034115

PHOSPHORUS

Lambda User Controlled Infrastructure for European Research

Integrated Project

Strategic objective:
Research Networking Testbeds

Deliverable reference number D0.4

Publishable Activity Report

Due date of deliverable: 2009-06-30
Actual submission date: 2009-06-30
Document code: Phosphorus-WP0-D0.4

Start date of project:
October 1, 2006

Duration:
33 Months

Organisation name of lead contractor for this deliverable:
PSNC

Project co-funded by the European Commission within the Sixth Framework Programme (2002-2006)		
Dissemination Level		
PU	Public	✓
PP	Restricted to other programme participants (including the Commission	
RE	Restricted to a group specified by the consortium (including the Commission	
CO	Confidential, only for members of the consortium (including the Commission Services)	

Publishable Activity Report

Abstract

This document contains a summary description of project objectives and achievements.

Project:	Phosphorus
Deliverable Number:	D0.4
Date of Issue:	30/06/2009
EC Contract No.:	034115
Document Code:	Phosphorus-WP0-D0.4

Table of Contents

Table of Figures

1 Project execution

1.1 Introduction

A new generation of scientific applications is emerging that couples scientific instruments, data and high-end computing resources distributed on a global scale. Developed by collaborative, virtual communities, many of these applications have requirements such as determinism (e.g. guaranteed QoS), shared data spaces, large data transfers, that are often achievable only through dedicated optical bandwidth.

There has been tremendous amount of research and development in the Grid community in terms of Grid services infrastructure and Grid application development. However, there has been very little work done in the area of using network as a first-class Grid resource. There is no existing implementation today that can demonstrate the power of exploiting the optical network as a first-class Grid resource and the challenges that arise in provisioning end-to-end light-paths across different management and control plane technologies spanning multiple administrative domains.

High capacity optical networking can satisfy bandwidth and latency requirements, but software tools and frameworks for end-to-end, on-demand provisioning of network services in coordination with other resources (CPU and storage) need to be developed.

In response to the above requirements, Phosphorus addressed some of the key technical challenges to enable on-demand e2e network services across multiple domains. The Phosphorus network concept and test-bed make applications aware of their complete Grid resources (computational and networking) environment and capabilities. Software and tools developed within the project allow to make dynamic, adaptive and optimized use of heterogeneous network infrastructures connecting various high-end resources.

PHOSPHORUS realises a service-centric infrastructure supporting the deployment of mission-critical applications on a global scale. PHOSPHORUS significantly enhances the capability of e-science applications, providing a unified network/Grid infrastructure that can flexibly adapt to the demands of applications having strong, combined requirements on CPU, memory and storage resources as well as on the communication network. Therefore, PHOSPHORUS applications can rely on a network infrastructure that adapts to the application, rather than having the application to adapt to the network, as proposed in other current EU-funded projects in the same area (e.g. MUPBED, EU-QOS).

