30.08.07

Deliverable DN3.0.5: Processes and Provision of Point-to-Point Services in GÉANT2

Deliverable DN3.0.5

Contractual Date:	31/05/07
Actual Date:	30/08/07
Contract Number:	511082
Instrument type:	Integrated Infrastructure Initiative (I3)
Activity:	NA3
Nature of Deliverable:	R (Report)
Dissemination Level	RE (Restricted)
Lead Partner	DANTE
Document Code	GN2-07-131v8

Authors: John Chevers (DANTE), Dale Robertson (DANTE), Michael Enrico (DANTE), Marian Garcia-Vidondo (DANTE), Xavier Martin-Rivas (DANTE)

Abstract

This document describes the point-to-point services deployed by the GÉANT2 project and the procedures put in place to facilitate their delivery. The background to the introduction of these services is discussed, both in terms of user demand, the available technologies, their capabilities and the economic circumstances which allowed their introduction. The different types of GÉANT2 service are described in detail, including the different technical variations offered and the geographic reach of the circuits. Importantly, the processes created to deliver these services, both within GÉANT2 and across multiple domains are outlined and the future plans for development of customer service are discussed.

Table of Contents

Table of Figures

0 Executive Summary

GÉANT2 is a hybrid network, combining the operation of a shared IP infrastructure with the ability to provide additional, switched, circuits reserved strictly for particular user groups. This dual role makes GÉANT2 unique: it is the first hybrid production network operating on an international scale.

The GÉANT2 point-to-point service is one of the most important developments of this, the seventh generation of European research and education backbone network. It reflects the evolution of technology and the telecommunications market and an ever-growing demand from the user community, not only for a higher volume network, but also for guarantees of the capacity available and of the network's performance. Point-to-point circuits are offered between NRENs in Europe where it has been possible to procure at an affordable price the necessary network infrastructure – usually dark (unlit) fibre optic cables on which circuits can be incrementally added as demand requires.

The GÉANT2 point-to-point service offers for the first time the opportunity to configure user-designated circuits of between 1Gbps and 10Gbps capacity between many GÉANT2 points of presence (PoPs) across Europe. Larger projects can request a series of such circuits to form an Optical Private Network (OPN) –effectively a stand-alone internet linking that user group's sites. A total of 22 GÉANT2 point-to-point circuits of this type were delivered between June 2006 and June 2007.

The availability of the infrastructure and equipment is not sufficient for the delivery of the point-to-point service: to efficiently deliver GÉANT2 point-to-point services to European NRENs, an effective administrative and implementation process was quickly recognised to be essential. To deliver an end-to-end service from user site to user site it is equally essential that effective coordination and communication exists between the domains[4] involved. Significant effort has gone into developing processes to coordinate the handling of requests and the management of circuits. Work is also underway to develop the service interface to improve communication and allow tracking of progress by the service requester.

The deliverable reports on the current state of the developments as of June 2007. In so much it is a snapshot of ongoing development and of achievements made, however much additional work has still to be performed in this field to realise the many potential benefits of these services.

[4] For the purposes of this deliverable, a domain is defined as a set of network elements (equipment and connectivity) deployed as one unified network infrastructure, administered under a single Autonomous System (AS) number and managed by a single entity such as DANTE, an NREN or an institution (usually a university) behind an NREN.

1 Introduction

The GÉANT2 network offers the European research and education community a new range of opportunities for international collaboration. In addition to a shared IP infrastructure, GÉANT2 can provide point-to-point connectivity between the network points of presence (PoPs) which it connects. Additionally, circuits may be provided over the alternative international routes provided by Cross Border Fibre (CBF) operated by GÉANT2 NREN partners. Dedicated point-to-point circuits, either complete 10Gbps wavelengths or sub-divisions thereof, are a new means of delivering high-bandwidth international network connections free from the constraints inherent in a shared, routed infrastructure.

Point-to-point circuits come in two designations:

- An NREN within the dark fibre "cloud"[5] receives, in addition to its IP connection, a GÉANT+ subscription. This service provides the NREN up to 10Gbps of pre-paid point-to-point capacity between the GÉANT2 point-of-presence (PoP) in its own country and the GÉANT2 PoPs connecting similarly-subscribing NRENs. This capacity can be used to provide connections dedicated to individual research and education projects. Because this capacity is pre-provisioned, circuits can be implemented or reconfigured at short notice and without incremental cost to the NREN. These circuits can also be extended across the Atlantic.

- Full 10Gbps wavelengths dedicated to a particular user group can be delivered between the GÉANT2 PoPs (or, indeed, over other routes if CBF is available) in various European NRENs. These international segments will in general be connected to national point-to-point infrastructures to create a dedicated end-to-end connection.

The portfolio of point-to-point services offered over the GÉANT2 dark-fibre footprint, and its background, are explained in detail in this document.

[5] Point-to-point services use dark fibre and therefore are generally only available to NRENs in those countries where dark fibre has been procured for GÉANT2. For more details see section 3.2.

2 Background to P2P Services

The research and education backbone network in Europe has now reached its seventh generation. All of these networks have provided an IP backbone connecting together NRENs using links with shared capacity.

The network requirements of the academic community have consistently exceeded those of the wider community, both in terms of capacity and the quality of the service provided. Thus the research and education networks have responded with infrastructures offering considerable advantages over the commercial, or commodity, Internet. Typically high quality services have been ensured by deliberate over-provisioning of the network. In this case, although the network is shared, there will be little competition for capacity between data flows and thus the quality of service is high. In addition it is possible, through functionality provided on the network routers, to provide quality of service (QoS) and multi-protocol label switching (MPLS). Whilst both of these technologies provide significant advantages over a standard IP service for a small number of demanding users, both have limitations.

QoS allows certain traffic to be expedited through the router's queuing system, meaning that jitter associated with network congestion is minimised. Quality of service does nothing, however, to address the jitter associated with the unpredictable rerouting of packets due to network failure. Most importantly, QoS can only be applied to between 10 and 20% of capacity on any network route, before queuing and congestion affect the performance offered. Thus, for increasing volumes of the highest quality traffic, ever more IP network capacity must be provided, which is highly inefficient if the additional 80% of non-QoS capacity is not required.

MPLS allows a predefined route to be chosen for certain traffic flows. This means that traffic engineering can take place to achieve a desired latency or avoid points of known congestion. When QoS is also enabled, a very well controlled service can be achieved. In the case of network failure, a secondary, preferred path may also be defined. MPLS does not, however, reserve or guarantee network capacity and should not be confused with dedicated point-to point circuits.

It is important to note that neither QoS nor MPLS, nor a combination of both technologies, actually increases the network capacity available. As such, they can only be used to improve the service offered to a limited number of users and do nothing to address the root causes of network congestion.

