Manual Multi-Domain Routing for Géant E2E Links with the I-SHARe Tool

Mark Yampolskiy[1,2,3], Wolfgang Fritz[1,2], Wolfgang Hommel[2,3], Gloria Vuagnin[4], Fulvio Galeazzi[4]
[1] *German Research Network (DFN), Alexanderplatz 1, 10170 Berlin, Germany*
[2] *Leibniz Supercomputing Centre (LRZ), Boltzmannstraße 1, 85748 Garching, Germany*
[3] *Munich Network Management (MNM) Team, Oettingenstraße 67, 80538 München, Germany*
[4] *Italian Research and Education Network (Consortium GARR), Via dei Tizii 6, I00186 Rome, Italy*
{*yampolskiy,fritz,hommel*}*@lrz.de,* {*gloria.vuagnin,fulvio.galeazzi*}*@garr.it*

***Abstract*—The term *routing* is usually associated with the fully automated routing in computer networks. Various routing techniques, protocols, and strategies have been established in LANs, WANs, Internet, and PSTN networks. However, all these routing approaches rely on a pre-installed, pre-configured, and well-maintained network infrastructure. The process of planning the network infrastructure remains the stronghold of manual planning. Whereas the planning within a single provider domain is a very common task for all network service providers (SPs), the planning of multi-domain backbone connections introduces several additional challenges, such as the coordination of connection options among multiple SPs' planning teams. In this paper we present the *I-SHARe* (Information Sharing across Heterogeneous Administrative Regions) tool, which has been developed in the pan-European collaboration Géant in order to foster inter-provider collaboration during the planning and operation of multi-domain backbone connections.**

***Keywords*-manual routing; backbone connections; management processes; tool-support**

I. INTRODUCTION

Géant is a collaboration of over 30 European National Research and Education Networks (NRENs). Whereas the purpose of every NREN is to provide network connections for national research and education institutions, the purpose of Géant is to interconnect NRENs and consequently to foster international research projects where participating organizations are connected to different NRENs. The portfolio of Géant includes various services among other conventional IP connections. However, such services cannot always fulfil all the challenging requirements of modern research collaborations. The Large Hadron Collider (LHC) project provides a very good example. Experiments performed in the CERN center near Geneva produce over 15 petabytes of raw data per year [1]. Mainly to store all this vast amount of data, it has to be transferred to 11 so called Tier-1 (T1) supercomputing centres spread across Europa and North America [2]. The analysis of the data is then performed in 160 T2 supercomputing centres spread around the entire world. CERN backups all the data on tape but has the local high speed storage capacity sufficient only to save experimental data of few days. Therefore, a bad network quality or long term outages of network connections between the T0 centre at CERN and the T1 centres might lead to an inability to process real time the data taken, with a negative impact on the analysis results. A bad quality of connections between T1 and T2 centres is not that critical, but nevertheless might lead as well to significant delays of the data analysis. Therefore, LHC needs permanent high-bandwidth and high availability connections between involved research organizations. Other good examples are Grids, e.g., WLCG [3] or EGEE [4], where the involved supercomputing centres require network connections as a means for job-transfers between the collaborating partners.

Realizing high-quality high-bandwidth connections in general purpose IP networks is very difficult, as communication flows can interfere with each other and therefore lead to bad connection quality. In order to cope with the challenging customer requirements, a novel service, End-to-End (E2E) Links, has been introduced in Géant. E2E Links are dedicated optical point-to-point connections realized at ISO/OSI levels 1 and 2, with connection segments provided by one or more NRENs [5]. Regarding their quality requirements, the used network technologies as well as the geographical dimensions, E2E Links are nothing else but multi-domain backbone connections, in which—in opposite to classical backbones—multiple network providers are involved and heterogeneous network technologies can be used.

The E2E Links service has been first introduced mid 2005. The speciality of the service is that a new connection can be ordered regardless of whether the required infrastructure is already installed or not. If new infrastructure is needed to fulfil the customer's request, it can be procured, installed, and configured according to the requirements to the new E2E Link. Consequently, all route planning procedures can only be done manually and require intensive interactions between the involved NRENs.