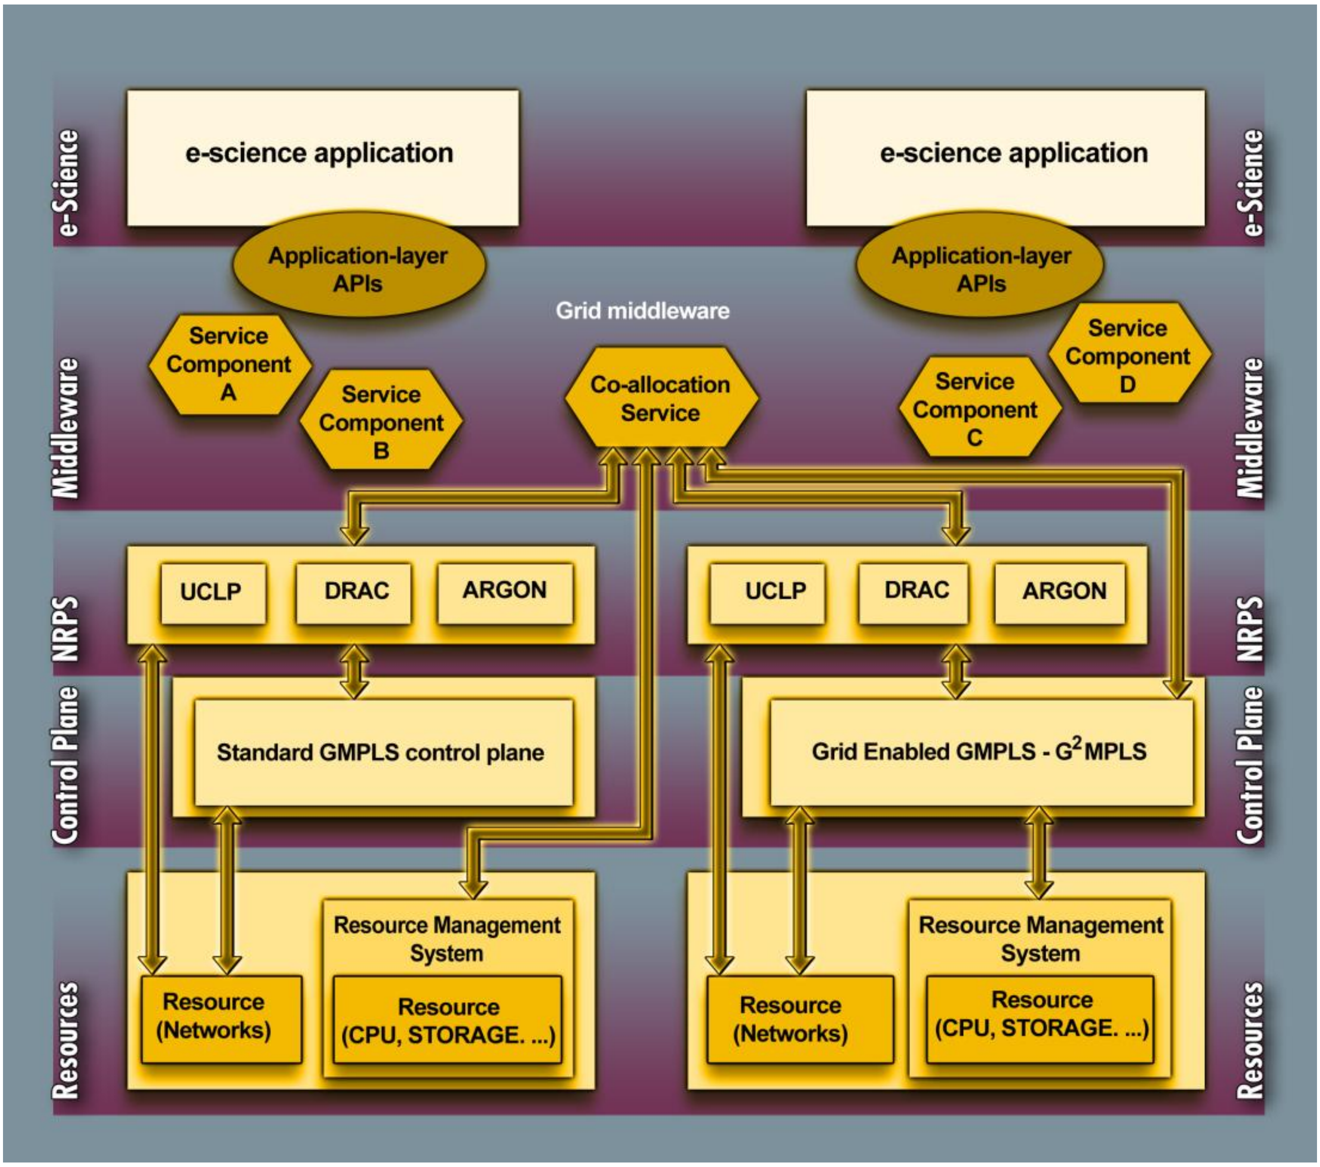

Figure 1-1: PHOSPHORUS architecture

The main innovation introduced by PHOSPHORUS is a network Service and Control Planes concept where the network (lightpath) and Grid (computational, storage) resources are provisioned in a single-step: network and Grid-specific resources are controlled and set-up at the same time and with the same priority, with a set of seamlessly integrated procedures. From a user's perspective, this results in a real, node-to-node deployment of on-demand Grid services.

The innovations introduced by PHOSPHORUS can be organized and analyzed according to the structure of the project's technical topics: Service Plane, Network Resource Provisioning Systems and Control Plane.

The Phosphorus assessment relied on experimental activities on a distributed test-bed interconnecting European and worldwide optical infrastructures. Specifically, the test-bed involved European NRENs and national test-beds, as well as international resources (GÉANT2, Internet2, Canarie, Cross Border Dark Fibre infrastructures and GLIF virtual facility). A set of highly demanding applications were adapted to prove the concept.

Phosphorus disseminated procedures, toolkits and middleware to the EU NRENs and their users, such as Supercomputing centres and the wider European and worldwide scientific users.