The telecommunications market in Europe has evolved to the extent that in recent years leasing or purchase of unlit fibre and transmission infrastructure has become an economically viable alternative to leased wavelengths. The GÉANT2 network is thus based on a mixture of leased circuits, leased wavelengths and leased dark (unlit)

fibre. On those links of the network for which dark fibre is used, additional incremental capacity can be brought into service relatively cheaply. This opened the way for the development of GÉANT2 as a hybrid network, offering the use of switched point-to-point circuits alongside the more traditional routed IP backbone.

The ready availability of additional infrastructure and the increasing demands for ever-greater network capacity from the most data-hungry of users, have led to the development of the GÉANT2 point-to-point service. Whilst still operating a shared IP network, the provision of user-dedicated layer 1 and layer 2 circuits on an international scale is an innovation for GÉANT2 and makes GÉANT2 the first international production network of its type.

3 Description of Services Offered

GÉANT2 is a backbone or transit network. As such it does not (with a special construct for CERN and the DEISA switch) provide connectivity directly to the end-sites of the research and education community which it serves. The NRENs are GÉANT2's direct "customers". It is therefore vital that effective communication exists between DANTE and the NRENs. This is especially so in the case of point-to-point connections where an end-to-end service can only be achieved through effective interconnection of the component circuits. Dedicated circuits are thus requested on behalf of user projects by the NRENs which connect the end-sites to the GÉANT2 backbone.

GÉANT2 offers two distinct classes of point-to-point services to NRENs requiring dedicated international circuits. The principal benefits of both are identical: they provide completely dedicated network capacity to the user group concerned - network capacity is not only available to them, but is guaranteed to be available, regardless of other network traffic; dedicated circuits, being free of routing equipment, can guarantee constant latency between end-sites and thus eliminate jitter; as no other users have access to the circuit, data security is maximised and the user is freed from any constraints imposed by protocols designed to be 'fair' to other network users.

The availability of point-to-point services is primarily dependent upon the use of dark fibre for the underlying infrastructure. For reasons of availability and/or economic viability of dark fibre it is not possible to offer point-to-point services to all GÉANT2-connected NRENs. This section lists the NRENs on the 'fibre cloud'. It is expected that new fibre-based NRENs will join this sub-group in the near future.

The following GÉANT2 PoPs provide access to GÉANT2 point-to-point services and can also connect point-to-point circuits from outside the GÉANT2 administrative area provided that acceptable use criteria and operational procedures can be agreed:

- Vienna (ACOnet; Austria)
- Ljubljana (ARNES; Slovenia)
- Brussels (BELNET; Belgium)
- Zagreb (CARNet; Croatia)

- Prague (CESNET; Czech Republic)
- Frankfurt (DFN; Germany)
- Milan (GARR; Italy)
- Dublin (HEAnet; Ireland)
- Budapest (NIIF; Hungary)
- Copenhagen (NORDUnet; Nordic region)
- Poznan (PSNC; Poland; currently GÉANT+ service only)
- Madrid (RedIRIS; Spain)
- Paris (RENATER; France)
- Bratislava (SANET; Slovakia)
- Amsterdam (SURFnet; Netherlands)
- Geneva (SWITCH; Switzerland)
- London (UKERNA; United Kingdom)

3.1 Classes of GÉANT2 Point-to-Point Service

The two classes of GÉANT2 point-to-point service are described in the following sections:

3.1.1 The GEANT+ Service

The GÉANT+ service consists of a 10G IP service plus access to as many light paths – with granularity above 1G - as one lambda can carry.[6]

All NRENs subscribing to GÉANT+ have a pre-paid quota of 10Gbps of GÉANT+ capacity. In practice this means that up to 9 x GEth circuits (or equivalent) can be requested at short notice to any combination of NRENs within the GÉANT2 fibre cloud (see definition later in this section). The standard interfaces provided for GÉANT+ are 10GEth or STM-64 (depending on the preference of the NREN).

[6] It is assumed that capacities lower than 1G can be accommodated on the IP network.

As these services are pre-subscribed there is no charge associated with GÉANT+ point-to-point orders. A GÉANT+ circuit can be requested by NRENs using the web-order form located on a password-protected page of the GÉANT2 stats portal:

http://stats.geant2.net/p2p/

See also Appendix A for a copy of the initial order-form spreadsheet (valid June 2006- August 2007) and associated instructions, although a recent development now provides a similar form, which can be submitted on-line, with improved functionality –as visible on the above link.

3.1.2 Full Wavelength (10Gbps) Services

Full capacity 10Gbps wavelengths can be ordered between any two fibre-cloud GÉANT2 points-of-presence (see next section for details). Such circuits can be configured as STM-64 or 10GEth and with either short-reach (1330nm) or long-reach (1550nm) optics specified at each NREN interface.

Should dissimilar interface types be required on each end of a single 10Gbps circuit, this can be accommodated using the GÉANT2 MCC switching equipment. See the 'Technologies' section of this document for details.

The form for ordering full wavelength services is the same as for GÉANT+ services described above.

3.2 Transatlantic Circuits

There are two 10Gbps transatlantic circuits dedicated to providing point-to-point circuits between GÉANT2 and the US research networks operated by Internet2 and ESNET. One connects Paris to New York (provided by GÉANT2), the other London to New York (reciprocally provided by Internet2). GEth circuits can be provided on these links. An additional charge, reflecting the cost to the project of these circuits, will be made to those GÉANT+ circuits which are extended across the Atlantic. Since the use of the transatlantic circuits means by definition that networks outside of the GN2 consortium are involved in the end-to-end circuit, a special engineering and coordination procedure is under development to facilitate the configuration of these circuits.

As with the GÉANT2 PoPs listed in section 3.2, these transatlantic circuits can be used to extend point-to-point connectivity globally using onward connectivity provided by cooperating networks worldwide.

4 P2P Services in the GÉANT2 Domain

4.1 Circuit Type

The GÉANT2 IP network is designed to be over-provisioned, allowing small-to-medium sized traffic flows an uncongested path. Unmanaged flows above 1Gbps are considered to run the risk of impacting other IP traffic and suffering congestion. The GÉANT2 point-to-point service offers circuits of between 1Gbps and 10Gbps which avoid congestion and offer un-contended service over the GÉANT2 domain.

4.2 Interface Type

4.2.1 GÉANT+

Because the transmission and switching equipment used for GÉANT+ circuits are pre-provisioned, this classification of service allows GÉANT2 to offer NRENs flexible, quickly configurable circuits between GÉANT2 PoPs across Europe. The GÉANT+ subscription cost includes a 10GEth or STM-64 interface on the GÉANT2 equipment, supporting up to 9 GEth circuits, although alternative capacity circuits (e.g. STM-16) are also possible. As described above, capacity granularities below 1Gbps are not offered.

4.2.2 10 Gbps Circuits

10 Gbps circuits can be provided with either 10GEth or STM-64 interfaces, depending on preference. Unless a media converted circuit is specified, interfaces at the A and B ends of the circuit must be of the same type. 10Gbps circuits must be planned several months ahead of the ready-for-service date, due to the long lead times associated with the optical transmission equipment needed to light the wavelength.

