05-10-2010

Deliverable DJ1.3.1: Architecture Considerations for Federated Backbone Networks Study

Deliverable DJ1.3.1

Contractual Date:	31-03-2010
Actual Date:	05-10-2010
Grant Agreement No.:	238875
Activity:	JRA1
Task Item:	T3
Nature of Deliverable:	R (Report)
Dissemination Level:	PU (Public)
Lead Partner:	SURFnet
Document Code:	GN3-09-250
Authors:	Bartosz Belter (PSNC), Maribel Cosin (RedIRIS), Paul van Daalen (SURFnet), Lars Fischer (NORDUnet), Ivana Golub (CARNet), Andreas Hanemann (DFN), Marijke Kaat (SURFnet), Milosz Przywecki (PSNC), Branko Radojevic (CARNet), Sue Tyley (DANTE), Srdjan Vukovojac (CARNet)

Abstract

This deliverable reports on an investigation into network federation, i.e. sharing resources among multiple independent networks. It defines key terms, benefits and challenges; assesses user demand for federated networks based on an analysis of current and future large-scale projects; evaluates existing GN3 resources, tools and services that can serve as building blocks for federated networks; introduces a set of generic federated network architecture models; and describes the GN3-related test cases by which the models will be verified and refined.

Table of Contents

Table of Figures

Table of Tables

Executive Summary

Federating networks means to share network resources among multiple independent, but collaborating networks in order to optimise the use of those resources, improve the quality of network-based services, and/or reduce costs for the provisioning of services. The purpose of this deliverable is to report on JRA1 T3's investigations into network federation. While the primary intended audience is the GN3 project partners, the results are of value and relevance not only to GÉANT and European NRENs but also to any special-purpose network and core networks in general.

The deliverable begins with a definition of federation and other key terms relating to federated networks, together with a clarification of the scope of the study (Section 2). After outlining the benefits and challenges of federation, the user demand for federated networks is assessed based on an analysis of current and future large-scale projects requiring international data transmission (Section 3). Section 4 describes and assesses a selection of existing GN3 tools and services that could serve as building blocks for federated networks. Section 5 introduces architecture models for building federated networks. The models propose layered architectures that are deliberately generic in order to be applicable to many scenarios. Section 6 describes the proposed GN3-related test cases by which the models will be verified and refined in future work.

The potential benefits of federation include cost savings, support for multi-domain services and improved user experience. The main challenges relate to the management, technological differences, missing standards, cost model and the federation-independent presentation of services.

The current large-scale projects analysed to assess user needs for federated networks included the Large Hadron Collider Optical Private Network (LHCOPN), Electronic Very Long Baseline Interferometry (e-VLBI), Enabling Grids for E-Science – Croatia (EGEE-III), and Distributed European Infrastructure for Supercomputing Applications (DEISA 2). Some of the future projects that will come under the European Strategy Forum on Research Infrastructures (ESFRI) umbrella were also included in the analysis. The key parameters for federated networks are geographic location, topology, type of end-to-end connection, data requirements, operations and services (i.e. work-flow and procedures, quality of service, performance monitoring and security) and cost model.

The existing GN3 resources, tools and services that were assessed for their potential usefulness for building a federated network were NREN networks, operational tools and supporting services, administration and procedures, and end-user services. The networks of 5 NRENs representative of European research infrastructures (DFN, PIONIER, RedIRIS, SURFnet and CARNet) were analysed for topology, coverage, utilisation and development. Particular attention was paid to three cross-border fibre (CBF) initiatives. The operational tools and supporting services assessed included those for performance troubleshooting (eduPERT), measurement support (perfSONAR), automated bandwidth provisioning (AutoBAHN, AMPS) and information storage (I-SHARe, c-NIS). The administrative and procedural principles, requirements and best practices considered include information sharing, agreements, collaboration between network operations centres, an operations model and a governance structure. The key points relating to end-user services concern how the services are composed and offered, i.e. whether by direct aggregation of similar services in the partners' networks or as a layer on top of the core network components provided by the federation partners; each method has fundamental implications for the network design.

The generic federated network architecture models, defined as a result of analysing current and future large-scale projects and existing GN3 resources, tools and services, are composed of three main layers: Infrastructure, Operations and Service. The lower, Infrastructure Layer consists of network infrastructure elements (the NRENs in the case of GÉANT); the middle, Operations Layer consists of the tools and support services (also known as intra-federated services) that are needed to provide support for services that are offered to end users (e.g. perfSONAR, AutoBAHN); the top, Service Layer contains the end-user services themselves. Elements within the same layer may be inter-related. Elements in the Service Layer are not permitted to interact directly with the Infrastructure Layer, but must go through elements of the Operations Layer. The two variant models, A and B, reflect the two ways communication is performed within and between the layers, with Model A being the simpler, restricting communication between neighbouring layers to the management component of each layer, and Model B being more complex, allowing direct relationships between individual elements in adjoining layers. While Model B is potentially the more scalable, it risks increasing the overhead, reducing robustness, and causing duplication.

Future work will verify and refine the models by applying them to GN3-related test cases. One approach is to analyse the building blocks and relationships of an existing collaboration from the perspective of the proposed federation model. A second approach is to analyse the implications of making a change in the federation, and which model is better suited to the new situation. The test cases proposed are the GÉANT network itself (with use cases including using CBF for GÉANT PoP-to-PoP connections, using CBF as an element in regional links, a federated GÉANT PoP and remote IP backup with CBF) and LHC Tier1-to-Tier2 connections. In addition, the federated models presented here will be compared with other models from the literature, e.g. those defined by the TM Forum, and with those used or proposed by other GN3 service or joint research activities.

1 Introduction

The improvement of services for the user community and the reduction of costs are important aims of the GÉANT3 (GN3) project consortium, for instance in Service Activity 1 Network Build and Operations.. The specific objective of Joint Research Activity 1, Future Network (JRA1) Task 3, Federated Network Architectures (T3) is to approach these aims by investigating possibilities related to federated networks, that is, the sharing of existing network resources among the domains in a highly structured and carefully managed manner.

To clarify the term further, and differentiate it from the similar “federation of networks”: the GN3 project is a federation of networks organisation. It has been successful in providing network services for research and education in Europe, addressing a potential digital divide by providing affordable network access for all members, and is successfully implementing multi-domain services. However, the GÉANT network is not a federated network because it uses its own dedicated resources and does not share resources with other networks. The work described in this report aims to investigate architectures for federated networks; work in Year 2 will investigate how this can be applied to the GÉANT network.

One driving force for investigating a federated approach has been the transformation of GÉANT and several other networks into hybrid networks. The term “hybrid” denotes a network that offers dedicated network bandwidth to users in addition to the plain IP service. The provisioning of dedicated bandwidth uses technologies such as Synchronous Digital Hierarchy (SDH), native Ethernet or Ethernet over Multi-Protocol Label Switching (EoMPLS) so that no IP routing is involved in the network except of the users’ own routers. The possibility of installing fibres to neighbouring countries at low prices has led to the establishment of a set of cross-border fibres (CBFs) that typically link two National Research and Education Networks (NRENs). The CBFs and NREN resources could be used to build an alternative path for connecting GÉANT Points of Presence (PoPs). Assuming a GÉANT PoP in (hybrid) network A and (hybrid) network B, as well as a CBF between A and B, the PoPs could be connected either with a carrier link or using the resources of A and B and the CBF (or, more precisely, one or several wavelengths based on the CBF). Such alternatives could be used for improving availability in the event of failure or as a way to reduce costs if only the connection with the lower costs is realised.

The purpose of this deliverable is to report on JRA1 T3’s investigations into network federation, this being a prerequisite for transforming GÉANT into such a network. While the primary intended audience is the GN3 project partners, the results are of value and relevance not only to GÉANT and European NRENs but also to any special-purpose network and core networks in general.

The deliverable begins with a definition of federation and other key terms relating to federated networks, together with a clarification of the scope of the study (Section 2). Then the benefits and challenges of federating networks are detailed, based on the definition and scope (Section 3). This section also provides an overview of current and future large-scale projects as the basis for investigating which kind of federations and services based on the federation users require.

Several resources and activities in the GN2 and GN3 projects provide building blocks that are assessed from the federated network point of view (Section 4). The aspects assessed are network topologies, operations and services. Contributions from the services and tools delivered by these activities serve as input for defining federated network architecture models.

The federated network architecture models, described in Section 5, are the main contribution of JRA1 T3. Two alternatives are proposed, which differ in the way the collaboration is carried out. Both architecture models use three layers and explicitly model the communication that is required to make a federation work.

The usefulness of the architecture models will be analysed in future work. To prepare for this analysis, JRA1 T3 has identified several test cases which refer to current open issues in the context of GÉANT and large-scale projects. These are described in Section 6.

2 Definition of Terms

This section introduces key terms used in the context of federated networks. Detailed definitions are provided because experience has shown that misunderstandings of terms can lead to major difficulties.

2.1 Definitions

A **federation** is traditionally an organisation composed of several autonomous members or partners. A federation therefore has a layered governance model and / or segmented spans of control. A key characteristic of a specific federation is the division of control and authority, i.e., which decisions are made and / or executed independently by the members, and which are made centrally and / or executed federation-wide. A federation is also characterised by some things being homogeneous (the same across all members of the federation), and others being heterogeneous (different for different members). The GN3 project is an example of a federation.

Federating is an approach to creating complex IT systems or structures. A federated IT system is composed of resources contributed by several autonomous (sub-)systems. Often, the autonomous sub-systems are owned and / or managed by autonomous organisations. The federated IT system is kept together by standard interfaces (protocols) for exchange of information and control between the autonomous systems, and by agreed procedures and workflows. This ensures that even though the autonomous systems may be technically different and under different control, interworking can be ensured.

Federated IT systems are often loosely coupled and based on a best-effort approach to service delivery. **Loosely coupled** means that there is no well-defined governance, no single view of the status of the federated system, no workflow coordination across the participating domains, etc. **Best effort** means that the partners contributing to the federated IT system make no service-level guarantees. The best-known example of a loosely-coupled best-effort federated IT system is the Internet.

For some use cases, however, best effort is insufficient for user needs. In such cases, a centralised approach is often used, with one organisation owning and operating the total system, providing governance and control, and issuing guarantees. However, we argue that a tightly coupled, agreement-based federated system can achieve the same result. In a **tightly coupled** system, a formal agreement is established between federation partners, providing a source of governance for the IT system; process, procedure and workflow coordination, as well as integrated operations support systems across domains, enable the system to make guarantees and enter into service level agreements. Examples of such systems are the Large Hadron Collider Optical Private Network (LHCOPN) [LHCOPN] and the Nordic DataGrid Facility [NDGF].

In the remainder of this document, when we talk of federations, unless otherwise noted we will be referring to tightly coupled, agreement-based federated IT systems.

The rationale for building a federated network is to share network resources between domains (defined below), for the benefit of all partners – for example, cost savings through reuse of resources owned and operated by the domains (see also Section 3.1 *Benefits of Federation* on page 9). Federating the network can take place on different technological levels. A domain can for example share single fibres with other domains, or separate wavelengths. It may use wavelength-based services, managed either by the domain itself or by a supplier. Parts of these wavelengths can be shared with other domains based on Layer 2 technologies such as Ethernet or Synchronous Digital Hierarchy (SDH). A domain may route IP traffic from other domains – for example, offering transit. Or it may virtualise its resources and offer them to other domains, e.g. virtual routers or virtual network paths. These resources can be used to build a federated network. In networks, the resources of each partner in the federated network belong to the domain of control of the autonomous system of the partner. The resources are combined as elements of the network architecture of the federated network.

A **domain** is described as a network entity with independent management and administration (e.g. a Campus Local Area Network (CLAN) or an NREN). This definition is similar to the definition of an Autonomous System provided by RFC1930:

An AS can also be referred to as a domain (as in the phrase “inter-domain routing”) and is usually managed by a single organisation as seen from the outside world. Inside an AS, parts of the network may be operated by other organisations, e.g., a campus network that is part of the NREN’s network. To other ASs it is the same domain as it adheres to the same routing policy. The concept of a domain is equally important at all network layers; a lightpath may extend across several domains. A federation is therefore a multi-domain system, with connections, physical or otherwise, between separate domains. This is what complicates operation, as it means the federation has to deal with independent spans of control.

Architecture can be defined in many ways. The IEEE 1471 standard [IEEE1471] states that architecture is: “The fundamental organization of a system embodied by its components, their relationships to each other and to the environment, and the principles guiding its design and evolution.” IEEE 1471 describes the architecture of a software-intensive system. This is not exactly the same as network architecture, but many parts and issues are very similar. The conceptual framework for an architecture description from this standard is very useful for describing network architectures as well.

The purpose of describing architectures is to be able to communicate more consistently about the system, to plan and manage the development, and to enable evaluation and comparison of different architectures. Expressing the system and its evolution in more or less standardised terms is essential to producing a useful description.

The following definitions from the IEEE Standard Glossary of Software Engineering Terminology [IEEE610.12], and descriptions from the conceptual framework described in IEEE 1471 [IEEE1471wiki] are useful for describing federated network architecture models:

- **Architectural Description** (AD): a collection of products (i.e. documents) that document an architecture.

- **System**: a collection of components organised to accomplish a specific function or set of functions. The term system encompasses individual applications, systems in the traditional sense, sub-systems, systems of systems, product lines, product families, whole enterprises, and other aggregations of interest.
- **System stakeholder**: an individual, team, or organisation (or classes thereof) with interests in, or concerns relative to, a system.
- **View**: a representation of a whole system from the perspective of a related set of concerns.
- **Viewpoint**: a specification of the conventions for constructing and using a view. A pattern or template from which to develop individual views by establishing the purposes and audience for a view and the techniques for its creation and analysis.
- Every system has an architecture, whether understood or not, whether recorded or conceptual. An architecture can be recorded by an architectural description (AD).
- A system has one or more stakeholders. Each stakeholder typically has interests in, or concerns relative to, that system.
- **Concerns** are those interests that pertain to the system's development, its operation or any other aspects that are critical or otherwise important to one or more stakeholders. Concerns include system considerations such as performance, reliability, security, distribution, and evolvability.
- An architectural description is organised into one or more constituents called (architectural) views. Each view addresses one or more of the concerns of the system stakeholders. A view is a partial expression of a system's architecture with respect to a particular viewpoint.
- An architectural description selects one or more viewpoints for use. The selection of viewpoints is typically based on consideration of the stakeholders to whom the AD is addressed and their concerns. A viewpoint definition may originate with an AD, or it may have been defined elsewhere.
- A viewpoint establishes the conventions by which a view is created, depicted and analysed. In this way, a view conforms to a viewpoint. The viewpoint determines the languages […] to be used to describe the view, and any associated modelling methods or analysis techniques to be applied to these representations of the view. These languages and techniques are used to yield results relevant to the concerns addressed by the viewpoint.
- A view may consist of one or more architectural models. Each such architectural model is developed using the methods established by its associated architectural viewpoint. An architectural model may participate in more than one view.

2.2 Views and Viewpoints

The IEEE 1471 standard describes a viewpoint as determining the set of concerns and the representations / modelling techniques, etc. used to describe the architecture to address those concerns. A view is defined as the result of applying such a viewpoint to a particular system. IEEE 1471 also states that architecture descriptions are inherently multi-view, since no single view adequately captures all stakeholder concerns. The purpose of views and viewpoints is to enable engineers to comprehend complex systems. Most complex-system specifications are too extensive for one single individual to understand all aspects. In a federated network environment there is by definition no one single person who can comprehend all details, as it deals with separate domains with their own technologies, rules and regulations. People will have different interests in the system, ranging from technical details, functional aspects, and security concerns to financial implications. A view gives a representation of the system from the perspective of a viewpoint. By focusing on a specific concern (or a related set of concerns), it suppresses certain details to provide a simplified model and focuses attention on certain aspects of the system. An architectural description, as stated above, selects one or more viewpoints for use. The products used to document an architecture are things like reports, analyses, (graphical) models, etc. These products provide a way to visualise the architecture data as graphical or textual representations.

Several view models and methodologies have been developed over the years (see [ViewModel] and [Sessions]). Each has its strengths and weaknesses and is more or less applicable to a certain kind of system. The best approach is to choose a set of views and viewpoints that apply to the architecture to be described.

In Section 5 of this document, federated network architecture models are described. Using the definitions and terminology of IEEE 1471 for each architecture model, the stakeholders should be identified and their concerns relevant to the architecture, the specifications of each viewpoint that was selected and the rationale for those selections should be determined; one or more architectural views should be defined and a rationale for selecting the architecture should be stated.

3 Benefits, Challenges and User Needs

This section explains the benefits that can be expected when networks are federated, as well as the challenges related to such a transformation process. It also presents the results of an analysis of a set of several research projects, undertaken to assess user needs for federated networks. The analysis has two aspects: investigation of support for projects by federated networks and identification of possibilities for the formation of a federated network within the project to serve its users.

3.1 Benefits of Federation

Building a federated network should benefit the participating domains (member networks). The potential benefits include:

- Cost savings, realised by sharing resources.
- Benefits for multi-domain services, as a result of managing the federation in an integrated manner. In particular, service provisioning can become easier and faster. This aspect is important for large projects that are supported by several domains within the federation.
- Improved user experience. The interface between the federated network and its users must abstract from the involved domains. This means that the users will not be aware which domain's resources they are using. From the user's point of view, this makes using the federated network much easier.