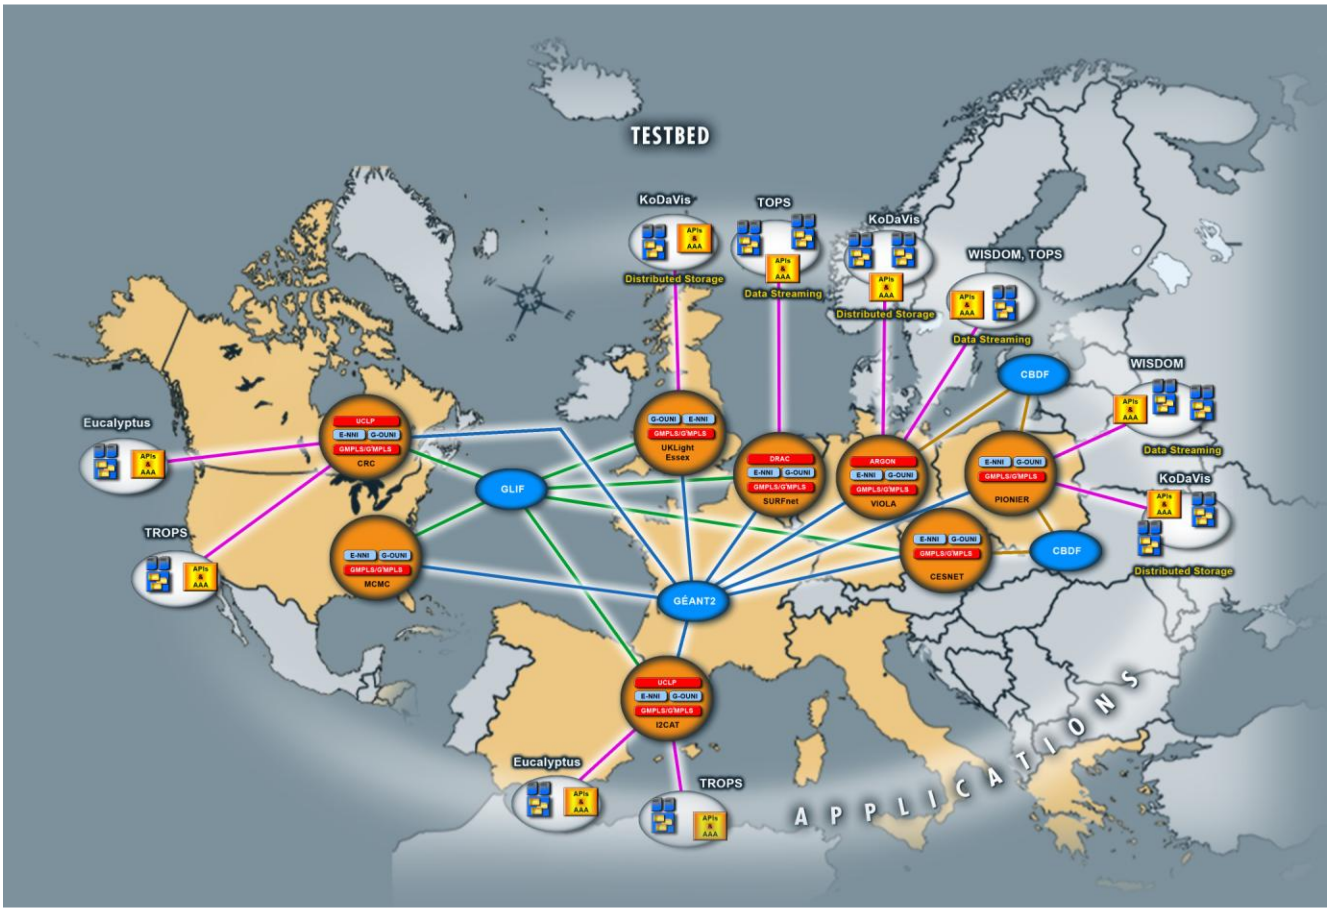

Figure 1-2: PHOSPHORUS project testbed and applications

The project lasted 33 months and was finished in June 2009.

1.2 Summary of the project objectives

The PHOSPHORUS project focused on delivering advanced network services to Grid users and applications interconnected by heterogeneous infrastructures. PHOSPHORUS enabled and tested dynamic, adaptive and optimised use of heterogeneous network infrastructures interconnecting various high-end resources. The ultimate goal of this project was to disseminate procedures, toolkits and middleware to EU NRENs and their users, such as supercomputing centres, to enable authorized end-to-end dynamic service provisioning across the European and worldwide heterogeneous network infrastructure. Furthermore, PHOSPHORUS is providing applications with the ability to treat the underlying network as first class Grid resource.