4.3 Media Conversion

In addition to the full-wavelength 10 GEth and STM-64 services offered by GÉANT2, it is also possible for the network to provide a circuit which transports 10GEth over the SDH network using GFP encapsulation. This brings an important advantage: that the GÉANT2 switches can then be used to allow the technology platform of the circuit to be different at each end of the circuit. Thus a circuit may be provided with mismatched interfaces (STM-64 to 10GEth and vice versa) if this is a requirement. This functionality is known as *Media Conversion*. Furthermore, a 10Gbps *media converted* circuit is able to be configured as multiple circuits of lower capacity as with a GÉANT+ subscription.

5 Development of P2P and E2E Services and Processes

5.1 Administrative Procedure for P2P Circuit Request and Delivery

Based on an NREN Policy Committee (PC) decision, a procedure was introduced defining internal project procedures for NRENs wishing to request DANTE, as the operator of the GÉANT2 network, to implement point-to-point circuits over GÉANT2.

To clearly define this procedure a flow chart was drawn up (see Figure 5.1) defining the way in which a user request would be handled and the responsibilities of the organisations and individuals involved.

To capture the technical and organisational details of each circuit request, a Service Request Form was created in June 2006 and placed on the GÉANT2 Intranet to be used by NREN staff to submit the point-to-point requests. This form and the associated submission procedure can be found as Appendix A to this document. A recent development of this form (August 2007) now allows the on-line submission and provides additional functionality allowing the user to track the progress of each request. Full details of the rationale behind these changes and the work carried out to achieve the improvements will be detailed in a subsequent deliverable.

5.2 Technical Procedure for End-to-End Circuit Implementation

5.2.1 Development of Methodology

Thus far, this document has concentrated on the procedures for provision of P2P circuits within the GÉANT2 network domain. Clearly, since all such circuits form only one part of any End-to-End (E2E) path, a method of coordinating and integrating the different domain segments is essential. The methodology applied to this problem is described in this section.

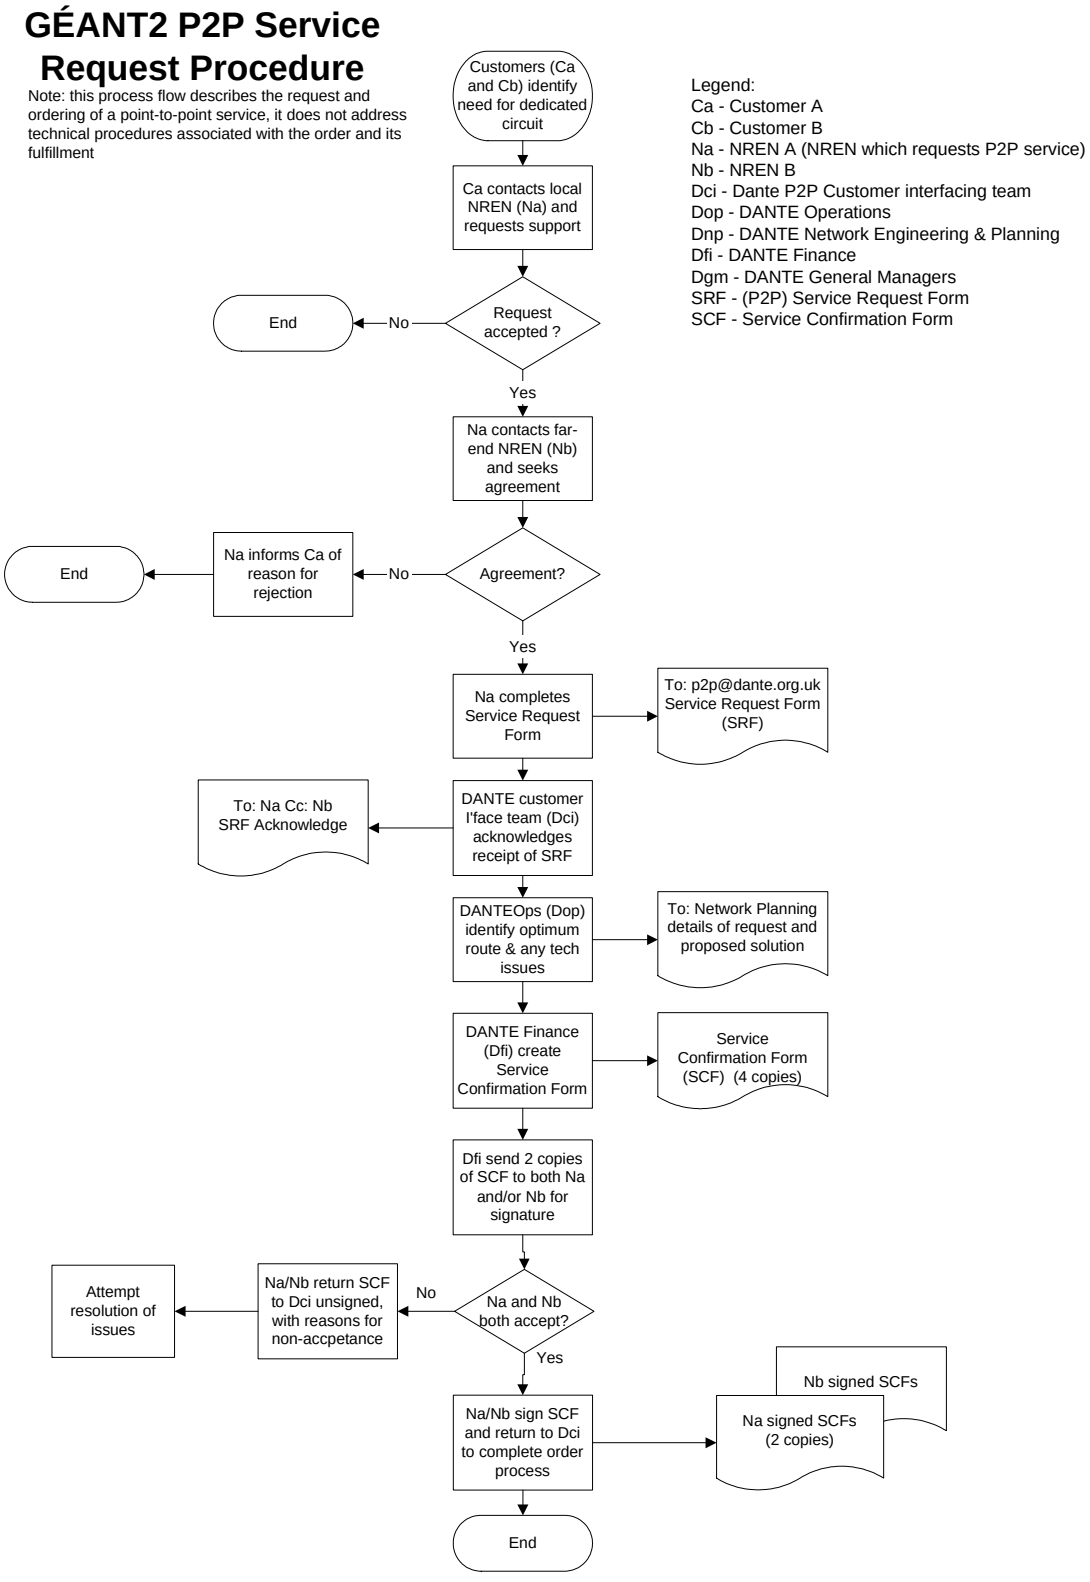

Figure 5.1: Process flow for GÉANT2 Point-to-Point orders

In recent years an increasing number of international and intercontinental E2E service instances have been requested by certain "power user groups" and the NRENs and DANTE in Europe and other R&E networks such as Internet2 in the US have endeavoured to provide these in a timely manner. This has still meant that the lead time for implementation of these has extended to many weeks and involved many hours of conference calls and the exchange of hundreds of e-mails. This very considerable workload can be attributed to a number of factors, including:

