order to add such a link it would also be necessary to upgrade Paris–Madrid to a fibre route (see Section 7.4.1 above).

7.5 Additional NREN Access Points

Traditionally, NRENs have accessed the backbone in only one physical location (typically this means via a single colocation site). For most, this has meant via the GÉANT PoP in their country. This does not mean to say that these access arrangements are grossly non-resilient. Many NRENs have backup IP accesses that

bypass the local GÉANT router and are transported to a router in the next-nearest GÉANT PoP (either switched over the GÉANT Plus provisioning platform or directly over the optical transmission platform). This arrangement ensures that the local GÉANT backbone router does not act as a single point of failure but the backup circuit does pass through the PoP and is hence not fully diverse from the primary access.

There are exceptions to this. For those that do not have an in-country GÉANT PoP, then usually two international unprotected but diversely routed access circuits have been used to provide access to two of the nearest GÉANT PoPs. These have, in the past, been used in a main and standby arrangement but in at least one case at the moment these two access circuits are used in a load-shared arrangement. A current example is the connection of ULAKNET (the Turkish NREN), which is facilitated by two 2.5 Gbps leased circuits – one to Sofia and the other to Bucharest. Recently there have been requests from some of the NRENs to have building-diverse second accesses to GÉANT.

These can be implemented using CBF solutions or by taking advantage of the fact that the existing GÉANT fibre footprint may actually pass through convenient locations that have hitherto only been used for inline amplification or may even be the location of a pair of simple “glass-through” (i.e. no active equipment) fibre splices.

7.5.1 Hamburg

The notion of using Hamburg as a location in which to provide second accesses to three or four NRENs has been discussed above in Section 7.4.2 (in the context of GÉANT infrastructure rationalisation).

7.5.2 Marseille

As indicated above in Section 7.3.1 *Diversity of Trunks into Geneva* on page 75, the Geneva–Madrid fibre trunk passes through Marseille. There is currently an amplifier in a facility (hut07.agu-gen.fr) near where a fibre route comes in from Milan (via Turin). Although GÉANT does not make use of this infrastructure for a fibre trunk, it does make use of the same infrastructure for a part of a managed wavelength (used as an IP layer resilience trunk) that connects Milan with Madrid. This can potentially be used to improve the meshing of the backbone as described previously, but it can also be used to facilitate access to the nearby MOLEN facility that RENATER have set up.

7.6 Adding Fibre Junction Flexibility Points

As described above there are three locations where parallel fibre lines diverge – namely Basel, Lyon and Dusseldorf. There is an interesting debate over how the use of these can be modified. Taking Geneva–Basel as an example, the way they are configured is as shown in Figure 7.6.

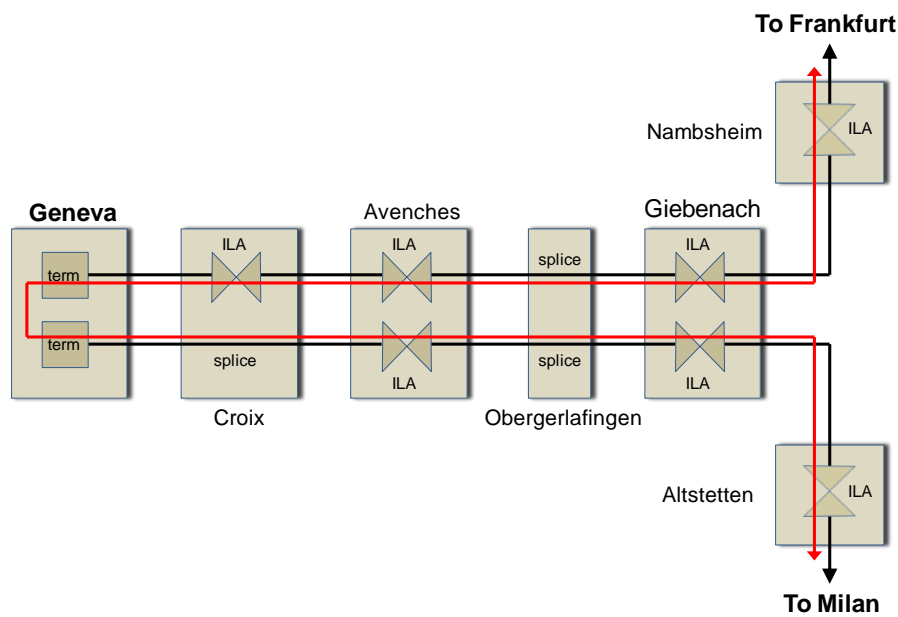


Figure 7.6: Current configuration of parallel fibre lines

One disadvantage of this is that circuits between the Frankfurt and Milan GÉANT PoPs need to transit the Geneva PoP and in so doing need to be regenerated (the red line in Figure 7.6). It is possible to place a 3-degree ROADM at the fibre junction (Basel/Giebenach in this case) as shown in Figure 7.7.

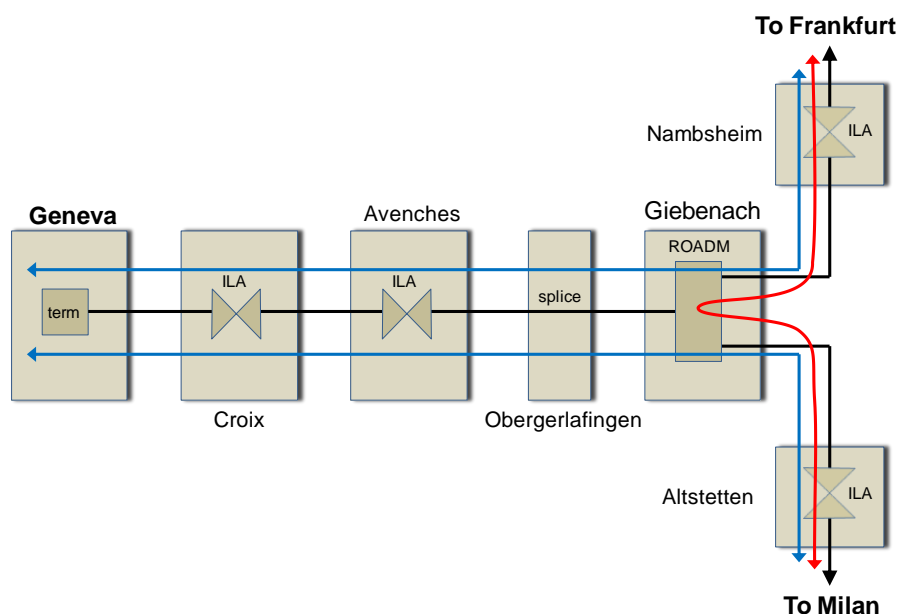


Figure 7.7: Placing a 3-degree ROADM at a fibre junction

Clearly, in this case the Frankfurt–Milan wavelength need not transit the Geneva–Basel section, or the Geneva PoP itself, with the following advantages:

- The wavelength is shorter (and hence the latency will be less).
- In many cases two transponders will suffice instead of four.
- The wavelength passes through fewer elements (in this case 5 fewer ILAs and 2 fewer terminals) and less transmission fibre so the availability should be higher (although this depends on the reliability of the ROADM that needs to be inserted at the fibre junction).

A disadvantage of the approach shown in Figure 7.7 is that there is now a greater shared risk between two wavelengths emerging from the Geneva PoP that diverge in Basel (shown by the blue lines in Figure 7.7) – they share transmission fibre and ILAs.

Another possibility is to retain the parallel transmission line (fibre and ILAs) and place a 4-degree ROADM in the fibre junction. This is illustrated in Figure 7.8.

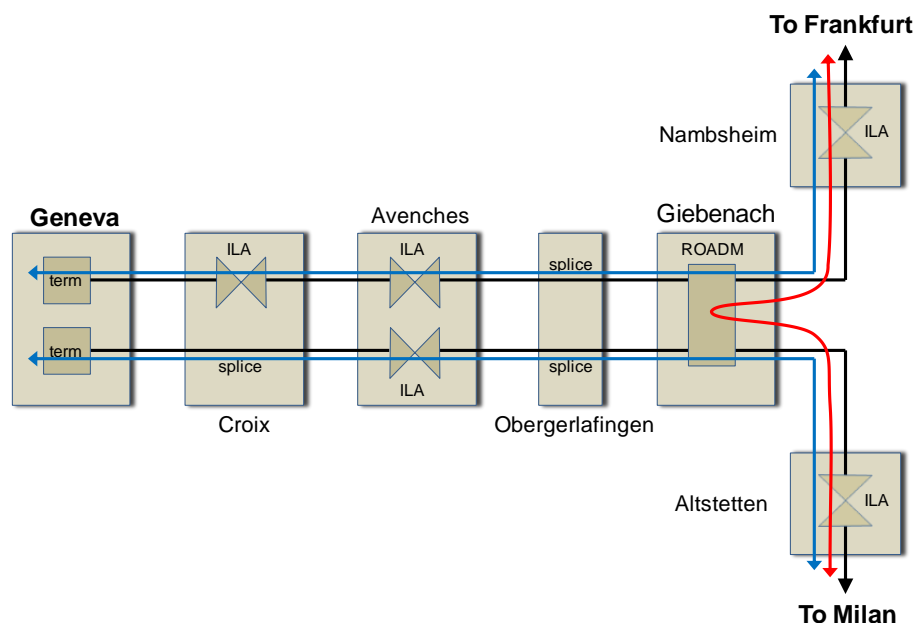


Figure 7.8: Placing a 4-degree ROADM at a fibre junction

The advantage of this is that the original level of “diversity” (of transmission fibre and amplifiers) along the commonly routed section is retained rather than being compromised more as it was in the setup shown in Figure 7.7.

In addition to the three fibre junction sites where parallel transmission lines diverge, there are two locations that currently house 3-way regenerator/terminal nodes. These are Finkenstein in Austria (where the fibre route coming up from Ljubljana meets up with the Milan–Vienna fibre route) and Bratislava. The former is just a junction (no services are added or dropped there), whereas the latter is a (“routerless”) GÉANT PoP. The reason these locations are 3-way regens is because degree- n ROADMs were not available when the current GÉANT transmission platform was rolled out five years ago. This situation is now much changed and these locations are also prime candidates for the deployment of degree-3 ROADMS (with add/drop in the case of Bratislava).

8 Architecture Options

This section puts together the building blocks described in previous sections and presents possible generic approaches to the GÉANT network architecture. It presents an overall framework for the architecture and the main architectural alternatives in terms of suitability for supporting the GÉANT services and meeting the associated quality parameters.

8.1 Introduction

The GÉANT network has a layered architecture, as shown in Figure 8.1 below. This shows the GÉANT services provided (or that can be potentially provided) by each layer (in green), the potential technologies that can be deployed to do so (in blue) and the nature of the applications that typically (or can be expected to) make use of services at this layer (in red).

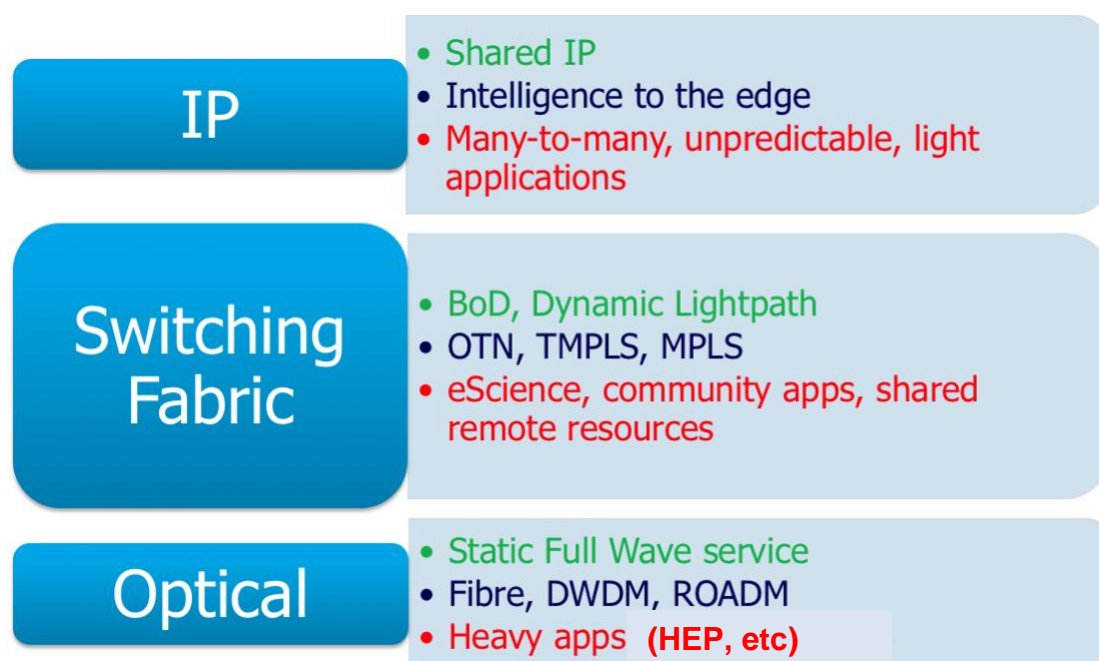


Figure 8.1: GÉANT network layered architecture

The three layers – optical transport, switching, and IP – are not new compared to previous generations of the GÉANT network. What is new is a greater possibility to make use of advances in both optical transport (in the DWDM transmission layer, which allows circuit switching of whole wavelengths) and sub-wavelength switching (be that based on TDM or packet-oriented approaches or both), thereby facilitating further evolution of the network.