To achieve this goal PHOSPHORUS had to enhance the level of integration between application middleware and the optical transport networks. It has been enabled by advanced interworking between heterogeneous network domains and their applications environments. Interfacing solutions facilitated vertical and horizontal communication between applications middleware, existing Network Resource Provisioning Systems and the Grid-GMPLS (G2MPLS) control plane. Integration of AAA mechanisms at various network and management layers was essential to ensure that stakeholder interests can be represented and enforced.

The PHOSPHORUS main goals has been broken down into the following, measurable, objectives:

Delivery of single-step on-demand services across multi-domain networks for e-science applications

The project aims to demonstrate on-demand service delivery across access-independent multi-domain/multi-vendor research network test-beds on a European and worldwide scale. The global test-bed in PHOSPHORUS project is composed of a number of local test-beds interconnected using multiple optical international networks. These include GÉANT2, CBDF, GLIF connections and NRENs. E-Science applications with extreme communication demands will be put into particular test-beds to demonstrate services delivery. The test-bed infrastructure will be available for all interested NRENs, end users and Research and Development projects.

Seamless way for Grid systems to access network resources and Grid middleware extensions to GMPLS protocol

The goal is to develop integration between applications, middleware and transport networks, based on three planes: service plane, NRPS plane and control plane. The service plane will consist of APIs specification for applications, services components exposing network and Grid resources in integrated fashion taking into account policy driven AAA mechanisms. Construction of NRPS plane assumes adaptation of existing NRPSs and full integration with middleware and control plane. The GMPLS control plane will be enriched with Grid extensions providing Grid middleware with access to optical network resources as first-class Grid resources.

Conduct accompanying studies to investigate and evaluate the further technological development of the project outcomes

Supporting studies will be carried out throughout the project in order to ensure a proper future outlook for the project results. Resource management and job scheduling algorithms will be studied, designed and finally tested in a simulation environment developed by the project. They will incorporate issues of network awareness, constraint based routing and advance reservation techniques. Recommendations for the design of an optical control plane will be analyzed and documented.

Disseminate the project experience and outcomes to the targeted actors: NRENs and research users

Partnership with NRENs and end users with highly demanding applications is very welcome. NRENs and their research end users from all over the world are invited in order to share the knowledge and results of PHOSPHORUS project.

Liaise with other European and Global Grid and Networking projects in order to understand and possibly integrate relevant developments

To disseminate ideas and developments the PHOSPHORUS consortium will strongly interact with other relevant programs, research activities and initiatives at the European and international level. Various network-oriented R&D projects are encouraged to share results and exchange ideas with PHOSPHORUS project.

1.3 List of contractors involved

Role		Name	Acronym	Country	Date Enter	Date Exit
1.	CO	Instytut Chemii Bioorganicznej PAN	PSNC	Poland	M1	M33
2.	CR	ADVA Optical Networking	ADVA	Germany	M1	M33
3.	CR	CESNET z.s.p.o.	CESNET	Czech Republic	M1	M33
4.	CR	Nextworks s.r.l.	NXW	Italy	M1	M33
5.	CR	Fraunhofer-Gesellschaft zur Förderung der angewandten Forschung e.V.	FHG	Germany	M1	M33
6.	CR	Fundació i2CAT. Internet i Innovació digital a Catalunya	I2CAT	Spain	M1	M33
7.	CR	Forschungszentrum Juelich GmbH	FZJ	Germany	M1	M33
8.	CR	Hitachi Europe SAS	HEL	France	M1	M33
9.	CR	Interdisciplinair Instituut voor Breedbandtechnologie VZW	IBBT	Belgium	M1	M33
10.	CR	Research Academic Computer Technology Institute	CTI	Greece	M1	M33
11.	CR	Research and Education Society in Information Technologies	AIT	Greece	M1	M30
12.	CR	Stichting Academisch Rekencentrum Amsterdam (SARA) Computing and Networking Services	SARA	Netherlands	M1	M33
13.	CR	SURFnet b.v.	SURFnet	Netherlands	M1	M33
14.	CR	Rheinische Friedrich-Wilhelms-Universitaet Bonn	UniBonn	Germany	M1	M33
15.	CR	University van Amsterdam	UvA	Netherlands	M1	M33
16.	CR	University of Essex	UESSEX	Great Britain	M1	M33
17.	CR	University of Wales Swansea	UWS	Great Britain	M1	M15

18.	**CR**	NORTEL Networks B.V.	NORTEL	USA	M1	M33
19.	**CR**	MCNC	MCNC	USA	M1	M15
20.	**CR**	Communications Research Centre	CRC	Canada	M1	M33
21	**CR**	University of Leeds	UNIVLEEDS	Great Britain	M16	M33