The experience made in the first years of the E2E Links service has revealed that information exchange and information management are key factors determining the time needed for manual connection planning and installation. Information exchange via email and planning via Excel sheets has proven to be error prone with a high probability

 ISBN: 978-1-61208-141-0

of losing necessary information or missing various events, e.g., the delivery of the procured infrastructure by neighbour NRENs. This results in a high fluctuation of the time needed to plan and install new connections. In order to improve the outlined situation, a tool supporting the information exchange among participating NRENs should be introduced for the E2E Links service.

The design and development of the *I-SHARe* (Information Sharing across Heterogeneous Administrative Regions) tool has been performed by an international team of researchers working for different NRENs. The *I-SHARe* tool covers information exchange for the whole life cycle of an E2E Link service instance, from its planning through installation and operation till decommissioning. In this paper, we focus on the tool support for the planning of a new E2E Link. We outline the most important challenges of such planning procedures in Section II. The manual routing procedure with *I-SHARe* is presented in Section III. In Section IV, we provide a brief history of the tool's design and development. In Section V we present our future plans. The presented paper aims to promote the gained knowledge about the *I-SHARe* tool, so that network operators facing similar challenges can use it as well.

II. Specific Challenges

Manual route planning has to overcome a combination of technical and organizational challenges. Both types are caused by the organizational independence of the involved NRENs.

Due to various domain-internal reasons, it is typical for organizations to have very restrictive information policies. This means that the amount of information, which is allowed to be shared with other project partners, is very limited. For instance, it is typically prohibited to share detailed information about the physical network topology or total capacity available on the already installed infrastructure. On the other hand, as NRENs have to collaborate with each other on service instance planning and realization, it is broadly acceptable to share service-instance-bound information.

Further, independent organizations tend to have different procurement policies and various preferences regarding hardware vendors and technologies. This is often caused by legal issues like procurement rules, by past experience, and contractual conditions with various hardware vendors as well as by the competence of organizational members with particular technology. Consequently, this results in a high level of heterogeneity of hardware and networking technologies used by different NRENs. The interconnection of different technologies is not an easy but very well understood task. This, however, requires the consideration of the compatibility of the used hardware, e.g., network interfaces, and network parameters like the maximal supported frame size. Sharing such information is essential in order to prevent problems caused by incompatibility or misconfiguration.

The planning of the network infrastructure is done manually in each NREN. It is inevitable that some infrastructure might have to be procured and information about both financial conditions and hardware properties should be first requested from the hardware vendor(s). This introduces unpredictable delays and an uncertainty of properties that NRENs will be able to provide, as for instance some hardware used in the past might become obsolete and is not supplied anymore. As all NRENs are independent organizations, changes of the planning conditions, e.g., the ordered infrastructure should be delivered at a certain time, are not automatically known to other involved partners. This raises the necessity to notify other involved partners about the completion of the own planning part and about the properties available for the connection.

In conjunction with the information exchange, a reduction of the information flood is also needed. This is especially important as the tool has to support manual processes. The reduction of information means, for instance, that only relevant partners should be involved and not all NRENs involved in the collaboration and the necessary information do not need to be resent many times. Furthermore, especially for the planning of interconnecting interfaces it is important to know the plans of the neighbour NREN for the particular connection instance, in opposite to all interfaces planned for all instances.

Last but not least, the coordination aspect has to be mentioned. As network planning teams of all involved NRENs operate independently, some sort of coordination is needed in order to achieve the common goal—planning of a new E2E Link. The lack of such information can result in unresolved deadlocks, if, for example, two neighbour NRENs have simultaneously planned incompatible infrastructure. Finally, the outlined information exchange has to be embedded in multi-domain operational procedures.