3.2 Challenges of Federation

A federated network leads to several challenges:

- Cost model. The sharing of resources has to result in the sharing of costs for these resources, which is a very challenging problem. The compensation that a domain receives for sharing its resources should compensate the costs that the domain incurs, but should not be too high, to ensure that sharing remains beneficial for the other involved domains. The issue of cost-sharing gets even more difficult if users receive on-demand services such as bandwidth on demand. For on-demand services, costs are incurred even if no user makes use of the service so that there is a financial risk of not being compensated when providing resources for such a service. In addition, certain levels of service-quality guarantees will also have cost implications so that they have to be selected related to that.

Agreement on a fair cost sharing model is necessary for the federation to be acceptable to the executives of the participating domains. In the GN3 project a special Cost Sharing Working Group as a subcommittee of the NREN PC is working on this topic.

- Management challenges. Joint service operation by the NRENs requires tighter network management collaboration with respect to the fault, configuration, accounting, performance and security (FCAPS) management areas. This issue is mainly an organisational challenge related to specifying processes and the responsibilities associated with them. Tools and other technical means are only used to support the processes.
    - Fault and performance. The fault and performance management processes and tools being used in the domains have to be integrated to resolve issues in federated services. In the GN3 project, the federated Performance Enhancement Response Team (PERT), perfSONAR and also I-SHARe (with its planned process management component) are already working towards this objective.
    - Configuration. Federated services must be appropriately configured with respect to their use of NREN resources. A distinction can be made between "loose coupling" and "tight coupling". Loose coupling can be based on a best-effort approach to service delivery, with rather informal agreements and without coupling provisioning systems for automated interaction, whereas tight coupling requires an agreement-based federated IT system.
    - Accounting. Depending on the cost-sharing and service models used, charging may be required. The charging can be used for federation-internal cost-sharing or for user accounting.
    - Security. Federated services should not lead to a lower level of security in comparison to non-federated services. Therefore, the whole federation needs to agree on common security standards.
    - Management systems today are built for single-domain purposes and are therefore not suitable for a multi-domain environment. The establishment of federated services also has to address this. GN3 already provides one building block for monitoring, i.e. the End-to-End Monitoring System (E2EMon).
- Technological differences and missing standards. Federating a network has to address the issue of technological differences on the layers below IP. Some domains are based on SDH, others on native Ethernet or Ethernet over MPLS. If federating occurs even on the optical level, the federation has to address missing standards for interworking on this level. The technological details are addressed in particular in JRA1 (Task 1 and Task 2) related to technical differences as well as in JRA2 (Multi-Domain Network Service Research, in particular Task 1) for management standards.
- Service offer. For the users, the fact that the service is provided using a federated network should be transparent. Rather, the service offer should be independent from it and simply specify a service with the needed features and guarantees. The offer should be no worse, in terms of service quality, than the offers of a single provider. For the details of how the service is provided to be invisible to the user, it is necessary to abstract the resources that are used from the network. Building this abstraction requires the technological and organisational differences to be addressed.

3.3 Overview of Current and Future Large-Scale Projects

A number of current large-scale projects within the GÉANT user community are examples of network services delivered by collaborating networks and feature network federation in some manner. They therefore provide models for typical users of a future federated network and so were analysed as part of our research in order to assess user needs for federated networks. The focus areas of the analysis included network topology, requirements for data delivery, workflow and cost models. These particular parameters were included in the analysis in order to gather information relevant to the federated network development process.

The GÉANT user community projects analysed include:

- Large Hadron Collider Optical Private Network (LHCOPN).
- Electronic Very Long Baseline Interferometry (e-VLBI).
- Enabling Grids for E-Science (EGEE-III).
- Distributed European Infrastructure for Supercomputing Applications (DEISA 2).
- CINEgrid.
- Low Frequency Array (LOFAR).

Some of the future projects that will come under the European Strategy Forum on Research Infrastructures (ESFRI) umbrella might have different connectivity and (federated) network support compared with that given to current projects, so these were also included in the analysis.

This section presents an overview of these projects and the results of our analysis, which were then used to develop potential models for a future federated network. The section ends with a summary and conclusion of our research.

3.3.1 Large Hadron Collider Optical Private Network (LHCOPN)

Description

The LHC project [LHCCERN] is organised in a tiered model with CERN, as the leader of the project and the site where experimental data will be produced, defined as Tier0. Tier1 participants have been defined in several countries and are connected directly to CERN through an Optical Private Network (OPN) provided by GÉANT. Some Tier1 participants are also connected with each other.

The LHC experiments generate enormous amounts of data which are distributed several times a day to 11 Tier1 computer centres for analysis by teams of physicists. This amount of information requires high-bandwidth capacity and a reliable network. The LHCOPN facilitates the efficient and seamless distribution of data. It consists of CERN (Tier0) and the Tier1 computer centres, enabled by the GÉANT network and its connections to other National Research and Education Networks (NRENs) around the world (see Figure 3.1 on page 12).

Topology

The topology consists of point-to-point links configured across different domains in the LHCOPN project. The architecture of the network is therefore heterogeneous. Every link has been constituted over RRENs (Regional Research and Education Networks), NRENs, GÉANT, CANARIE and Internet2 infrastructures, which are all based on dark fibre or leased lines. In some cases, dedicated infrastructure is used, such as dark fibre between some Tier1s or the LHC network in the US. However, the main characteristic is that every connection is a 10 G transparent link.

The connections used for the LHCOPN project have been established for a long time and are static – as required by the project users.

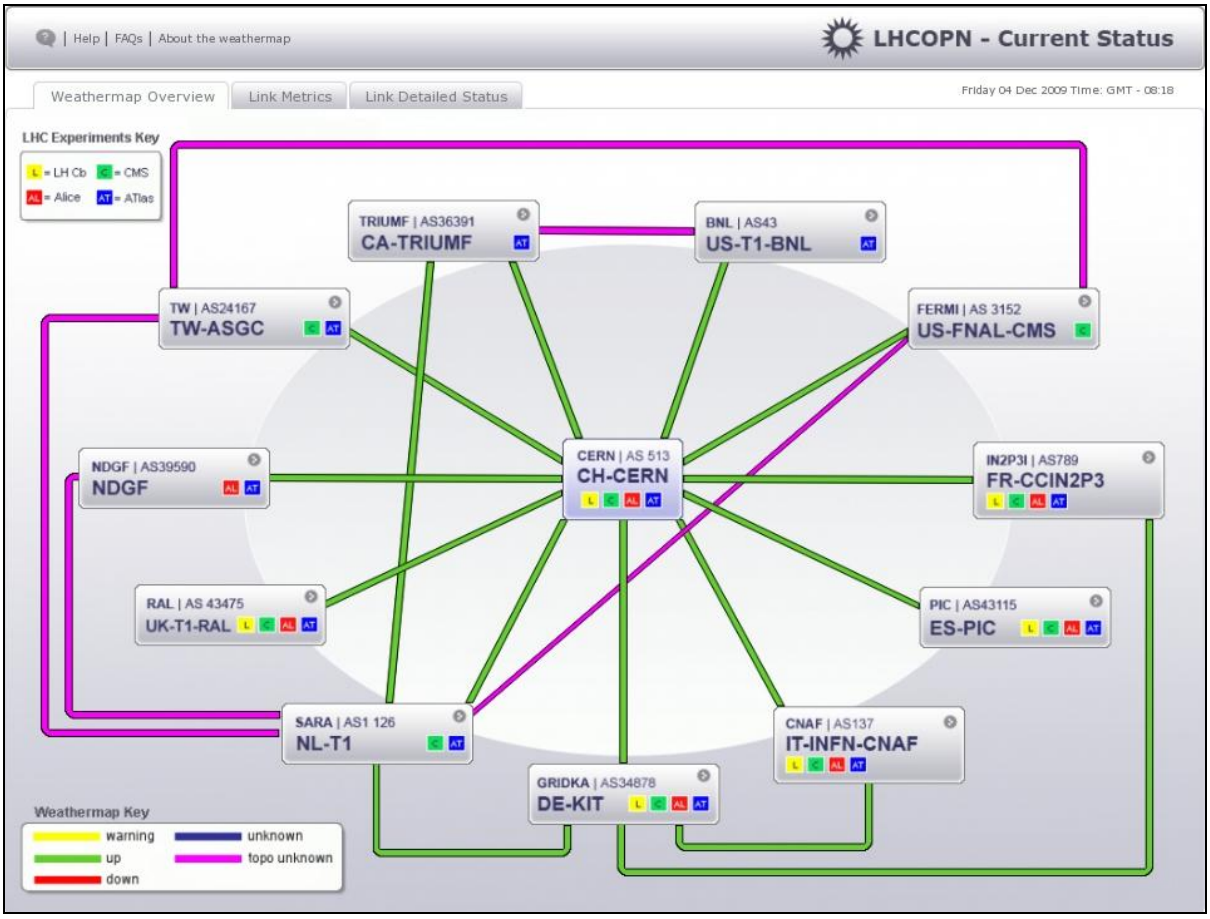

Figure 3.1: LHCOPN deployment map

Data Requirements

At the time of writing, we don't have any firm information about data-flow characteristics. This is due to the fact that the experiments are not yet producing real data. We plan to obtain and analyse further information, including data-flow estimates produced by the CERN IT department, as part of our future work.

Operations and Services

The LHCOPN project has defined its own operational model with an SLA involving several entities.

No Quality of Service (QoS) mechanisms for IP traffic are used in this project at the moment. However, some tests have been done, because traffic between Tier0 and Tier1 sites needs to be prioritised over other, lower priority traffic. These mechanisms will be deployed in the near future.

As LHCOPN is a private network, security procedures are not implemented. As far as we know, there are no security monitoring tools.

Cost Model

Usually, every Tier1 has to cover all the costs related to the connection with the Tier0 and another Tier1. This is specifically related to the cost of the network, not of the whole project, as in the other examples.

3.3.2 electronic Very Long Baseline Interferometry (e-VLBI)

Description

In e-VLBI [EVLBI] the data from distant radio telescopes is streamed to the central processor through dedicated optical fibres, and correlated in real-time. The greatest advantage is that real-time processing enables more flexible observation planning.

Location and Topology

The correlation centre and 14 radio telescopes are located in Europe. The rest of the infrastructure is located in South America, Africa and Asia.

The connectivity between remote sites is realised in the form of static and long-term connections. Because the correlation process performed on data streamed from radio telescopes to the Joint Institute for VLBI in Europe (JIVE) is very sensitive and vulnerable to any anomalies in a network, the connectivity is heavily tested several weeks before the actual experiment takes place. The role of JIVE, which is located in Dwingeloo, Netherlands, is to operate the data correlator for the European VLBI Network (EVN) and provide support to the world-wide user community.

Data Requirements

VLBI data needs to be transferred at a specific rate, which may be at least 16 Mbps and up to 1024 Mbps. Although Transport Control Protocol (TCP) is the usual method for transporting VLBI data, it is possible to use User Datagram Protocol (UDP). EVN’s experience has been that there is no significant improvement using UDP (from the end-user’s point of view).

The e-VLBI data transmission places the following requirements on the underlying network:

- Fast transport (>0.5 Gbps) through long, large capacity networks. The network may be shared with other users or an Optical Private Network (OPN) might be used.
- Fairly reliable transmission. To maintain synchronisation of the data from the two telescopes, the correlator depends on receiving time stamps from headers in the application data stream.
- Constant transfer rate. The network transfer rates are controlled by the application. The data arrives from the telescope at a constant bit rate.
- Predictable latency. One-way delay for link (jitter) should remain constant, ensuring timely delivery of data.
- Bit-wise correct data. Incoming data should be bit-wise correct.
- Lost packets must be detected by the receiver and reported to the application.
- Low re-ordering in the network such that data can be placed in the correct order in a simple manner by the application in the receiver.

- Packet duplication must be minimal.
- Distributed correlation on grid systems requires the same portion of data from a given telescope to be sent to different remote nodes. For example, a given compute node may only correlate one pair of telescopes.

Operations and Services

- It must be possible to monitor and test the network before any astronomical experiment takes place, to ensure good connectivity during the experiment.
- No formal or informal workflows or procedures have been implemented in the EVN environment.
- Currently, the connections are maintained by the NRENs and available for the experiments (which are performed several times a year).
- Network failures are directly reported to the NRENs involved and/or GÉANT NOC.
- Connectivity performance problems are either reported to the NRENs involved or raised with the Performance Enhancement and Response Team (PERT).

The radio astronomers have already conducted some initial work together with the GN2 project partners in the area of QoS. Recent developments are the result of a collaboration between the GN2 AutoBAHN team and the Software Correlator Architecture Research and Implementation for e-VLBI (SCARIe) project.

e-VLBI monitors the quality of data coming from the correlator and the stations during an e-VLBI experiment. The software is called Data Status Monitor (DSM). DSM provides a quick overview of the quality of data and allows further investigation of possible problems. However, this monitoring system does not perform any measurements on the underlying network. It just helps the operators (end users) to rate the quality of the experiment's results.

With regard to "pure" network performance monitoring, it is worth mentioning the Express Production Real-time e-VLBI Service (EXPRES) project, where the development of the monitoring infrastructure for e-VLBI takes place. In EXPRES, the project partners deliver network monitoring support for workflow management in e-VLBI experiments with a software module that makes use of the existing and newly installed perfSONAR monitoring and measurement nodes. The open architecture of perfSONAR allows external applications to use data collected by the perfSONAR system through the NMWG XML schema. The measurement framework collects performance data (bandwidth utilisation and Round-Trip Time (RTT)) using perfSONAR's Command Line Measurement Point (CL MP) service installed in each involved domain and covering network paths between the resources. EXPRES has already deployed the first version of the monitoring infrastructure in several locations and is now finalising the validation and testing phase.

The VLBI real-time data does not have any security requirements.

Cost Model

The project is financed by the European Commission and participating countries.

3.3.3 Enabling Grids for E-sciencE (EGEE-III)

Description

Enabling Grids for E-sciencE (EGEE-III) is the world's largest computing and storage distribution infrastructure, consisting of more than 260 sites in 55 countries. The global infrastructure schema is available on the EGEE-III website [EGEE-III].

Topology

The project deploys homogeneous services on top of a Red Hat Enterprise Linux 4 (RHEL4) operating system. Each resource centre serving a grid site deploys services that have a set of public interfaces. Users and other grid sites access the services via the public interfaces. Authentication and authorisation of users on all grid sites is performed by using standard grid protocols based on X.509 certificates: Grid Security Infrastructure and Virtual Organisations.

Data Requirements

A grid infrastructure is inherently distributed, therefore all services in a grid utilise the network. The two most common use types are user job execution and large data transfers. Job execution consists of a small amount of network communication, needed for services to initiate and monitor user jobs. Large data transfers are mainly performed using the GridFTP protocol and this utilises bandwidth heavily.

Operations and Services

All workflows are performed completely automatically. Services spanning multiple domains automatically communicate with each other constantly. As mentioned above, security between domains is based on grid security protocols, which don't require the administrators' interaction. A grid provides information services (e.g. BDII) that provide other services (e.g. services for job executions or data transfer) with up-to-date information about all available resources.

Individual resource centres can provide a defined set of resources (e.g. CPUs, storage space). Reservations are achieved by using standard cluster batching systems (e.g. SGE, Torque) in the case of CPUs and grid storage systems (e.g. dCache, DPM) in the case of storage space.

There is no centralised network performance monitoring in EGEE-III. Activity EGEE Network Operations Centres (ENOCs) perform grid site and service availability monitoring and provide network failure information to relevant grid administration personnel. Besides ENOC there are several non-network-specific monitoring tools, which are aimed mainly at monitoring service availability (e.g. SAM, Nagios monitoring, BDII). Individual grid sites monitor their local networks using various tools.

There is no specific network traffic security monitoring tool.

Cost Model

The project is financed by the European Commission and participating countries.

3.3.4 Distributed European Infrastructure for Supercomputing Applications (DEISA 2)

Description

DEISA2 [DEISA2] is a consortium of national supercomputing centres that currently deploys and operates a continuous, production-quality, distributed supercomputing environment with continental scope. The purpose of this FP7-funded research infrastructure is to enable scientific discovery across a broad spectrum of science and technology, by enhancing and reinforcing European capabilities in the area of high-performance computing. This becomes possible through an integration of existing national high-end computing platforms, tightly coupled by a dedicated network and supported by innovative information systems and grid software.

DEISA2 involves seven European countries with large supercomputing centres.

Topology

DEISA2 operates a heterogeneous High-Performance Computing (HPC) infrastructure, currently formed by eleven European national supercomputing centres that are tightly interconnected by a dedicated high-performance network. The term “heterogeneous” refers to the variety of HPC system architectures, operating systems, batch schedulers and local file systems provided by the DEISA2 supercomputing centres and which is typical for an HPC environment.

The interconnections used by the DEISA2 network are established on a long-term basis and are static.

Data Requirements

DEISA2 also runs a Global File System (GFS) where each supercomputing centre has a common data space, which is a shared file system transparently accessible from every compute node in the cluster at one site. Such a cluster-wide shared file system offers a single system view to compute jobs that are running locally on the cluster.

From a networking point of view, maintaining such a file system may be vulnerable to jitter and delay.

Operations and Services

The application support service is provided on different levels. Early adopters of the DEISA2 infrastructure from different scientific communities (e.g. Materials Science, Cosmology, Plasma Physics, Life Sciences, Engineering and Industry) are individually supported.