- relative inexperience of many of the operational staff in provisioning and operating circuit-based services;
- complexity of provisioning these services in the highly technologically and administratively heterogeneous, multi-domain environment that is the R&E networking world;
- paucity of tried and tested provisioning processes;
- configuration errors;
- time zone differences between NOCs.

Efforts have been made over the last year or so to improve and streamline the provisioning process by making it shorter, less labour intensive and less prone to configuration errors.

The first such refinement was to adopt the approach whereby a single person took on the role of a "coordinator" for the provisioning of each new service instance (or a set of instances where they are closely related in some way such as linking the same end points or being requested by the same user group). This person has responsibility for managing the implementation process – organising conference calls, preparing documentation (such as end-to-end circuit implementation diagrams), setting milestones and generally driving the process forward.

Appointing a coordinator gives the implementation process a certain momentum but it is still labour intensive and lengthy (>1 week) and, although manageable whilst End to End (E2E) service request volumes remain low (say at the level of <5 per month), it is clearly not going to scale up to levels in excess of 10 per month.

There is also the question of whether this centralised approach to service operations is suitable, in a global context, in a predominantly multi-domain environment that has historically (and remains) more at home operating in a highly decentralised, peer-centric manner.

To address these issues an attempt has been made to develop and trial tools that may help to optimise a more decentralised approach to handling E2E service requests and provisioning. The most significant of these has been the design of a service request and provisioning record form.

The original intention was that this form could be used by the NRENs (almost always in close collaboration with their respective end user(s)) to make an initial request and for the same form to then act as a vehicle by which

a less centrally coordinated provisioning process could proceed as efficiently as possible. To this end the form was developed (based on an Excel spreadsheet). This can be found in Appendix B.

This spreadsheet consists of multiple worksheets the first of which is based on the DANTE point-to-point service request form but with the GÉANT2-specific features removed this will be filled in by (or on behalf of) the domains hosting the end-sites of the circuit. The next worksheet is designed to capture the inter-domain routing of the E2E service instance being requested. Although this is arguably a centralised part of the process (someone decides on the domain level routing of the service early on in the provisioning process), it is thought that it could nevertheless also be used in a decentralised manner. In the case of the latter mode of using the form, each domain decides which is the next domain in the path and passes the form on to the appropriate person within that domain.

The remaining worksheets in the spreadsheet are there to facilitate the negotiation and capture of the configuration parameters necessary to "hand-off" an individual E2E service instance between two neighbouring domains. Given that this can be very complex in a technologically heterogeneous environment (which is almost always the norm amongst the worlds' R&E network operators) then this aspect of provisioning is often referred to as "technology stitching" or simply "stitching". Within GÉANT2's JRA3 activity quite some considerable effort has been expended in trying to rigorously parameterise this process since it is of critical importance when considering a fully automated provisioning of E2E services in the context of "bandwidth-on-demand" services where the target is to achieve zero-touch provisioning times of the order of minutes for the whole E2E service instance. More information on this can be found in the GÉANT2 deliverable DJ3.5.3 "Report on Testing of Technology Stitching". The problem with this very formal and rigorous approach to the parameterisation of stitching is that it is quite difficult to understand – especially by operations personnel who may only be familiar with the capabilities of the particular equipment used within their domain and, more importantly, the paradigm adopted by the designers of the management system they use on a day-to-day basis. Therefore, the "hand-off" worksheets that can be found in Appendix B are essentially simplified versions of the JRA3 efforts to parameterise stitching. Each worksheet is used to describe each one of the successive hand-offs and they are labelled as "NNIs" (Network-Network Interfaces). Given that each NNI consists of two sides (the neighbouring borders of the successive domains) then a negotiation needs to be undertaken by the operators of the neighbouring domains to ensure that the hand-off is correctly configured. The relevant NNI worksheet then acts as a repository of the final configuration for that given hand-off.

It is worth noting in passing that there is another way in which this chain of configuration information can be captured across the entire form. This is to have each worksheet associated with each successive domain. The advantage of this approach is that only a single worksheet needs to be completed (and updated as and when necessary) by the relevant (technical) people in each domain. However, the two stitching functions necessary at the edges of each domain mean that the negotiation that will be needed between the operators of neighbouring domains will require the updating of two worksheets. It is likely that only experience with using the form-mediated provisioning processes as described above will provide any indication as to which is the most optimal.

In many respects, this way of using the E2E form can be looked upon as being a kind of manual, process-based analogue of the way in which the routing and RSVP-mediated signalling of circuit provisioning operates in an ASON/GMPLS capable transport network. Whereas software processes running on network elements would undertake the necessary roles (communication, negotiation, maintenance of state, etc) in a true

ASON/GMPLS environment, these are instead performed by operational staff as a part of a more streamlined manual provisioning process.

Of course, this whole approach has to be tested and at the time of writing this deliverable the first attempts to trial the process(es) described above are being undertaken.

5.2.2 Operational Configuration of the End-to-End Path

This section describes, using examples of actual circuits, the procedures followed in coordinating the implementation of a multi-domain E2E path over NRENs and GÉANT2.

Firstly this section will describe the processes to be followed by each domain at the time of provisioning each individual *DomainLink* of the E2E circuit. Secondly this section describes the involvement of a coordinator whose main role is to manage the set up of the entire E2E circuit.

Establishment of Individual *DomainLinks*

Once the request has gone through the GÉANT2 administrative procedures (described in section 5.1), the configuration of each *DomainLink*[7] is taken over by the technical implementation teams of each domain. These implementation teams will ensure that the set-up requested does not present any technical issues in their respective domains. Checks will ensure that the technology to be used is appropriate, for example GEth versus SDH, and the capacity required is available in their networks. They will also make sure that the interface type requested, for example SR versus LR, is compatible with the interfaces implemented in the neighbouring domains. If the network does not have spare capacity, a purchase order for the hardware will be raised. The technical implementation team will then inform all the parties involved of the expected *DomainLink*'s delivery dates.

In order to confirm the spare capacity in each domain it is not only necessary to confirm the availability of the ingress and egress node, but also the internal path that the circuit will use in each of the domains. Once the path is defined and the hardware has been purchased as necessary, it is vital for each domain to keep a record describing the solution. These records will provide an overview of network developments and will avoid any over-booking of the same hardware for different projects.