1.4 Results

1.4.1 Network Resource Provisioning Systems for Grid Network Services

Phosphorus Work Package 1 activity aimed to design an architecture and implement a set of interfaces that allow interoperability in a seamless environment between different Network Resource Provisioning Systems (NRPSs) towards the Grid Middleware and the standard GMPLS control plane of the Phosphorus project. In Phase 2, the G2MPLS control plane was also introduced. This has been possible thanks to the implementation of the Network Service Plane (NSP), which is the component responsible for dealing with the NRPSs in order to provide end-to-end paths, manage AAI issues, keep track of the resource usage and to coordinate the different actions done. Morevoer, the Network Service Plane has been enhanced to increase its performance and scalability and it has been integrated with the Internet 2 IDC and GÉANT2 JRA 3 systems by means of making them interoperable.

The partners involved in this work package were: i2CAT (leader), FHG, SURFnet, UniBonn, UvA, UESSEX, NORTEL, and CRC.

NRPSs involved in PHOSPHORUS WP1

ARGON (Allocation and Reservation in Grid-enabled Optic Networks) was developed to manage resources of advanced network equipment as it is present in the German VIOLA test-bed. The advance reservation service of ARGON is able to operate on the GMPLS as well as on the MPLS level. It guarantees the requested level of QoS for applications for the requested time interval. This feature enables a Meta-Scheduling Service to seamlessly integrate the network resources into a Grid environment.

Nortel's DRAC (Dynamic Resource Allocation Controller) was a commercial-grade network abstraction and mediation middleware platform, acting as an agent for network clients (users, applications, compute resource managers) to negotiate and reserve appropriate network resources on their behalf. DRAC uses client's QoS requirements and pre-defined policies to negotiate end-to-end connectivity across heterogeneous domains in support of just-in-time or scheduled computing workflows.

Argia/UCLPv2 provides a network virtualization framework upon which communities of users can build their own services or applications. Articulated Private Networks (APNs) are presented as the first services. APNs can be considered as a next generation Virtual Private Network where a user can create a complex, multi-domain network topology by binding together network resources, time slices, switching nodes and virtual or real routing services.

1.4.1.1 *Key Points and Objectives*

The work done under this work package has focused on the definition, design, and implementation of a new set of east/west southbound/northbound interfaces. The main problems addressed in this work package encompassed:

- Definition, design and implementation of NRPS and GMPLS control plane boundaries, in terms of functionalities and capabilities.
- Development of east/west interfaces for NRPS interoperability, under the implementation of a Network Service Plane
- Development of southbound interfaces between different NRPSs, and the GMPLS control plane.
- Implementation of the Network Service Plane within the Service Layer of the Middleware, which represents/exposes the network resources and grid resources in a seamless environment within the middleware.
- Ability to create point-to-point or point to multi-point connections using resources from several domains in a transparent way. The solution implemented speeds up the creation of complex connections with advance reservations features involving several systems by making them interoperable.
- Simplification of AAI management.

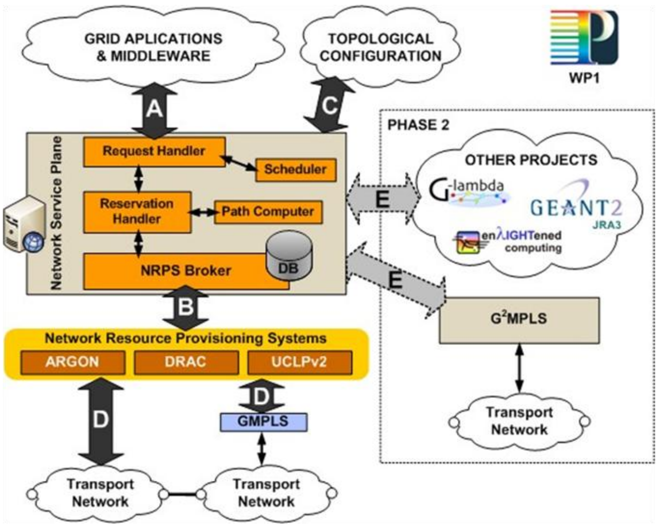

Figure 1-3: Communication interfaces designed and implemented under WS schema.