DEISA2 is using connections with dedicated bandwidths, so no other QoS procedures were implemented.

DEISA2 is running local network performance tools and is now investigating the possibility of using perfSONAR.

As DEISA2 is running a dedicated (private) network, security procedures are not implemented, and there are no security monitoring tools.

Cost Model

DEISA2 is funded by FP7.

3.3.5 European Strategy Forum on Research Infrastructures (ESFRI)

ESFRI [ESFRI] is not a project. It is a body concerned with the study of the infrastructure requirements of future projects to be conducted in Europe. The networking requirements of some of the ESFRI projects are described below. The choice of the projects has been the result of an analysis of all projects described in the ESFRI roadmap and by selecting those where large amounts of data are likely to be involved so that a federated project network may be relevant.

ESFRI's description of projects is focused on the research that is carried out. It is often hard to estimate how much data will be generated and whether it is planned to process this data in a distributed manner. Out of the listed projects, the following are likely to generate a large amount of data and require a data-distribution infrastructure.

Description

- **LifeWatch** [LifeWatch]. The project aims to bring together research on biodiversity from all over Europe by joining forces on data mining and simulation. The project explicitly refers to itself as an e-infrastructure. The summary of the activities on the website says that contact with EGEE has been established, while stating only that the collaboration should happen "over the Internet". The timeline of the project foresees a preparation phase from 2008 until 2011 and a construction phase from 2010 until 2018.

- **European Biobanking and Biomolecular Resources Research Infrastructure** (BBMRI). The aim of the project is to link data centres that hold biological information (e.g. patient data such as DNA, tissues and blood) across Europe, with nearly all countries participating. The website says that the biological databases currently contain 10 million samples. It is unclear whether a special requirement for the infrastructure exists or whether it can be handled over the commodity Internet. The timeline of the project foresees a preparation phase from 2008 until 2010 and a construction phase from 2010 until 2013.

- **Infrafrontier** [Infrafrontier]. This project is going to link 15 laboratories that conduct experiments with mice and which have collected large experimental databases over time. As with BBMRI it is unclear whether the exchange of data requires a special infrastructure. The timeline of the project foresees a preparation phase from 2008 until 2011 and a construction phase from 2011 until 2020.

- **X-Ray Free-Electron Laser (European XFEL)** [XFEL]. For this project a new facility for X-ray laser research will be constructed in Hamburg. It will serve purposes like the detailed analysis of chemical reactions and the analysis of material structures. The project documentation says that many external researchers will come to the site to conduct their experiments. Even though the project description is already quite detailed, it does not provide information about computing facilities, so that it is unclear whether experimental data will be analysed on-site or will be distributed. The accelerator and first six experimental stations will be commissioned starting in 2014.

- **European Extremely Large Telescope (E-ELT)** [E-ELT]. This project will deal with a new generation of telescopes (42 meters in diameter) where the decision about the construction site will be taken in

2010. The construction is planned for 2010 until 2017. There is no information available on how much data will be generated and distributed, but it seems likely that the amount of data will be large.

- **Facility for Antiproton and Ion Research (FAIR)** [FAIR]. This project deals with a large experimental facility that will be constructed in Darmstadt, Germany. It will allow experiments with high-energy beams of ions for which already more than 2,500 researches have declared interest. The construction starts in 2009 and the start of the experiments is planned for 2015. So far information on a required networking infrastructure cannot be found, and it is not clear whether a network like the LHCOPN will be required.

- **Partnership for Advanced Computing in Europe (PRACE)** [PRACE]. The PRACE project wants to link supercomputing centres in Europe (the principal partners being GENCI/France, Gauss Center/Germany, CINECA/Italy, NCF/Netherlands, BSC/Spain and EPSRC/UK) and intends to build a structure with several tiers for it. The project is not linked directly to DEISA even though DEISA does also link European supercomputing centres.

- **European Incoherent Scatter (EISCAT)** [EISCAT]. EISCAT is an international research organisation operating three incoherent scatter radar systems, at 931 MHz, 224 MHz and 500 MHz, in Northern Scandinavia. It studies the interaction between the Sun and the Earth as revealed by disturbances in the magnetosphere and the ionised parts of the atmosphere (these interactions also give rise to the spectacular aurora, or Northern Lights). EISCAT 3D is the next generation EISCAT radar system due to go into production in 2015.

Location

The projects have different geographic distribution scenarios. LifeWatch, BBMRI and Infrafrontier link research facilities in many countries. PRACE is similar, but may be limited to a few countries only. In E-ELT the telescopes may be located outside Europe and only send data for processing to several European locations. XFEL and FAIR will have their centres in Germany and may distribute the data to several other countries, while EISCAT 3D will have several radars in Scandinavia.

Topology

LifeWatch, BBMRI, EICSAT 3D and Infrafrontier have participants who act as peers. A computing grid comparable to EGEE therefore seems to be the right approach. PRACE intends to form a hierarchy of the supercomputing centres, based on their computing power. Within each tier a partial mesh can be established. While these projects would clearly use point-to-point services only, multi-cast can be relevant for XFEL, FAIR and E-ELT, where data from a central generation point is likely to be distributed to several computation points. For these projects a hierarchy of tiers is also likely.

For LifeWatch, BBMRI and Infrafrontier it is not obvious whether data-exchange patterns are going to be quite stable, which will certainly influence how dynamic the exchange possibilities should be. For PRACE, XFEL, FAIR, E-ELT, and EICSAT 3D it is likely that a permanent topology will be useful in the core network, but more dynamic configurations may be preferable for the edge or lower tiers. It is interesting to compare this situation with the solution that has to be found for the Tier1-to-Tier2 connection in the LHC network.

Data Requirements

It is too early to give a statement about data characteristics. For E-ELT and EICSAT 3D it is likely to be a continuous stream of data, while bursty traffic is more likely for LifeWatch, BBMRI and Infrafrontier.

Operations and Services

There is not enough data regarding QoS mechanisms and performance monitoring to give a statement on what will be required.

The security of data may be particularly important when dealing with medical data in the BBMRI project.

Cost Model

These projects will be funded under the rules of under FP7/8.

3.4 Project Information Mapped to Federated Network Parameters

Based on currently available information on networking parameters for existing and future projects, we believe that we have a valid representation of the services that could be (or are) users of the federated network model(s).

Table 3.1 below summarises the four major current projects using the parameters that we consider to be fundamental to our research.

Parameter	LHC	e-VLBI	EGEE-III	DEISA 2
Geographic locations	1 Tier0 site (CERN) 11 Tier1 sites (7 EU, 4 global)	6 EU sites, 4 global sites	260+ global sites	11 EU sites
Topology	Heterogeneous Multi-domain Private network	Private network	Public network	Heterogeneous Layer 2 Private network
Type of E2E connections	Static / long term	Static / long term (interest in short-term solutions)	N/A	Static / long term
Data characteristics	No information at this time Up to 10 Gb/s during the runs	Fast transport Large amount Constant stream Can recover from some data loss Limited jitter	Fast transport Large amount Occasionally Can recover from data loss and jitter/delay	GFS Fast transport Large amount Constant connectivity Constant jitter/delay requirements

Parameter	LHC	e-VLBI	EGEE-III	DEISA 2
Workflow and procedures	Has its own operational and SLA with several parties involved	None (NRENs/DANTE/ PERT)	Automatic	Formal Automatic
QoS	Optical protection (dedicated network)	Tried AutoBAHN	N/A	None
Performance monitoring	perfSONAR	DSM / EXPRES perfSONAR	No centralised monitoring ENOC performs some	Local Investigating perfSONAR
Security model	Not required	Not required	N/A	Not required
Cost model	Tier1 pays for the network connection to Tier0 and another Tier1. This is the cost of the network, not the cost of the LHC project as a whole.	EU and participating countries	EU and participating countries	EU/FP7

Table 3.1: Summary of current projects using key parameters for federated networks

It is clear that the projects are not tied to a single network service provider and that their characteristics and requirements vary greatly. Some projects do not have a very strong concept of what the network should be like and have opted for fairly simple – even sub-optimal – network designs.

Most projects have a strong concept of performance monitoring and incident resolution at the level of their application. Some have incorporated end-to-end monitoring of network connectivity, or have even introduced perfSONAR. It is clear that the ability to verify the quality and integrity of network services is important. Some projects have a dedicated Network Operations Centre (NOC) function that provides a centralised approach to incident resolution for the project.

All projects have some concept of how to interact with the network service providers. Projects tend to deal directly with the local NREN, sometimes facilitated by the central site, often in collaboration with GÉANT. However, this does not resolve the possible issues relating to the functioning of the service on a global scale that might arise.

4 Building Blocks

The continuous development of network technologies in National Research and Education Networks (NRENs), complemented by a new approach to interconnecting neighbouring NRENs with fibre infrastructure (cross-border fibre (CBF) development), opens new perspectives for federating network infrastructures at the “technology” layer and brings benefits not only to the parties involved (the NRENs), but, most importantly, to end users. However, the advances in technology have not been fully reflected at the operational level. This was identified at an early stage of the GN2 project and addressed during the development phase. The work done by the GN2 project partners resulted in a number of tools to support end-to-end requests from the end users of research networks. These tools supporting multi-domain workflows present an opportunity to analyse how a future federated network in Europe might be built.

This section introduces the three basic building blocks of future federated networks: network technologies, operations and services. Each component is described in detail, with particular focus on existing elements that are potentially useful for building a federated network in Europe.

4.1 Networks

This section gives brief summaries of the infrastructures of some European NRENs. It is based on a survey conducted among the NRENs participating in GN3 JRA1 T3. The aim of the survey was to get an overview of the current status of networks in a number of countries and investigate the issues involved in a federated GÉANT network. The results of the survey show the diversity of the research networks in Europe, which can be seen in many aspects, from the network coverage to the technologies used in the NRENs.

In order to investigate a potential federated research network in Europe, the emphasis was on the following aspects:

- Network topology –how NRENs have built their networks and what kind of technology is used.
- Network coverage – the current coverage of NRENs in particular countries.
- Network utilisation –the utilisation of links in NRENs.
- Network development – future plans for network development in the countries participating in the survey.

The following NRENs took part in the survey:

- DFN – German Research Network (Deutsches Forschungsnetz).
- PIONIER – Polish Optical Internet (Polski Internet Optyczny).
- RedIRIS – Spanish academic and research network (Red de Inteconexión de Recursos InformáticoS).
- SURFnet – Dutch Research and Education Network.
- CARNet – Croatian Academic and Research Network (Hrvatska akademska i istrazivacka mreza).

4.1.1 Network Technologies in Research and Education Networks

This section summarises the network technologies used in the NRENs taking part in the survey.

- According to the results of the survey questionnaire, some NRENs lease or own dark fibre links, based on Dense Wavelength Division Multiplexing (DWDM) and Multi-Protocol Label Switching (MPLS) technologies in the core. On the other hand, there are also core networks based on leased capacities.
- Some NRENs have cross-border fibres (CBFs) to neighbouring NRENs, and also some point-to-point links dedicated to special projects.
- NRENs have one or more connections to the commodity Internet, via the GÉANT network and via other providers.
- Bandwidth in the core ranges from a few Mbps up to 40 Gbps, depending on the NREN and on the part of that country.
- There are differences between countries regarding vendors of the networking equipment used. The vendors mainly mentioned in the survey are Cisco, Juniper, Foundry Networks, Huawei and Nortel.

4.1.2 Network Coverage and Utilisation

There are differences among NRENs regarding how much of their network is already built and the coverage of their network. Some NRENs already connect all their member institutions, but some connect less than 60% of them.

In addition, some NRENs' member institutions are at the university level (universities, research institutions, corporate R&D departments, scientific libraries, teaching hospitals), while others' members also include primary and secondary schools and related educational institutions, hospitals, libraries and other government institutions.

In the countries of the NRENs that completed the survey questionnaire [JRA1S], all institutions that are participating in EU projects are already connected to their network.

The most demanding users in those five countries are projects like LHC, DEISA, e-VLBI or LOw Frequency ARray (LOFAR) that are being carried out on a national level as well as on an international level. Four of the NRENs identified those projects as their most demanding users. Apart from projects, NRENs usually have a few member institutions that are more active than others. These are universities or member institutions that are part of a university.

According to the responses, network link utilisation varies from less than 10% to up to 90% of available bandwidth. For a federated network that might mean that not all links can be treated and counted on equally, otherwise network saturation might result.

For links that are used for EU projects, it was not easy to say how much bandwidth is already used, or how much available bandwidth can be counted on for a federated network. Some NRENs do not monitor performance on those dedicated links, while some do not monitor the activities of those projects. However, in one of the NRENs surveyed, such links have a load of 10% – 40% [JRA1S].

4.1.3 Cross-Border Fibre (CBF) Initiatives in Europe

The current direction of development within European NRENs is mainly related to the expansion of optical networks based on leased or owned fibre infrastructure. Optical networks guarantee high capacity and flexibility both in network growth and in the choice of technology used. After covering their own country with optical networks, NRENs started to connect directly to each other. Cross-border fibres, in addition to standard GÉANT services like GÉANT Plus, have become a complementary way of providing international connectivity. Figure 4.1 gives an overview of current and planned cross-border fibre between European NRENs. (The map is taken from the TERENA compendium [TCNREN] and presents the status of CBF in Europe in 2009. It is based on information provided by the NRENs themselves, which in some cases may not have been sufficiently detailed for the map to reflect with complete accuracy the current status of the CBF initiative.)

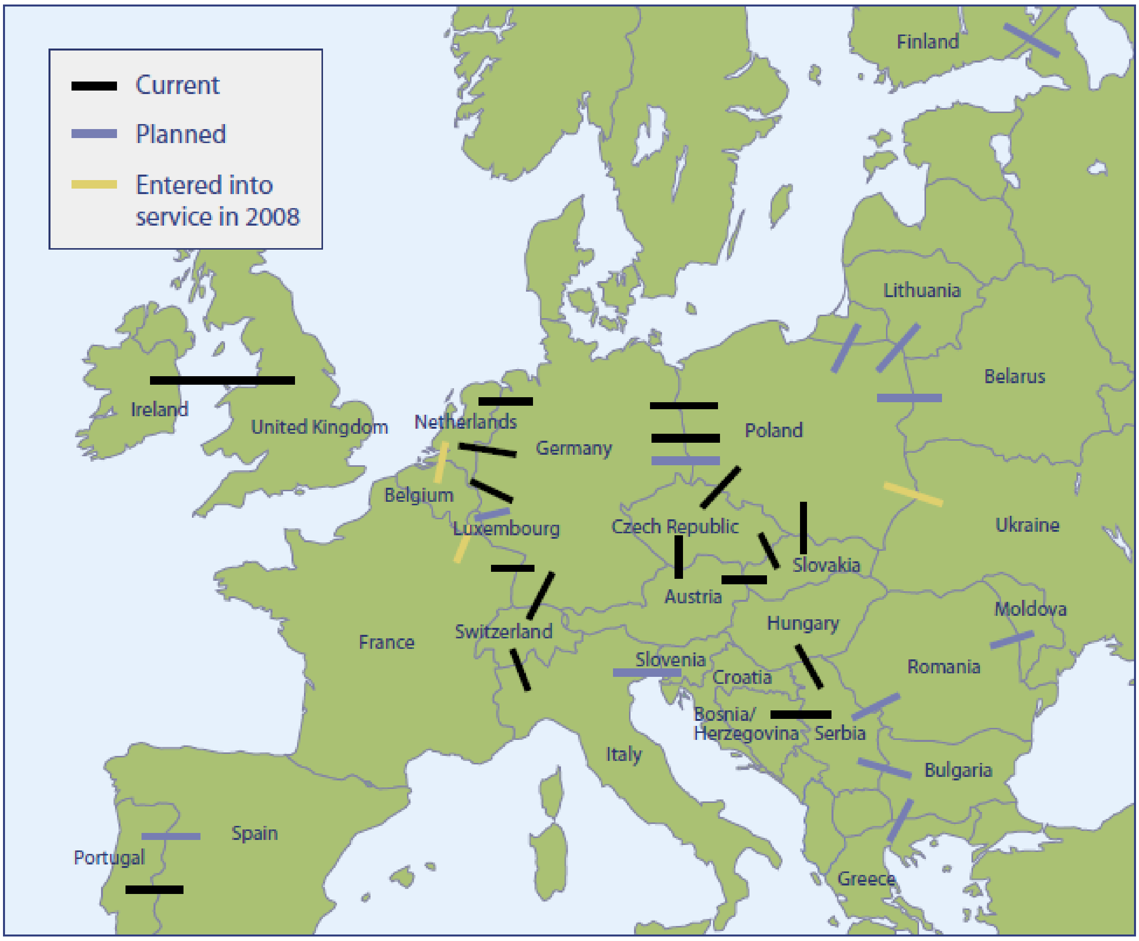

Figure 4.1: Overview of cross-border fibres between European NRENs, TERENA Compendium 2009 [TCNREN]

As Figure 4.1 shows, the level of CBF development varies across Europe, with a visible concentration of cross-border fibres in the centre of Europe and a relatively low number of CBFs in Western, Southern and Eastern Europe. The lower number of cross-border fibre connections in those regions may indicate that there are no requirements for creating CBF points or it may be a consequence of the level of development of the NRENs, of regulations that limit the possibilities for owning fibre or creating international connections, or of difficulties with acquiring dark fibre from network operators.