For the GÉANT2 domain, DANTE has developed a database where all these items are registered. This database gives an accurate and up-to-date view of the network status. To achieve this, the database must act as an asset database, recording all relevant items in each PoP (Point of Presence). These items will include new hardware purchased (optical transponders etc) together with its specifications, for example transponder type and reach. This database will also record which equipment the hardware is installed in, for example transponders installed in optical switching equipment.

[7] An E2E link can be composed of several parts connecting different optical equipment. The parts that belong to only one domain are called DomainLinks. Also see section 5.3.

The second aim of the database used by DANTE is to provide a network topology and a view of the network capacity available. In order to do so, the network is represented at different layers (Layer 1 or optical layer, Layer 2 or switching layer and Layer 3 or IP layer).

Information will be introduced onto the database at the time when the hardware is ordered. This allows DANTE to take a proactive approach to the provisioning of new circuits by being able to predict future network requirements.

The current database used at DANTE does not have the capability of viewing the complete end to end circuit to be provisioned, being restricted solely to the GÉANT2 domain. It is expected that the forthcoming CNIS (Common Network Information System) facility will provide answers to this and other related questions.

Once the ordered hardware is delivered, each domain will use their NMS (Network Management System) to implement each *DomainLink* and include them in their respective monitoring tools.

The Establishment of the E2E Working group and Appointment of the Circuit Coordinator

At the time of implementing the E2E circuit, a coordinator is appointed by the domains involved on the path. The main role of the coordinator will be to facilitate communication between the different domains and ensure delivery of the E2E circuit in a timely manner.

To achieve this objective, the first task of the appointed coordinator is to contact all the stakeholders. The coordinator will thus gather the contact information of the implementation teams' representatives in each network domain. The second task for the coordinator will be to collect specific technical details of the set-up, giving special attention to the *InterdomainLink* and inter-domain issues such us demarcation points. The form created to manage this information was described in 5.2.1 and is available as Appendix B. The coordinator will be responsible for forwarding these forms to the technical representatives of each domain. Once retrieved, this information is stored in a place accessible to all the parties. The information is currently stored in the GÉANT2 wiki page with guaranteed access to all the parties involved in the project. In the future a common Information System is planned to be developed.

With all the necessary inter-domain information at hand, and after discussion with all the technical representatives, the coordinator is able to identify any potential issues likely to impact the implementation. In order to feedback all this information back to each of the domains, a mailing list is created to facilitate day-to-day information flow. If more critical problems arise, the coordinator may arrange conference calls with the appropriate parties.

E2E Implementation: A Practical Example

DANTE put these procedures into practice during the set up of the PHOSPHORUS circuit between PSNC in Poland and CANARIE in Canada. In this case the coordinator appointed was a representative of the GÉANT2 network. The coordinator initiated contact with the main entities involved (PSNC, CANARIE and GÉANT2). He collected the contact information of the technical representatives involved and managed the completion of the E2E form for each domain.

The same procedures were followed for the set up of the IN2P3-Fermi circuits. These circuits involved several domains (RENATER, GÉANT2, Abilene MANLAN and ESnet). For this case a mailing list was created [in2p3-fermi-2ge@dante.org.uk]. Experience has shown that a well organised mailing list helps to speed-up the whole set-up process and problem resolution. The experience gathered in this and other projects will feed back to provide improvements in set up procedures.

Once the circuit has been established, the coordinator will then organise basic testing to ensure that the circuit functions end-to-end. This step may not be trivial since most segments of end-to-end circuits are optical wavelengths or GE circuits. To conduct these tests, IP equipment along the path or some other type of testing equipment is required. The coordination of these tests is an important part of the implementation process. Every domain must agree on a thorough testing and debugging process. In the case of the PSNC-CRC-PHOSPHORUS-001 circuit the testing process revealed technical issues such as MTU limitation which were identified and solved.

5.3 End-to-End Circuit Monitoring

Each domain Network Operation Centre (NOC)) with an established capability to oversee and operate optical equipment[8] monitors its own physical infrastructure but to ensure an optimal End-to-End operation a wider monitoring is required. The role of the coordinator is to introduce to all the parties the latest developments of the JRA1-PerfSONAR and JRA4 activities and to pursue a global view of the E2E link. For the PSNC-CRC-PHOSPHORUS-001 circuit the coordinator introduced CANARIE to the end-to-end monitoring technologies developed in Europe and requested all domains install and populate local Measurement Points (MP). At the moment the CANARIE and GEANT2 Domain Links for this PSNC-CRC E2E link can be viewed on the end-to-end visualization tool. Several NREN network domains have already followed this initiative (CESNET, DFN, GARR, PSNC, RedIRIS, RENATER and SWITCH plus a number of non-European partners). The long-term goal is for a significant majority of the GN2 partners' domains to provide status information to E2EMon.

The real-time monitoring of layer-1 and layer-2 circuits that traverse more than one administrative domain is a considerable challenge. JRA4 has been working hard to address the problem and in so doing has developed a solution that has become known as "E2Emon". This involves the notion of breaking down a given E2E circuit into a number of concatenated *Domainlinks* which are themselves interconnected by *InterDomainLinks* (or *IDLinks*) with the responsibility for the real-time monitoring of these component parts being distributed to the relevant administrative domains along the length the E2E circuit. In principle the monitoring can include operational and administrative link status (up/down) and performance monitoring (monitoring of bit or frame errors). The latter is a substantial challenge so for the time being efforts have been concentrated on basic link status.

The E2Emon management information model is illustrated in Figure 5.2. It shows the subtlety of monitoring the IDLinks mentioned above. Typically these IDLinks represent the back-to-back patching of collocated network transmission equipment (e.g. DWDM terminal multiplexers) belonging to and operated by neighbouring administrative domains and no single domain can easily monitor the integrity of both directions of transmission

[8] This capability is sometimes referred to as a 'T-NOC' function. A T-NOC function can be performed either within an existing NOC or by a separate, specialist NOC organisation; this varies depending on the particular technical solutions deployed in each domain.

at this point. To address this, the E2Emon model allows IDLinks to be represented as two (unidirectional) IDPartialLinks the status of which can be monitored via the receiving "client" interfaces on the equipment belonging to the respective neighbouring domains. This is shown in Figure 5.2..