III. Planning a new route with I-SHARe

If a project like LHC requires a new E2E Link, for example, between CERN and Brookhaven National Laboratory (BNL) in USA, a corresponding request can be submitted to one of the NRENs to which end-points are connected. The project has to specify two end-points and the required properties of a new connection. If the contacted NREN approves the request, network planning team(s) start to work on the planning of a new connection and consequently on its installation. The *I-SHARe* tool is dedicated to support this work.

In this section we first briefly describe the system architecture of *I-SHARe*. After that, we outline the support of the service instance life cycle. Finally, we present the usage functionality of *I-SHARe* by planning and setting up a new connection.

 ISBN: 978-1-61208-141-0

A. Data separation in I-SHARe's system architecture

The handling of single-domain and multi-domain information is clearly distinguished in the *I-SHARe* system architecture (see Figure 1). Information such as operational groups and group members, their responsibility areas and contact data are handled in the *domain part*. This information can be maintained in the *I-SHARe* domain part, or in an NREN's domestic management tool. In both cases the single-domain information is propagated to *I-SHARe* via the *I-SHARe Domain Interface*. The *I-SHARe Central Server* stores the copy of the provided information, so that it can be incorporated in the supported processes.

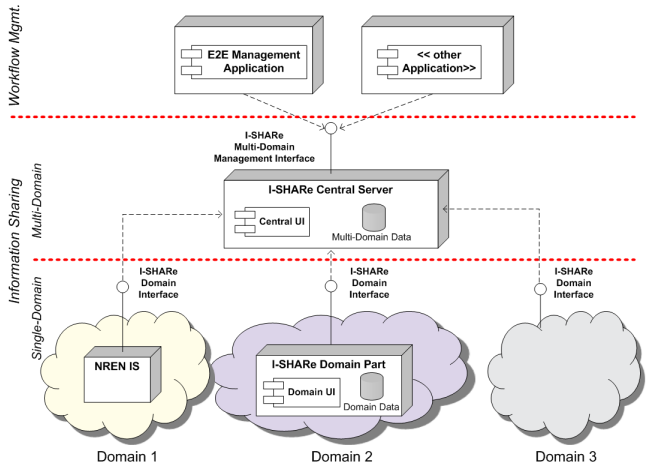

Figure 1. System architecture of *I-SHARe* [6]

Multi-domain information like the route of an E2E Link through NRENs, interfaces of the adjacent connection parts provided by neighbor NRENs, and states of various operations are stored directly in the *I-SHARe Central Server*. This information can be accessed and edited via a web-based GUI (see Sections III-C and III-D for the detailed description).

All information stored in the *I-SHARe Central Server* can be accessed from other applications through the *I-SHARe Multi-Domain Interface*. This north-bound interface provides the means for integrating *I-SHARe* with other tools, e.g., with workflow management or analysis tools.

B. Life cycle coverage in I-SHARe

I-SHARe is designed to support the multi-domain manual management processes during the whole life cycle of E2E Links. According to the specifics of the E2E Links service, one has to distinguish between four phases: (i) *Ordering* of a new E2E Link, (ii) *Setting up* of the ordered E2E Link, (iii) *Operation* of E2E Links in-service, and (iv) *Decommissioning* of no longer needed E2E Links. E2E Links in different phases can be accessed through different views (top-level tabs in the GUI) of the *I-SHARe Central Server*. The distinguishing between views is needed, as the tasks in the corresponding phases require different knowledge, skills, and competences, which generally are provided by different teams. Furthermore, also the information needed in various phases overlaps only partially.

The GUI views and the corresponding tasks are defined as follows:

- The **Ordered** view contains all links that have been requested. During the first phase the general feasibility of the requested E2E Link with the specified quality parameters is investigated. Among other tasks, this includes the selection of NRENs, which should participate in the realization of E2E Link segments.
- The **Set Up** view incorporates information about all ordered E2E Links, whose general feasibility has been approved. The purpose of this phase is to oversee and coordinate all steps needed for the establishing of the planned connection, i.e., the installation and configuration of the needed equipment, the interconnection of all adjacent connection segments, and the accomplishment of all integration tests needed for the allowance for service.
- The **Operational** view provides access to information about all E2E Links, which are delivered as a service to the end-customers. This view incorporates all plans and details elaborated in the previous phases. The changes in this phase can be used, e.g., in order to plan an upgrade of used infrastructure.
- The **Decommissioned** view provides access to all formerly operational, now obsolete connections. This view can be used, e.g., in order to reuse solutions elaborated for E2E Links, which are not in service anymore.