Three examples of the successful implementation of cross-border fibres by European NRENs are described below.

4.1.3.1 *PIONIER*

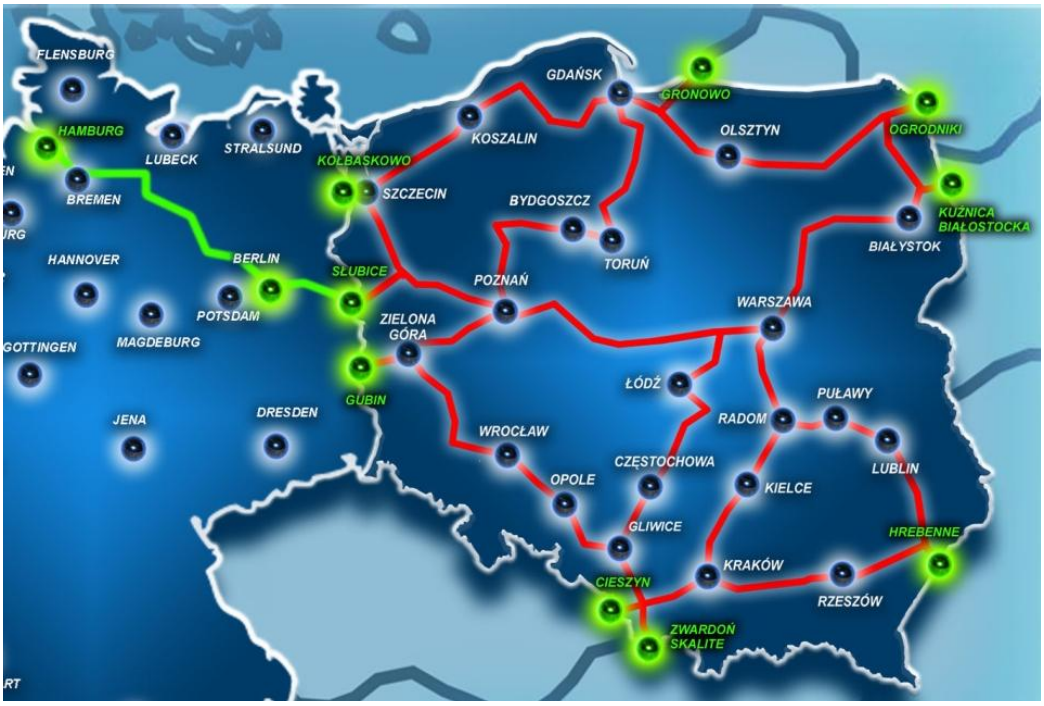

Figure 4.2: PIONIER fibre network and CBFs

The PIONIER network, which is based on optical fibres, has 11 cross-border fibre connection points. These make it possible to connect to each of Poland's neighbours: Russia, Lithuania, Belarus, Ukraine, Slovakia, Czech Republic and Germany. Moreover, thanks to new fibre between Słubice and Hamburg, as well as additional connection points to Germany, new possibilities for connecting to the Netherlands and the Nordic countries have arisen. The full list of CBF points is as follows:

- Germany (Hamburg, Berlin, Kołbaskowo, Słubice, Gubin).
- Czech Republic (Cieszyn).
- Lithuania (Ogrodniki).
- Belarus (Kuźnica Białostocka).
- Ukraine (Hrebenne).
- Slovakia (Zwardoń).
- Russia (Gronowo) – to be ready in 2010.

4.1.3.2 *SURFnet*

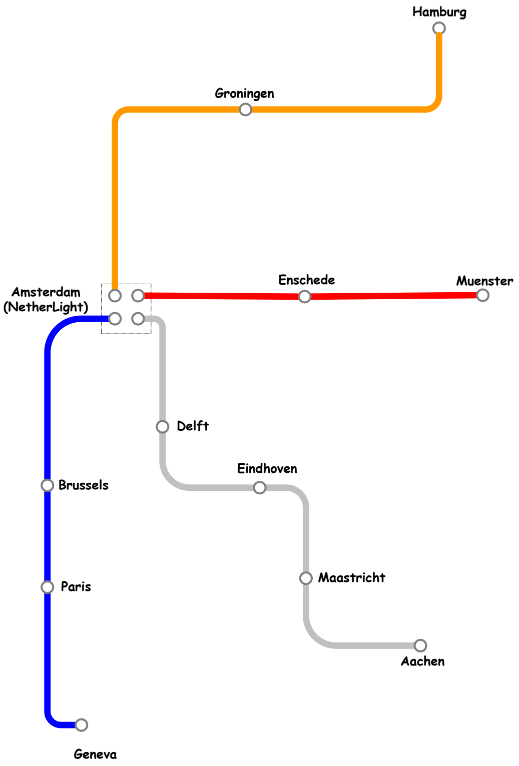

Figure 4.3: SURFnet cross-border fibres

Currently, SURFnet has four CBFs: three to Germany (Hamburg, Muenster, Aachen) and one that connects Amsterdam to Geneva through Brussels and Paris (see Figure 4.3).

4.1.3.3 *CBF Triangle – Czech Republic, Austria and Slovakia*

Cross-border fibres are also used by NRENs for increasing the reliability of connections. For example, there is a fibre triangle between the Czech Republic, Austria and Slovakia running at 10 Gbps, which allows the quick (around 60 ms) redirection of traffic in the event of failure of the CBF connection between two countries. It is achieved using Layer 2 protocols [SSK]. The CBF triangle is shown in Figure 4.4.

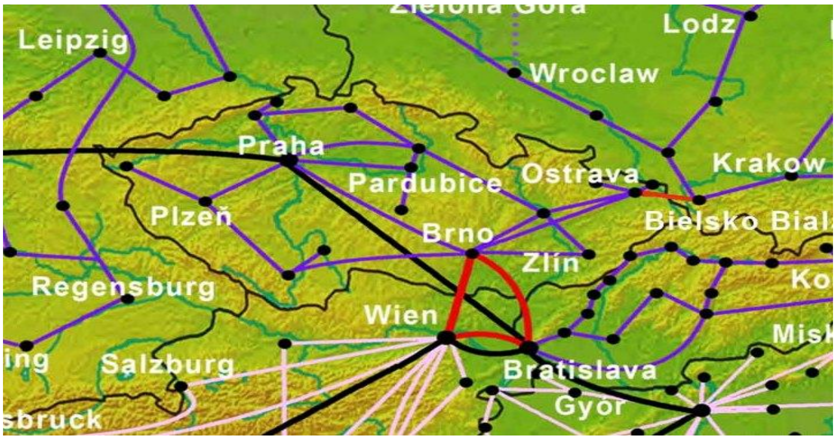

Figure 4.4: CESNET, SANET and ACONET CBF triangle [SSAW]

4.1.3.4 *CBF Configuration Options for a Federated Network*

With regard to the technological aspects of cross-border fibre connections between NRENs, different configurations for interconnections can be considered, some of which are shown in Figure 4.5.

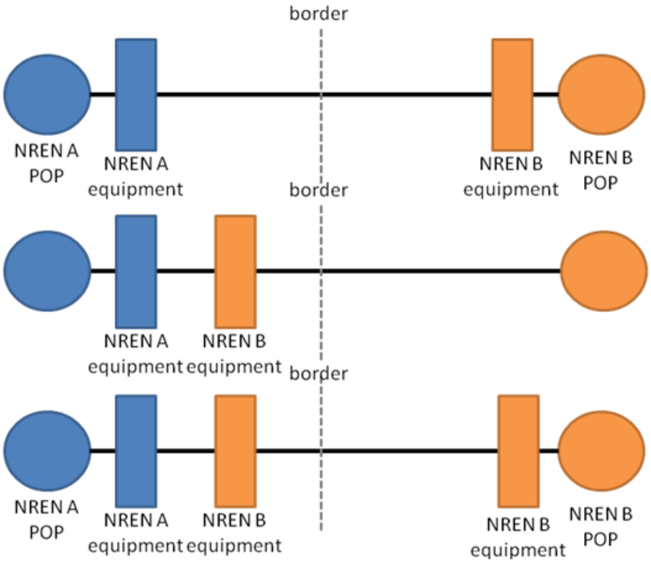

Figure 4.5: Examples of cross-border connections between NRENs

Figure 4.5 above presents some possibilities for configuring cross-border fibre interconnections between two NRENs. In practice, configuration of the real connection depends heavily on many factors, such as availability of fibre, PoP locations, availability of tele-housing, the optical transmission system used, etc., so there is no universal rule for connecting two NRENs. In each case an agreement describing each NREN’s area of responsibility is necessary, as well as an agreement on cost sharing. The agreement needs to specify which NREN is responsible for which part of the infrastructure, who will service which part of the infrastructure, and who will pay for the servicing and maintenance.

4.2 Operation

This section summarises the existing tools developed in GN2 and which continue to be developed as part of GN3. The requirement to support multi-domain workflows and procedures is one of the tools' basic design assumptions. eduPERT [eduPERT] is the first example of a potentially valuable element that could be re-used when building a federated environment. In contrast to the other elements, eduPERT is a service. Moreover, it is a production service, which started at the end of GN2. It is mentioned here because it is helpful for a general discussion on federation; although eduPERT is not explicitly concerned with federating networks, its findings and results are nonetheless relevant to this activity.

The section starts with a general description of GÉANT tools and supporting services (Section 4.2.1), then summarises the features that are most important for a federated network (Section 4.2.1.5).

4.2.1 Tools and Supporting Services

This section describes the following groups of tools and supporting services developed by GN2 and GN3:

- Performance troubleshooting in federated networks.
- Measurement-support services for a federated environment.
- Automated bandwidth provisioning tools.
- Multi- and single-domain information systems.

4.2.1.1 *Performance Troubleshooting in Federated Networks*

One example of a service that supports end users in their day-to-day operations in a multi-domain environment is eduPERT [eduPERT]. eduPERT shows how a federated network can be supported with additional services, helping end users to get optimal performance from their networked applications.

Originally, the Performance Enhancement Response Team (PERT) was set up as an informal group of network experts, who expressed their interest in solving difficult performance issues. The complex nature of such problems, especially in the context of end-to-end investigations, resulted in changes to the PERT's organisation. In consequence, the PERT was formally established during the GN2 project, to investigate pure end-to-end network paths. A strong need for the decentralisation of some GN2 services, including the PERT, led to the subsequent migration towards a federated PERT service (eduPERT).

Started in September 2008, eduPERT is a federated PERT that combines the independent PERTs (the GÉANT PERT and the national, local and project PERTs) with a portfolio of central services to aid them in their network investigations. Importantly, eduPERT is not limited to NRENs. Universities, other organisations and international projects are encouraged to join, by setting up “local” or "project" PERTs and becoming part of the eduPERT Knowledge Base. Once registered, a PERT may apply to become accredited. Accredited PERTs commit to providing a minimum level of service and guarantee to respond to requests for help within a selfimposed maximum time limit of no more than two working days. Accredited PERTs are guaranteed access to the equivalent information from other accredited PERTs [eduPERT].

Administrative Structure of eduPERT

eduPERT implements a model that seeks to combine the strengths (and minimise the weaknesses) of the fully centralised and fully decentralised PERT models. These models have been thoroughly studied and described by the GN2 project [BBD]. Concentrating for now on Europe, the eduPERT consists of a well-funded, well-resourced central PERT (the GÉANT PERT) and all the European PERTs (national and (if they exist) regional and campus). Formally, there are four types of PERT:

- Central GÉANT PERT.
- National PERTs.
- Local PERTs.
- Project PERTs.

A detailed description of these types may be found in [TRP]. In principle, the GÉANT PERT is managed by DANTE, while national PERTs are managed by NRENs. Local PERTs are PERTs established on university campuses or within other similar institutions. Project PERTs are those set up by international projects that have a centralised networking function, and in which no single country is involved.

A fundamental issue for eduPERT is communication. In order to ensure the efficient and smooth investigation of PERT cases, it should be clearly stated *who* has to be contacted and *how* to contact them. This directly implies a clear and simple organisation structure. Figure 4.6 shows the administrative structure of eduPERT.

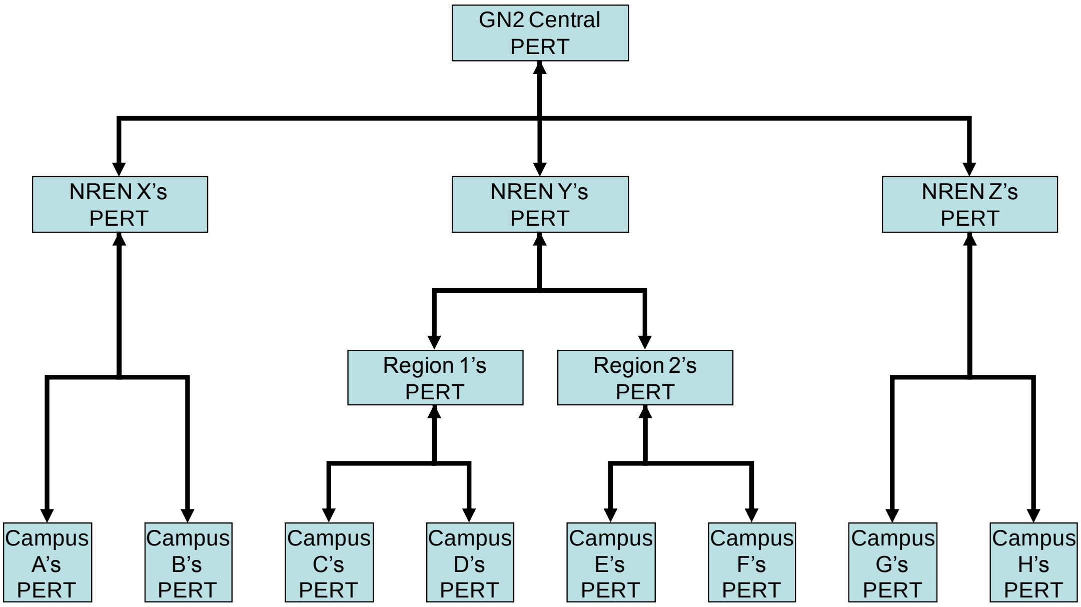

Figure 4.6: Hierarchical structure of eduPERT

Each parent PERT must maintain a register of any “children” that it has, including their statement of capabilities (hours of operation, response time, contact information, etc.). A regional PERT should keep a register of any campus or institutional PERTs under it. The PERT registers are available to the whole PERT community and speed up the process of contacting the right person in the course of a PERT investigation [SOP].

eduPERT should be considered as adding value to a federated network, as it actually achieves federation, but at a slightly different level: while it does not address any of the specific challenges involved in federating networks, it does solve a number of general issues related to federation. Any consideration of a future federated research network in Europe should therefore take into account how federation has been achieved in eduPERT. For example, a decision was made to keep a central database for eduPERT services, to allow the efficient management of performance investigations across Europe. [BBD] describes how a fully decentralised model imposed too many constraints on the overall case-investigation process to be realised efficiently in a live environment. In consequence, eduPERT has been implemented as a compromise between fully centralised and decentralised models.

4.2.1.2 *Measurement Support Services for a Multi-Domain Environment*

Mechanisms for monitoring the federated network architectures are one of the most important elements enabling trouble-free operation of vast networking infrastructures. Monitoring the most important network characteristics within a single domain allows quick problem resolution. The problem arises when dealing with services spanning multiple domains. Typically, each domain gathers its own monitoring data and defines specific policies for obtaining such data. Differences in the range of network characteristics gathered within domains and in access policies present significant obstacles to tracking issues in a multi-domain network environment. Defining and agreeing a common set of details to be provided by each domain allows technical users to access the necessary network performance data and diagnose the causes of problems, and is consequently one of the most important aspects of a multi-domain monitoring service.

One way to address these challenges effectively is perfSONAR (Performance Service-Orientated Network-Monitoring Architecture) [perfSONAR], a multi-domain monitoring service, created and maintained as part of GN2 and continuing to be developed in GN3. The perfSONAR project is carried out by a consortium of organisations (including GÉANT 2/3, ESnet and Internet 2 in the USA, and RNP in Brazil) and aims to design and build network performance middleware (the perfSONAR web services) and visualisation tools that assist in the analysis of local or remote interconnected network domains. One of the project’s main goals is to make it easier to solve end-to-end performance problems on paths crossing several networks. It allows consistent, standardised performance data to be accessed from any of the participating research networks. It is also able to carry out network measurements in each of these different networks [PIVIR, PMM].

perfSONAR allows flexible use of monitoring tools within a domain; depending on the requirements, an organisation can deploy a different set of monitoring tools. Access to the data from other domains is possible because an agreement has been reached between networks on unified data formats and the tools developed for perfSONAR. In this way, an efficient monitoring of paths crossing multiple different domains is possible.

The latest version of perfSONAR is 3.2. Work on perfSONAR is continuing in GN3 in JRA2 Task 3. and SA2 Task 3, which maintains the software produced in GN2. JRA2 Task 3, on the other hand, carries out research on new perfSONAR features (e.g. monitoring of lower layers, etc).

There is also a special version prepared for the Large Hadron Collider Optical Private Network (LHCOPN). perfSONAR MDM (Multi-Domain Monitoring) provides LHCOPN users with access to network measurement data from multiple network domains. This service is used to monitor the IP and circuit operations through probes placed at the 12 nodes of the network and providing users with access to network measurement data through specific weather maps and diagnosis tools.