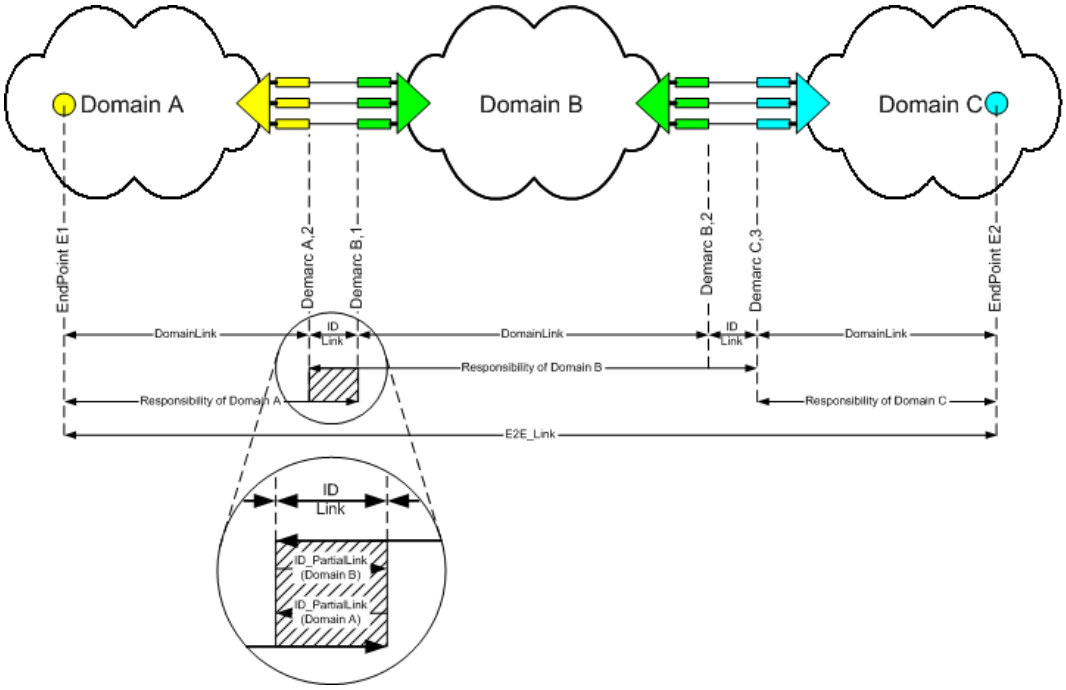

Figure 5.2: E2Emon management information model

The "value-added" of the E2Emon approach is that it introduces an automated and near real-time mechanism for the concatenation of the status monitoring drawn from the different administrative domains.

The domains themselves are required to assemble the monitoring information for all those parts of all E2E circuits for which they are responsible. In the E2Emon jargon this means the *DomainLinks* and *IDLinks* (or more likely *IDPartialLinks*). Each domain does this by synthesising whatever vendor- and/or technology-specific information they have into a commonly agreed abstract representation based on the model described above and encoded in an XML schema that has been integrated into the PerfSONAR schema. Each domain then publishes this monitoring information (updating it in near real-time) via a perfSONAR compliant Measurement Point or Archive. The latter MA allows the maintenance by the domain of local copies of historical records which can prove useful for retrospective debugging of the E2Emon system or resolution of the sources of false alarms.

A centrally operated component of the E2Emon system (in fact, there can be many instances of this central component) periodically polls all the distributed MPs (or MAs), assembles the information for each E2Ecircuit and stores it in a database. The central component also facilitates the publishing of the near real-time E2E status monitoring information via a web interface which exports an intuitive graphical representation of the information (to which access can be controlled according to a number of criteria). It also facilitates integration with other umbrella network management systems. In particular, an SNMP (Simple Network Management Protocol) agent has been incorporated into the central component along with a specially designed Management

Information Base (MIB) which together allow alarms to be asynchronously exported (via SNMP traps) to the umbrella management system. This in turn means that operational personnel (e.g. those in the E2ECU – see below) can be informed of problems with an E2E circuit in a timely manner. All of this is illustrated in Figure 5.3..

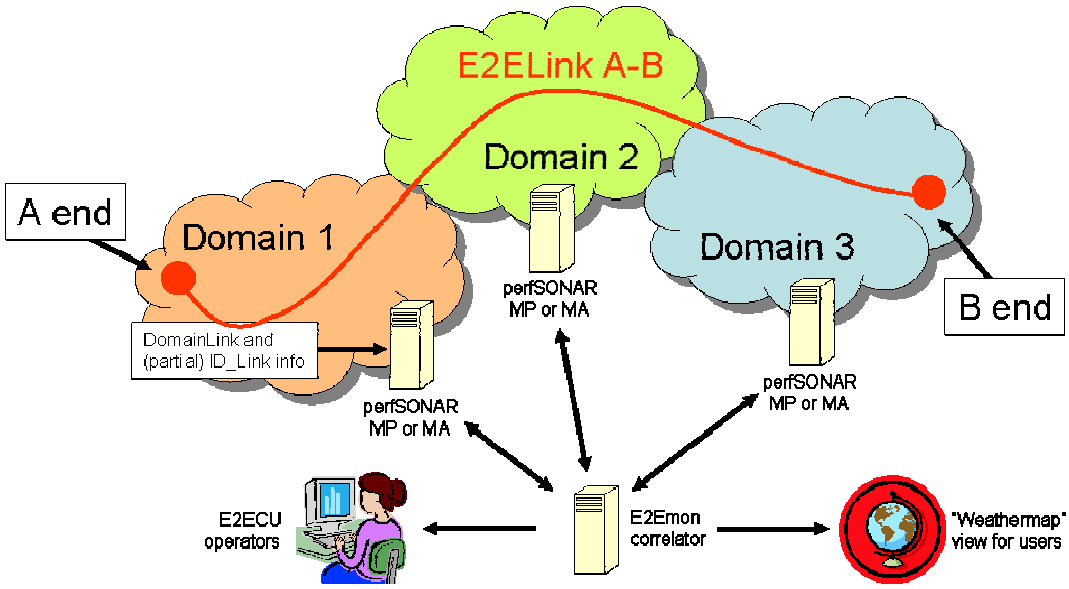

Figure 5.3: E2Emon system architecture

The E2Emon system is undergoing continuous development (bug fixing, feature addition, etc) within JRA1 and JRA4 and adoption of the system by GÉANT2 NRENs has been good with most NRENs that support E2E circuit services now having in place working E2Emon MPs or MAs.

The following Domains have MPs or MAs actively responding to queries from the E2Emon for a variety of project links: CANARIE, CERN, CESNET, DFN, ESNET, FERMI, GARR, GEANT2, IN2P3, NETHERLIGHT, PSNC, REDIRIS, RENATER, SWITCH and USLHCNET.

For global monitoring of E2E circuits, GÉANT2 introduced the E2ECU (End-to-End Coordination Unit). The E2ECU's main role is to coordinate information flow between the different domains that are part of the E2E circuit. This coordination will be focused in making sure that:

- All the domains are aware of the method to install the different MPs/MAs (Measurement Points/ Measurement Archives)
- All the domains populate the MPs/MAs with the correct data
- All the E2E circuits are named uniquely and each relevant Domain is informed of the name
- All E2E circuits appear in the global visualization tool
- Trouble Tickets (TTs) are opened when a fault occurs in a domain that affects the E2E circuit

- TTs related to any fault affecting an E2E circuit are updated and forwarded to all the domains involved.

The E2ECU is currently monitoring 21 E2E circuits belonging to 2 different projects. It is expected that in the future the E2ECU can act as coordinator at the time of setting up the E2E circuits.