Options that are envisaged and explored further in this chapter are centred on making use of the strides that have been made in recent years in (i) the introduction of great flexibility (agility) in DWDM transmission systems (itself based mainly on the widespread availability of tunable and wavelength selective switching components) and (ii) the area of digital switching. In addition, advances in optical transmission capabilities mean that 100 Gbps (on a single DWDM channel) is now accessible and this is almost certain to be a key part of the evolution of GÉANT that takes place during GN3.

Another option is to increase the emphasis on using the switching layer for the support of a modified architecture for the IP network compared to the current situation where the IP network is almost all based on the underlying optical transmission infrastructure (and managed capacity where the GÉANT fibre footprint is not present). One such architecture envisages pushing IP layer intelligence towards the network edge (this will be investigated in work that will be done following up on this deliverable).

Several implementation alternatives are defined for the switching layer. These are the main architectural choices for the GN3 network design. The switching layer can be realised with an OTN switching capability, with a “carrier Ethernet” capability (as defined below in the context of this deliverable), or with an MPLS-focused IP capability.

The GÉANT network architectural components and possible implementation alternatives are therefore as follows:

1. IP component: edge IP capability, BPG4 routing and interfacing, MPLS interfacing.
2. Switching component:
 - **Carrier Ethernet (cE)** implemented by equipment supporting: (i) plain old Ethernet enhanced with OAM, QoS and >10 G (as per section 3.2 in [DJ1.1.1]) perhaps with Layer 2 routing (SPB, etc.) or (ii) MPLS-TP or (iii) PBB-TE.
 - **OTN/NG-OTN** implemented by equipment capable of ODU-x switching. Note that the other functionalities of OTN/NG-OTN, e.g. error correction, are viewed as an integral part of the optical transmission component.
 - **MPLS** implemented with equipment supporting MPLS switching and EoMPLS services on high-bandwidth trunks, offering a network-wide MPLS switching fabric.
3. Optical transmission component:
 - Fixed (P2P) transmission.
 - Agile transmission (allowing circuit switching): implemented by optical transmission equipment including ROADMs with colourless and directionless features.

Note that the switching layer has several functions:

- It can potentially be the basis for the IP service layer, allowing edge equipment, BGP4 capability, and other IP service equipment to interconnect at high bandwidth. Specifically, it allows NREN edge routers to connect to remote GÉANT routers rather than to local ones.
- It provides the infrastructure for the GÉANT Plus lightpath service. With a high-bandwidth switching fabric (100 G), a GÉANT Plus service up to 10 G can be provided.

An important issue with regard to the technological categorisation above is the interaction between the three layers, in particular between the IP and switching components; there is a clear tendency among the transport vendors to promote an architectural model according to which the pieces of the IP component are interconnected via a network of switching equipment. The target is to minimise the number of IP hops between a source–destination pair of IP nodes, thus reducing the number of expensive IP ports (router off-load) and increasing the number of cheaper switching ports, and to enable the switching fabric to facilitate additional meshing of the IP layer.

On the other hand, some IP equipment vendors claim that the profit from router off-loading is cancelled out by the added cost of the switching equipment. Router off-load makes limited sense if combined with a 100 Gbps routing capability at all nodes. For instance, for a 2-degree IP node, 2 x 100 Gbps ports is expected to be adequate to manage all IP traffic instead of $N \times 10$ Gbps per direction – hence it is not possible to further reduce (off-load) the number of IP ports. To achieve the full benefit of an extended switching capability, the number of nodes with full routing capability would have to be reduced. In this design, the IP layer would have a reduced number of routers, and NREN routers would connect to the IP layers using switching services rather than connect to a local GÉANT router. In any case, this is an architectural issue that needs to be separately examined for the GÉANT case so as to fully understand the cost and functionality trade-offs.

It is worth noting that the current GÉANT IP layer provides a fully functional MPLS switching fabric. If the MPLS architecture is chosen, this can be used to initially implement part of the switching layers. Doing so would initially have the IP layer equipment also implement part of the switching layer, moving to separate switching equipment as the edge router equipment is upgraded. The choice of approach and phasing of implementation depends on cost, including cost and requirement for 100 G service.

Control plane interactions among the defined technological components need to be investigated separately. These bring additional complexity and a survey of relevant standardisation activities as well as vendors' proprietary solutions needs to be carried out in order to understand the complexity, depth and the maturity of the proposed solutions.

It should be mentioned that, unless stated otherwise, each technological component is assumed to be implemented by discrete equipment.

Finally, each architectural alternative is evaluated overall in terms of:

- Reliability.
- User-network separation.
- Maturity of involved technologies.
- Provisioning time.
- Multi-domain deployment.
- Management flexibility.

8.2 Fixed (P2P) transmission + carrier Ethernet + IP/MPLS

8.2.1 Service Delivery

In this scenario, the GÉANT services portfolio would be delivered as follows:

- GÉANT IP: provided by IP/MPLS equipment (as today).
- GÉANT Plus: provided by “carrier Ethernet” (cE) equipment and perhaps using incumbent technology (IP/MPLS) as an alternative for cE protocols not yet fully implemented in commercially available equipment.
- GÉANT Lambda: provided by fixed (P2P) transmission equipment (i.e. without wavelength selective switching capability).

8.2.2 Pros and Cons for GÉANT Plus Implementation

Pros	Cons
<ul style="list-style-type: none"> • GÉANT Plus is implemented with a packet-based technology, which is expected to be more compatible with NREN circuit provisioning layers. • Non-transparent implementation implies the ability to provide detailed traffic statistics per service instance, e.g. number of forwarded frames. • OAM tools per service instance for troubleshooting and performance monitoring are either available or under development by the relevant standardisation bodies (IEEE, ITU-T, IETF). • Tools for protection and restoration are either available or under development by the relevant standardisation bodies (IEEE, ITU-T, IETF). • E-LINE, E-TREE and E-LAN services are possible. However, the E-LAN service for PBB-TE has not been standardised yet. 	<ul style="list-style-type: none"> • A partial mesh of wavelengths among the cE equipment needs to be pre-installed. Hence, this alternative bears the risk of having wavelengths severely under-utilised – especially where there is the desire to always ensure short lead times for provisioning new service instances. • Since statistical multiplexing is possible, admission control and policing/shaping functionality is required at the edges of the network. This means that traffic loss for bursty traffic is possible over heavily utilised links serving a high number of service instances. This issue could be alleviated by over-provisioning the cE network using (potentially costly) 40 Gbps / 100 Gbps wavelengths.

Table 8.1: Pros and cons for GÉANT Plus implementation

8.2.3 Pros and Cons for GÉANT Lambda Implementation

Pros	Cons
<ul style="list-style-type: none"> • It is a well-known mode of operation, having been 	<ul style="list-style-type: none"> • The wavelength path provisioning methodology is

Pros	Cons
in place for almost 5 years.	<p>inflexible, requiring manual and on-site intervention from the vendor's personnel. This implies a lead time (a few days or sometimes weeks) for wavelength path provisioning.</p> <ul style="list-style-type: none"> • It would not be possible to deliver “photonic services” within a reasonable time frame. • Even in mesh topologies, 1+1 wavelength path protection is the only possible resiliency strategy that can be followed (assuming that resilience is implemented in the transmission layer alone).

Table 8.2: Pros and cons for GÉANT Lambda implementation

8.2.4 Overall Evaluation

- **Reliability:** cE technologies include protection and restoration mechanisms. Especially for PBB-TE standardised head-end path restoration mechanisms are still under development. At the optical layer O-SNCP can provide a 50 ms protection if needed and is the only available resiliency tool.
- **User-network separation:** cE technologies guarantee full isolation of the core nodes; however edge nodes need to maintain customer-specific information. At the optical layer, user-network separation is implemented by means of a back-to-back interconnection among client and network DWDM platforms; at this point regeneration of the optical signal is implemented on the network DWDM platform providing adequate visibility on received/transmitted optical signals status.
- **Maturity of involved technologies:** cE technologies (including MPLS-TP and PBB-TE) are relatively new technologies and standardisation of some features is still ongoing. At the optical level, fixed (P2P) transmission is a well-known technology, commercially available for over a decade.
- **Provisioning time:** cE technologies guarantee minimal time (e.g. a few seconds) in service provisioning. However, at the optical layer, fixed (P2P) transmission usually requires some working days since it involves on-site visits and re-cabling along with adjustments of the optical signal. (In both cases, idle trunk capacity has to be present *a priori*.)
- **Multi-domain deployment:** PBB-TE was developed as a single domain technology. On the other hand, MPLS-TP can seamlessly integrate with customer MPLS networks. On the optical layer, multi-domain deployment without reductions in optical signal monitoring capabilities is feasible as long as regeneration is implemented at the edges of the network.
- **Management flexibility:** Usually cE technologies and optical transmission systems require a NMS system.

8.3 Fixed (P2P) transmission + OTN + IP/MPLS

8.3.1 Service Delivery

In this scenario, the GÉANT services portfolio will be delivered as follows:

- GÉANT IP: provided by IP/MPLS equipment (as today).
- GÉANT Plus: provided by OTN equipment (more precisely, equipment implementing the digital switching and multiplexing parts of the G.709 recommendation from the ITU-T).
- GÉANT Lambda: provided by fixed (P2P) transmission equipment.

8.3.2 Pros and Cons for GÉANT Plus Implementation

Pros	Cons
<ul style="list-style-type: none"> • GÉANT Plus is implemented in a transparent way using dedicated capacity per service instance. Hence, no packet loss is expected if client traffic is mapped to an appropriately sized ODU structure. • Tools to detect errors at the ODU level (switching level) are available. • Tools for protection and restoration are available. • Line, Tree and LAN topologies can be provided. However, in order to provide the relevant services (E-LINE, E-TREE and E-LAN), another switching layer is required (e.g. existing IP/MPLS equipment could be used, implying that OTN switches interconnect with local IP equipment) or the OTN switching equipment should support sub-port-level grooming capabilities. 	<ul style="list-style-type: none"> • A partial mesh of wavelengths among the OTN equipment needs to be pre-installed. Hence, this alternative bears the risk of having wavelengths severely under-utilised – especially where there is the desire to always ensure short lead times for provisioning new service instances. • Since OTN is a transparent transport service, no OAM visibility into the actual users' traffic is possible. Arguably this is undesirable depending on how service monitoring (through use of OAM) will be done in the multi-domain environment.

Table 8.3: Pros and cons for GÉANT Plus implementation

8.3.3 Pros and Cons for GÉANT Lambda Implementation

Same as in Section 8.2.3.

8.3.4 Overall Evaluation

- **Reliability:** Digital OTN technologies include protection and restoration mechanisms, where the control plane is based on GMPLS or similar. At the optical layer O-SNCP can provide a 50 ms protection if needed and is the only available resiliency tool.
- **User-network separation:** OTN is a transparent technology and it will not maintain any customer information. At the optical layer, user-network separation is implemented by means of a back-to-back interconnection among client and network DWDM platforms; at this point regeneration of the optical signal is implemented on the network DWDM platform providing adequate visibility on received/transmitted optical signals status.
- **Maturity of involved technologies:** OTN technology was initially developed almost ten years ago. However, the features that make it attractive for deployment in the GÉANT network (e.g. ODUflex) are currently in the initial general availability phase for many vendors. At the optical level, fixed (P2P) transmission is a well-known technology, commercially available for over a decade.
- **Provisioning time:** OTN guarantees minimal time (e.g. a few seconds) in service provisioning. However, at the optical layer, fixed (P2P) transmission usually requires some working days since it involves on-site visits and re-cabling along with adjustments of the optical signal. (In both cases, idle trunk capacity has to be present *a priori*.)
- **Multi-domain deployment:** OTN can seamlessly operate with other client OTN networks implementing the carriers' carrier scenario. On the optical layer, multi-domain deployment without reductions in optical signal monitoring capabilities is feasible as long as regeneration is implemented at the edges of the network.
- **Management flexibility:** Usually OTN technologies and optical transmission system requires a NMS system.

8.3.5 Special Case: ODU Switching

It should be noted that there is at least one vendor today whose optical transmission platform supports fixed (P2P) transmission and also integrated flexible ODU switching: all incoming wavelengths transit an ODU-based switching matrix at every OADM node. This implementation can be considered as a special case of the scenario outlined above and has several implications that are highlighted below.

Pros	Cons
<ul style="list-style-type: none"> • Unified management of GÉANT Lambda and GÉANT Plus services. Actually, the boundaries between the two services are not quite discrete, since in operational terms there is no difference between managing a GÉANT Lambda service (an ODU-2/3/4 signal) and a GÉANT Plus service 	<ul style="list-style-type: none"> • Some argue that there is increased cost, since all wavelengths are regenerated at every OADM node even if, in principle, the highly integrated nature of the regeneration leads to it being less of an issue than if it were based on discrete regenerator components.

Pros	Cons
<p>(e.g. 1 x ODU-0 signal).</p> <ul style="list-style-type: none"> The GÉANT Lambda service becomes more flexible and blocking probability is significantly reduced; e.g. a wavelength path can use different frequencies along its route. The GÉANT Lambda service can include restoration characteristics. Simpler optical engineering, since the range of optical domains transparency is limited. 	<ul style="list-style-type: none"> Since digital processing is involved at each OADM node, there is an RTT penalty for the GÉANT Lambda service due to unnecessary processing at intermediate OADM nodes – although this is usually much smaller than transmission delays over typical GÉANT inter-PoP distances. Currently lambdas are implemented in bundles of 10x10G, and concerns regarding upfront costs can be associated with this approach. The overall power consumption for this platform, having an integrated ODU switch, might be higher than for a traditional DWDM platform. This should be studied further.