Work Package 1 work has impacted in overall Phosphorus objectives, defined as following

- **Objective 1:** ”To demonstrate on demand service delivery across access-independent multi-domain / multi-vendor research network test-beds”
- **Objective 2:** “To develop integration between application middleware and transport networks”.
- **Objective 3:** “Supporting studies”
- **Objective 4:** “Dissemination and standardization”
- **Objective 5:** “Liaise with other European and global Grid- and Networking projects”

WP1 Research Objective 1

WP1 main outcome has been the Harmony system. The Harmony system is a multi-domain, multi-vendor, and multi-technology network resource broker with advance reservation features. The definition, design and implementation of the Harmony system, demonstrated in several international events, has contribute to this Phosphorus objective by means of demonstrating the feasibility to provide services across multi-domain, and multi-vendor research network test-beds.

Moreover, Harmony’s test-bed has involved up to ten independent domains, proving the multi-domain capabilities of the system and the NRPS utilized, since each domain involved in the test-bed was composed of different physical equipment.

WP1 Research Objective 2

The Harmony system, apart from controlling multi-domain scenarios, enables the Network Resources in the Grid by means of the Harmony Service Interface. The interface developed has allowed the integration of the Grid middleware and the transport networks, since it is the component of the Harmony system responsible for offering the network resources to the applications in a seamless way. The Harmony system implements a resource co-allocation and scheduling capability (reservation service), able to reduce the probability of resource blocking, and providing inter-domain topology awareness services (topology service).

WP1 Research Objective 3

WP1 has contributed to the objective 3, entitled Supporting studies, by means of the close collaboration maintained with WP5 during the whole project. The main outcome to this objective has been the simulator of the Network Service Plane developed by WP5 jointly with WP1, which enables the simulations of the Network Service Plane in order to simulate their performance and scalability under different scenarios.

WP1 Research Objective 4

WP1 has contributed to this overall objective by means of the presence of the Harmony system in several international events (conferences and workshops). Thus, the Harmony system has been presented in several international events. The complete list of events can be found in Phosphorus Deliverable D7.1.2. Moreover, the Harmony system was demonstrated also in several conferences.

Regarding the standardisation efforts, Harmony has been (and is) present in two main standardisation groups:

- Open Grid Forum Network Service Interface (OGF-NSI) working group.
- General Network Interface (GNI) working group within the GLIF community.

WP1 Research Objective 5

WP1 has contributed pro-actively in order to achieve this overall objective. In this sense, WP1 has launched cooperation agreements between the PHOSPHORUS consortium and the Korea Institute of Science and Technology Information as well as the EU FP7 FEDERICA project.

1.4.1.2 *Starting point of work*

Different available and independent NRPS (Network Resource Provisioning System) system implementations, already operational in Europa and Canada, were considered as starting point (table 1.1) for the WP1 developments. These systems, developed under the scope of International projects, were provided by the project partners. A NRPS is a system that is able to accept e2e reservation requests and establishes intradomain paths between 2 end points.

One of the extra functionalities that Phosphorus required to deal with Grid users was the advance reservation functionality across heterogeneous domains. After some partner implemented by themselves this functionality within their NRPS architecture, since the project proposal Phosphorus did not involve any change within the architectures of the NRPS but on its external interfaces, we started the design of the service layer. ARGIA was one of the modified systems, since initially it did not support this functionality. ARGON already supported it and DRAC could not been modified, so the advance reservation service was provide by the HARMONY system itself. As the NRPS systems work at the intra-domain level, a new high level system capable to perform the inter-domain connections and reservations across different NRPS was needed (the service layer). For this purpose the HARMONY system was designed and built. However, and as one of the goals of HARMONY was its future interconnection with the G^2MPLS system developed in WP2, another NRPS was added and supported by HARMONY. These new NRPS was a standard GMPLS Control Plane. For that purpose the needed interfaces were developed, and the GMPLS control plane that was already deployed along the VIOLA test-bed in Germany was used. This development was very useful towards start preparing HARMONY to deal with the future WP2 outcomes.

When the system architecture of HARMONY was designed, only a prototype of the AUTHOBAN system developed within GN2 was ready. Although AUTOBAHN had some similarities with HARMONY, it did not support advance reservation and was not integrated with the Grid layer to allow the middleware to set up or book network resources automatically. Moreover, AUTOBAHN was a system more complex since it was performing most of the actions that in HARMONY are performed by the NRPS's.