5.3.1 Monitoring Procedure

1. A Domain, such as an NREN or GÉANT2, ensures that their Measurement Point will be given information about up/down alerts on the End-to-End circuits that traverse or terminate in their network. Since the Domains use a variety of hardware in their network, each needs to write software to interact between their Network Monitoring System or actual hardware, and the PerfSONAR Measurement Point or Archive.

2. The central End-to-End Monitoring System (E2EMon) polls the individual Domain MPs and MAs every five minutes to gather information about the constituent Domain and Inter-Domain Links. Since each is tagged as belonging to a particular End-to-End link and names its neighbour Domains, the central software concatenates these smaller Links together to form the End-to-End link on a graphical display that can be viewed with a web browser.

3. Any errors relating to the population of the XML files used by the MPs and MAs are listed on the central E2EMon Domain View and the E2ECU sends an email to the relevant Domains and assists them in interpreting the errors.

4. The engineers at the E2ECU have written plug-ins for their proprietary monitoring system which receives alerts from E2EMon whenever an End-to-End circuit has an outage on any of the constituent parts.

5. The E2ECU engineers are watching their monitoring screen from 06:00 to 22:00 CE(s)T during the week and from 09:00 to 18:00 CE(s)T at weekends.

6. When an outage occurs, a ticket is raised containing information such as the names of the Domain Link or Inter-Domain Links affected, the name(s) of the Domain(s), the name of the Project, such as LHC-OPN, that is affected and the time of the outage.

7. The E2ECU then contact the relevant Domains to request information regarding the outage; in the case of an Inter-Domain Link they will contact both Domains involved.

8. Any updates regarding the outage are distributed to all partners in the project that includes the affected End-to-End link.

A similar procedure exists for planned maintenance, whereby the Domain planning the work contacts the E2ECU who will open a ticket and distribute information to all partners of the relevant project.

The E2ECU sends monthly reports to DANTE including availability statistics of the various End-to-End Links and a list of End-to-End Links recently added to the central Monitoring System.

DANTE Operations are overseeing the procedures that the E2ECU uses in their day-to-day operation and are also in close co-operation with the software developers of the End-to-End Monitoring System in JRA4, to aid in the continual refinement of the monitoring tools available.

6 Refining P2P and E2E Delivery in GÉANT2

The delivery of GÉANT2 point-to-point circuits has now become a regular occurrence and a total of 22 circuits were delivered between June 2006 and June 2007. However, it is hoped that by further refinement of the procedures the customer experience can be improved and unnecessary workload reduced.

The initial point-to-point Service Request Form (Appendix A) was accessed on a password-protected page of the GÉANT2 intranet, accessible to the project's NREN partners. The form could then be downloaded as an Excel spreadsheet, completed and submitted to a dedicated P2P email address. With a further development of this form, it is now possible to fill in all details of the request directly on-line after logging–in to the GÉANT2 website. This will reduce the work involved for the submitting NREN and allow additional features to be added to improve communication. The new on-line process is designed to be able to:

- Generate an automatic email to indicate that a request has been received
- Allow the request to be modified up to the point at which equipment has been purchased
- Provide the ability to track the progress of the delivery process remotely by the requester
- Offer access to the stored technical and administrative details of the request via a web interface[9].

It is thus hoped that the process for requesting point-to-point circuits will be greatly improved for both the requester and the staff providing the service. A future development of this form is envisaged to be to expand its features to cover the full E2E process described in 5.2. A report detailing the development of the process workflow, request status milestones and the necessary software to manage point-to-point requests will be included in a future SA2 deliverable.

It is important to remember that GÉANT2 also includes a research activity dedicated to the realisation of zero-touch provisioning of E2E circuits, essentially transferring the manual process described in section 5.2 into a fully automated process. Whilst this considerable undertaking is still some way from production-level service, it is already proving useful insights for operational procedures, for example the stitching framework developed by JRA3 engineers. Thus it is intended that both operational and research activities will benefit from an exchange of experiences and ideas in this important work area.

Looking at E2E customer service beyond the processing of the order and circuit delivery, it is envisaged that two new features will benefit the circuits' users:

[9] Currently at http://stats.geant2.net/p2p/ (login required)

- The multi-domain monitoring tools under development by the GÉANT2 JRA1 and JRA4 activities will be of significant benefit in assessing the performance of the individual circuits and in ensuring optimal operation of the service.

- The GÉANT2 end-to-end coordination unit will bring around-the-clock multi-domain operational assistance to international point-to-point services. This function will ensure that the complex issues associated with circuits crossing several network domains (in some cases including CBF) and spanning diverse time zones are adequately addressed.

Although of direct relevance to point-to-point users, both of these services are beyond the scope of this document and will be discussed in other GÉANT2 reports.

7 Conclusions

The GÉANT2 network is the first international production hybrid network. The point-to-point services offered represent a new era in network services which, whilst providing the opportunity to deliver unprecedented levels of service to the research and education community, require significant coordination effort to manage. The availability of the infrastructure and equipment is therefore not sufficient for the efficient delivery of the point-to-point service: an effective set of procedures associated with the service is also essential to ensure the success of this initiative.

This document has described the services offered and the benefits that they bring to end users. It has shown that effective procedures are in place to deliver these services and that a significant number of such circuits are now operational. Building on the experience gained work is well underway to develop and improve the point-to-point and end-to-end procedures to improve the customer service still further.

8 Acronyms

P2P –Point-to-point (a dedicated circuit configured between two points over a single network domain)

E2E –End-to-end (a dedicated path over multiple network domains each segment comprising a P2P circuit)

NREN -National research and education network

Gbps – Gigabit per second

PoP –Point of presence

OPN –Optical private network

GEth – Gigabit Ethernet

SDH –Synchronous Digital Hierarchy

JRA –Joint research activity

CNIS –Common Network Information System

NMS -Network Management System

NOC – Network Operation Centre

E2ECU – End to End Coordination Unit

TT – Trouble Ticket

MA – Measurement Archive

MP – Measurement Point

LHC – Large Hadron Collider

ID – InterDomain

IDLink – InterDomainLink

SNMP – Simple Network Management Protocol

MIB – Management Information Base

Appendix A Point-to-Point Service Request Form

GÉANT2 POINT-TO-POINT SERVICE REQUEST FORM

NOTE: NREN requesting the service (NREN A) - This person should complete the full form covering both sites A and B

This Service Request Form will be used to prepare a Service Confirmation Form (SCF).
The SCF will form the basis of a legally binding contract between the organisations named below and DANTE

A. Details of NRENs		
	NREN A	**NREN B (details of person contacted who agreed to this service request**
Contact Name:		
Organisation		
Position		
Telephone		
E-mail		
B. Commercial Contacts (if different from above)		
Site Name		
Contact Name:		
Telephone		
E-mail		
C. Project		
Project Name:		Field of education or research:
Project Description:		
D. End-sites to be connected		
	Site A	**Site B**
Site Name:		
Contact Name:		
Telephone No:		
E-mail:		
E. Date Required		
Required Ready for Service Date:		Termination Date:
Further point-to-point services required for this project?		