Table 8.4: Pros and cons for ODU switching

8.4 Fixed (P2P) transmission + IP/MPLS

8.4.1 Service Delivery

In this scenario, the GÉANT services portfolio will be delivered as follows:

- GÉANT IP: provided by IP/MPLS equipment (as today).
- GÉANT Plus: provided by IP/MPLS equipment. Note that the cost of using the existing IP/MPLS routers for implementing the GÉANT Plus service is not negligible and it may make sense, in economic terms, to deploy lower capability MPLS-enabled routers at existing IP nodes and at the GÉANT nodes that do not currently host an IP router. This option will be evaluated during the forthcoming equipment RFP.
- GÉANT Lambda: provided by fixed (P2P) transmission equipment.

8.4.2 Pros and Cons for GÉANT Plus Implementation

Pros	Cons
<ul style="list-style-type: none"> Unified management of GÉANT IP and GÉANT Plus services, using existing tools and methodologies. Cost reduction, because the need to maintain a separate wavelength network for providing the GÉANT Plus service is eliminated; GÉANT IP traffic is “statistically multiplexed” with GÉANT Plus traffic (although, in practice, overbooking is 	<ul style="list-style-type: none"> Since statistical multiplexing is involved, appropriate admission control and policing/shaping functionalities are required at the edges of the network to avoid losses. It should be noted that this issue could be alleviated by using 40 Gbps / 100 Gbps wavelengths.

Pros	Cons
<p>unlikely to be done).</p> <ul style="list-style-type: none"> • GÉANT Plus is implemented with a packet-based technology, which is likely to be more compatible with NREN circuit-provisioning layers. • Non-transparent implementation implies the ability to provide detailed traffic statistics per service instance, e.g. packet/frame statistics. • OAM tools (per service instance) for troubleshooting and performance monitoring are either available or under development by the relevant standardisation bodies (ITU-T, IETF). • Tools for protection and restoration are available. • E-LINE, E-TREE and E-LAN services are possible. 	

Table 8.5: Pros and cons for GÉANT Plus implementation

8.4.3 Pros and Cons for GÉANT Lambda Implementation

Same as in Section 8.2.3.

8.4.4 Overall Evaluation

- **Reliability:** IP/MPLS technology includes lots of resiliency mechanisms. At the optical layer O-SNCP can provide a 50 ms protection if needed and is the only available resiliency tool.
- **User-network separation:** IP/MPLS technology guarantees full isolation of the core nodes; however edge nodes need to maintain customer-specific information. At the optical layer, user-network separation is implemented by means of a back-to-back interconnection among client and network DWDM platforms; at this point regeneration of the optical signal is implemented on the network DWDM platform providing adequate visibility on received/transmitted optical signals status.
- **Maturity of involved technologies:** IP/MPLS technology has been deployed in data networks and GÉANT for around ten years. At the optical level, fixed (P2P) transmission is a well-known technology, commercially available over a decade.
- **Provisioning time:** IP/MPLS technology guarantees minimal time (e.g. a few seconds) in service provisioning. However, at the optical layer, fixed (P2P) transmission usually requires some working days since it involves on-site visits and re-cabling along with adjustments of the optical signal.
- **Multi-domain deployment:** Initially MPLS was developed for single-domain deployment but it was gradually equipped with multi-domain features; this is an area of active development for this technology.

On the optical layer, multi-domain deployment without reductions in optical signal monitoring capabilities is feasible as long as regeneration is implemented at the edges of the network.

- **Management flexibility:** IP/MPLS technology can be fully managed by DANTE's existing management tools. The optical transmission system needs a separate NMS.

8.5 Agile transmission + carrier Ethernet + IP/MPLS

8.5.1 Service Delivery

In this scenario, the GÉANT services portfolio will be delivered as follows:

- GÉANT IP: provided by IP/MPLS equipment (as today).
- GÉANT Plus: provided by "carrier Ethernet" (cE) equipment and perhaps using incumbent technology (IP/MPLS) as an alternative for cE protocols not yet fully implemented in commercially available equipment.
- GÉANT Lambda: provided by agile transmission equipment (meaning with wavelength selective switching capability).

8.5.2 Pros and cons for GÉANT Plus Implementation

Pros	Cons
<ul style="list-style-type: none"> • GÉANT Plus is implemented with a packet-based technology, which is likely to be more compatible with NREN circuit-provisioning layers. • Non-transparent implementation implies the ability to provide detailed traffic statistics per service instance, e.g. packet/frame statistics. • OAM tools per service instance for troubleshooting and performance monitoring are either available or under development by the relevant standardisation bodies (IEEE, ITU-T, IETF). • Tools for protection and restoration are either available or under development by the relevant standardisation bodies (IEEE, ITU-T, IETF). • E-LINE, E-TREE and E-LAN services are possible. However, the E-LAN service for PBB-TE has not been standardised yet. 	<ul style="list-style-type: none"> • A partial mesh of wavelengths among the cE equipment needs to be pre-configured. Hence, this alternative bears the risk of having wavelengths severely under-utilised • Since statistical multiplexing is involved, admission control and policing/shaping functionality is required at the edges of the network. This means that traffic loss for bursty traffic is possible over heavily utilised links serving a high number of service instances. It should be noted that this issue could be alleviated by over-provisioning the cE network using potentially costly 40 Gbps / 100 Gbps wavelengths.

Table 8.6: Pros and cons for GÉANT Plus implementation

8.5.3 Pros and Cons for GÉANT Lambda Implementation

Pros	Cons
<ul style="list-style-type: none"> Reduced lead time (a few minutes) in provisioning a new wavelength path, assuming that some spare transponders are pre-installed at strategic points of the network. Increased flexibility leading to significant reduction in wavelength blocking probability. More specifically, a wavelength path can dynamically adapt its transmission frequency along its route so as to utilise regions of the spectrum that are not currently used by another wavelength path. Wavelength path restoration is possible either via an advanced control plane or via a management application. It will be possible to deliver “Photonic services” within reasonable time frame. 	<ul style="list-style-type: none"> The upfront cost associated with the installation of all the requisite ROADMs components (mainly WSSs) could be quite substantial.

Table 8.7: Pros and cons for GÉANT Lambda implementation

8.5.4 Overall Evaluation

- Reliability:** cE technologies include protection and restoration mechanisms. Especially for PBB-TE standardised head-end path restoration mechanisms is under development. At the optical layer, protection as well as dynamic restoration mechanisms (possibly coupled with an advanced control plane) could be deployed.
- User-network separation:** cE technologies guarantee full isolation of the core nodes; however edge nodes need to maintain customer-specific information. At the optical layer, user-network separation is implemented by means of a back-to-back interconnection among client and network DWDM platforms; at this point regeneration of the optical signal is implemented on the network DWDM platform providing adequate visibility on received/transmitted optical signals status.
- Maturity of involved technologies:** cE technologies (including MPLS-TP and PBB-TE) are relatively new technologies and standardisation of some features is still ongoing. At the optical level, ROADM filters based on Wavelength Selective Switches (WSSs) is considered a mature technology that is currently commercially available from most of the vendors. However, the colourless and directionless feature of the ROADMs is a new technology that is only now becoming generally available.
- Provisioning time:** cE technologies guarantee minimal time (e.g. a few seconds) in service provisioning. At the optical layer, colourless and directionless ROADMs guarantee service (wavelength path) delivery in a few seconds.

- **Multi-domain deployment:** PBB-TE was developed as a single-domain technology. On the other hand, MPLS-TP can seamlessly integrate with customer MPLS networks. On the optical layer, multi-domain deployment without reductions in optical signal monitoring capabilities is feasible as long as regeneration is implemented at the edges of the network.
- **Management flexibility:** Usually cE technologies and optical transmission systems require a NMS system.

8.6 Agile transmission + OTN + IP/MPLS

8.6.1 Service Delivery

In this scenario, the GÉANT services portfolio will be delivered as follows:

- GÉANT IP: provided by IP/MPLS equipment (as today).
- GÉANT Plus: provided by OTN equipment.
- GÉANT Lambda: provided by Agile transmission equipment.

8.6.2 Pros and Cons for GÉANT Plus Implementation

Pros	Cons
<ul style="list-style-type: none"> • GÉANT Plus is implemented in a transparent way using dedicated capacity per service instance. Hence, no packet loss is expected if client traffic is mapped to an appropriately sized ODU structure. It should be noted that although OTN is a circuit-based technology, it includes mechanisms for mapping Ethernet client signals in an efficient way. • Tools to detect errors at the ODU level (switching level) are available. • Tools for protection and restoration are available. • Line, Tree and LAN topologies can be provided. However, in order to provide the relevant services (E-LINE, E-TREE and E-LAN), another switching layer is required (e.g. existing IP/MPLS equipment could be used, implying that OTN switches interconnect with local IP equipment) or the OTN switching equipment should support sub-port-level grooming capabilities. 	<ul style="list-style-type: none"> • A partial mesh of wavelengths among the OTN switching equipment needs to be pre-configured. Hence, this alternative bears the risk of having wavelengths severely under-utilised • Since OTN is a transparent transport service, no OAM visibility into the actual user's frames is possible.

Table 8.8: Pros and cons for GÉANT Plus implementation

8.6.3 Pros and Cons for GÉANT Lambda Implementation

Same as in Section 8.5.3.

8.6.4 Overall Evaluation

- **Reliability:** Digital OTN technologies include protection and restoration mechanisms, where the control plane is based on GMPLS. At the optical layer O-SNCP can provide a 50 ms protection if needed and is the only available resiliency tool.
- **User-network separation:** OTN is a transparent technology and it will not maintain any customer information. At the optical layer, user-network separation is implemented by means of a back-to-back interconnection among client and network DWDM platforms; at this point regeneration of the optical signal is implemented on the network DWDM platform providing adequate visibility on received/transmitted optical signals status.
- **Maturity of involved technologies:** OTN technology was initially developed almost ten years ago. However, the features that make it attractive for deployment in the GÉANT network (e.g. ODUflex) are in the initial general availability phase for some vendors. At the optical level, ROADM filters based on Wavelength Selective Switches (WSSs) is considered a mature technology that is currently commercially available from most of the vendors. However, the colourless and directionless feature of the ROADMs is a new technology that is only now becoming generally available.
- **Provisioning time:** OTN guarantees minimal time (e.g. a few seconds) in service provisioning. At the optical layer, colourless and directionless ROADMs guarantee service (wavelength path) delivery in a few seconds.
- **Multi-domain deployment:** OTN can seamlessly operate with other client OTN networks implementing the carriers' carrier scenario. On the optical layer, multi-domain deployment without reductions in optical signal monitoring capabilities is feasible as long as regeneration is implemented at the edges of the network.
- **Management flexibility:** Usually an OTN network requires the deployment of a NMS system. The same holds for the optical transmission system.

8.6.5 Special Case: ODU Switching

A special case for this scenario is that the optical transmission platform supports Agile transmission and also integrated flexible ODU switching: all incoming wavelengths are transited to an ODU-based switching matrix at every OADM node. The implications of this case are highlighted below.

Pros	Cons
<ul style="list-style-type: none"> • Unified management of GÉANT Lambda and 	<ul style="list-style-type: none"> • Optical transmission agility makes limited sense

Pros	Cons
<p>GÉANT Plus services. Actually, the boundaries between the two services are not quite discrete, since in operational terms there is no difference between managing a GÉANT Lambda service (an ODU-2/3/4 signal) and a GÉANT Plus service (e.g. 1xODU-0 signal).</p> <ul style="list-style-type: none"> • The GÉANT Lambda service becomes more flexible; e.g. a wavelength path can use different frequencies along its route. • The GÉANT Lambda service can include restoration characteristics. • Simpler optical engineering, since the range of optical domains transparency is limited. 	<p>since incoming wavelengths are digitally processed and switched; directionless and colourless features are implemented via the OTN matrix.</p> <ul style="list-style-type: none"> • Increased cost, since all wavelengths are regenerated at every OADM node. • Since digital processing is involved at each OADM node, there is an RTT penalty for the GÉANT Lambda service due to unnecessary processing at intermediate OADM nodes. • Current lambdas are implemented in bundles of 10x10G, and concerns regarding costs should be associated to this approach. This statement is only valid for this special case. • The overall power consumption for this platform, having an integrated ODU switch, might be higher than for a traditional DWDM platform. This should be studied further

Table 8.9: Pros and cons for ODU switching

8.7 Agile transmission + IP/MPLS

8.7.1 Service Delivery

In this scenario, the GÉANT services portfolio will be delivered as follows:

- GÉANT IP: provided by IP/MPLS equipment.
- GÉANT Plus: provided by the IP/MPLS equipment. Note that the cost of using existing IP/MPLS routers for implementing the GÉANT Plus service is not negligible and it may make sense, in economic terms, to deploy lower capability MPLS-enabled routers at existing IP nodes and at the GÉANT nodes that do not currently host an IP router. This option will be evaluated at the forthcoming equipment RFP.
- GÉANT Lambda: provided by the Agile transmission equipment.

8.7.2 Pros and Cons for GÉANT Plus Implementation

Pros	Cons
<ul style="list-style-type: none"> • Unified management of GÉANT IP and GÉANT Plus services. 	<ul style="list-style-type: none"> • Since statistical multiplexing is involved, appropriate admission control and policing/shaping functionalities are required at the

Pros	Cons
<ul style="list-style-type: none"> • Cost reduction, because the need to maintain a separate wavelength network for providing the GÉANT Plus service is eliminated. Production IP traffic is statistically multiplexed with GÉANT Plus traffic. • GÉANT Plus is implemented with a packet-based technology, which is more compatible with NREN circuit-provisioning layers. • Non-transparent implementation implies the ability to provide detailed traffic statistics per service instance. • OAM tools per service instance for troubleshooting and performance monitoring are either available or under development by the relevant standardisation bodies (ITU-T, IETF). • Tools for protection and restoration are available. • E-LINE, E-TREE and E-LAN services are possible. 	<p>edges of the network to avoid losses. It should be noted that this issue could be alleviated by using 40 Gbps / 100 Gbps wavelengths.</p>

Table 8.10: Pros and cons for GÉANT Plus implementation

8.7.3 Pros and Cons for GÉANT Lambda Implementation

Same as in Section 8.5.3.