The perfSONAR MDM support tool is vital to the efficient operation of LHCOPN. It enables optimum use of the available bandwidth and eases the burden on technical support staff.

Any consideration of a federated network should include a monitoring service. We therefore recommend re-using the existing monitoring tools deployed in the GÉANT environment (i.e. GÉANT plus the NRENs) to provide a monitoring facility in a federation.

4.2.1.3 *Automated Bandwidth Provisioning Tools*

Automated provisioning tools remain the focus of research activities in Europe. As part of the GN2 project, two prototypes were developed and deployed in the GN2 test environment: AMPS and AutoBAHN. These tools allow end users to request end-to-end paths for their demanding applications. This section presents the multi-domain tools in the context of their use in a federated network in Europe.

AMPS

The Advanced Multi-domain Provisioning System (AMPS) [PAAMPS] was developed in GN2 to deal with Premium IP (PIP) requests in multi-domain environments. The Premium IP service is defined by certain QoS parameters, which are guaranteed for the traffic flow between service endpoints. The parameters include bandwidth, maximum one-way delay (latency), negligible packet loss, negligible packet reordering and guaranteed maximum IP delay variation (jitter). An authenticated and authorised user can reserve an end-to-end path across multiple federated domains [PAAMPS].

AMPS enables authorised users to reserve PIP capacity across the GÉANT2 network and adjoining PIP-capable domains. Each PIP-capable domain must deploy its own instance of AMPS. These instances operate in a federated (rather than hierarchical) manner; the federation of PIP-capable domains is called the extended AMPS domain.

Currently, the only administrative domains participating in the AMPS pilot are GÉANT and GRNET. The plans for future expansion are not clear. Currently, no other domain has expressed interest in deploying AMPS.

AutoBAHN

The Automated Bandwidth Allocation across Heterogeneous Networks (AutoBAHN) [AutoBAHN] activity was started in GN2 under the JRA3 activity. AutoBAHN is a service that enables automatic bandwidth provisioning across heterogeneous NREN domains. It supports multiple technologies (e.g. SDH and Ethernet, with GMPLS planned for GN3). The AutoBAHN architecture involves two main entities: Domain Manager (DM), which focuses on single-domain activities, and Inter-Domain Manager (IDM), which performs activities related to inter-domain aspects of bandwidth reservation. Both entities have to be deployed within a domain to allow bandwidth reservations.

AutoBAHN creates a new set of interfaces at the Inter-Domain Manager level to exchange information between involved administrative domains. It shares the topology information exchange (routing) and exposes a dedicated interface for triggering network actions (signalling). Information about the topology in particular domains is abstracted and made available to other domains for potential use (e.g. for multi-domain pathfinding).

The current work plan of the GN3 SA2 T2 activity assumes the AutoBAHN pilot starts March/April 2010 and will last for a couple of months, resulting in Operational Guidelines and Procedures to be developed by June 2010. It is expected that at least 5 NRENs will participate in the pilot with their own network infrastructure. AutoBAHN will also be one of the building blocks of the GN3 Bandwidth-on-Demand service, which is being defined within GN3 SA2 Task 1. The first rollout of the BoD service is expected in the middle of GN3 Y2.

4.2.1.4 *Information Systems*

I-SHARe

Information Sharing across Heterogeneous Administrative Regions (I-SHARe) [EMUC] was started in SA3 of the GN2 project. I-SHARe, previously known as version 2 of cNIS [CNISW], aims to support multiple domains (while cNIS builds a database for use within a single domain (NREN or GÉANT2)). The goal of the I-SHARe activity is to design and develop an information and workflow-tracking system to support the operations teams within the multi-domain environment of the GÉANT project [EMUC].

In general, I-SHARe focuses on multi-domain services, collecting and formalising the following operational procedures:

- Sharing contact information.
- Ordering an E2E link.
- Setting up an E2E link.
- Managing E2E link faults.
- Managing E2E link maintenance.
- Change management for an E2E link.
- Troubleshooting unresolved E2E link faults.
- Escalation.
- Decommissioning an E2E link.

A detailed description of these procedures can be found in [EMUC].

I-SHARe uses a central server to collect information about the different domains. Each domain can provide information about their links through the I-SHARe Domain Interface. Currently, a fully working prototype of I-SHARe is being deployed in a number of participating NRENs and DANTE [GCPST].

c-NIS

The Common Network Information Service (cNIS) is a unified repository of all relevant network information about a single administrative domain. cNIS has been developed (together with AMPS, I-SHARe, AutoBAHN) as a component of the unified platform of services supporting multi-domain workflows in the GÉANT network and the NRENs. Within GN3, many of the applications and services developed in GN2 will be able to use cNIS in place of their internal topology information storage systems [CNISW].

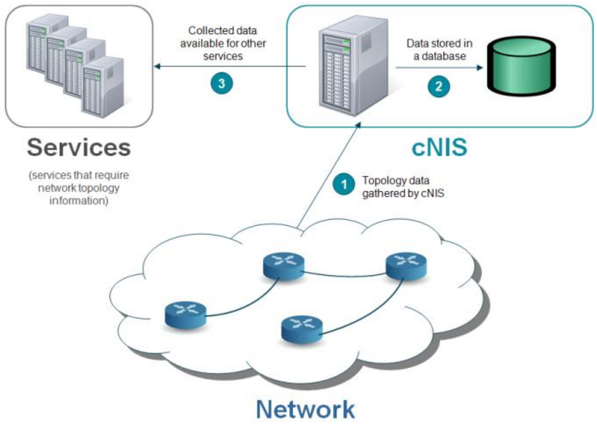
Collected data available for other services
3
Data stored in a database
2
Services
(services that require network topology information)
cNIS
1
Topology data gathered by cNIS
Network

Figure 4.7: cNIS overview [CNISW]

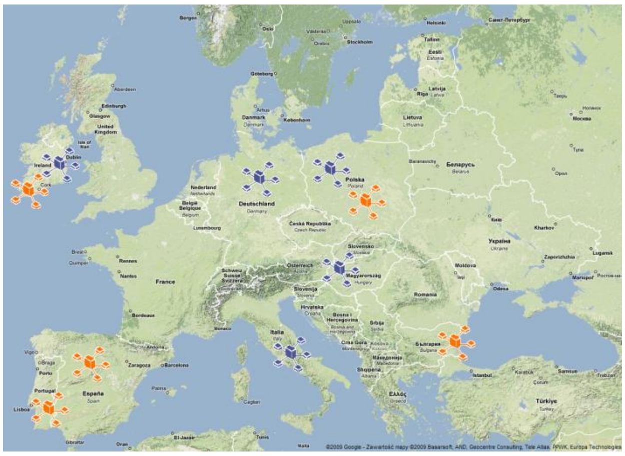

Figure 4.8: cNIS deployments [CNISW]

cNIS is a "single point of storage", but in fact it is more than just a database. Apart from the internal functionality required for populating, validating and updating the database, it is equipped with modules for analysing the topology data and presenting the data in a client-specified format. An example is the Path Finder module, which is able to calculate the routed path across a domain, and the Topology Service, which presents the topology data in the XML format specified by the Global Grid Forum's (GGF's) Network Measurement Working Group. cNIS is able to store topology data not just for a domain's IP infrastructure, but also for other network technologies (Ethernet, SDH). Automatic population is a significant feature of cNIS, since it simplifies the work of network administrators (NRENs and others). As a result, cNIS delivers a mechanism for the automatic discovery of IP, Ethernet and SDH technologies [CNISW].

Since the first release of cNIS, the service has been deployed in selected European NRENs. At the end of GN2, four test instances were running. At the beginning of GN3, five operational instances were launched, in BREN, HEAnet, PIONIER, FCCN and RedIRIS [CNISW].

4.2.1.5 *Tools and Services Summary*

Table 4.1 presents a summary of the most important features of the existing GÉANT tools and services, taking into account their readiness to support future federated environments in Europe.

Type	GN2/GN3 Tool/Service	Main features
Performance troubleshooting	eduPERT	• eduPERT is a **service** offered to end users, not a tool • (partly) **decentralised** • a central database • **multi-domain** issues processing supported • a **tree-like** administrative structure
A measurement support service	perfSONAR	• **multi-domain** monitoring supported • **decentralised** database • unified formats of data and interfaces to various tools
Automated bandwidth provisioning	AMPS	• **multi-domain** QoS provisioning (Layer 3) • **decentralised** system • **transactional** system
	AutoBAHN	• **multi-domain** bandwidth provisioning (Layer 1-2) • **decentralised** system • topology abstraction • **inter-domain pathfinding**
Network Information Systems	I-SHARe	• **multi-domain** information processing • **central** server/database • unified operational procedures
	cNIS	• supports **single administrative domains** only • supports **automated population** of database with information about networks

Table 4.1: Summary of operational tool and service features important to federated networks

4.2.2 Administration and Procedures

Currently, establishing E2E paths for projects requires manual management of incoming requests for establishing connectivity, manual configuration of the devices used for providing the connection, and manual provisioning of resources. For example, to establish an E2E link it is necessary to exchange a lot of information describing technical details (hardware and interfaces) and contact details (locations, full addresses, phone numbers, email addresses, contact personnel). Usually, therefore, establishing E2E paths is a costly process in terms of time and manpower.

Work in GN2 to facilitate E2E path creation and management is being continued in GN3. As already described in this section, this was the aim of the I-SHARe initiative, begun in GN2 SA3. The main goal of I-SHARe was to define common interfaces for information exchange across network domains within GN2 to help the NOC engineers to manage E2E paths. For a federated network, such tools for information sharing between domains are crucial. Service provisioning in a federated network is possible only with shared workflows across the domain and across the partner organisations.

As mentioned in Section 2 *Definition of Terms* on page 5, a federation may be loosely coupled or tightly coupled. In a loosely coupled federation, each federation partner can handle operations and service provisioning autonomously. All that is needed for a loosely coupled federation is agreement on standard interfaces for traffic and information exchange. However, in a tightly coupled federation, the operations of the federation partners' networks become directly interdependent. It is therefore necessary to have an agreed set of known procedures and workflows, and to back these with a set of service level agreements (SLAs).

Of crucial importance is the collaboration between the Network Operations Centres (NOCs) of the partners in the federation. The NOCs must agree on procedures for incident recovery, on who is responsible for creating fault tickets, who is responsible for taking action on tickets, for following up, etc. As errors may affect more than one network domain, mechanisms for passing responsibility must exist. Procedures must be detailed, clear, accurate, predictable, and efficient, and must be well-understood by all partners. Likewise, agreed procedures must be followed for information distribution, ticket forwarding, etc.

For workflows and procedures to be well known, agreed by all partners, and efficient, they must be described in an operations model that details the interaction between the partners, the data sources used, and how data is updated. The operations model should bring together the actions and procedures of all partners. An example of such an operations model is the one used for the LHCOPN [LHCOPNOPS].

In addition to an operations model, a federation will need intra-federation service level agreements. A tightly coupled federation will be able to offer SLAs, covering, for example, traffic capacities, incident reaction and recovery times, etc., to end users. Such external SLAs can only be honoured if partners enter into intra-federation SLAs which guarantee the services each partner provides to the federation. These SLAs are similar to those required from sub-contractors in order to deliver end-to-end services in a non-federated network.

In a traditional network, with a central organisation owning and operating the network resources and offering the services, relations between the network owner and users, and between the network owner and suppliers, can be based on contractual and financial agreements. Procedures for handling these relations are well known. In a federated network, both user-network and supplier-network relations will become relations between partners. Such intra-federation relations are often political in nature. In order for intra-federation relations to be manageable, a governance structure must be agreed. A federation or consortium agreement must be defined, and steering bodies, chains of command, and mechanisms for resolving disputes must be established. The governance structure should be both efficient and fair; it may, for example, assign voting rights and agree cost-sharing principles, as well as defined the overall strategy and principles for the federation.

4.3 Services

The goal of a federated network is to provide a range of network services to end users. Resources owned and operated by the federation partners are used to create, provision, and operate end-user network services. For some services, this can be done by simply combining the federation partners' network services directly, while in other situations federated network services must be built as a layer on top of the core network components provided by the federation partners.

4.3.1 Types of Services

The essential network services that are required from a federated network are end-to-end circuit services, shared IP transport, and network virtualisation services.

- End-to-end circuit services can be provided on a number of network layers. Typical services requested include end-to-end wavelength services, end-to-end transport services (based on SONET, OTN or Ethernet, for example), and end-to-end connectivity across a shared IP network (based on MPLS, for example). In all cases, the objective is to connect end-user sites directly, using a circuit with known properties and QoS parameters.
- From an end-user perspective, shared IP transport is usually provided as an inclusive Internet connectivity. However, for federation partners, shared IP services can include peering, mutual IP transit, joint upstream access, etc.
- Finally, network virtualisation services can again be provided on several network layers. Virtualisation services can include virtualisation of basic network building blocks (links, switches) in an infrastructure-as-a-service framework such as UCLP; it can mean virtualisation of IP network capabilities (virtual routers, virtual IP networks) such as in MANTICORE; or it can mean simple VPN-type services across shared IP networks.

It is important to note that in a federated approach to networks, network services have a dual role: they are both services for end users (e.g. an end-to-end SDH circuit connecting an instrument to a computational site) and building blocks provided by federation partners to the federated network (e.g. an SDH link between two PoPs provided by a federation partner). This is true for both simple network links and for network virtualisation services.

4.3.2 Building Services in a Federated Network

Services in a federated network can in general be built either by direct aggregation of similar services in the partners' networks, or by using resources provided by partners as building blocks to construct a core network on top of which network services are offered.

Direct aggregation is the approach often taken for shared IP services, where peering and IP transit are used to aggregate the IP services of partner networks into a coherent IP service for end users and to share access to external parties. Likewise, end-to-end services can be provisioned by stitching together the end-to-end services of individual partners, providing end-to-end services by composing the services of individual networks as is done in the GLIF collaboration and by European NRENs in combination with GÉANT end-to-end services, for example. However, for such aggregation to be successful, interfaces must be compatible and the control planes used by partners must inter-operate.

Creating a core network from partners' resources allows the federated network to offer a wider range of services, and makes it better able to offer end-to-end guarantees. For example, federation partners might provide transport-level network services to the federation; the federation can then use these services to build and operate an inter-federation transport transit network (or a federation-wide transport network), offering a range of network services, including end-to-end services, shared IP services, and virtualisation.

The choice of how services are offered leads to quite different network designs for the federation. In a core network, operated as a single entity, service provision is no different from that of a traditionally designed network. With this approach, the federated network architecture is not directly visible at the service-delivery level. With the aggregated service approach, on the other hand, the service delivery itself becomes federated, offering services made of composable network elements. This choice is therefore fundamental to the network design.

5 Federated Network Architecture Models

The main purpose of this section is to propose generic models for the federated network architecture, developed as a result of analysing existing projects, processes and services that are either potential users of a federated network or that feature some elements of federation. It defines the structure, relationships and dependencies between particular elements of the architecture. Two models are proposed and are elaborated in detail. Both models are based on the assumption that the federated network architecture is composed of three main layers. Inter-communication procedures between these three layers, supported by internal procedures within each layer (intra-communication procedures), make the whole federation possible. Each model reflects one of the two ways communication is performed within and between layers.

Although some references are made to particular existing tools and services, the architecture models are very generic. The practical implementation of these models in the GÉANT/NREN environment will be the subject of study in Y2 of the GN3 project and is introduced briefly in Section 6 *Future Work and Recommendations for Test Cases* on page 50. Furthermore, as part of the future work, the architecture models can be placed in context with other models from the literature as well as with models used or proposed by other GN3 service and joint research activities.

In order for a federated network to work, it is necessary to have a common understanding of the general collaboration within the federation. A brief overview of the agreements necessary for participating in a federation is therefore given before the specifications of the federated network architecture models.

5.1 Agreement for Participation in a Federation

A federated network is about sharing existing resources that are owned and maintained by multiple domains. In order to be able to do that efficiently, reasonably and for a long time, all participants should commit to the same rules and regulations. It is obvious that different users will have different needs;. the rules should therefore be such that they support a wide range of needs, i.e. not be too strict or rigid. However, it is important to define basic rules, common for all situations and every participant, and also to specify as many variations as necessary to accommodate different uses and users.

The basic agreement should first define the constitution of the federation, its purpose, members, federation services, and resources. Conversely, it should also define who or what cannot be part of the federated network, and what services are not part of the federation. In addition, the agreement will have to define financial issues such as the costs of providing and using the federation and its services, and financial dependencies among participating members.

It should then define how to become part of the federation (how to join), how to "behave" within the federation (what are the rules, roles and procedures within the federation), and how to leave the federation. Also, mandatory and optional requirements should be clearly distinguished.

By signing such an agreement, participants confirm their commitment and accept the mandatory requirements. Later on, less strict, non-mandatory options can be chosen for each particular situation, user or service.

5.2 Generic Federated Network Architecture Model

The architecture model for a federated network consists of three main layers: Infrastructure, Operations and Service. Inside each of these layers there are different elements, with connections between them. Figure 5.1 shows the generic model for a federated network architecture.