In the remainder of this section we will present, how the *I-SHARe* tool can support the manual work performed in the first two phases of the service instance life cycle.

C. The ordering process

After the contacted domain approved the request for a new E2E Link, one of the domain's experts needs to login to *I-SHARe* and open a new link request. First, he specifies the end points given by the request and both connecting domains (see Figure 2). The rectangles in the figure represent organizational domains. If domains connecting the end-points are not neighbors, one or more transit-domains can be inserted by clicking the "+"-button on the right hand of the vertically arranged route. The acronyms of the domains can be selected from a drop-down box in the middle of the rectangle. Furthermore, for each domain the *Point of Presence* (PoP) can be specified, at which it should be connected to the neighbor domain. The connection at the end-point site is not specified, as it is the end-customer's responsibility. In the GUI, the domain-specific PoP list is represented as a drop-down box at the top or bottom edge of the rectangle. The list is provided to the *I-SHARe Central Server* by each domain via the *I-SHARe Domain Interface* (see Section III-A).

 ISBN: 978-1-61208-141-0

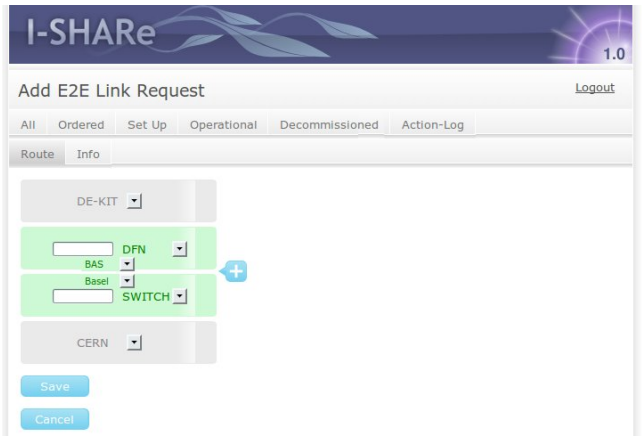

Figure 2. Define end points and connecting domains [7]

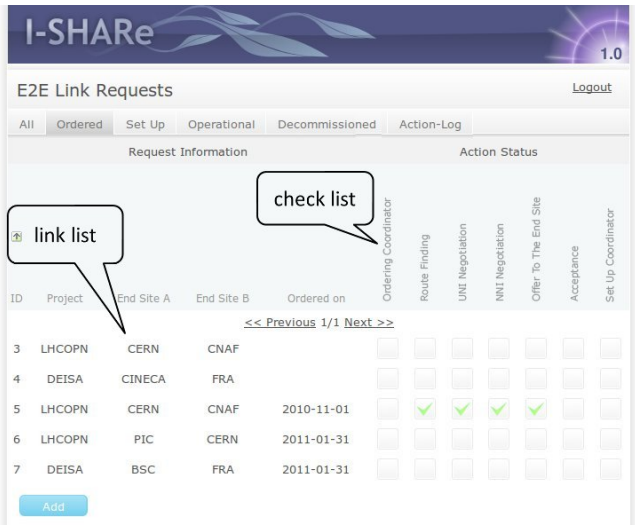

Figure 3. I-SHARe's list of ordered links [7]