8.7.4 Overall Evaluation

- **Reliability:** IP/MPLS technology includes lots of resiliency mechanisms. At the optical layer, the protection as well as dynamic restoration mechanisms (possibly coupled with an advanced control plane) could be deployed.
- **User-network separation:** IP/MPLS technology guarantees full isolation of the core nodes; however edge nodes need to maintain customer-specific information. At the optical layer, user-network separation is implemented by means of a back-to-back interconnection among client and network DWDM platforms; at this point regeneration of the optical signal is implemented on the network DWDM platform providing adequate visibility on received/transmitted optical signals status.
- **Maturity of involved technologies:** IP/MPLS technology has been deployed in data networks and GÉANT for around ten years. At the optical level, ROADM filters based on Wavelength Selective Switches (WSSs) is considered a mature technology that is currently commercially available from most of the vendors. However, the colourless and directionless feature of the ROADMs is a new technology that is only now becoming generally available.

- **Provisioning time:** IP/MPLS technology guarantees minimal time (e.g. a few seconds) in service provisioning. At the optical layer, colourless and directionless ROADMs guarantee service (wavelength path) delivery in a few tenths of seconds.
- **Multi-domain deployment:** Initially, MPLS was developed for single-domain deployment but it was gradually equipped with multi-domain features; this is an area of active development for this technology. On the optical layer, multi-domain deployment without reductions in optical signal monitoring capabilities is feasible as long as regeneration is implemented at the edges of the network.
- **Management flexibility:** IP/MPLS technology can be fully managed by DANTE's existing management tools. The optical transmission system needs a separate NMS.

9 Summary and Next Steps

The information collected and analysed during the backbone architecture study indicates that the use of the GÉANT infrastructure follows a path of constant growth in the amount of IP traffic and number of high-capacity circuits dedicated to projects. In addition, increasingly advanced services and functionalities in the areas of authorisation and authentication, security, monitoring, and dynamic provisioning are being requested. The technology has evolved since the implementation of the GÉANT2 network at the end of 2005, offering new optical equipment capabilities and switching platforms, and marking the decline of SDH/SONET and a ubiquitous acceptance of the Ethernet protocol. The increasing importance of data transmission for research and for organisations' operations and daily activities is placing greater importance on the resiliency and redundancy of the services. The requirement impacts the whole infrastructure, from ensuring diverse physical routes to diverse fibres to the logical topology of the IP network.

The information now collected allows a new generation of the GÉANT network to be engineered that is not just an upgrade of the current infrastructure, and which is well placed to meet the new challenges and requirements as stated in the GÉANT3 white paper [GN3 white paper].

9.1 An Architectural Model for the Next Generation

The study and operational experience have confirmed that the hybrid infrastructure at the core of the GÉANT network is a valid building block and provides the correct foundations for the next-generation infrastructure. This will be based on the fibre available to GÉANT and the NRENs, with the most appropriate switching layer at the packet and frame level added on top of it.

Figure 9.1 below shows a high-level representation of the basic layers of the new architecture; the common functions of monitoring and authentication and authorisation are part of each layer and are depicted vertically for clarity and to show the required integration. Each layer has its own control and management planes (not shown); their integration between layers is subject to technological choices and ongoing research and development. The examples given in Section 8 *Architecture Options* on page 88 demonstrate the wide variety of possible overlaps and integration.

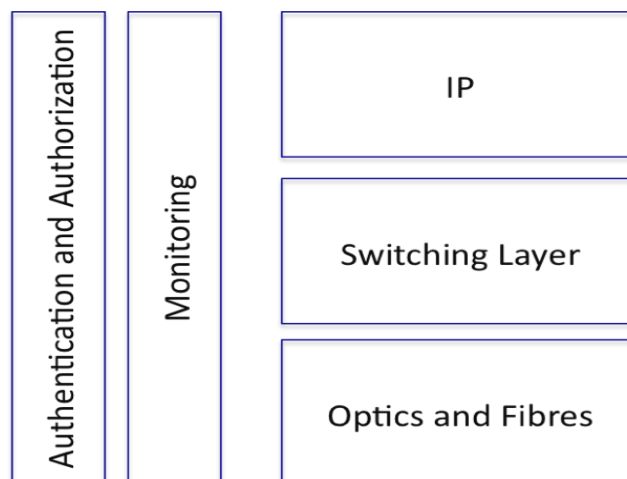


Figure 9.1: Representation of the next-generation infrastructure architectural model

In the architecture shown in Figure 9.1, the fibre provides the fundamental foundation for offering the raw capacity and flexibility of the hybrid infrastructure, its physical resilience, and its ability to outreach to federate with similar networks in Europe and in other continents.

The switching layer provides the increased capabilities for frames and packets handling currently available and which are independent of the Internet protocols. As briefly summarised in Section 8, this layer may contain a wide variety of technologies. The switching layer's functionalities allow it to be decoupled logically from both the optical layer below and the IP layer above, and to provide its own services, such as a point-to-point Ethernet circuit or a virtual local area network that spans multiple optical domains. This layer enhances current services and provides new ones, and may offer them more dynamically, logically, and with much lower cost and faster provisioning than in the current architecture.

In the new architecture the IP layer is more independent from the physical and switching topology and may be engineered to provide more rapid connectivity and virtual networking across the whole domain.

The implementation of this model will be subject to a detailed analysis based on the technical information presented in this document, the information provided by vendors, cost-benefit considerations, and the views of the NRENs.

The following section presents further technical considerations, conclusions and recommendations which will be used in the engineering and tendering processes.

9.2 Technical Considerations and Recommendations

A number of important issues are common to all layers of the architecture and require, in addition to analysis at the level of each individual layer, a solution that takes into account the interaction of all the layers. These are: resiliency and robustness to failures; fast recovery from failures; ease and speed of reconfiguration. In addition, the infrastructure should be transparent to the users and allow innovation. Additional considerations and recommendations relate to:

- Upgrading the current optical layer.
- Enhancing the physical topology.
- Switching layer.
- IP layer.
- Monitoring and authentication, authorisation and accounting.

Each of these is discussed below.

9.2.1 Optical Layer

The current GÉANT optical transmission capability needs to be upgraded in order to cope with emerging requirements. It should be noted that there is a minimum of five years remaining on the fibre contracts for their very low-cost right of use. In several cases the right is indefinite. “Upgrade” in this context does not necessarily imply only an incremental upgrade of the existing GÉANT optical transmission platform but also includes the possibility of platform replacement.

Optical Transport Network (OTN) technology features in the roadmaps of many vendors and it may provide a significant improvement in the capabilities of the optical layer. It should be noted, however, that:

- To realise the full advantages of OTN switching, the whole path should be OTN compliant. This implies deployment in all NRENs and also in the end users’ sites.
- Mature OTN products will not be commercially available in the near future.
- The cost of the equipment may be significantly higher.

The new generation equipment has to be capable of cost-effectively lighting up 100 Gbps; all the indications from the transport technology Request for Information (RFI) are that this goal should be easily achievable. There is also some evidence that fully optimising the optical layer to support transmission of 100 Gbps wavelengths may not be compatible with retaining support for discrete 10 Gbps and 40 Gbps wavelength transmission in the same fibres and optical equipment. The new-generation network topology must therefore be carefully planned to ensure the most cost-optimised migration/upgrade path, taking into account expected future service demands. The final decision can only be reached by conducting a transmission platform procurement, which will include a topology layout and detailed fibre specifications. The tender for additional fibres will then provide preliminary information to be used in the optical equipment tender.

The transport technology RFI has also shown that it is now feasible to construct and operate an “agile” (i.e. remotely reconfigurable) optical transmission layer, even taking into account the very-long-haul and geographically dispersed nature of the GÉANT PoPs. The possibility is confirmed when taken in conjunction with the introduction of the latest generation of (long-haul, high-capacity) “coherent” DWDM transponders.

Introducing agility into the GÉANT optical transmission layer is achieved by introducing photonic switching capability (most likely in the form of Reconfigurable Optical Add-Drop Multiplexers or ROADMs). There is a spectrum of agility that can be considered, depending on the number and functional flexibility of the ROADMs that are introduced. At one end of the spectrum, ROADMs are only added to a few fibre junctions (between the main PoPs, e.g. in Basel and Lyon), resulting in some optimisation of the provision and performance of some

wavelength services. Further along the spectrum, the introduction of basic ROADMs (with limited add/drop flexibility) can be considered for all GÉANT PoP sites. Although this represents an upfront investment, the main advantage is that it will yield future rewards when the provisioning of long-haul, high-capacity wavelength services will rely on less intermediate regeneration and therefore be more cost-effective and rapid. At the furthest end of the agility spectrum (and representing the largest upfront investment), each ROADM is designed to be fully flexible with respect to the add/drop of channels (the technical term is that the ROADMs become “directionless and colourless”). The main advantage of this approach is that the provision of resilient wavelength services becomes feasible and more cost-effective, although this is further improved by adding to the meshing of the fibre footprint (see Section 9.2.2 below). The term resilient here encompasses wavelengths that are 1+1 protected and those that can be automatically restored. Additionally, through the strategic pre-placement of a number of DWDM transponders, it will be possible to speed up the fulfilment of many wavelength service provisioning requests.

It should be noted that there is also a so-called “digital optical” approach to delivering the agile optical transmission platform described above, in which photonic integrated circuits are extensively used. This approach achieves all of the benefits described above, but has the possible disadvantage that it may not easily lend itself to the support of the new research and scientific applications that rely on all-optical transmission paths as described in Section 3.3 on page 26.

The recommendation is to provide these new functionalities in the optical transmission layer such that it can support cost-effective 100 Gbps transmission and agility at the photonic layer. The target for the latter should be the fully agile option, to be validated by a cost/benefit analysis. The validation should be based on a long-term view, and should ensure that GÉANT is not dependent on the development roadmap of a single vendor.

9.2.2 Physical Topology

A number of enhancements to the GÉANT physical topology are also recommended. The motivations are the need to improve overall network resilience, consolidate on network efficiency, and allow greater and more resilient access to the GÉANT backbone. Specific recommendations are:

- Increase the meshing of the GÉANT fibre footprint. The denser the meshing the more it helps with the ability to realise the benefits of an agile optical transmission layer.
- Ensure that the main connections run on physically diverse trunk paths. For a fully agile optical transmission layer, it is preferable that such connections are fibre links or dedicated wave bands rather than based on managed wavelength services (whether sourced from commercial providers or via the CBF route).
- Consider having more than one PoP per country where appropriate (e.g. Germany). This scenario increases the capability of the infrastructure to optimise cost and maximise the usefulness of the physical topology. This can be facilitated by the deployment of agile transmission technologies such as ROADMs.
- As owned fibres are now becoming common in the NRENs, it is recommended that the best synergy between NRENs’ fibres and GÉANT fibres is reached.

9.2.3 Switching Layer

The switching layer will complement and enhance the functionalities of the other two layers, in particular for circuit provisioning at varying speeds, with or without capacity guarantees. Three broad categories of technological approaches, identified and analysed in the study, can be considered viable: using the digital switching and multiplexing capabilities of next-generation OTN (G.709) (see Section 6.3 on page 54); using “carrier Ethernet” (in the broadest sense of the definition, which encompasses a wide range of technologies and protocols to overcome some of the limitations in monitoring and control of standard Ethernet – see Section 6.5 on page 65); and Ethernet over MPLS (see Section 6.4 on page 57). NG-SDH has not been considered as an alternative, due to its cost, development plans and vendors’ roadmaps.

With respect to EoMPLS, there is again a spectrum of options that can be considered. This ranges from the installation of a full set of devices (one per PoP) that act as a one-for-one replacement of the platform currently used to provide GÉANT Plus (17 Alcatel-Lucent 1678 MCCs), possibly eliminating the majority of SDH interfaces, to implementing a service to provide circuits on demand using Ethernet over MPLS on the existing GÉANT IP/MPLS platform. It is expected that on a per-port (or per-Gbps) basis the cost of the latter will prove significantly higher.

For any of these technologies, it is clear that Ethernet will be the preferred data-link technology and enhancements are needed in the areas of monitoring, control and management to operate what originated from a local area environment in a wide area environment. These enhancements will also facilitate the provision of multi-domain services, by making easier the communication of network-related information. The use of the basic Ethernet protocol will, in any case, simplify the provision of some multi-domain services (e.g. virtual LANs) and traffic engineering currently based on the IP protocol. MPLS-based solutions require instead a tighter integration of layers (at least the IP and switching layers). It is expected that the two solutions will coexist in the network.

The project will be cautious about adopting the complexities of these technologies, and will also consider capital and operational costs. Given the environments the NRENs have to serve, an agile infrastructure with simpler control planes is key: it will be better suited to fast provisioning and changes in technologies, and will ensure more transparency to researchers, an equally important consideration.

9.2.4 IP Layer

The IP layer is a long-established, well-known infrastructure component and is the fundamental asset of any transmission network. Nonetheless the analysis has indicated areas of challenge. The requests for increased capacity and the recommendation for higher meshing point to the need for more powerful IP layer equipment, which must at the same time be less expensive. In addition to the capacity increase, it is expected that the IPv6 protocol will increase its share of capacity, adding to existing usage rather than replacing it, and placing more demand on hardware platforms. Part of this strain may be alleviated by the switching and optical layer through circuits bypassing routers.