Network Resource Provisioning Systems (NRPS)	
ARGON	The Allocation and Reservation in Grid-enabled Optic Networks system was developed to manage resources of advanced network equipment as it is present in the German VIOLA test-bed. The advance reservation service of ARGON is able to operate on the GMPLS as well as on the MPLS level. It guarantees a certain QoS for applications for the requested time interval. This feature enables a Meta-Scheduling Service to seamlessly integrate the network resources into a Grid environment.
DRAC	The Dynamic Resource Allocation Controller system was developed by NORTEL and it is a commercial-grade network abstraction and mediation middleware platform, acting as an agent for network clients (users, applications, compute resource managers) to negotiate and reserve appropriate network resources on their behalf. DRAC uses client's QoS requirements and pre-defined policies to negotiate end-to-end connectivity across heterogeneous in support of just-in-time or scheduled computing workflows.
UCLP	The User Controlled LightPaths system was developed by CRC, Inocybe, i2CAT and UofO under the CANARIE support. It provides a network virtualization framework upon which communities of users can build their own services or applications. Articulated Private Networks (APN) are presented as the first services. The APN can be considered as a next generation Virtual Private Network where a user can create a complex, multi-domain topology by binding together network resources, time slices, switching nodes and virtual/real routing services.

Table1.1. NRPS systems

Globally, at that time there was no system already integrated with grid applications to dynamically request network path with advance reservation features. Some ideas were also being developed in US and in Japan. These ideas were under the projects enlightened and G-Lambda, which whom Phosphorus has established strong links. Moreover, we also considered the work done under Internet2, so we identified the IDC protocol, a system to provide on demand connections services. Thus, and as the Autobahn system was being defining the interface for interoperability with IDC, the HARMONY system extended its interface to support also IDC and therefore be able to interoperate with Autobahn. This has been achieved by the end of the project, so IDC has become another NRPS supported. It is also important to advertise that not all the NRPS have the same functionalities, so HARMONY accepts a minimum set of the functionalities performed by all system, and needed for Grid applications.

1.4.1.3 *Methodology to reach the objectives*

WP1 was structured in tasks as follows in order to achieve the objectives of the work package:

- Task 1.1 Heterogeneous NRPS interoperability
- Task 1.2 Interoperability of NRPS and GMPLS control plane
- Task 1.3 Integration of the Network Service Plane with the Service Layer of the Middleware
- Task 1.4 Interoperability between NRPS, GMPLS control plane and the service layer
- Task 1.5 Interoperability between the NSP and the G2MPLS Control Plane
- Task 1.6 NSP architecture enhancements
- Task 1.7 Interoperability of the NSP with GÉANT2 JRA 3 and other related projects (Internet2, G-Lambda, EnLIGHTened)

Each task was divided in activities in order to easily achieve the objectives of each task. These activities were leaded by one partner, and there were contributions also from involved partners in each activity. The communication between the members of the work package has been really fluent and there have been around 100 executive video conferences hosted weekly during the whole project, as well as 50 technical video conferences in order to discuss and solve the technical issues occurred during the project. At the same time, communication with the other PHOSPHORUS work packages has been also fluent and often-occurring. So WP1 has contributed pro-actively to the joint activities with all the other PHOSPHORUS work packages (WP2, WP3, WP4, WP5, WP6, and WP7).

The major results achieved by WP1 during the three years of the period can be summarized in:

First year key achievements:

a. Definition, design and development of the common Web Service interface.
b. Definition, design and development of the Adapter architecture (common and NRPS dependent parts).
c. Implementation of the East-West interfaces and the NRPS Adapters for ARGON, DRAC and UCLP.
d. Implementation of the communication between NRPSs and the NSP.
e. .Implementation of a "Thin NRPS" module and a GMPLS driver for integration of GMPLS domains into the WP1 architecture.
f. Implementation of advance reservations in the NRPS (where necessary). Definition, design and implementation of the Northbound Web Service interface and its operations.

g. Complete design and implementation of the internal architecture of the NSP: topology and reservation web services, database, request handlers, path computer, NRPS manager and other modules.

h. Basic AuthN/AuthZ schema.

i. Pre-design of a distributed communication schema for future developments and interoperation with other projects.

j. Demonstration of the full operation of the NSP and the NRPSs

Second year key achievements:

k. Adoption of a branding name for the system: Harmony. Dissemination efforts focused on this name.

l. Release of a stable prototype under the new branding name. Internal entities also have received a branding name for easier dissemination in public events.

m. Enhancement of the performance in overall in the Harmony system, in comparison with the first prototypes released in year 1.

n. Release of administration tools for the Harmony system.

o. Enhanced integration with application middleware from WP3.

p. Enhanced architecture for easing the interoperability with other systems. Enhancement of the Harmony NRPS Adapter architecture (divided now in interface, common NRPS part and NRPS-dependent part).

q. Enhancement of the Thin NRPS for allowing a higher degree of integration between a GMPLS control plane and the Harmony system.

r. Integration of the NRPSs and the GMPLS control plane with the NSP and the Application Middleware (service layer).