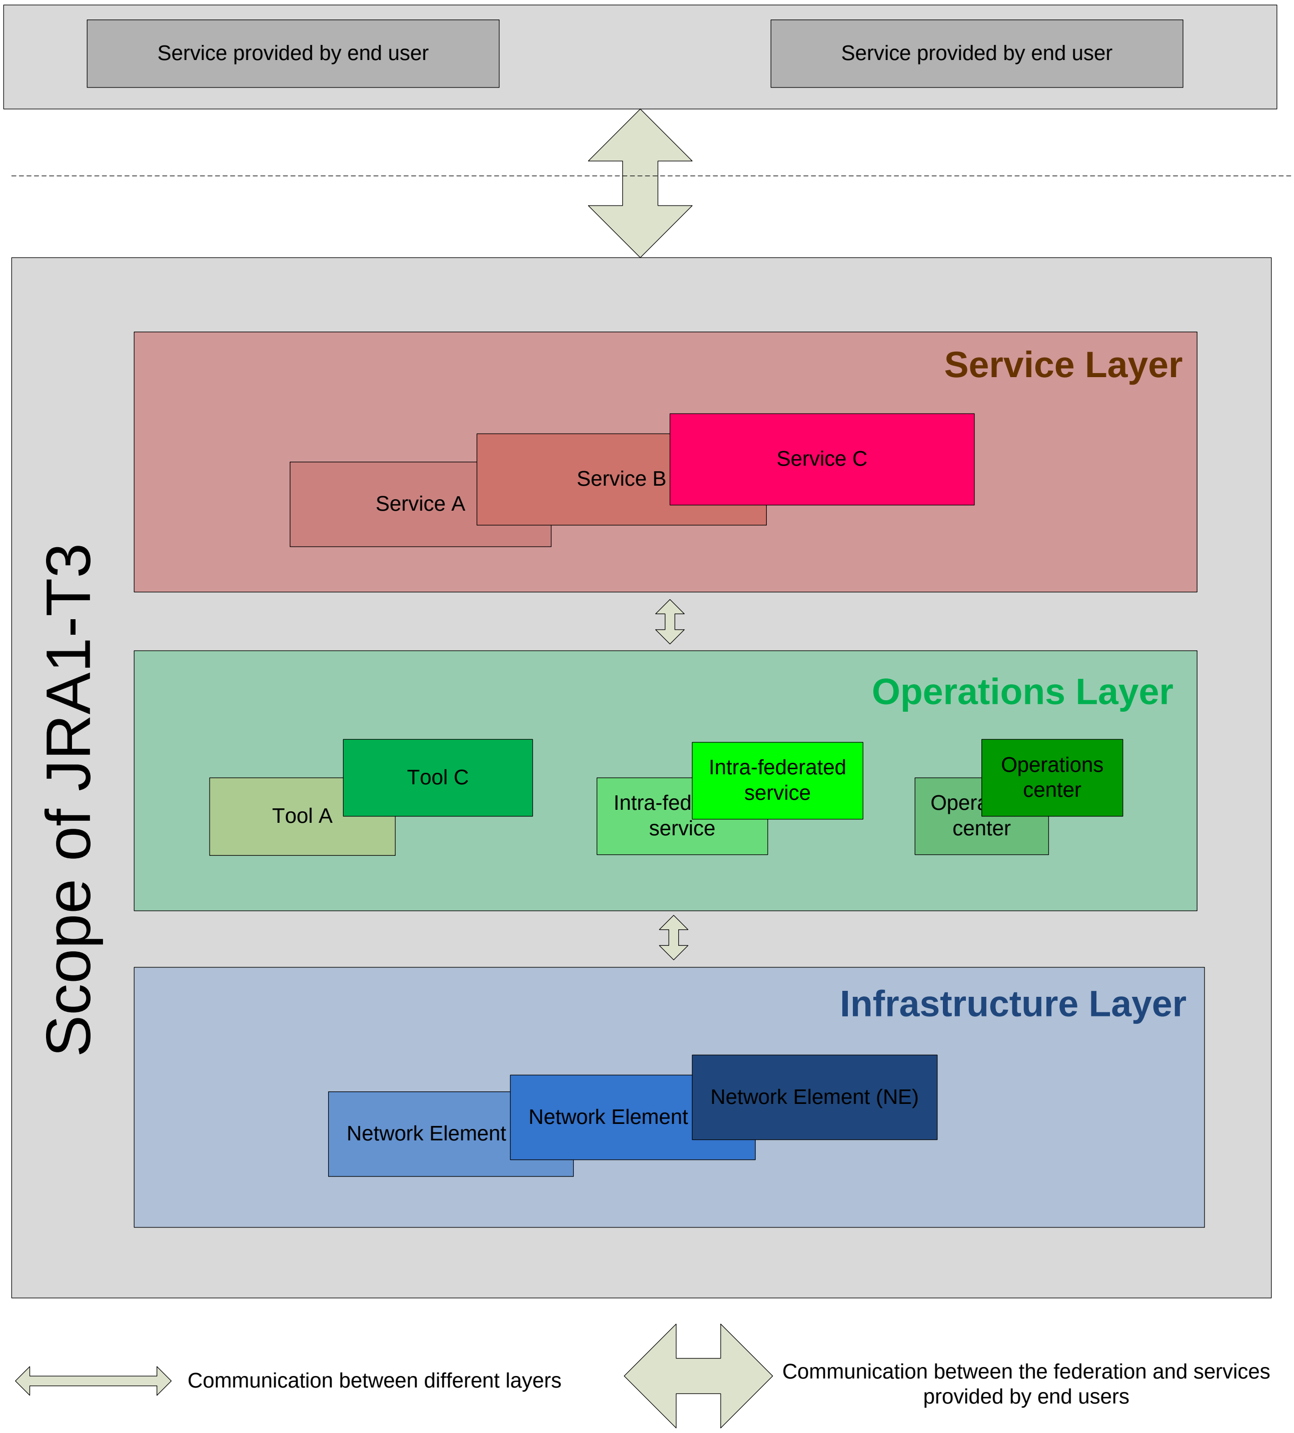

Figure 5.1: Generic federated network architecture model

The Infrastructure Layer, situated at the bottom of the architecture, includes the network infrastructure elements, essential to support end users and their needs. The Infrastructure Layer consists of several elements; all of them together compose the first layer of the architecture model. In federations with many participants, the individual elements of the Infrastructure Layer may differ significantly. This should be taken into account when a decision about new interconnections is being made to create a stable, sustainable, solid and scalable base for the federation.

For example, GÉANT interconnects a number of NRENs to form a unique European research and education network. Individual NRENs have their own specifics regarding topology, network coverage, equipment, technologies at lower layers, etc. However, the same services can be provided for the whole GÉANT community. So, in the case of GÉANT, the Infrastructure Layer has NRENs as individual elements that are interconnected and whose differences should be taken into account in order to build and operate the Infrastructure Layer of a federated network architecture model.

In Figure 5.2 (page 43) and Figure 5.3 (page 46), the linking lines between elements within the Infrastructure Layer (represented also as the “Internal communication” box) might represent physical connections (e.g. cross-border fibres, GÉANT links) or might represent some logical relationship between elements, like intensive collaboration (e.g. NRENs that have heavy traffic exchange for different projects, so they need a high-capacity, high-availability connection between them).

The middle layer of the model is called the Operations Layer. This layer includes tools and services that are needed to provide support for services that are offered to end users. Since these are the services needed to make the federation work, and, as such, should not be visible to users outside the federation, they are also referred to as intra-federated services. This layer binds together the Infrastructure Layer and the Service Layer. The building blocks of this layer are tools (such as perfSONAR or AutoBAHN from the GÉANT perspective), operations centres and intra-federated services (such as eduPERT in GÉANT). It is expected that different tools within the layer will be interconnected (or, in some cases, integrated) to enable or enhance operations; for example, operations centres may rely on some tools or intra-federated services. Some services might exist more or less independently within the layer, being independent in the nature of their work, but needed to support the overall function of the layer. In Figure 5.2 and Figure 5.3, the interaction between the components of this layer is shown through the Internal communication box. It is not appropriate to show the significance and meaning of each link on the overall architectural blueprint because it is implementation-specific. The common property for each link is that it defines a procedure or process between two elements, relating to the operation of those elements. However, the exact meaning for each link depends on the specific relationship between the actual two elements that it connects. The subsequent detailed practical implementation of the model in a specific environment will address the question of the importance of some interconnections, relative to others in the same layer.

The uppermost layer is called the Service Layer. Services that are offered to end users will be identified and detailed at this layer. Examples of these services are bandwidth-on-demand services, static wavelength-based connectivity services, or a service like eduroam. Steps must be taken to ensure that services provided by the end users are capable of interacting with services offered by the federation. There are also certain services that are established, maintained and operated by end users, which use the federated infrastructure. Examples of such services can be those provided by any of the large-scale projects described in Section 3.3. Although the provisioning, management or operation of those services might not be in the scope of the federation itself, the services will use the federated network in a legitimate way, so steps must be taken to ensure that the services provided by the end users are capable of interacting with those offered by the federation.

Following the principle of layered architecture and inter-layer communication, the Service Layer must communicate only with the layers it is attached to. In this case, depending on the model used, one or more interconnections may exist between the Service and Operations Layers. The communication may be realised either via the management component of each layer (Model A) or between individual elements of both layers (Model B). Elements of the Service Layer are not permitted to interact directly with the Infrastructure Layer. This must be realised through elements of the Operations Layer.

In Figure 5.1. communication between two neighbouring layers is presented by a bi-directional arrow. The two models described in the following sections differ primarily in the way communication is performed between the layers and elements. The generic model therefore presents only the main layers and elements within the layers; the communication elements – the source of the models' differences – are presented in detail in the respective variant models, in Figure 5.2 and Figure 5.3. From that point of view, the generic model presented in Figure 5.1 does not contain all the elements that the two model variants contain (e.g. the Layer Management box and the Internal communication box).

For a federated network architecture model it is important to define exact operational procedures for each building block within each architecture layer, and to define who will actually perform the individual tasks specified for that interface.

The practical implementation of a federation should take into account how the existing building blocks of the federation have been prepared. A distinction should be made between elements that, at least partially, support the federation, and elements that are completely unaware of any aspects of federation. With regard to the first option, all federation-related tasks will be performed as a standard activity of these elements. Communication between the elements will be realised in the form of direct communication. The element itself will be capable of handling the whole procedure to support a federation. The second option assumes that communication between non-federation-capable elements will be realised through an adapter element. In practice the adapter may take the form of a person appointed to perform all tasks related to a federation (manual or semi-automatic operation) or of a separate component that allows bilateral communication between independent components of the architecture that do not inherently support a federation (automatic operation).

5.3 Federated Network Architecture Model A

In the federated network architecture Model A, shown in Figure 5.2, there are connections between the different building blocks within the same layer. In order to simplify the figure, an additional "Internal communication" box has been introduced, which reflects these intra-layer connections. In addition, there is a single inter-connection between neighbouring layers of the model. The model assumes that direct connection between the Infrastructure Layer and Service Layer is not allowed. All communication between these layers must be realised through the Operations Layer. Communication between neighbouring layers is carried out through a special element responsible for the management of the layer (Service Layer Management, Operations Layer Management and Infrastructure Layer Management elements). Model A is simpler and more restricted than Model B because it has fewer communication channels: communication between layers has to be performed through the Layer Managements boxes of neighbouring layers; communication between individual elements from neighbouring layers is not allowed.

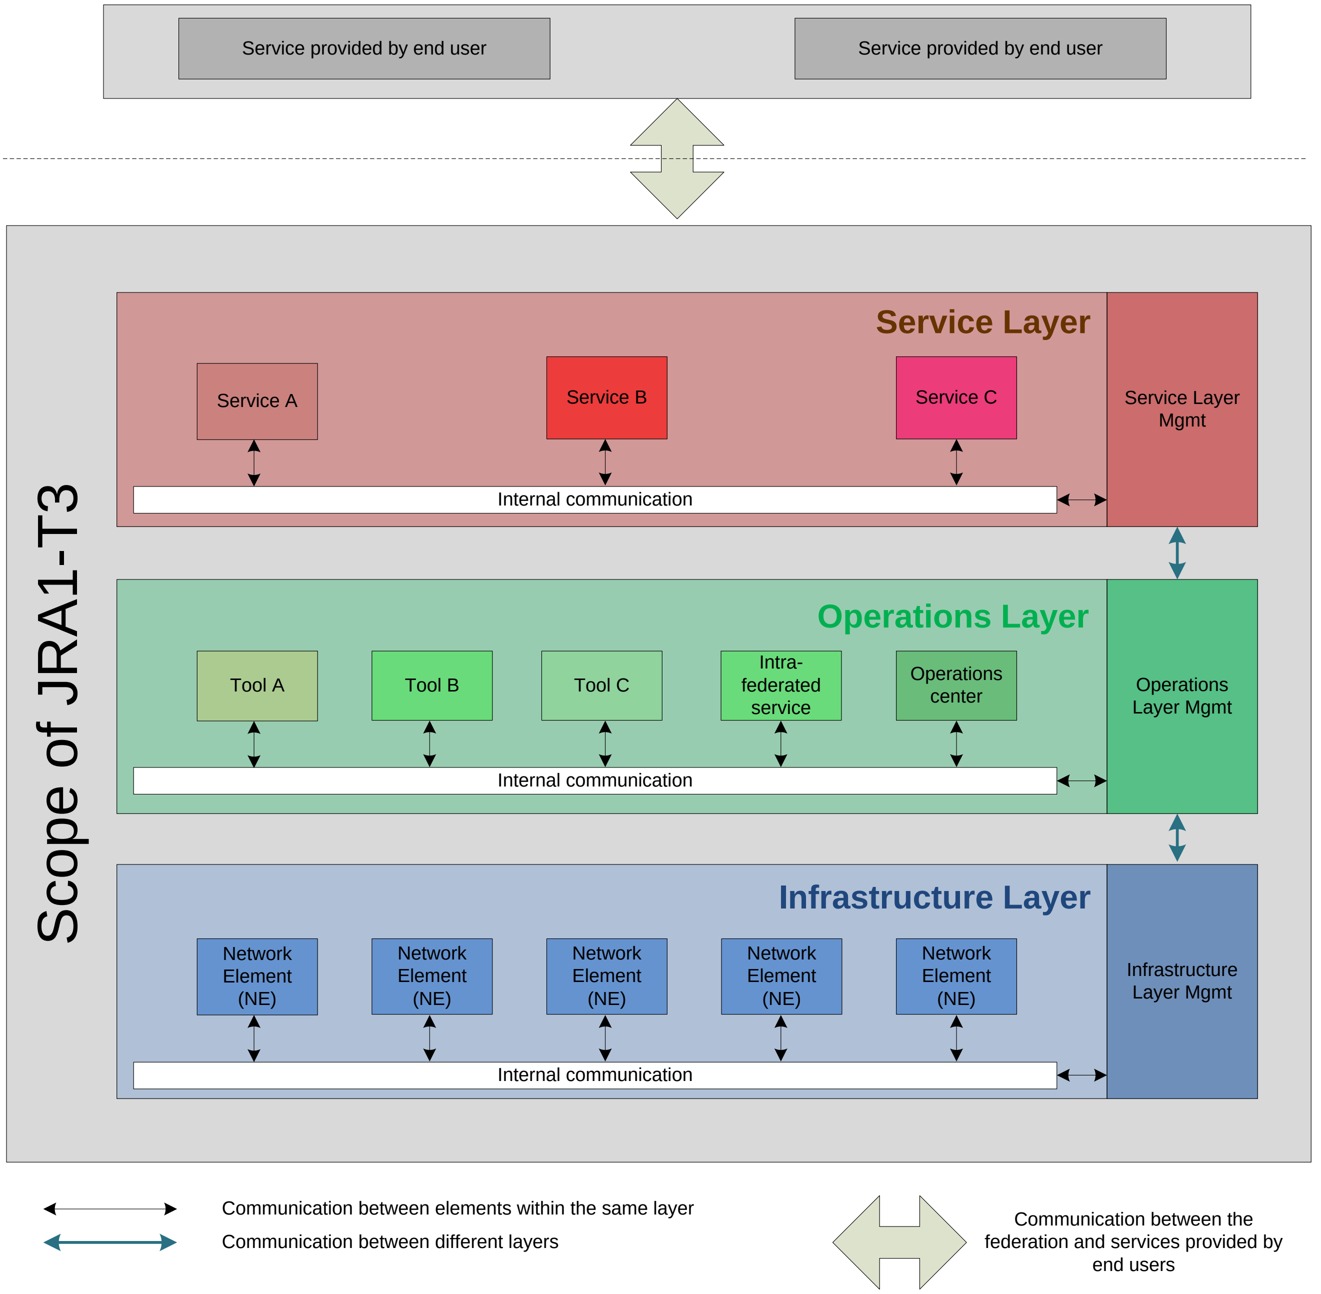

Figure 5.2: Federated network architecture model – Model A

All links between building blocks in each layer, as well as links between federated network architecture layers, are supported by operational procedures that describe all possible interactions between the elements of the architecture. The procedures include communication channels in the control plane and data-exchange in the data plane.

5.3.1 Federated Network Topology

In Figure 5.2, the Infrastructure Layer consists of a number of network elements (NEs) that are connected to other elements within the layer by one link (represented as a link to the Internal communication box). This link represents one relationship defined with adequate procedures.

For example, the relationship represented by the link might be network connections between neighbouring NRENs. The operational procedures would define establishment, monitoring and maintenance of those infrastructure links.