9.2.5 Monitoring and Authentication, Authorisation and Accounting

Monitoring is essential for all GÉANT activities (operation, development and services) and must be appropriately reflected in both the equipment capabilities and operational aspects. The monitoring function has been modelled vertically in Figure 9.1 to show its presence at all layers and the need for tight integration of information gathered at each of them. The information collected and its correlation between layers is vital to rapid recovery from failures, service quality validation and user support (e.g. Performance Enhancement Response Teams (PERTs) and hosted projects like Partnership for Advanced Computing in Europe (PRACE)). The monitoring capabilities of all the equipment, their compliance to standards, and whether operation should be centralised or decentralised have to be evaluated for their impact on the infrastructure and on the services' operations.

The same considerations are valid for the Authentication, Authorisation and Accounting (AAA) functionalities of the whole infrastructure and services offered. The emphasis is on simplifying access, whilst maintaining fine-granularity security levels.

9.3 Final Considerations and Next Steps

The information collected and the analysis performed allow an architecture for the next-generation GÉANT network to be defined. The fibre has an enhanced role in the new architecture as a fundamental asset and the current hybrid infrastructure has been confirmed as a sound foundation on which to build the new network.

The next-generation network will be strengthened, at all layers, in the areas of:

- **Resilience.** Improved resilience is required to provide IP and circuit services that comply with the users' increased requirements. For each layer, and for the combined infrastructure, the study has identified clear strategies (some of which are already being implemented).
- **Agility and timely configurability.** These capabilities are now available at the optical layer and at the switching layer, providing greater flexibility in infrastructure design and the ability to satisfy user requirements in a timely manner.
- **Capacity.** Increasing the total capacity of the infrastructure will first be achieved at the fibre and optical layer. A combination of more meshing, wavelength density and increased single-circuit capacity (up to 100 Gbps) will create the capability to support the projected increase. Careful pre-provisioning will also ensure that capacity will be available where it is needed in a short timeframe (i.e. days).
- **Interoperability.** Interoperability at all layers with the NRENs and international infrastructures will facilitate the deployment of services and operations.

From the technical studies conducted so far there are clear preferences for the future GÉANT network. As already stated, these will be subject to further analysis. The preferences are:

- Availability of an agile transmission platform based on ROADMs, to facilitate the resilience improvements needed, ensure the more efficient use of the topology and infrastructure, and facilitate additional access points.

- Availability of a logically separate switching layer, using Ethernet and based on one of the identified approaches (EoMPLS, CcE or OTN).
- Given the developments possible at the transmission and switching layer, there is now also the opportunity to review and optimise the IP layer

The next steps are to compare the technical information and plans with vendors' contractually available solutions and reliable cost data. Further planning is required to devise an appropriate schedule for the staged approach(es) necessary to arrive at recommendations for solutions that may be implemented. This will include an assessment of the need for further Request for Proposal (RFP) work and/or commencement of some initial tendering phases, a typical part of the dialogue phase in a Competitive Dialogue process. Work on developing this schedule is now considered a priority.

During this process the current implementations of NRENs' and international peering networks will be carefully considered to ensure that the largest number of services (including monitoring), may be seamlessly implemented.

The costs resulting from the need to migrate seamlessly from the current backbone architecture to whatever alternative is chosen and procured are expected to be not insignificant, and will play an important part in the final procurement decisions. Some of these migration costs have already been estimated – for example, costs associated with short-term parallel fibre pair leases to facilitate gradual transmission platform migration as opposed to the alternative “big bang” (overnight) approach. However, the estimation process is not yet complete, and in some areas the available costing information on which to base the estimates has not been reliable. Migration estimates have therefore not been reported in this document but will be refined and considered during the forthcoming procurement activity.

It is stressed again that the project will be cautious with regard to the possible complexities arising from novel technologies and it will ensure that the technologies selected involve low capital and operational costs, while maintaining the broadest possible compatibility and inter-operability with peering networks at all layers. Consideration will also be given to openness (which enables the user to participate in the design and development of services) and transparency of equipment in all three layers, which are important requirements of the R&E community.

The project will monitor closely the needs of users such as those identified in Appendix A, to ensure that the new architecture is able to meet them in terms of both capacity and service provision.

The availability of a greater number of fibres and wavelengths (a combination of GÉANT's own fibres and those provided by partners of the consortium) will help to keep complexity low, provide simpler solutions to resiliency, and enrich the services' capabilities.

Appendix A **User Requirements**

This appendix reproduces the text of the document “User Requirement Input to the Architecture Supervisory Committee” (version date: 5 July 2010), prepared by Richard Hughes-Jones, GÉANT Technical Customer Support Manager, for the meeting of the SA1 Supervisory Committee that took place 7–8 July 2010.

A.1 Introduction

This note summarises some of the networking requirements coming from projects and the user communities. Please note that it is very much work in progress; many of the emerging user groups and projects are only just considering their computing requirements, the storage that might be needed, and the relative locations of their users. The values noted are only rough and early estimates, and are expected to change as the user groups develop.

The following sections discuss the input gained from meetings with the ESFRI subject areas, meetings with individual emerging potential users, and the user requirements gathering meetings organised by the NRENs.

A summary table is included at the end of the document. It is intended to update this document at regular intervals.

A.2 Campus Issues Concerning Moving Data

The physical capacity of between 1–10 Gbit/s is usually sufficient for the current use of a “big science” department, but users find it hard to get transfers to work at the expected rates across these links, especially for long distances. This applies to all connectivity services routed IP, point to point circuits, and lambdas. There is a requirement for help with advice and best practice for:

- The incorporation of circuits & lambdas connectivity into Campus Policy.
- Consideration of IP addressing and routing.
- Security policy.
- Location of the bottleneck or resource problem in a multi-domain environment.
- Tuning applications, TCP, and hosts to obtain the expected and desired performance.

A.3 Distributed Computing Environments

With requirements for distributed computing and storage, latency is important as well as bandwidth. Given that the speed of light is a major factor contributing to the propagation time between sites, this implies that the network topology is important as well as minimising the router count. Consideration should be given to the locations of the data stores and the required compute.

A.4 Data Traffic Patterns

Users report that for various reasons data transfers occur in a bursty nature not as a regular pattern. This implies that the simple approach of taking the number of TBytes per day under estimates the required bandwidth. Note that even transient bottlenecks can reduce the observed performance dramatically for long distance transfers. Bio-Informatics, Climate and weather modelling, and HEP are examples of users with bursty flows.

Some workflow tools work so well from the user point of view that without realising it the users could soon be a big impact on the network traffic. The astronomy Virtual Observatory tool is a good example and it is hard to get a traffic pattern in these cases.

There are estimates that the aggregated effect from a few researchers in an active department may be transfers requiring several 100 Mbit/s for several hours of a working day. This implies the need to capacity and flexibility; perhaps automatic network provisioning would help in the future.

A.5 Integrating R&E with Community and Public Sectors

There are discussions on providing infrastructure for government and public sectors such as health, work and pensions, etc. This and the ESFRI subject areas of Bio-Informatics and Social Science and Humanities could have implications on architectural issues including:

- Carrying data at higher assurance levels
- Design of network security
- Access management e.g. smart cards to control access to the network and applications.

A.6 Connection Services

Some users are starting to examine the costs. If they need large data transfers to and from a few data points around the world and if they have BW from the NREN then they can use routed IP. But if the economics are better – they could prefer circuits or lambdas.

There are some projects, like the synchronisation of atomic clocks, and photonic protocol research, that require special facilities such as carrying light from 3rd party equipment – perhaps as an all optical alien wave. It might also be useful to dedicate a region of the DWDM wavelengths for network research projects.

A.7 Dynamic Networks

This could in principle include VPN, MPLS, point to point circuits, and full lambdas. There is current interest from some projects such as NEXPreS, HEP, connecting real-time visualisation & haptics to supercomputers, and high-definition multi-media. Through DICE ESNET and Internet2 indicate increasing usage, but this could be focused more on provisioning MPLS than physical infrastructure. Several NRENs see the potential for provisioning and managing their connection services.

The user communities see dynamic networks as an important way to obtain high bandwidth or specialised connectivity only when they need it and at a cost much reduced from a permanent connection. At then moment it seems we have a few users requiring long connection times, not the traditional telephone case of many users with short duration connections.

A.8 PRACE & HPC

HPC in this context is taken to mean specialised systems such as MPPs, thin/fat node-clusters, hybrid systems, vector systems, not large blade farms.

PRACE is part of an initiative to integrate world class HPC-facilities into an European e-Infrastructure. It will include:

- Tier-0: 4-5 European Centres (formed from PRACE funding)
- Tier-1: National Centres
- Tier-2: Regional/University Centres

PRACE will coordinate procurement so that there are systems of different architectures needed to best solve the different types of computing challenges. Currently the first Tier-0 is JUGENE IBM BlueGene/P Jülich and the second Tier-0 will follow in France during 2011. There are also 6 prototypes installed at other sites.

Networking requirements:

- All sites connected at 10 Gbit/s to the DEISA network.
This provides the strong links to national supercomputers
- High bandwidth is required between Tier-0 sites
- Large data transfers foreseen to & from user home sites
This would use the routed IP service.
- Connection to SC in other world regions

The exact detailed requirements are still being discussed, but it would appear that each PRACE site would require something like a 10 Gigabit connection to the routed IP infrastructure and one or more lambdas or dynamic network services. Given the increase in the amount of data being moved between the user sites and the PRACE and other supercomputer sites there are implications for the country access links.

A.9 Square Kilometre Array

SKA is a truly global collaboration with 21 countries including South Africa, Australia, New Zealand and Brazil. The instrument could be located in South Africa along with 7 other countries, or Australia and New Zealand.

Current thinking is for images to be stored at a small number of Regional Repositories around the world and there will be a constant flow of images 24/7. The Regional Repositories will also provide user access to the data. The requirement is for high bandwidth on a world scale.

Phase 1 with 250 dishes will produce image data at ~40 Gbit/s. Phase 2 (2200 dishes) image output rate is 44 GBytes/s or 352 Gbit/s.

A.10 High Definition Multi-media

The data rates of 4k multi-media are impressive at 15.2 Gbit/s uncompress and 764 Mbit/s compressed (even at 382Mbit/s for 24 frames/sec). The work of projects like CineGrid or that in the Networked Media Laboratory at University of Essex is demonstrating the potential of this technology, for example in:

- Digital Cinema & movie industry
- Tele-medicine both live video of surgery to remote experts, diagnostics and teaching
- Performing arts – linking nations in real time
- Humanities and digital cultural heritage - documents; archaeological sites
- Research into display and visualisation techniques
- Visualisation of data & simulations along with supercomputing, Grids and cloud computing.

A.11 HEP

Currently the multi-10 Gigabit LHCOPN provides HEP with a push model where data flows from CERN to the Tier-1s making it available so users can do their analyses in as many places as possible. The raw data are processed (at least once and may be re-processed with the best available calibrations) and the “interesting” sub-set of the data flagged for Physics analysis. To help user access to this data they are discussing adding a pull model and caching the interesting data when users access it. The big change is expected to be in the data paths and access patterns with Tier-2 to Tier-2 becoming of a similar scale as the current Tier-1 to Tier-2 traffic.

A guesstimate is that a Tier-2 might need 10 Gbit/s but it is under discussion if this is routed IP, a set of 1G circuits, 10Gigabit on demand or a mixture.

Please note that there are many on-going discussions between HEP and the Networking community; the above is only a simple outline.

A.12 Bio-Informatics

It is reported that data volumes are growing very fast (in the Pbytes scale), that the cost of genome sequencing machines is such that many labs can now afford them, and the number of applications is growing. Also more imaging is being used.

Recent discussions with labs in China indicated 24/7 rates to Europe of ~100 Mbit in 2009 and a predicted need of ~460 Mbit/s by the end of 2010.

Recent traffic measurements of Bio-Informatics data in the UK showed that after the firewall hardware had been upgraded, rates of over 2 Gbit/s were sustained to one remote facility over a period of 2 days. Also the general traffic pattern now has more peaks of 1-2 Gbit/s. These data tend to support the anticipated growth of data transfers.

A.13 Tabular Summary

Health warning: the values noted are only rough and early estimates, and are expected to change as the user groups develop their work models.

Project, group or area	Requirement	Expected bandwidth	Impact, implication, issue for investigation
Astronomy HEP Bio-Informatics Climate & weather modelling	Data is bursty and on routed IP Need for advice with high-bandwidth transfers	100s Mbit/s Gigabits Gigabits	Load on country access links & backbones
Network development research Metrology Mobile network research FIRE	Specialised access Isolation from production flows The need for testbeds	Lambdas Lambdas	Interlink research islands
Dynamic Networking	Interest by several subject areas Need for cost advantage Ease of use in campus		Delivery to NREN Delivery to Campus/Lab
PRACE	Connect Tier-0 to national SC Tier-0 to Tier-0 Significant data to/from user sites	10G Lambdas DEISA 10G routed IP	Load on country access links & backbones
SKA	Phase 1 40 Gbit/s to world sites. User access to world regional sites	5-10 Gbit/s routed IP	World-wide Lambdas Load on country access links & backbones
HEP	LHCOPN for the data push Data pull and cache by Tier-2	Multiple 10G Lambdas 10G / Tier-2	How to deliver; what topology; possible use of dynamic networks
ESFRI project plans to use clouds or NGIs	User campus access to HPC, grids / NGI, clouds and storage facilities	Bandwidth of individual services unknown, but expected to be many	Load on country access links & backbones
HD and 4k TV/multi-media	Real-time users-supercomputers Tele-medicine links	380/760 Mbit/s compres 7.6/15.2 Gbit/s raw	Possible use of dynamic circuits Possible use of dynamic lambdas

Table A.1: Summary of user requirements

Appendix B Gap Analysis

As mentioned in Section 1.3 *NREN Input to Architecture Planning* on page 13, in order to ensure that the workplan for GÉANT architecture planning considered in an appropriate measure the different elements involved, members of the SA1 Supervisory Committee (SC) (appointed by the GN3 management team to oversee this work and provide strategic direction and advice) each presented their detailed view on what aspects the SA1 team needed to work on towards developing the next GÉANT architecture. This has resulted in a gap analysis, which explains how each aspect put forward by an SC member needs to be addressed by the SA1 team. The gap analysis covers the following topics: user requirements, topology issues, services, technology, virtualisation, service quality, use of NREN resources, federation, overlays, peerings, and L2 interconnections.