F. Technical Details of Request							
	NREN A			**NREN B**			
1. Access Point							
2. Technical Contact Person Name							
Telephone No							
E-mail address							
3. Total Capacity required (Gbps)							
4. *New Interface type facing GÉANT2 equipment	1 or more GEth			1 or more GEth			
	10 GEth full rate			10 GEth full rate			
	10 GEth with VLAN functionality			10 GEth with VLAN functionality			
	VLAN in already existing 10 GEth interface			VLAN in already existing 10 GEth interface			
	STM-16			STM-16			
	STM-64			STM-64			
5. Type of Optics	SDH			SDH			
	GEth			GEth			
	Other (please specify)			Other (please specify)			
6. *In case of 10 GEth full rate interface or STM-64 interface the Service should be provided via							
7. In case of VLAN request, specify the VLAN number to be used							
8. In case of VLAN on a pre-existing 10 GEth interface, are the previously configured VLANs used for a different project(s) (If Yes, please specify)							
9. If STM-16 or STM-64 is required into the TDM switch, which time slots should be used for this project?	From (1-64)		To (1-64)	From (1-64)		To (1-64)	
10. Is a redundant circuit required for protection? (Note: this will be treated as an additional point-to-point request and charged as such)							
G. Payment							
Please indicate how the service is to be invoiced to NREN A and B			NREN A - %		NREN B - %		

* Please note that question is not applicable if the requested service will use GÉANT+

I agree that the above request is a true reflection of the requirements of my organisation and that it reflects the spirity of the GÉANT2 usage policy. I expect to be billed accordingly for the services provided, in line with the pricing structure agreed by the NREN PC

FOR DANTE USE ONLY	
Request Reference	

NOTES BELOW ARE PROVIDED TO ACCOMPANY THE SERVICE REQUEST FORM SPREADSHEET (ABOVE) AND ARE INTENDED TO ASSIST NREN STAFF FILLING IN THE FORM

Summary of the GÉANT2 Point-to-Point Service Request Procedure

The service request form contained within this spreadsheet is designed to allow NRENs to register the technical requirements for provision of the following dedicated connection technologies:

- 1 or more GEth
- 10 GEth full rate
- 10 GEth with VLAN functionality
- VLAN in already existing 10 GEth interface
- STM-16
- STM-64

The process to request such services is as follows:

1) NREN A should establish that the requested service meets the requirement of the requesting user group and is supported by NRENs A and B, both technically and financially (where appropriate)

2) NREN A should complete the service request form and return by email to the following address:

p2p@dante.org.uk

NREN B should be copied on the mail and be thus informed of the formal request

3) DANTE will check the form for completeness and respond to confirm receipt

4) DANTE will consider the response from a technical, financial and policy perspective and, if appropriate, issue a service confirmation form to both NREN A and NREN B. This document will detail:

i) The technical details of the circuit which will be provided

ii) The target delivery date for that circuit

iii) The financial cost to the NRENs concerned

5) NREN A and NREN B will each return two copies of the signature page of the service request form to DANTE signed by the appropriate authority

6) DANTE will sign and return one original service confirmation form to NREN A and NREN B, retaining one copy on file

NOTE: NREN patch cards must match the following specification to interface with GÉANT2 equipment - Connector type:

- E2000/APC (Single mode)
- E2000/PC (Multi mode)

Appendix B E2E Coordination Spreadsheet

GENERIC END-TO-END SERVICE IMPLEMENTATION RECORD

NOTE: Cells in **yellow** **have drop-down menus**
grey **are not used**

A. Details of END SITES (or USER DOMAINS)		
	SITE A	**SITE B**
*Site Name:		
Organisation:		
Organisation type:		
Contact Name (requesters):		
Position:		
Telephone No:		
E-mail:		

B. Project			
Project Name:		Field of education or research:	
Project Description:			

C. Date Required			
Requested Ready for Service Date:		Service End Date:	
Further point-to-point services required for this project?			

D. Technical Details of E2E Service		
	SITE A (UNI-A)	**SITE B (UNI-B)**
1. Access Point		
2. Site Technical Contact Name:		
Telephone No:		
E-mail address:		
3. E2E Service type (UNI-A to UNI-B):		
4. Total capacity (Gbps):		
5. New or existing physical interface?		
(skip this row)		
6. Interface type:		
7. Interface transceiver:		
(skip this row)		
(skip this row)		
9. Requested MTU? (EPL service only)		
10. Protection level required?		

NOT REQUIRED	Wavelength range [nm]		
	TX max launch [dBm]		
	TX min launch [dBm]		
	TX min extinction ratio [dB]		
	RX operating range [nm]		
	RX sensitivity [dBm]		
	RX saturation [dBm]		
	RX max reflectance [dB]		

ROUTE PLAN:

SITE/User Domain A	
Type:	
Name:	

? **UNI-A**

Provider Domain 1	
Type:	
Name:	

? **NNI-1** [,]

Provider Domain 2	
Type:	
Name:	
Last provider domain?	

? **NNI-2** [Form: ,]

Provider Domain 3	
Type:	
Name:	MANLAN
Last provider domain?	**No**

? **NNI-3** [Form: ,]

Provider Domain 4	
Type:	
Name:	
Last provider domain?	

? **NNI-4** [Form: ,]

Provider Domain 5	
Type:	
Name:	
Last provider domain?	

? **NNI-5** [Form: ,]

Provider Domain 6	
Type:	
Name:	
Last provider domain?	**Yes**

?

SITE/User Domain B	
Type:	
Name:	

GENERIC POINT-TO-POINT SERVICE NNI CONFIGURATION REQUEST & RECORD

NOTE: For Network Operator Use ONLY

This form need only be used at a domain boundary and is used by operational personnel for the configuration of "NNIs" between domains. Use one of these worksheets per NNI along the route specified on the "Route Plan" sheet ensuring they are named "NNI-1", "NNI-2", etc in the direction from SITE A towards SITE B.

A. Details of provider domains neighbouring NNI		
	Provider Domain 1	Provider Domain 2
Organisation		
Tech contact		
Position		
Telephone		
E-mail		

B. Link to be used		
Link ID		
Link owner		
Link type		
	Provider Domain 1	Provider Domain 2
Interface ID		
(status)		
(allocated?)		
(configured?)		

C. Link parameters		
C.1 Link parameters for SONET/SDH-based NNI		
Structure		
No of VCG members		
Positioning within multiplex		
At or starting from time slot…		(use STS-1 equivalent time slot notation so x=1…192)
(status)		
(allocated?)		
(configured?)		
C.2 Link parameters for Ethernet-based NNI		
	Provider Domain 1	Provider Domain 2
VLAN trunking?		
VLAN tag		
MTU		
	(FOR INFO: smallest MTU along path is 1518 bytes)	
(status)		
(allocated?)		
(configured?)		