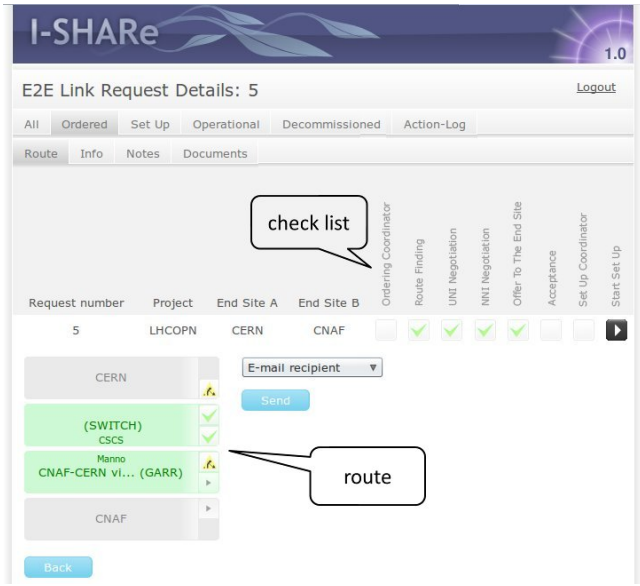

Figure 4. Detailed view of an ordered link [7]

At this stage, the *I-SHARe* tool also asks for further relevant information, such as the assigned project's name and the customer requirements, e.g., the guaranteed bandwidth. Other relevant entries, like the request date, are added automatically. Besides the information already published by the domain part, more contact details may be specified then as well. Furthermore, already at this stage it is possible to assign an *Ordering Coordinator* (OC)—a special role responsible for coordinating the actual ordering process among different steps, persons, and institutions. The OC makes sure everyone takes the right actions and keeps track of the overall progress. The OC is selected from the list of all network specialists (reported via the domain interface), whose domain designated him as qualified to take the responsibility for this role. Usually, the OC is selected from members of the connecting domain.

After a new link request is saved, it will appear in the *Ordered* view (top-level tab in the GUI). This view is divided into two parts—a link list showing the pending link requests and a check list indicating the progress (see Figure 3).

By clicking at an ordered link, its detailed—alike divided—view can be accessed (see Figure 4). The route part shows all involved end sites and domains, whose detailed (contact) information people can access by clicking at them.

The check list contains the state of all steps that the experts have to take before the link may be "promoted" to the next life cycle phase. It covers the selection of an *ordering coordinator*, *route finding*, *UNI negotiation*, *NNI negotiation*, *offer to the end site*, *acceptance*, and selection of *set up coordinator*. First, all involved domains need to agree on an OC, which is the first action. After that, they have to work together to find a feasible route from the requesting customer to the destination. The OC and the experts are supposed to set the individual states so that other people can easily keep track of the link negotiation's progress. During both, *UNI and NNI negotiation*, the experts provide general and technical information about their interfaces, such as its capacity, hardware, and interface type, where it is located, transceiver and media types or even fine grained parameters such as the MTU size or the used wavelength. Figure 5 outlines a typical view during NNI negotiation. In this view, a member of one domain can edit various parameters and at the same time see the corresponding parameters its neighbor specified.

In order to notify neighboring domains about different events, *I-SHARe* provides user friendly email functionalities, too. The tool supports the selection of recipients relevant for the particular connection. Further, at any time notes and documents can be stored in the system, for example, plans for patch panel interconnection.

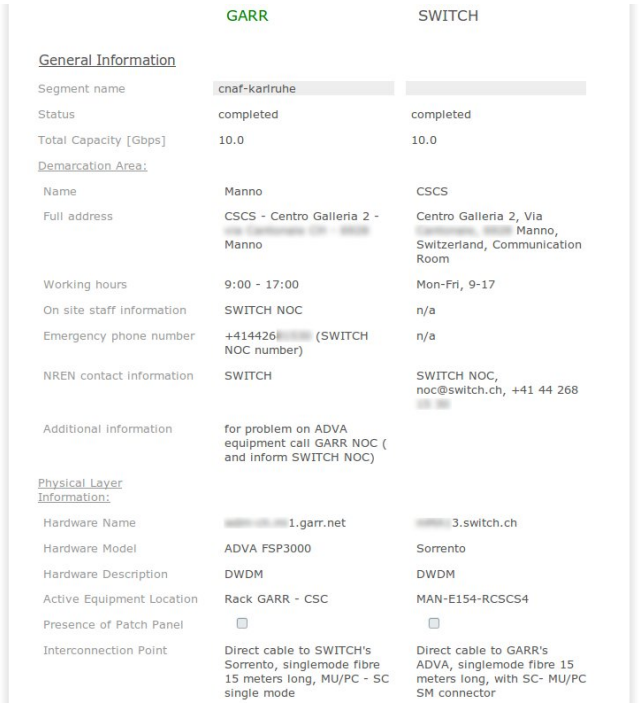

Figure 5. NNI negotiation page [7]