5.3.2 Operations Topology

In Model A, the Operations Layer is connected to the Infrastructure Layer, and also to the Service Layer. All requirements or responses from users for the infrastructure, and vice versa, go through the Operations Layer. The tools and services within this layer support the operations of the federation and federated services, as well as fulfilling requirements from the layers above and below the Operations Layer.

Procedures have to be defined between the building blocks in order to perform the operations. Also, the correlation between elements within this layer should be clearly defined, e.g. which operations centre uses which tools and intra-federated services, which intra-federated services use which tools and other intra-federated services, etc.

5.3.2.1 *Tools*

It is very important to identify the tools that are needed to perform all the necessary operations, then to identify the tools that already exist and can be re-used, those that are missing and how they can be provided. Also, it is important to draw a usage map for each of the tools, showing which services will need which tool, and how it will be used. The usage map is closely related to the operational procedures for each building block.

In Model A, tools can be used by any other building block in the same layer. If such a relationship exists, it is represented by an arrow. In real-life examples, arrows can be categorised in order to distinguish dependencies between building blocks, e.g. to show who is a provider of the tool (and responsible for its usability and availability), and who is a user of the tool. It is also possible to have two-way relationships (represented by bi-directional arrows), for example when a building block uses the tool, but also provides data or other input for the tool.

5.3.2.2 *Administration and Procedures*

Operational procedures have to be defined for each connection between building blocks within the layer, represented in the model with arrows. It is important to define rules, roles and policies and to ensure that they are accepted and well known by all the participants involved. Relationships can be represented by a one- or two-headed arrow, which indicates the direction of communication, data flow or sequence of execution. The correct meaning of the line should be defined in procedures, and depends on the specific element and neighbouring blocks involved.

5.3.3 Service Portfolio

The elements within the Service Layer are the services provided by the federation. In Model A, relationships exist between different services within this layer; each of these should be supported with well-defined procedures. All requirements for lower layers are propagated through the Operations Layer.

5.3.4 Pros and Cons

Model A is very simple, with very few connections between building elements and a small number of communication channels. All communication between layers or elements has to go through pre-defined channels; alternative methods are not allowed. On the one hand, such a strict communication procedure brings more control because the path for the information (and data) flow for each communication is defined in advance. On the other hand, communication channels must be of high capacity with enough resources to support the communication. All participants should know the communication procedures and paths, and follow them without an exception. The area of responsibility of each building block is very well defined.

However, it is possible that each element of the architecture layer will be very complex, in order to be able to meet all the requirements and challenges, especially the part of the element that performs the communication with another element. It is also possible that processing will not be as fast as needed.

Model A’s scalability depends on the optimisation of processes within an element and procedures for the elements' communication. It is also important to predict and provide enough resources, again especially for those parts that perform common services for multiple elements.

5.4 Federated Network Architecture Model B

In Model B, connections and relationships exist between elements within a layer, but also between elements in different architecture layers. There might be more than one connection between two elements. As with Model A, this model assumes that a direct connection between the Infrastructure Layer and Service Layer is not allowed. All communication between these layers must be realised through the Operations Layer. Communication between neighbouring layers is carried out either through a special element responsible for the management of the layer (Service Layer Management, Operations Layer Management and Infrastructure Layer Management elements) or directly between elements located in the neighbouring layers. Model B is shown in Figure 5.3.

Model B is much more complex than Model A. Multiple relationships between building blocks demand well-defined processes and procedures in order to avoid extra overhead. However, such a complex model might be able to embrace more services and be more flexible for future applications, projects and services.

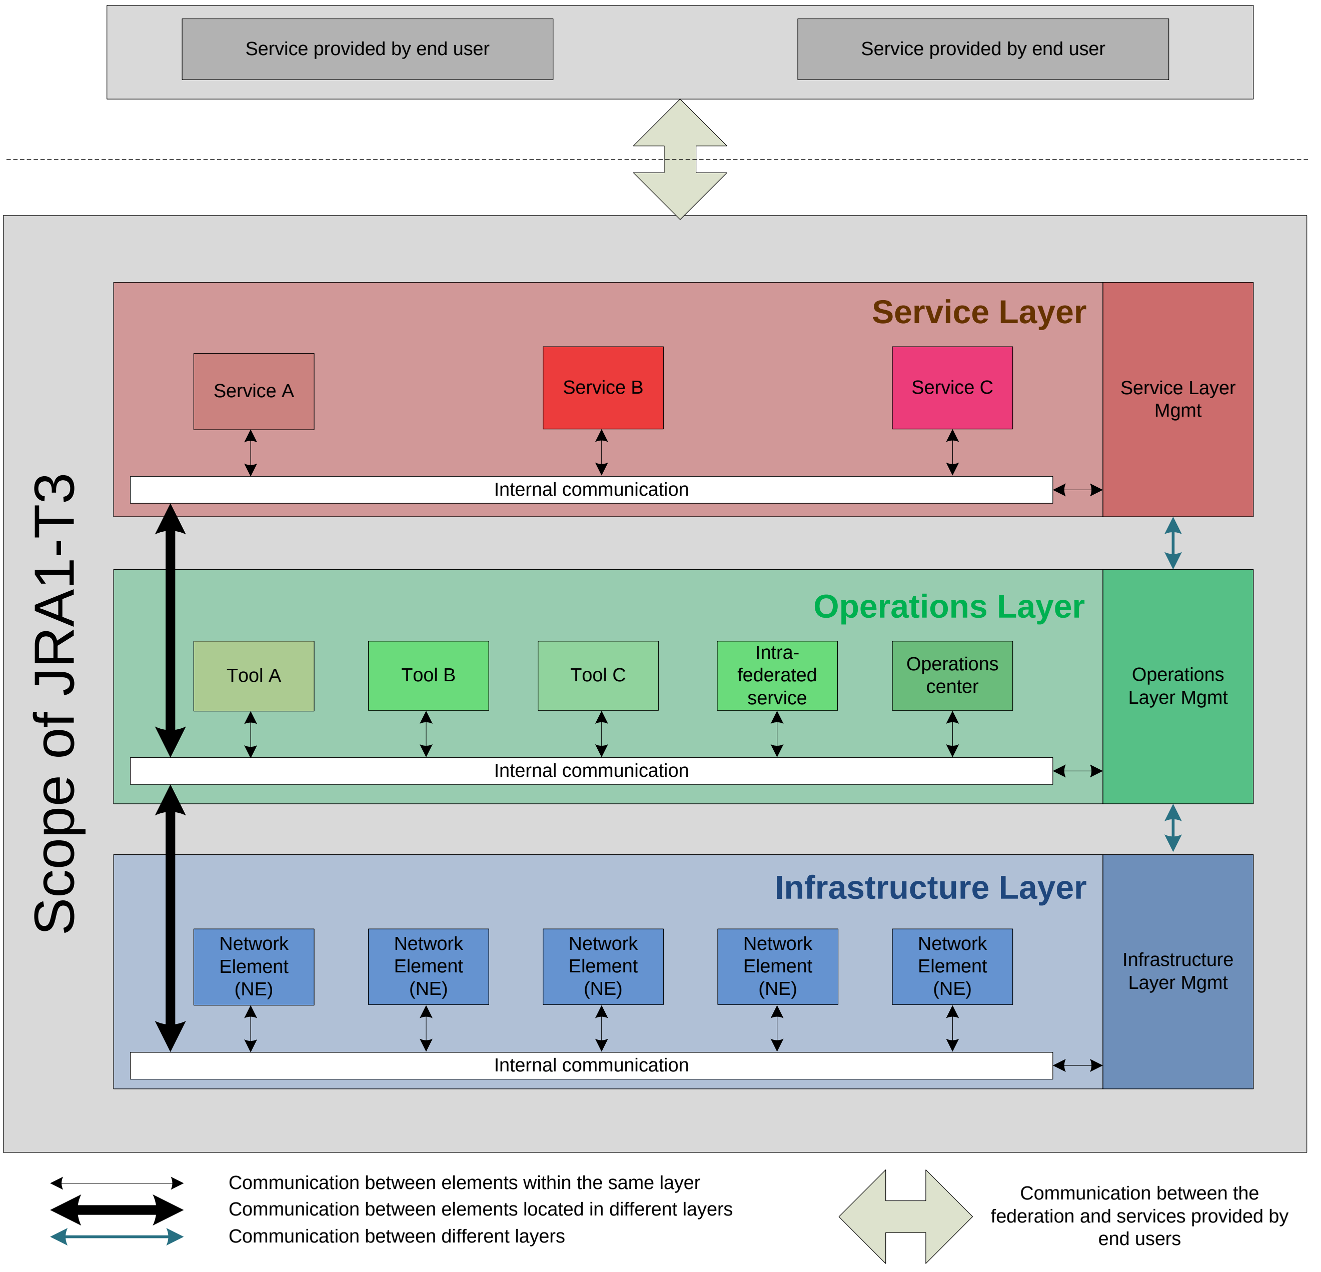

Figure 5.3: Federated network architecture model – Model B

5.4.1 Federated Network Topology

In Model B, each element in the Infrastructure Layer might have more than one connection to other elements within the layer. In the diagram, this can be represented with more than one line between two elements, or with a thicker line. Connections can be of one kind, for example, a physical link between two countries. Alternatively, connections can have multiple meanings, including physical connectivity, intensity of traffic flow, dependencies in important projects, and so on.

Links from the Operations Layer to the elements in other layers (represented in Figure 5.3 as links between the Internal communication boxes of each layer) represent processes that are performed between those specific elements. For example, one service from the Service Layer can use a specific tool from the Operations Layer through a direct communication between those elements. The processes can be performed in such a way that no other elements in any of the three architecture layers have any knowledge or awareness of them. In Figure 5.3, it is not explicitly shown, but is only implicitly represented with a path from an element within the Operations Layer and a service in the Services Layer (passing also through the 'Internal communication' boxes in each of these layers and linking the layers).

A third type of connection that might originate from an element in the Infrastructure Layer is a connection to the Infrastructure Layer Management element. Where there is a virtual organisation that is coordinating the activities of the federation, such connections represent the relationship between the element and the virtual organisation. In such circumstances, all elements should have such a connection. However, it is possible that some links will be stronger than others. It is also possible that some elements will use other elements as a proxy, in which case they will not have a connection to the Infrastructure Layer Management itself.

5.4.2 Operations Topology

As in Model A, the Operations Layer includes tools and intra-federated services. Relationships may exist between the elements, defined with appropriate processes and procedures, and represented in Figure 5.3 by an arrow to the Internal communication box. A bi-directional arrow means that the relationship is two-way. A uni-directional arrow means that it is one-way, and indicates the direction of the communication (or data flow).

In Model B, it is possible to have more than one connection between two elements, meaning that there is more than one process or relationship between those two elements. As in the Infrastructure Layer, a connection from an element to the Layer Management entity implies a dependency on a virtual organisation. That organisation might be the same as for the Infrastructure Layer, or it might not.

Tools and services from this layer can be used by elements in the other two architecture layers. Communication can go directly from one element in the Operations Layer to another element in either the Infrastructure or Service Layer. This communication involves the Internal communication box in each of those layers as well as communication between elements in different layers. Tools and services in the Operations Layer can serve as a mediator between the Infrastructure and Service Layers. Communication can also be performed through the Layer Management entity (Infrastructure Layer Management, Operations Layer Management and Service Layer Management).

5.4.2.1 *Tools*

Tools in the Operations Layer imply the programs and systems needed to perform the set of tasks of the Operations Layer, either programmed to perform the tasks themselves, or used by operations centres or intra-federated services. They can also be used by elements in the other two layers.

Within the tool set in the Operations Layer some tools can have overlapping features with other tools, at the same time having unique functionalities that are important for the overall operation, which make them an important part of the set. It is important to optimise the tool set to provide the necessary and sufficient number of tools that cover the required functionalities.

5.4.2.2 *Administration and Procedures*

Being the middle layer of the federation model that performs all the inter-layer communication, the Operations Layer has procedures that are specific to this layer and also procedures that describe the interfaces to the Infrastructure and Service Layers. Therefore, definition of the communication procedures of the Operations Layer requires the involvement of elements in the other two layers, and possibly involvement of multiple elements within the Operations Layer. From this perspective, the administrative and operational procedures of this layer might be more demanding to define (not necessarily to perform) than those in the other two layers.

Elements that are participating in the federation and have relationships with neighbouring elements share some administrative and operational procedures. Those procedures are a part of their everyday work included in process descriptions and workflows. They should be defined by both participating elements, optimised to use as little resource as possible, and well-known to workers in the participating elements. The same applies to any pair of elements no matter which layer they are in, and whether they all belong to the same layer, or they cross different layers. If procedures and processes are not well-defined and well-known to the participating entities, the overall functionality of the federation might be sub-optimal or even put at risk.

Since a federation is a changing structure, it is possible to add or remove elements within any layer. For each addition, all aspects of communication should be analysed in order to define the right place for the new element within the architecture. Also, operational and administrative procedures should be reviewed with any change, and redefined if necessary.

5.4.3 Service Portfolio

As in Model A, the elements within the Service Layer are the services provided by the federation. In Model B, relationships exist between different services within this layer. The difference between Models A and B is in the links originating from each element in the Service Layer. Links can exist between elements within the Service Layer as well as between an element in the Service Layer and an element in the Operations Layer. All requirements for the lower Infrastructure Layer are propagated through the Operations Layer.

There can be more than one link between two elements, regardless of which layer they are in. Each link represents a process that exists between two elements.

Federation Model B should also enable the coexistence of services provided by an end user that have to use the federation, such as those mentioned in Section 3.3. Although provisioning of such services is not in the scope of the federation, the definition of the relationship between any such service and an element of the federation should be within the scope of the federation.

5.4.4 Pros and Cons

Model B is much more complex than Model A. There are two main differences: Model B has direct connections between elements located in neighbouring layers and the number of links between elements is not strictly defined. This means that the federation can offer more functionality and services than Model A. It can be assumed that, if well organised, Model B can be more scalable that Model A.

The disadvantage is that, if operational processes are not well defined, there is a possible danger of much more overhead and less robustness than in Model A. There is also a risk of process or operations duplication, limited awareness of similar processes and functions, etc. In order to avoid this, it would be necessary to analyse and optimise all processes.

6 Future Work and Recommendations for Test Cases

The federated network architecture model was presented in the previous section, together with its two variants: the simple, sequential Model A and the more complex Model B. However, these two models appear abstract if not mapped to existing multi-domain structures that are candidates for federation. One approach for future work is to choose one (or more) existing collaborations and analyse their building blocks and relationships from the perspective of the proposed federation model.

A second approach for future work is to analyse the implications of making a change in the federation, and which model is better suited to the new situation. For example, what would be the implications for the federation of adding an element to any of the architecture layers? How would the federation behave if the number of elements was reduced? How would an architecture layer perform if one element were excluded, or joined with another in the same layer, and so on?

In addition, the federated model presented here should be compared with the methods of federating systems services and data defined by the TM Forum: Service Level Federation (SLF) and Repository Level Federation (RLF).

6.1 Analytical Verification and Testing

In order to perform verification and testing of the federated network architecture model, it is necessary to identify and list all the elements and links within the scope of the test case. After that, all relevant processes for those elements and links must be identified, named and analysed from the perspective of communication, data exchange, and data flow. Those analyses will determine how the links are presented in the model, e.g. the size and thickness of the connecting lines between any two elements, the number and direction of arrows on the line, and the type of link-end.

Where an element is to be added or removed, the existing links should be thoroughly reviewed in order to determine which ones should be changed, added or deleted.

6.2 Federated Network Designs for GÉANT

To start the verification process, the GÉANT network and services can be used as a test case for the model(s). NRENs including CARNet, PIONIER, SURFnet, DFN, RedIRIS and NORDUnet will be elements in the Infrastructure Layer; tools such as perfSONAR and services like eduPERT will form the Operations Layer; services like eduroam and eduGAIN, and projects like EGEE, LHCOPN and DEISA 2 will be the elements of the Service Layer. Links should be analysed for all elements to identify and describe the existing processes among them. After that, the federation should be examined to determine which model suits it better – Model A or Model B. For each architecture layer, it should be established whether there are any missing processes and/or services that should be defined, or whether some activities be optimised to improve the overall organisation of the federation.

It is proposed that several test cases be carried out involving making changes to the existing structure and analysing the implications of the change for the behaviour of the federation. The objective of these test cases is to demonstrate and verify the federated network models by applying them to a number of practical use cases within the context of GN3. We have selected three use cases that a) have different levels of complexity, and b) demonstrate different aspects of federated networking.