The table on the following pages groups the input received from NRENs into themes and then for each theme shows the individual items put forward (in the **DESCRIPTION** column), which NREN made the suggestion, and an assessment as to whether the topic is already part of the SA1 work plan. The **NREN Comments** column provides further clarification of the suggestion raised by the NREN, and is followed by an assessment of whether more work is needed. Finally, the **SC discussion** column summarises the discussion held in the SC on each topic.

Theme	DESCRIPTION	NREN	Item in SA1?	NREN Comments	Follow-up needed?	SC discussion
User requirements -> general	User needs	DFN, PIONEER, RENATER	NO	need a process to include user requirements. Some work has been done re NREN requirements but more needed especially to include e.g ESFRI. RENATER view to consider "common" users as well	yes, with NA4	The planned Gn3 service portfolio offers all the elements that are required to fulfill the projected needs of user projects. NA4 and SA1 need to work together to gain more precise user requirements and disseminate the GN3
	Expanding GEANT's Userbase	JANET(UK)		Widening the userbase of GEANT beyond R&E networks to public sector uses such as health and meteorology (and perhaps metrology!) may have architectural impacts. This should be considered for any future GEANT architecture.	YES, with SC and NA4	continue dialogue with these sectors via NA4 and the NRENs to ensure SA1 continues to have the right services and architecture to serve these emerging needs. It can only be seen as a continuous iterative process
	interactions with other sectors such as government and health	PIONEER	NO	mainly a policy matter, there may be some implications on architecture	policy/strategic issue for exec/nrenpc	as above and has implications on architecture and security. Not for SA1 T1 team at this stage, but more high level

Theme	DESCRIPTION	NREN	Item in SA1?	NREN Comments	Follow-up needed?	SC discussion
user requirements -> topology	IP traffic matrix analysis	DFN	YES	based on netflow data for IP. The demand pattern can determine the traffic matrix for static circuits and this can be used for planning purposes which in turn affects topology	NO	WIP, revisit when content peering traffic on-line. Do not repeat NREN questionnaire but look at history of traffic patterns. revisit topology on a continuous basis taking into account resilience and flexibly enable additional access points. Examine how topology enhancements can increase levels of resilience for services
	Topology for future needs	SURFnet	PARTLY	The current topology and distribution of traffic cannot be taken for granted when designing the future network. This needs	YES, within SA1 T1	
	Large data centre needs, topology to consider user sites	GARR	PARTLY	the topology studies and additional network access points cover this point in part. We need to include this aspect from GARR as one of the reasons to consider additional access points. Google is a	YES, within SC	
	Capacity planning for the interconnect backbone	SURFnet	PARTLY	Capacity planning for the interconnect backbone should be established in cooperation with all the other partners, based on a distribution of traffic across the facilities contributed by each partner.	YES, within SA1 T1	This is related to evaluation of use of NREN resources and is WIP re CBF

Theme	DESCRIPTION	NREN	Item in SA1?	NREN Comments	Follow-up needed?	SC discussion
user requirements -> flexible network	Coping With Increasing Traffic Levels and New Demands	JANET(UK)	PARTLY	The output of this task would be a document that describes the requirements for a network that will scale to meet the demands of its users for the next five to ten years. This will cover IP, wavelengths and sub-lambda circuits.	YES, as next step in GN	The planned Gn3 service portfolio offers all the elements that are required to fulfill the projected needs of user projects. NA4 and SA1 need to work together to gain more precise user requirements and disseminate the GN3
Services	Enable new types of network applications for Research and Education Community	CESNET	NO	Main purpose of GEANT is not to do service instead of ISPs, more advanced types of applications should be enabled (for example real time applications). All-optical sub-networks in GÉANT (clouds without OEO conversions) should be as large as possible. Wide-area dark fibre Experimental Facilities (testbeds) are	YES, within SA1 T1 and consultation with SC	use cases available. It is related to testbeds (see PSNC). Will need to prioritise taking costs into account
	Testbed	PIONIER	NO	some of these aspects are covered in JRA2 T5. We ought to formulate clear testing requirements. The utilisation of GN2 and GN1 testbeds has been very low from the community, limited to backbone operational tests, ipv6 and autobahn.	Yes, discussion in SC. There is also work ongoing in JRA1 but it appears not very coherent at present.	as above
	Development of advanced and innovative services	SURFnet	PARTLY	There is too little focus on the development of advanced and innovative services in the existing workplan. The network services will have to provide both static and dynamic connections between a large number of locations, many of which are as yet unknown.	YES, within SA1 T1 for physical and transmission layers and with collaboration with SA2 for layers above	ensure the services are accessible everywhere (PIONEER made same point) where it is requested (exp. GEANTPlus in Athens)
	all nrens should be able to access all gn3 services	PIONIER	NO	how can we offer i.e. Lambda services where we do not avail of fibre ?	SC discussion	as above

Theme	DESCRIPTION	NREN	Item in SA1?	NREN Comments	Follow-up needed?	SC discussion
Services -> QoS	Architectural constraints of end-to-end services	JANET(UK)		Whilst we're getting better at provisioning end to end services such as lightpaths within GEANT, we should consider any architectural requirements that would assist in doing that better, and on a wider scale.	Yes, with SA2	x-activity team to focus on the SA1 services and define their quality parameters (SA1-SA2-JRA1-JRA2)
	availability of GEANT IP	DFN	PARTLY	in network service resilience studies for the GEANT network part and also in additional access points item. DFN make the point about multidomain availability of IP with focus on operations. If this is to be tackled, it belongs to SA2.	yes, in SA1 T1	as above
	general QoS parameters for IP and L2 services	DFN	PARTLY	In network service resilience studies. There has been little work on identifying QoS parameters for L2 services, apart from	yes, in SA1 T1 with initial discussion in SC	as above
	granularity of BoD services Mbps - nx10Gbps	PIONEER	Partly	relates to SA2 work. It affects architecture and technology esp re granularity	yes in consultation with SA2	as above
	Near real-time configuration of multi-domain circuits	SURFnet	NO	The emerging need for rapid establishment of circuits across domains demands an architecture supporting near real-time configuration of multi-domain circuits. Once this is in place, such an architecture can be extended to include the required user-controlled and application controlled other ICT services.	YES, within SA1 T1	this is about having pre-provisioned capacity in the backbone. How much is needed and how to interface to nrens ? (i.e, today GEANTPlus has a 10G backbone)Dynamic circuit services is WIP. Start with 1Lambda
	set up time for e2e link	DFN	N/A	this is an operational development, which is addressed in SA2 on one hand. On the other hand, it is also a technology aspect in that some technologies facilitate faster provisioning of circuits	SA2	technology paramter to consider in equipment spec

Theme	DESCRIPTION	NREN	Item in SA1?	NREN Comments	Follow-up needed?	SC discussion
Services -> QoS ctd	SLA at European scale	GARR	N/A	this is an operational development issue, which ought to be tackled in SA2. It has started to be discussed in the multidomain services workshop 24-25 June. There need to clearly defined per-domain SLAs and e2e SLAs. The commercial market can provide SLAs, so	SA1	SA1 has to specify the service quality of GEANT services and define how they are monitored
	QoS for specific flows using statistical multiplexing technology	GARR	PARTLY	GEANT IP has Premium IP. For a future non-TDM based GEANTPlus service, network design rules are required. This part is missing	yes, in SA1 T1	in conjunction with specifying quality parameters for the various services, we need to clarify which technologies are suitable and how do we implement the quality guarantees. Capacity planning is
Services -> virtualisation	computing systems and use of virtualisation	GARR	NO	Specific needs have not yet been set out. This is a study item in JRA1 and JRA2	Ensure the architecture can accommodate other devices. Keep dialogue with JRA1 and JRA2 and SA2	technology choice should not preclude virtualisation. SA1 and JRA1 should work together to define which virtualisation services can be offered. What are the benefits to the users ?
	GN3 should facilitate network virtualization	NORDUnet	PARTLY	GN3 network architecture should enable virtualization and support for virtualized infrastructures.	YES, by ensuring in architecture specification	

Theme	DESCRIPTION	NREN	Item in SA1?	NREN Comments	Follow-up needed?	SC discussion
CBF/use of NREN resources	Rules for procurement of Lambdas over CBFs should be defined and agreed	CESNET	NO	Guidelines for NRENs preparing offer are needed (especially how to do cost claims)	YES, within SA1 T2	This is WIP. These rules should address aspects of cost as well as ability of a CBF solution to support lambda services. Emphasis on service reliability and economics. long term use of CBF should be compared to normal market taking into account past costs re installation and setup
	Use of NREN-owned fibre	JANET(UK)	NO	All forms of fibre acquirement used by NRENs should be considered (ownership, lease, IRU). Relation/differences to CBF based lambda services should be specified. Economical analysis and service reliability analysis are needed (building of reliable services by less	YES, as pilot experiment and preparation of strategic document for GN3 successor	
	An architecture encompassing NREN owned and operated resources	SURFnet	PARTLY	An architecture encompassing NREN owned and operated resources, including Cross Border Fibers and Open Exchanges, should be developed. Currently no architecture work with NREN owned and operated resources at the base is	YES, within SA1 T1	
	Cost-recovery model	SURFnet	NO	Out of the box cost-recovery model development, e.g. for NREN owned and operated resources, needs to be worked	YES, within SA1 T1	
	on use of CBF and nren resources	PIONEER, RENATER	YES	work on criteria for using CBFs is ongoing in SA1, looking at cost and operations	the output of the SA1 work to date needs to be discussed within the SC, and if needed more different work should be undertaken	

Appendix B Gap Analysis



Theme	DESCRIPTION	NREN	Item in SA1?	NREN Comments	Follow-up needed?	SC discussion
Architecture-> general	Outlook for a future NREN and GÉANT architecture based on JRA1 results	CESNET	NO	JRA1 deliverable is new and SA1 planning was done before	YES, within SA1 T1 and consultation with SC	should define which result of JRA1 will bring benefits to SA1. Should be done by SA1
	Open/libre architecture of future GÉANT transmission system	CESNET	NO	Freedom to provide services to Research and Education Community cannot be limited by development roadmap of one equipment vendor. Open multivendor	YES, within SA1 T1 and consultation with SC	need a through analysis of pros&cons of multivendor solutions conducted in parallel to a technical field trial
	NEW : availability of circuit services in other world regions	PIONIER	NO	how to make circuit services available i.e. in Asia	yes, general architecture work. At the moment clear plan to deliver in North America only	develop a plan to deliver circuit services in other world regions
Architecture -> Federation	GN3 should pursue federated networking	NORDUnet	NO	We have the prerequisites in place to include NREN owned-and-operated resources in the core GÉANT network on a substantial scale. We must ensure that the network architecture is built to include CBF, NREN-provided connectivity to external partners, and access to open exchanges, and that support systems and workflows are built to integrate resources and staff from several organizations.	YES, by planning transition from centralized to federated GEANT architecture	Use of CBF is covered in existing work. The architecture does not preclude more general federated networking concepts, but the definition, benefits (for example reduced cost, improved quality, better user experience, simplified operations), developments and business case for these should be done in JRA1 and follow the project hierarchy.

Theme	DESCRIPTION	NREN	Item in SA1?	NREN Comments	Follow-up needed?	SC discussion
Architecture -> Federation ctd	The development of new operational concepts in a federated environment	SURFnet	NO	The development of new operational concepts, in a federated environment, is lacking from the workplan.	YES, within SA1 T1 for physical and transmission layers and with collaboration with SA2 for layers above	
	Multi-domain federated networking.	SURFnet	PARTLY	There is currently hardly any view on multi-domain, federated networking. This should be added, and the work from JRA1 needs to be adopted already now.	YES, within SA1 T1	
	federated services	PIONEER, RENATER	N/A	extend scope of SA3, with easy-win topics such as GDS, VoIP, DNS mirrors, and extend dialogue with other co-ordination	very general matter, needs wider discussion in GN3	no impact, apart from housing equipment
architecture-> overlays	Design of network architecture with different topology of layers	SURFnet	PARTLY	Based on the definition of advanced and innovative services, we have to architect each layer of the network separately (L0, L1, L2 and L3) and in relation to each other. These layers and their PoPs can be optimized separately, and should not be constrained by the current fiber layout.	YES, within SA1 T1	yes more detailed work is required on this aspect
	overlay topologies, IP, ethernet	PIONEER	partly	the logical topologies of the overlays and the overall architecture are being looked at, but not thoroughly enough. Some limited work has been done by the team	YES, within SA1 T1 and consultation with SC	as above. Consider tailoring architecture of each PoP to local NREN needs ? Or one size fits all ?