Figure 6. Install and configure the network infrastructure [7]

After all planning steps have been completed successfully and the customer accepted the particular offer, the link may be put to the next life cycle phase, the set-up phase. Prior to that a *Setup Coordinator* (SC) has to be selected— a role similar to OC, but responsible for the coordination of all activities related to the link set up. If the domains could not find a feasible route for a particular link request, they will reject it and inform the customer.

D. The set up process

Similar to the ordering phase, *I-SHARe* provides a separated list for each link that is currently in the process of being established. By clicking at a link, the experts can access a more detailed view that is alike divided. On the one hand, there is a more detailed description of the selected route—now not only containing sites, but also interconnecting links. It is itself divided into two parts—one with detailed check boxes, where the actual states can be set depending on the type of the entry and a global section that computes the overall states automatically, based on the segments' states (as indicated in Figure 6). The top part presents a check list (containing *E2E link ID assigned*, *set up request sent*, *infrastructure ready*, *connection tested*, ready for monitoring, *set up completed*—see Figure 6). The first thing *I-SHARe* will ask for is a unique name for the link (until now, the identifier was just a number). After that, the SC asks all involved domains to start the actual hardware installation and configuration (apply connectors, fibres, etc.) and indicate that by setting *infrastructure ready* properly (done, work in progress or delayed, see Figure 6). This step can be done by various domains simultaneously to save time. As soon as all involved NRENs have completed that step successfully, they need to test the new connection. This may not only include local test, but also tests in cooperation with some or all of the other partners to guarantee the whole link is working. Last but not least, the experts have to include that new link in their local monitoring systems and export information about it to the Géant multi-domain monitoring tool. This allows the experts to keep track of the link's current status. After all previous steps have been completed, the SC marks the last step, *set up completed*, as done, and then declares the link operational. That means the whole process was successful and the customer's request lead to a new E2E link, which can be then used by the end users.

IV. Short History of I-SHARe

As the *I-SHARe* tool had to be developed within and provided to an international collaboration of different NRENs, the whole *I-SHARe* team has also been assembled from members of the participating NRENs. In order to separate responsibilities for key aspects and hereby avoid conflict of interests, the *I-SHARe* team consisted of a designer and a development team from the beginning.

The design of the *I-SHARe* tool has been led by the Italian and the German NRENs (GARR and DFN respectively) with a strong participation of the Swiss NREN (SWITCH). During different phases members of RENATER (French NREN) and DANTE (operator of the Géant network) have also participated in the efforts of the designer team. Only NREN

members with experience in requirements analysis, system design and other key project management tasks have been assigned to the designer team. The *I-SHARe* development team consists of members of the Polish supercomputing centre PSNC (working for the Polish NREN PIONIER), who have proven to be experienced in the development of web-based applications.

The major difficulty of *I-SHARe's* design was to gather and analyse the operations' requirements. The interviews with network planning and operational teams have begun in the mid of 2008 and revealed different and sometimes even controversial needs of these teams. The de-facto multi-domain process was defined by the *I-SHARe* designer team in the Géant deliverable DS3.16 [8]. After the NRENs had approved it, this deliverable has been used to identify necessary information that have to be exchanged during the different steps planning, setting up, and operation of E2E links. In order to evaluate gathered requirements, an *I-SHARe* prototype with reduced functionality has been designed. Implemented by the *I-SHARe* development team, this prototype has been evaluated by the NRENs' operational teams. The received feedback has been used to improve the *I-SHARe* system specification, which has been finished in summer 2009. The implementation of the first fully functional version of *I-SHARe* tool has been finished in summer 2010.