Examples of test cases are:

1. *Using CBF for GÉANT PoP-to-PoP connection.* For GÉANT PoPs in neighbouring countries, it might be possible to replace dedicated GÉANT capacity (dark fibre or leased line) with resources owned and operated by NRENs. The aim is to reduce cost. A potential candidate would be to analyze the Poznan-Frankfurt/Main and Poznan- Prague link.
2. *Using CBF as an element in regional links*. NREN-owned and -operated network resources might be used for (part of) long-haul connections between GÉANT PoPs, or for connections from GÉANT PoPs to local NRENs. The aim is to reduce cost. A candidate is connectivity to the DANTE PoPs in Russia and the Baltic states using NORDUnet connectivity from Copenhagen to Helsinki and CBF to St Petersburg.
3. *Federated GÉANT PoP*. Traditionally, there is a GÉANT PoP in each country that has a connected NREN, with the PoP connecting only one NREN. However, where NREN CBF is available, it will be possible to share a PoP between several NRENs and use NREN owned- and -operated transmission resources for lambda and IP backhaul connections. The PoP can be shared for both L2 (lightpath) and L3 (IP) connections. The goal is to reduce cost by reducing fibre and equipment needs. In addition, making the PoP an exchange point for traffic and connections between several NRENs can increase flexibility. A candidate for this use case is the NORDUnet – SURFnet – PSNC CBF connection in Hamburg. By turning this facility into a DANTE PoP, three (or four) NRENs can be served in one PoP.
4. *Remote IP backup with CBF.* Traditionally, NREN IP connections to GÉANT take the form of a primary connection to a router at the GÉANT PoP in the same country, and a backup connection using a GÉANT-provided lambda to a router at a GÉANT PoP in a different country. An alternative is to use CBF (when available) to connect the NREN directly to a GÉANT PoP in another country. The goal is to increase resilience by providing site redundancy. A candidate is backup IP connectivity for NORDUnet and SURFnet, using NREN-owned and -operated CBF for transmission capacity. A draft design for this case has already been proposed and can be implemented within the scope of the activity.

> Note: DFN and RENATER have implemented such a backup already. DFN is connected via its own site in Erlangen to Paris and from there to the GÉANT PoP. A backup for RENATER has also been configured.

For all four use cases, the approach will be to select one scenario where appropriate CBF resources are available, do a technical design, propose a plan for implementing the design in the operational GÉANT network, taking into account both technical and operations requirements. The cost, reliability, and operations impacts of the solution will then be studied.

The test cases are expressed above in terms of their infrastructure elements, notably CBF. However, in testing, the full interlayer communication must be tested; operations and service delivery should be considered for each case.

From the point of view of the operational and service layers, there are other implications. In both cases 1 and 2, NRENs will act as carriers and DANTE will be the customer. Detailed operational procedures should be written with clear points of contact and a Service Level Agreement agreed by both sides. On the other hand, in cases 3 and 4, other actors are involved. Other NRENs will be the users of this infrastructure. So, in addition to the procedures that must be defined as in cases 1 and 2, other kinds of procedures must be written, clarifying who manage what, and the steps that these NRENs must follow in case of incidents, programmed works or any kind of query.

A detailed design must be carried out to deploy L2 and L3 services. Since they will be configured through NREN infrastructure, they must be well differentiated from the NREN's own services. These L2 and L3 services could be deployed both for basic IP connectivity and also for projects, so there will be an important group of users that will make use of these services. It means that a major effort in the definition of operational procedures should be done. I-Share can provide useful input to this work.

6.3 Federated Network Design for LHC Tier1-to-Tier2 Connections

For distributing the results of the LHC experiments, an Optical Private Network (OPN) has been built, linking CERN as the Tier0 centre to 11 Tier1 centres around the world (see Figure 3.1 on page 12). The traffic that will be distributed is regarded as stable over time because it is clear that CERN itself cannot store the huge amount of data that will be generated by the experiments.

One outstanding issue is the connection of Tier1 centres to Tier2 centres for the further distribution of data. The overall topology of these connections is unknown to a large extent. For example, the following aspects remain unclear:

- Which topology is required? In particular, it is not yet known how many Tier1 centres may be contacted by a Tier2 centre. It is likely that each Tier2 centre will retrieve data from more than one Tier1 centre.
- Which are the major sources and sinks of data? This is currently unknown because there is considerable uncertainty about the computation model.

- Which paths will the data flows take? This is related to the previous point and will be less uncertain once the sources and sinks are clear.
- Will these paths be more or less the same over time?
- What volumes are the data streams going to have?
- What are the quality expectations in terms of availability and monitoring?

The uncertainty impacts the way a federated network should look in order to provide the best support for the project. There are several alternatives for how the GÉANT consortium might support such connections:

- As with the LHCOPN, static connections can be established between Tier1 and Tier2 centres.
  - These are useful for flows with large data volumes that are always exchanged between the same end sites and where quality expectations (concerning availability, jitter, security) are high.
- A bandwidth-on-demand solution may be provided.
  - BoD is useful for flows with large data volumes where the communicating end sites change. This solution should fulfil high quality expectations (concerning availability, jitter, security) once the service is in operation.
  - The implementation of this solution is dependent on technical support from the NRENs, because BoD is a new technology and initial problems due to the newly developed software have to be taken into account. It may therefore be considered as too high a risk for this very important project.
- No additional infrastructure may be required if plain IP is regarded as sufficient.
  - This option is suitable for flows with low data volumes and no special quality demands.

Note that it is not necessary to adopt the same solution for every Tier1-to-Tier2-centre connection. It would even be possible to combine all three.

In Y2, JRA1 T3 is planning to model the situation using the architecture models in order to try to help assist in decision-making. The experience of instantiating the models with LHCOPN will also serve as a feedback mechanism for improving them.

7 Conclusions

In Y1 of the GN3 project the JRA1 T3 members have investigated the benefits and challenges of, and user requirements for federated networks with respect to current and future research projects. Existing tools and services have been analysed to identify possible contributions towards federated networks and to highlight areas where additional efforts are required. Based on these investigations, architecture models for federated networks have been developed.

The proposed model consists of three main layers: Infrastructure, Operations and Service Layer. Each contains elements which are building blocks for that federation layer: networking elements (e.g. NRENs in GÉANT) in the Infrastructure Layer; tools, operations centres and supporting intra-federation services in the Operations Layer; and various services provided to the end users by the federation in the Service Layer. Within an architecture layer, pairs of elements are inter-related, a relationship which might represent communication, data-transfer flows or some logical relationship. In addition, similar relationships exist between neighbouring layers of the architecture, i.e. between the Infrastructure and Operations Layers, as well as between the Operations and Service Layer. It is also recognised that a number of services provided by the end users will use the federation, and that the federation should therefore support these. However, since providing such services is not within the scope of the federation (although the requirements of those services are), they are represented within an outer box in the model.

Two variants of the federated network architecture model have been presented, Model A and Model B, reflecting the two ways communication is performed within and between layers, In Model A, all aspects of communication between layers are carried out only through the Layer Management entities of neighbouring layers. In addition, all data and information exchange between the Service and Infrastructure Layers has to go through and be performed by the Operations Layer. Communication within a layer is performed via internal communication. Model A is more restricted than Model B. Model B allows more communication and information exchange channels. In Model B, individual elements within an architecture layer can communicate with another element in a different architecture layer not only via the Layer Management block but also directly. Such direct communication can speed up processes and operations. However, if processes and procedures are not well defined, an increased number of communication channels can create additional overhead, or poor information flow, both of which should be avoided.

In both models, links between elements can have different meanings, depending on the type of relationship, including physical connectivity, intensity of traffic flows, directions of data and information flows, dependencies in specific projects, and so on. Since each federation has its own specific elements and element connections, it is not useful to limit the definition of the term "relationship" between elements. Any potential federation should be analysed case by case to determine the existing procedures behind each of the links that will be represented in the architecture model. In addition, the links in both models represent specific processes and procedures with their respective work- and information flows; these should be very well defined and well known to the elements involved. At each different stage of federation – establishment, the addition or removal of an element, or federation restructuring – the relevant processes should be examined for opportunities for enhancement and optimisation.

As future work, the federated network architecture models can be placed in context with other models from the literature as well as with models used or proposed by other GN3 service and joint research activities. In addition, the models will be verified and analysed using a number of test cases, including GÉANT and LHC Tier1-to-Tier2 connections.

References

[AutoBAHN] http://www.geant2.net/server/show/ConWebDoc.2544

[BBD] B. Belter, S. Leinen, T. Rodwell, M. Sotos, “GN2 Deliverable DS3.12.1: Description of a Decentralised PERT” http://intranet.geant2.net/upload/pdf/GN2-06-310v5-DS3-12-1_Description_of_a_Decentralised_PERT.pdf

[BBMRI] www.bbmri.eu

[CNISW] http://cnis.psnc.pl

[DEISA2] http://www.deisa.eu

[E-ELT] www.eso.org/projects/e-elt

[eduPERT] http://edupert.geant.net/

[EGEE-III] http://www.eu-egee.org/

[EISCAT] www.eiscat.se

[EMUC] E. De Marinis, M. K. Hamm, A. Hanemann, G. Vuagnin, M. Yampolskiy, G. Cesaroni, S.-M. Thomas, “Deliverable DS3.16.1: Use Cases and Requirements Analysis for I-SHARe”

[ESFRI] http://cordis.europa.eu/esfri/roadmap.htm

[EVLBI] http://www.evlbi.org/

[FAIR] www.gsi.de/fair/index_e.html

[GCPST] G. Cesaroni, M. K. Hamm, M. Labedzki, G. Vuagnin, M. Wolski, M. Yampolskiy, “I-SHARe – a Process Support Tool for Multi-Domain Services”, TNC 2009, Malaga, 8-11 June 2009

[GLIF] http://www.glif.is

[GN2AW] http://www.geant2.net/server/show/ConWebDoc.2544

[GN3-08-034] http://www.geant2.net/upload/pdf/GN3-08-034-GN3-White-Paper_20080808173508.pdf

[IEEE1471] ANSI/IEEE 1471-2000, “Recommended Practice for Architecture Description of Software-Intensive Systems” http://www.iso-architecture.org/ieee-1471/

[IEEE1471wiki] http://en.wikipedia.org/wiki/Ieee_1471

[IEEE610.12] “IEEE Std 610.12-1990 IEEE Standard Glossary of Software Engineering Terminology – Description” http://standards.ieee.org/reading/ieee/std_public/description/se/610.12-1990_desc.html

[Infrafrontier] www.infrafrontier.eu

[JRA1S] GN3 Intranet: GÉANT Research Activities > JRA1: Future Network > T3 Federated Network Architectures > Documents > Studies (access restricted to GN3 participants) https://intranet.geant.net/sites/Research/JRA1/T3/Documents/Forms/AllItems.aspx?RootFolder=%2fsites%2fResearch%2fJRA1%2fT3%2fDocuments%2fStudies&F

olderCTID=0x012000912272C74D96A04DB1A6DE8EBA2B4A95&View=%7b0AD97875%2dF1CD%2d49C0%2dBC26%2dD057C2B48AC8%7d

[LHCCERN] http://lhc.web.cern.ch/lhc/

[LHCOPN] https://twiki.cern.ch/twiki/bin/view/LHCOPN/WebHome

[LHCOPNOPS] G. Cessieux, “Proposed LHCOPN Operational Model”
https://twiki.cern.ch/twiki/bin/view/LHCOPN/OperationalModel

[LifeWatch] www.lifewatch.eu

[NDGF] http://www.ndgf.org/ndgfweb/home.html

[PAAMPS] A Patil, B. Belter, A. Polyrakis, T. Rodwell, m. Przybylski, M. Grammatikou, “GEANT2 Advance Multi-domain provisioning system”, TERENA Networking Conference 2006, Catania, Italy

[perfSONAR] http://www.perfsonar.net/

[PIVIR] GÉANT2 Deliverable DJ1.3.4 JRA1 Phase IV Implementation Report
http://www.geant2.net/upload/pdf/GN2-08-057v2-DJ1-3-4_JRA1_Phase_IV_Implementation_Report_20080401092109.pdf

[PMM] http://www.geant2.net/server/show/nav.1801

[PRACE] www.prace-project.eu

[RedIRIS] http://www.rediris.es/20aniversario/sobre_rediris.html

[RFC1930] J. Hawkinson, T. Bates, “RFC1930 – Guidelines for creation, selection, and registration of an Autonomous System (AS)”, Category: Best Current Practice, March 1996
http://www.faqs.org/rfcs/rfc1930.html

[Sessions] R. Sessions, “A Comparison of the Top Four Enterprise-Architecture Methodologies”, May 2007
http://msdn.microsoft.com/en-us/library/bb466232.aspx

[SOP] “Setting up and operating a PERT”, GN2 PERT Workshop,
http://wiki.geant2.net/pub/SA3/Sa3PertTraining/per_sl_modules_1_to_5_version_1.0.ppt

[SSAW] S. Sima, L. Altmannova, “Use of NREN resources as network elements of the GÉANT”, Fourth Workshop on the Architecture of the Future GÉANT Network, Rome, 19-20 October, 2009
http://wiki.geant.net/pub/Main/ArchWshop4/SS-_Use_of_NREN_resources_as_network_elements_of_the_GANT.ppt [access restricted to GN2/3 participants]

[SSK] http://www.porta-optica.org/files/kiev/05_Kiev_Sima.pdf

[TCNREN] TERENA compendium of National Research and Education Networks in Europe 2009 Edition,
http://www.terena.org/activities/compendium/2009/pdf/TERENA-Compendium-2009.pdf

[TRP] T. Rodwell, S. Thomas ”GN2 Deliverable DS3.12.4: Policy for a Federated Performance Enhancement Response Team (PERT)
http://intranet.geant2.net/upload/pdf/GN2-08-101-DS3-12-4_Policy_for_a_Federated_PERT.pdf

[ViewModel] http://en.wikipedia.org/wiki/View_model

[XFEL] http://www.xfel.eu

Glossary

AD	Architectural Description
AMPS	Advance Multi-domain Provisioning System
AutoBAHN	Automated Bandwidth Allocation across Heterogeneous Networks
BDII	Berkeley Database Information Index
CBF	Cross-Border Fibre
CERN	European Organisation for Nuclear Research
CLAN	Campus Local Area Network
CL MP	Command Line Measurement Point
cNIS	Common Network Information Service
CPU	Central Processing Unit
DEISA 2	Distributed European Infrastructure for Supercomputing Applications
DM	Domain Manager
DPM	Disk Pool Manager
DRAC GUI	Dynamic Resource Allocation Controller graphical user interface
DSM	Data Status Monitor
DWDM	Dense Wavelength Division Multiplexing
E-ELT	European Extremely Large Telescope
e-VLBI	Electronic Very Long Baseline Interferometry
E2E	End-to-End
E2EMon	End-to-End Monitoring System
ENOC	EGEE Network Operations Centre
EoMPLS	Ethernet over Multi-Protocol Label Switching
EGEE-III	Enabling Grids for E-sciencE – Croatia
EISCAT	European Incoherent Scatter
ESFRI	European Strategy Forum on Research Infrastructures
EVN	European VLBI Network
EXPRES	Express Production Real-time e-VLBI Service
FAIR	Facility for Antiproton and Ion Research
FCAPS	Fault, Configuration, Accounting, Performance and Security
FEDERICA	Federated E-infrastructure Dedicated to European Researchers Innovating in Computing network Architectures
FP7	European Union’s Seventh Framework Programme for Research and Technological Development
FTP	File Transport Protocol
GFS	Global File System
GGF	Global Grid Forum

GLIF	Global Lambda Integrated Facility
HPC	High-Performance Computing
IDM	Inter-Domain Manager
IEEE	Institute of Electrical and Electronics Engineers
IETF	Internet Engineering Task Force
I-SHARe	Information Sharing across Heterogeneous Administrative Regions
IXP	Internet Exchange Point
JIVE	Joint Institute for VLBI in Europe
JRA	Joint Research Activity
JRA1 T3	Joint Research Activity 1: Future Network; Task 3: Federated Network Architectures
JRA2 T3	Joint Research Activity 2: Multi-Domain Network Service Research; Task 3: Monitoring
L1	Layer 1
L2	Layer 2
LAN	Local Area Network
LHC	Large Hadron Collider
LHCOPN	Large Hadron Collider Optical Private Network
LOFAR	LOw Frequency ARray
MAN	Metropolitan Area Network
MPLS	Multi-Protocol Label Switching
MDM	Multi-Domain Monitoring
NE	Network Element
NMWG	Network Measurement Working Group
NOC	Network Operations Centre
NREN	National Research and Educatoon Network
NZDSF	Non-Zero Dispersion-Shifted Fibre
OGF	Open Grid Forum
OPN	Optical Private Network
PIP	Premium IP
PRACE	Partnership for Advanced Computing in Europe
perfSONAR	Performance Service-Orientated Network-Monitoring Architecture
PERT	Performance Enhancement Response Team
PoP	Point of Presence
PRACE	Partnership for Advanced Computing in Europe
QoS	Quality of Service
R&D	Research and Development
RHEL4	Red Hat Enterprise Linux 4
RLF	Repository Level Federation
ROADM	Reconfigurable Optical Add-Drop Multiplexer
RREN	Regional Research and Education Network
RTT	Round-Trip Time
SA2 T1	Service Activity 2: Multi-Domain Network Services; Task 1: Multi-Domain Network Connectivity Services Development
SA2 T2	Service Activity 2: Multi-Domain Network Services; Task 2: Multi-Domain Service Coordination & Operations
SA2 T3	Service Activity 2: Multi-Domain Network Services; Task 3: Monitoring and Performance

SAM	Service Availability Mapping
SCARIe	Software Correlator Architecture Research and Implementation for e-VLBI
SDH	Synchronous Digital Hierarchy
SGE	Sun Grid Engine
SLA	Service Level Agreement
SLF	Service Level Federation
TCP	Transport Control Protocol
UCLP	Universal Commerce Language and Protocol
UDP	User Datagram Protocol
VPLS	Virtual Private LAN Service
VPN	Virtual Private Network
XFEL	(European) X-Ray Free-Electron Laser
XML	Extensible Markup Language