Appendix B Gap Analysis



Theme	DESCRIPTION	NREN	Item in SA1?	NREN Comments	Follow-up needed?	SC discussion
architecture-> peerings	Exchange of Traffic with Commercial Peers	JANET(UK)	PARTLY	The SC should reflect on experiences gained through that and describe a coherent peering policy and service description for the commercial peering service beyond the short-term experiment.	YES, as next step in GN3 SA1 T1	The IP peerings is WIP. We should evaluate the experience. Consider operators as well as content providers ?
	Development of new IP peering fabric	SURFnet	NO	An IP peering fabric needs to be developed, not necessarily on L3.	YES, in collaboration of SA1 T1 with above SA	
	access to commercial peerings and open exchanges	PIONEER, RENATER	YES	access to commercial peerings is ongoing work. Access to open exchanges is also in place today via NRENs. This is mainly a policy matter though, as technically and architecturally this is already possible. RENATER view to facilitate for NRENs where commodity IP is difficult by	SC to advise if more work is needed	

Theme	DESCRIPTION	NREN	Item in SA1?	NREN Comments	Follow-up needed?	SC discussion
architecture -> I2 interconnect	GEANT node as L2 interconnect	DFN	YES	they are already L2 interconnects	Need to ensure L2 interconnects continue to allow to interconnect circuit services to other regions	Need. To clarify what L2 means. For the provision of point to point (ethernet) circuits, the GEANT nodes are already L2 interconnects. From a technical perspective no further work is needed apart from normal continuous evolution. Note that the "open" is very much a policy matter that should be solved at exec or nrenpc level
	Future GEANT Network should contain Open Lightpath Exchanges as Architectural Element	SURFnet	NO	The world-wide implemented and proven model of Open Lightpath Exchanges should be adopted as an architectural element in the Future GEANT Network. The "open" means that this switch fabric is policy free. Lightpath coordination should be undertaken on a global scale,	YES, within SA1 T1	

Theme	DESCRIPTION	NREN	Item in SA1?	NREN Comments	Follow-up needed?	SC discussion
technology	GN3 should pursue OTN technology	NORDUnet	PARTLY	For future high-end end-to-end circuit services, a true switching capability is required. Without such a capability, end-to-end services will remain on the	YES, by ensuring in architecture specification	these are candidate technologies that should be kept into account as possible solutions.
	GN3 should pursue an EoMPLS backbone	NORDUnet	PARTLY	European EoMPLS switching capability will be ideal for hosting a future GÉANT IP service, for acting as a European peering fabric allowing NRENs to peer directly and to reach a number of European peering points, and as the backbone for end-to-end circuit services	YES, by ensuring in architecture specification	
	MPLS technology for multipoint services	PIONEER	NO	PIONEER's experiences a high demand for multipoint. DEISA is also an example. In multidomain environment, the most important to place to start with is GEANT	in SA1 and SA2	SA2 will address in 2011, need to ensure technology choice does not exclude multipoint
	Monitoring Infrastructure	GARR	NO	I don't believe this materially affects the GEANT network architecture. It is a required add-on and of course monitoring and management capabilities of equipment needs to be considered in any procurement activity. It is essential for	Yes, as ongoing operational monitoring and management developments	no impact, apart from housing equipment

Theme	DESCRIPTION	NREN	Item in SA1?	NREN Comments	Follow-up needed?	SC discussion
High Level Planning	Dark fibre footprint for Research and Education Community is necessary	CESNET	PARTLY	Optimal leasing duration is very long (e.g. 10 years). Dark fibre footprint procurement should be started without additional delay (and repeated yearly for lines, where will be unsuccessful). At least two physically diverse last mile fibre routes are needed for GÉANT PoPs.	YES, set as priority subtask	SA1 should do
	involvement in ESFRI projects etc	PIONEER	NO	user workshops are not enough, GN3 engineers should be part of ESFRI projects in order to understand better their requirements. A bit like Federica, so we also understand better their	Not from SA1, it is an NRENPC matter	No consensus on this in SC meeting on 7 July. We should at least co-ordinate involvement of NRENs/DANTE with projects
	Projects involving China, Japan, India	RENATER	NO	Need to ensure that connectivity to these regions can fulfil project needs.	Yes, in SC and SA1	
	Bridging the 'Digital Divide.'	JANET(UK)	PARTLY	The output of this task would be a document that highlights what work should be done with the fibre owners and cable layers to improve matters for research and education.	YES, concurrently with fibre procurement and in collaboration with NA4 (liaison to industry)	keep an eye open to opportunities for accessing fibre in the less favourable areas.
	Dark fibre footprint topology should be largely independent on today knowledge of traffic and lambda topology	CESNET	PARTLY	Lighting equipment will be changed 2-3 times in period of fibre lease. Traffic maps are relevant for lambda topology, fibre topology could be different. Light speed in fibre and wave switching enables connection of PoPs by lambdas as needed.	YES, GÉANT Architecture will be layered (fibre topology, all-optical transmission topology, OTN transmission topology, digital services topology, IP services topology...)	
	NREN having fibres into GEANT PoP will participate in procurement process .	CESNET	NO	To avoid pointless duplication of fibre footprint segments and to achieve evaluation of the best solution.	YES, within SA1 T2	SA1 work

Appendix C Third-Party Connectivity

As part of the architecture planning, SA1 T1 has held ongoing discussions with NRENs about their ability to provide cross-border fibre (CBF) solutions. A Connectivity Resources Questionnaire [NREN_CBF_QAIRE] was issued in September 2009. Its aim was to understand the optical fibre and wavelength resources, popularly known as cross-border fibre (CBF) (though, more precisely, lambdas provided across NRENs' lit fibre infrastructure), that NRENs are in a position to contribute to the GÉANT infrastructure. It is possible that, in some areas, the current network topology, which is built on dark fibre and wavelength connectivity between GÉANT Points of Presence (PoPs) leased from various commercial companies, could be augmented and/or some parts replaced with connectivity resources provided by the NRENs themselves, with the aim of reducing the total cost of running the GÉANT network.

The questionnaire identified some possibilities for CBF-based “half-circuit” solutions, where an NREN uses its own resources to provide the part of the circuit within their country up to the border, and some possibilities for potential full end-to-end CBF solutions.

As a follow-up to the questionnaire, and after a series of meetings and consultations between members of GN3 management, SA1 Tasks 1 and 2 and a process analyst, a process has been defined for identifying and evaluating where resources offered by NRENs (CBF) can be integrated into the network where it makes technical and economic sense to do so. Known as the Capacity Acquisition Process, it was reviewed by the SA1 Supervisory Committee on 21 June 2010. The next stage is for the SA1 Supervisory Committee to finalise and agree the two documents that form part of the process, namely, the Managed Connectivity Services and Service Level Requirements questionnaire [CBF_Reqts], and related scoring document, which were piloted by NORDUnet for the link to Helsinki from Copenhagen, and to obtain Executive Committee approval for the process, questionnaire, and scoring document.

The primary objectives of the capacity acquisition process are to:

- Ensure independent and impartial procurement of new capacity.
- Deliver the most cost-effective connectivity solution.

Evaluation criteria for the resources offered by NRENs include cost, quality, timeframe, operations and the actual need for the resources within GÉANT.

There are borders in Europe and its neighbours, where NRENs have not yet built CBF and therefore GN3 received no offers to the questionnaire. On the other hand, procurement of such CBFs could bring further possibilities of cost-effective solutions for GÉANT. This looks to be very important from a long-term strategic

view (a way to improve cost-effectiveness and network service reliability). We should investigate, for example, possibilities of CBF solutions for the connection of “off-fibre” countries and possibilities of mesh enhancement using CBF for “on-fibre” countries (but with only one line connection at present). An overview of fibre footprints used by NRENs is available in the TERENA Compendium and in the proceedings of the CEF Networks workshops.

Figure C.2 presents an overview of the capacity acquisition process. The full process is documented in “Capacity Acquisition Process” [CAP].

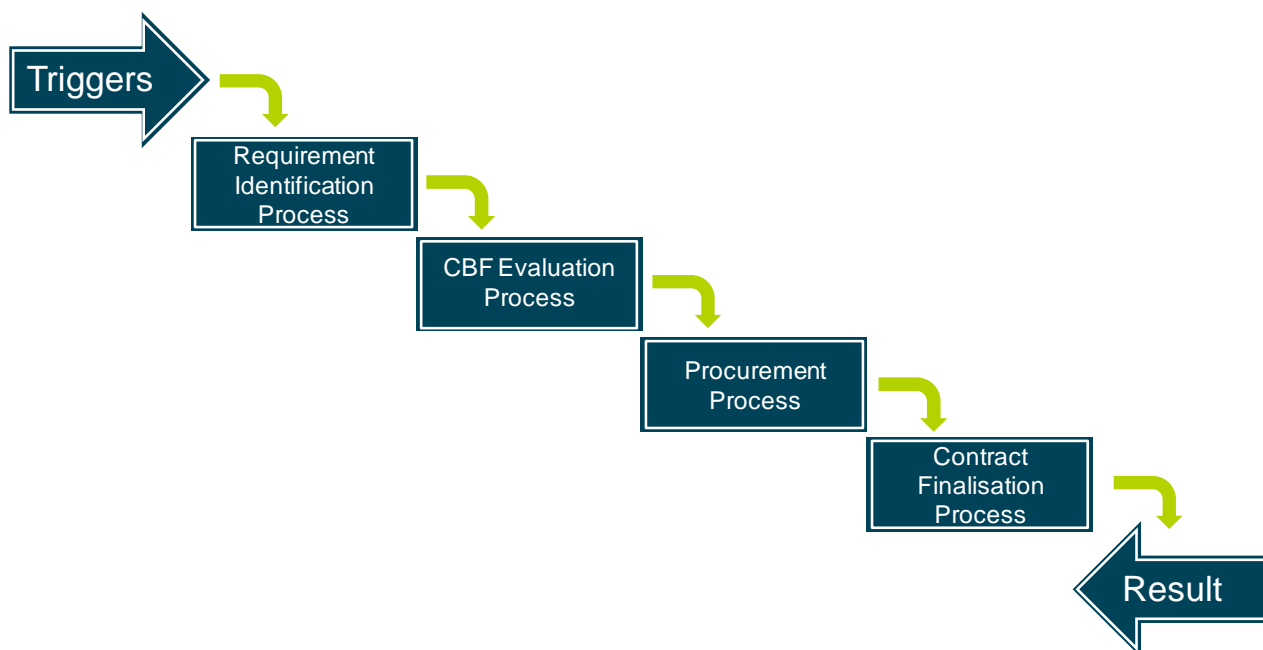


Figure C.2: Capacity acquisition process overview

References

- [CAP]** “Capacity Acquisition Process”
<https://intranet.geant.net/sites/Services/SA1/T2/Documents/GN3-10-312v1%20a%20Capacity%20Acquisition%20Process.doc> [access limited to GN3 project partners]
- [CBF_Reqs]** “Managed Connectivity Services and Service Level Requirements for GÉANT Network”
<https://intranet.geant.net/sites/Services/SA1/T2/Documents/GN3-10-313v1.4%20Managed%20Services%20Requirements%20Document.doc> [access limited to GN3 project partners]
- [CEFWorkshop]** 6th Customer Empowered Fibre (CEF) Networks Workshop, Prague, 13 – 14 September 2010
<http://www.ces.net/events/2010/cef/>
- [DJ1.1.1]** GN3-09-224 “Deliverable DJ1.1.1: Transport Network Technologies Study”
http://www.geant.net/Media_Centre/Media_Library/Media%20Library/GN3-09-224-DJ1-1-1v1-0_Transport_Network_Technologies_Study_Read_Only.doc
- [DJ1.2.1]** GN3-10-122 “Deliverable DJ1.2.1: State-of-the-Art Photonic Switching Technologies”
http://www.geant.net/Media_Centre/Media_Library/Media%20Library/GN3-10-122%20DJ1%202%201v1%200%20State%20of%20the%20Art%20Photonic%20Switching%20Technologies_Read%20Only.doc
- [DJ1.3.1]** GN3-09-250 “Deliverable DJ1.3.1: Architecture Considerations for Federated Backbone Networks – Study”
http://www.geant.net/Media_Centre/Media_Library/Media%20Library/GN3-09-250%20DJ1.3.1v1.0%20Architecture%20Considerations%20for%20Federated%20Backbone%20Networks%20Study.pdf
- [DS1.1.1]** GN3-09-227 “Deliverable DS1.1.1 Report on the Backbone Architecture Study”
http://www.geant.net/Media_Centre/Media_Library/Media%20Library/GN3-09-227%20Deliverable%20DS1-1-1%20Report%20on%20Backbone%20Arch%20Study%20v1.6.pdf
- [DS1.1.2]** GN3-09-228 “Deliverable DS1.1.2: Procurement Plan”
<https://intranet.geant.net/sites/Services/SA1/Deliverables/DS112/Documents/GN3-09-228%20DS1.1.2%20Procurement%20Plan%20v1.4.pdf> [access limited to GN3 project partners]
- [DS1.2.2,1]** GN3-10-230 “Deliverable DS1.2.2,1: Follow-Up on Short-Term Procurement Plans”
<https://intranet.geant.net/sites/Services/SA1/Deliverables/DS1.2.2,1/Documents/GN3-10-230v1.0%20DS1.2.2,1%20Follow-Up%20on%20Short-Term%20Procurement%20Plans%20-%20final%20version.pdf> [access limited to GN3 project partners]