After the quality assurance performed by the designer team, *I-SHARe* v 1.0 was approved for the 6 month long pilot phase. During this phase, the operations of selected NRENs evaluate the suitability of the developed tool for their daily work regarding the planning and the management of E2E Links. For the pilot phase the *I-SHARe* installation is hosted at Leibniz Supercomputing Centre (LRZ), a tight partner of DFN. The pilot phase started at the end of 2010. NRENs participating in the pilot phase are (in alphabetical order) DANTE, DFN, GARR, PIONIER, REDIRIS and SWITCH. In order to introduce *I-SHARe*, an online training course has been delivered to the network operation teams of these NRENs.

V. First Experiences and Further Steps

During the requirement analysis, system design, development, and quality assurance stages only one team (either designers or developers) was in charge of a particular task at a time. Interactions with potential users and among these teams took place with clear responsibilities and were of rather simple nature. The start of the pilot phase has introduced the necessity of communication not only among these two teams, but also with end-users and the hosting provider. In order to overcome possible misunderstandings and deadlocks, the designer team is now in charge for developing a proposal for operational procedures covering the whole life cycle of *I-SHARe*. Among other, procedures are about to be defined for the treatment of user-feedback, planning new releases, rollout by the hosting provider, and Incident & Problem management during the *I-SHARe* operation. These procedures will define roles, their rights and responsibilities as well as the way of interaction in different situations.

Another development is planned after the *I-SHARe* operation is settled and broadly used by NRENs. The main goal of *I-SHARe* is to support the information exchange between manually performed operational processes. In case these processes become settled and may be even standardized, the development of another tool for workflow management is planned. This tool should reuse *I-SHARe* as an information exchange platform and access it via the "north bound" interface already implemented in the tool.

Acknowledgment

The authors would like to thank all organizations participating in the design, development and operation of *I-SHARe*, especially DFN, GARR, PIONIER, SWITCH, LRZ and PSNC. The authors also wish to thank the members of the Munich Network Management (MNM) Team for valuable comments on previous versions of this paper. The MNM Team directed by Prof. Dr. Dieter Kranzlmüller and Prof. Dr. Heinz-Gerd Hegering is a group of researchers at Ludwig-Maximilians-Universität München (LMU), Technische Universität München (TUM), the University of the Federal Armed Forces and the Leibniz Supercomputing Centre (LRZ) of the Bavarian Academy of Science.

References

[1] Knobloch, J. and Robertson, L., "LHC computing Grid, Technical design report," CERN, http://lcg.web.cern.ch/LCG/tdr/, Tech. Rep., 2006 (last accessed 2011/04/13).

[2] "LHC website," http://public.web.cern.ch, (last accessed 2011/04/13).

[3] "WLCG website," http://lcg.web.cern.ch/LCG/public/, (last accessed 2011/04/13).

[4] "EGEE website," http://www.eu-egee.org/, (last accessed 2011/04/13).

[5] Schauerhammer, K. and Ullmann, K., "Operational Model for E2E links in the NREN/GÉANT2 and NREN/Cross-Border-Fibre supplied optical platform," Géant, Tech. Rep., 2006.

[6] Cesaroni, G. and Hamm, M. and Simon, F. and Vuagnin, G. and Yampolskiy, M. and Labedzki, M. and Wolski, M., "I-SHARe: Prototype specification," Géant, Tech. Rep., 2008.

[7] "I-SHARe pilot installation," cs.ishare.geant.net/, (last accessed 2011/04/13).

[8] De Marinis, E. and Hamm, M. and Hanemann, A. and Vuagnin, G. and Yampolskiy, M. and Cesaroni, G. and Thomas, S.-M., "Deliverable DS3.16: Use Cases and Requirements Analysis for I-SHARe," Géant, Tech. Rep., 2008.

 ISBN: 978-1-61208-141-0