- [GÉANT]** GÉANT Backbone Topology Map
http://www.GÉANT.net/Media_Centre/Media_Library/Media%20Library/1081%20GN%20Topology%2011.09_5_High%20Res.pdf
- [GÉANT_Perf]** “GÉANT Performance Overview” presentation
- [GÉANT_Svc_Proposals]** GN3-09-256 “GÉANT Backbone Service Proposals”
- [GN3 white paper]** GN3 Proposal Review Board, “GÉANT3 White Paper”, 15 July 2008
<http://wiki.geant.net/pub/PD/DocumentsMain/GN3-08-034-GN3-White-Paper.pdf> [access limited to GN3 project partners]
- [IEEE 802.1ad]** IEEE 802.1ad, “Provider Bridges”
<http://www.ieee802.org/1/pages/802.1ad.html>
- [IEEE 802.1ag]** IEEE 802.1ag, “Connectivity Fault Management”
<http://www.ieee802.org/1/pages/802.1ag.html>
- [IEEE 802.1ah]** IEEE 802.1ah, “Provider Backbone Bridges”
<http://www.ieee802.org/1/pages/802.1ah.html>
- [IEEE 802.1aq]** [IEEE 802.1aq, “Shortest Path Bridging”
<http://www.ieee802.org/1/pages/802.1aq.html>
- [IEEE 802.1D-2004]** IEEE 802.1D, “MAC Bridges”
<http://www.ieee802.org/1/pages/802.1D-2003.html>
- [IEEE 802.1q]** IEEE 802.1q, “Virtual LANs”
<http://www.ieee802.org/1/pages/802.1Q.html>
- [IEEE 802.1Qat]** IEEE 802.1Qat, “Stream Reservation Protocol”
<http://www.ieee802.org/1/pages/802.1at.html>
- [IEEE 802.1Qav]** IEEE 802.1Qav, “Forwarding and Queuing Enhancements for Time-Sensitive Streams”
<http://www.ieee802.org/1/pages/802.1at.html>
- [IEEE 802.1Qay]** IEEE 802.1Qay, “Provider Backbone Bridge Traffic Engineering”
<http://www.ieee802.org/1/pages/802.1ay.html>
- [IEEE 802.3ah]** IEEE 802.3ah, “Ethernet in the First Mile Task Force”
<http://grouper.ieee.org/groups/802/3/efm/public/index.html>
- [IETF SPB IS-IS]** P. Ashwood-Smith, D. Fedyk, D. Allan, J. Chiabaut, N. Bragg, “Shortest Path Bridging and Backbone Bridging with IS-IS”, Internet Draft, June 2009
<http://tools.ietf.org/html/draft-fedyk-isis-spb-00>
- [IETF TRILL]** P. Ashwood-Smith, D. Fedyk, D. Allan, J. Chiabaut, N. Bragg, “Shortest Path Bridging and Backbone Bridging with IS-IS”, Internet Draft, June 2009
 R. Perlman, D. Eastlake, D. G. Dutt, S. Gai, A. Ghanwani, “RBridges: Base Protocol Specification”, Internet Draft, January 2010
<http://tools.ietf.org/html/draft-ietf-trill-rbridge-protocol-15>
- [ITU G.709]** ITU G.709, “Interfaces for the Optical Transport Network (OTN)”
<http://www.itu.int/rec/T-REC-G.709/e>
- [ITU-T G.8011.1]** ITU-T G.8011.1, “Ethernet Private Line Service”
<http://www.itu.int/itudoc/itu-t/aap/sg15aap/history/g8011.1/index.html>
- [ITU-T G.8011.2]** ITU-T G.8011.2, “Ethernet Private Line Service”
<http://www.itu.int/itudoc/itu-t/aap/sg15aap/history/g8011.2/index.html>
- [ITU-T Y.1731]** ITU-T Y.1731, “OAM functions and mechanisms for Ethernet based networks”
<http://www.itu.int/itudoc/itu-t/aap/sg13aap/recaap/y1731/>
- [IXP_Usage]** “Commodity IP and IXP usage in GN3”: presentation by M. Enrico and O. Kreiter to Third Architecture Workshop 7–8 July 2009. Access restricted to GN3 project participants.

- http://wiki.geant.net/pub/Main/ArchWshop3/NEP-09-042v6_Commodity_IP_and_Peering_in_GN3.pdf
- [MEF 16]** MEF 16, "Ethernet Local Management Interface"
http://metroethernetforum.org/PDF_Documents/technical-specifications/MEF16.pdf
- [Nortel PLSB]** D. Fedyk, "Provider Link State Bridging"
<http://www.ieee802.org/1/files/public/docs2007/aq-fedyk-plsb-present-0107.pdf>
- [Noutsios]** P. Noutsios et al., "Optical latency on a core network", Journal of Optical Networking, Vol. 3 Issue 4, pp. 242–246 (2004).
- [NREN_CBF_QAIRE]** GN3-09-197v2 NREN Connectivity Resources Questionnaire
[restricted access – GN3 internal use only]
- [Phosphorus]** www.ist-phosphorus.eu
- [RFC 2547bis]** E. C. Rosen, Y. Rekhter, RFC 2547bis, "BGP/MPLS IP VPNs", Internet Draft, October 2004, [draft-ietf-l3vpn-rfc2547bis-03.txt](http://tools.ietf.org/html/draft-ietf-l3vpn-rfc2547bis-03)
<http://tools.ietf.org/html/draft-ietf-l3vpn-rfc2547bis-03>
- [ROADM_Sandesha]** K. Sandesha, "Colorless, Directionless and Contentionless multi-degree ROADM architecture for mesh optical networks"
<http://www.cedt.iisc.ernet.in/people/kuri/Comsnets/Presentations/Colorless-Devarajan.pdf>

Glossary

10GE	10 Gigabit Ethernet
3R	Re-amplifying, Reshaping, Retiming
ALU 1626 LM	Alcatel-Lucent 1626 Light Manager (DWDM equipment)
ALU 1678 MCC	Alcatel-Lucent 1678 Metro Core Connect (optoelectrical switching equipment)
ATM	Asynchronous Transfer Mode
AUP	Acceptable Use Policy
AutoBAHN	Automated Bandwidth Allocation across Heterogeneous Networks
B-DA	Backbone Destination Address in 802.1ah MAC-in-MAC header
BER	Bit Error Rate
BFD	Bi-directional Forwarding Detection
BGP	Border Gateway Protocol
BoD	Bandwidth on Demand
BOL	Beginning of Life
B-VID	Backbone VLAN ID in 802.1ah MAC-in-MAC header
BW	Bandwidth
CAPEX	Capital Expenditure
CBF	Cross-Border Fibre
CBS	Committed Burst Size
CTTI	Carrier Class Transport Infrastructure
CCTNT	Carrier Class Transport Network Technology
CD	Chromatic Dispersion
cE	carrier Ethernet
CET	carrier Ethernet Transport
CFM	Connectivity Fault Management
CFP	100 Gigabit Small Form Factor Pluggable
CIR	Committed Information Rate
CMD	Composite Material Dispersion
CO-PS	Connection Oriented – Packet Switched
CoS	Class of Service
CP	Control Plane
CPE	Customer Premises Equipment
DCM	Dispersion Compensation Module
DRAC	Dynamic Resource Allocation Controller
DSL	Digital Subscriber Line
DSP	Digital Signal Processing

DWDM	Dense Wavelength-Division Multiplexing/Multiplexed
E2E	End to End
EDFA	Erbium Doped Fibre Amplifier
E-FEC	Extended FEC
E-FTTx	Ethernet Fibre to the x
E-LMI	Ethernet Local Management Interface
EIR	Excess Information Rate
E-LAN	Ethernet LAN (Service)
E-Line	Ethernet Line (Service)
E-OAM	Ethernet OAM
EOL	End of Life
EoMPLS	Ethernet over MPLS
EPL	Ethernet Private Line
ESFRI	European Strategy Forum on Research Infrastructures
E-TREE	Ethernet Tree (Service)
EVPL	Ethernet Virtual Private Line
FCAPS	Fault, Configuration, Accounting, Performance and Security
FEC	Forward Error Correction
FIB	Forward Information Base
FPGA	Field-Programmable Gate Array
FRR	Fast Reroute
GE	Gigabit Ethernet
GFF	Gain-Flattening Filtering
GFP-F	Generic Framing Protocol – Framed
GLIF	Global Lambda Integrated Facility
GMPLS	Generalised Multi-Protocol Label Switching
GOLE	GLIF Open Lightpath Exchange
GUI	Graphical User Interface
H-QoS	Hierarchical QoS
H-VPLS	Hierarchical Virtual Private LAN Service
IEEE	Institute of Electrical and Electronics Engineers
IETF	Internet Engineering Task Force
IGP	Interior Gateway Protocol
ILA	In-Line Amplification/In-Line Amplifiers
IP	Internet Protocol
IPoDWDM	IP over DWDM
IPTV	Internet Protocol Television
IS-IS	Intermediate System to Intermediate System
ITU	International Telecommunication Union
ITU-T	ITU – Telecommunication Standardisation Sector
IXP	Internet Exchange Point
JRA1	GN3 Joint Research Activity 1, Future Network
JRA1 T1	JRA1 Task 1, Core Network Technologies
JRA1 T2	JRA1 Task 2, Photonics
JRA1 T3	JRA1 Task 3, Federated Network Architectures
JRA2	GN3 Joint Research Activity 2, Multi-Domain Network Service Research
L1	Layer 1

L2	Layer 2
L3	Layer 3
LAN	Local Area Network
LCAS	Link Capacity Adjustment Scheme
LSP	Label Switched Path
LSR	Label Switch Router
MAC	Media Access Control
MBGP	Multiprotocol Border Gateway Protocol
MEF	Metro Ethernet Forum
MPLS	Multi-Protocol Label Switching
MPLS-TP	Multi-Protocol Label Switching Transport Profile
MSPP	Multi-Service Provisioning Platform
MSTP	Multi-Service Transport Platforms
MTBF	Mean Time Between Failures
MTTR	Mean Time to Recovery
NG	Next Generation
NGN	Next-Generation Network
NG-SDH	Next-Generaton Synchronous Digital Hierarchy
NMS	Network Management System
NREN	National Research and Education Network
NRZ	Non Return to Zero
OAM	Operations, Administration and Maintenance
ODU	Optical Channel Data Unit
OEO	Optical-Electrical-Optical
OOO	Optical-Optical-Optical (optical input, optical switching, optical output)
OPEX	Operating Expenditure
OPN	Optical Private Network
OPUk	Optical Channel Payload Unit-k
OSNR	Optical Signal-to-Noise Ratio
OSPF	Open Shortest Path First
OSS	Operational Support System
OTDR	Optical Time-Domain Reflectometer
OTN	Optical Transport Network
OWD	One-Way Delay
P2MP	Point to Multipoint
P2P	Point to Point
PB	Provider Bridges
PBB	Provider Backbone Bridges
PBB-TE	Provider Backbone Bridge Traffic Engineering
PBT	Provider Backbone Transport
PCE	Path Computation Element
PDH	Plesiochronous Digital Hierarchy
PDL	Polarisation-Dependent Loss
PIM-SSM	Protocol Independent Multicast – Source Specific Multicast
PLSB	Provider Link State Bridging
PMD	Polarisation Mode Dispersion
PM-QPSK	Polarisation Modulation Quadrature Phase Shift Keying

PON	Passive Optical Network
PoP	Point of Presence
PoS	Packet over SDH
POTS	Packet Optical Transport Service
PW	Pseudowire
PWE3	Pseudowire Emulation Edge to Edge
PXC	Photonic Cross Connect
QFP	Queuing and Forwarding Protocol
QoE	Quality of Experience
QoS	Quality of Service
R&E	Research and Education
RFC	Request for Comment
RFI	Request for Information
RFP	Request for Proposal
RFQ	Request for Quotation
ROADM	Reconfigurable Optical Add-Drop Multiplexer
RSTP	Rapid Spanning Tree Protocol
RSVP-TE	Resource Reservation Protocol – Traffic Engineering
RTT	Round-Trip Time
SA1	GN3 Service Activity 1, Network Build and Operations
SA1 T1	SA1 Task 1, Network Planning and Procurement Preparation
SA2	GN3 Service Activity 2, Multi-Domain Network Services
SC	Supervisory Committee
SDH	Synchronous Digital Hierarchy
SE	Synchronous Ethernet
SFP	Small Form Factor Pluggable
SG-15	ITU-T Study Group 15
SLA	Service Level Agreement
SONET	Synchronous Optical Networking
SPB	Shortest Path Bridging
SPBM	SPB MAC
SPBV	SPB VID
SPoF	Single Point of Failure
SRG	Shared Risk Group
SRP	Stream Reservation Protocol
STM	Synchronous Transport Module
STP	Spanning Tree Protocol
TCM	Tandem Connection Monitoring
TCO	Total Cost of Ownership
TCP	Transmission Control Protocol
TDM	Time-Division Multiplexing / Time-Division Multiplexed
TE	Traffic Engineering
TRILL	Transparent Interconnection of Lots of Links
T-ROADM	Tuneable ROADM
UNI	User-to-Network Interface
VC	Virtual Circuit
VCAT	Virtual Concatenation

Glossary

VID	VLAN Identifier
VLAN	Virtual Local Area Network
VOA	Variable Optical Attenuator
VoD	Video on Demand
VoIP	Voice over IP
VPLS	Virtual Private LAN Service
VPMS	Virtual Private Multicast Service
VPN	Virtual Private Network
VPWS	Virtual Private Wire Service
WAN	Wide Area Network
WDM	Wavelength-Division Multiplexing/Multiplexed
WRED	Weighted Random Early Detection
WSO	Wavelength Switched Optical Network
WSS	Wavelength Selective Switching
XFP	10 Gigabit Small Form Factor Pluggable