s. Migration from the early version of the UCLP NRPS to a stable version called Argia 1.4, provided by I2CAT, Inocybe Technologies (Canada) and CRC.

t. Definition, design and implementation of the advance reservation mechanisms for Argia/UCLP NRPS.

u. Definition, design and implementation of a port Ethernet support in Argia/UCLP NRPS.

v. Definition, design and implementation of the support for Calient DiamondWave FiberConnect switches in Argia/UCLP NRPS.

w. Modularisation of the internals of the Network Service Plane and enhanced code architecture.

x. Implementation of the hierarchical architecture for the NSP. Originally, the NSP was composed only by an Inter-Domain Broker entity. In the system available in year 2, the NSP can have several IDB entities which communicate each other using a hierarchy.

y. Implementation of the distributed architecture for the NSP. Several IDB entities can be configured to act as peers and flood each other’s information about the underlying network for creating an end-to-end path. This configuration eases the communication with external systems that operate in a distributed mode (Géant2’s AutoBAHN, Internet2’s IDC, etc.).

z. Pre-release implementation of the authentication and authorisation infrastructure (AAI) for Harmony under WP4 guidelines using Generalised AAA Toolkit from UvA.

aa. Interface definition to support interoperability between the Harmony system and G^2MPLS from WP2.

bb. Demonstrations of the full interoperation of the NSP and the NRPS/Thin-NRPS-for-GMPLS using the project test-bed in several networking events such as ONDM2008 and TNC2008.

cc. Initiated active collaboration lines with GNI/GUSI initiative from Glif, NSI initiative from Open Grid Forum and the Korean Institute of Science and Technology Information (KISTI).

Third year achievements:

dd. The Harmony system, branding name adopted during the second year, has been disseminated with the efforts focused on this name in several conferences.

ee. The Harmony test-bed has been extended with the integration of three new international domains.

ff. Release of administration tools for the Harmony system.

gg. Cooperation agreement signed with EU FP7 FEDERICA.

hh. Cooperation agreement signed with Korea Institute of Science and Technology Information (KISTI).

ii. Standardisation efforts within the Open Grid Forum Network Service Interface (OGF-NSI) working group.

jj. Standardisation efforts within the Generic Network Interface (GNI) of the GLIF community.

kk. Definition, design and implementation of the support for low-cost Allied Tellesys switches in Argia/UCLP NRPS.

ll. Definition, design and implementation of the multi-technology support in Argia/UCLP and ARGON NRPS. Enhanced integration between Harmony and G^2MPLS.

mm. Performance and Scalability Analysis of the NSP

nn. Simulation of the Network Service Plane.

oo. Support for malleable reservations further tested.

pp. Interoperability with GÉANT2 JRA 3 (AutoBAHN) system.

qq. Interoperability with Internet2 IDC.

rr. Enhanced communication with the Grid Middleware. Security modules have been added in order to achieve secure communication. Agreed licensing scheme for the Harmony prototype. LGPL is the chosen license.

ss. Integration of the Harmony system with the Health Services Virtual Organization (HSVO) / SAVOIR software.

tt. Tested the AAI modules over the virtual test-bed infrastructure.

uu. A. Willner, C. Barz, J.A. García-Espín, J. Ferrer, S. Figuerola, and P. Martini: *Harmony: Advance Reservations in Heterogeneous Multi-domain Environments*. IFIP/TC6 Networking 09.

vv. S. Figuerola, J. A. García-Espín, J. Ferrer, and A. Willner: *Scalability Analysis and Evaluation of the Multi-domain, Optical Network Service Plane in Harmony*. ECOC 09.

ww. S. Figuerola, J.A. García-Espín, J. Ferrer, and A. Willner: *Performance Analysis of Harmony, an Optical, Multi-domain Network Resource Broker*. ICTON 09.

xx. S. Figuerola, J.A. García-Espín, J. Ferrer, H. Zhang, and M. Savoie: *Enabling Network Resources in the Grid: Functionalities and Services in Harmony*. GridNets 09

yy. Harmony *prototype* demonstrated at several international events (SC08, TNC09, or ICT08).

Task 1.1 Heterogeneous NRPS interoperability

The main goal of this task was to perform an analysis of commonalities between different NRPSs, defining requirements for interoperability, and implementing the required interfaces. The task was sub-divided into four activities. The issues addressed within this task consisted mainly in identifying the common features between the heterogeneous NRPS, the study of the interoperability requirements, and finally the definition, design and implementation of a set of east/west interfaces for NRPS interoperability. Moreover, the task involved one demonstration activity. Thus, the outcome of this task was demonstrated at Super Computing 07 event (SC07),

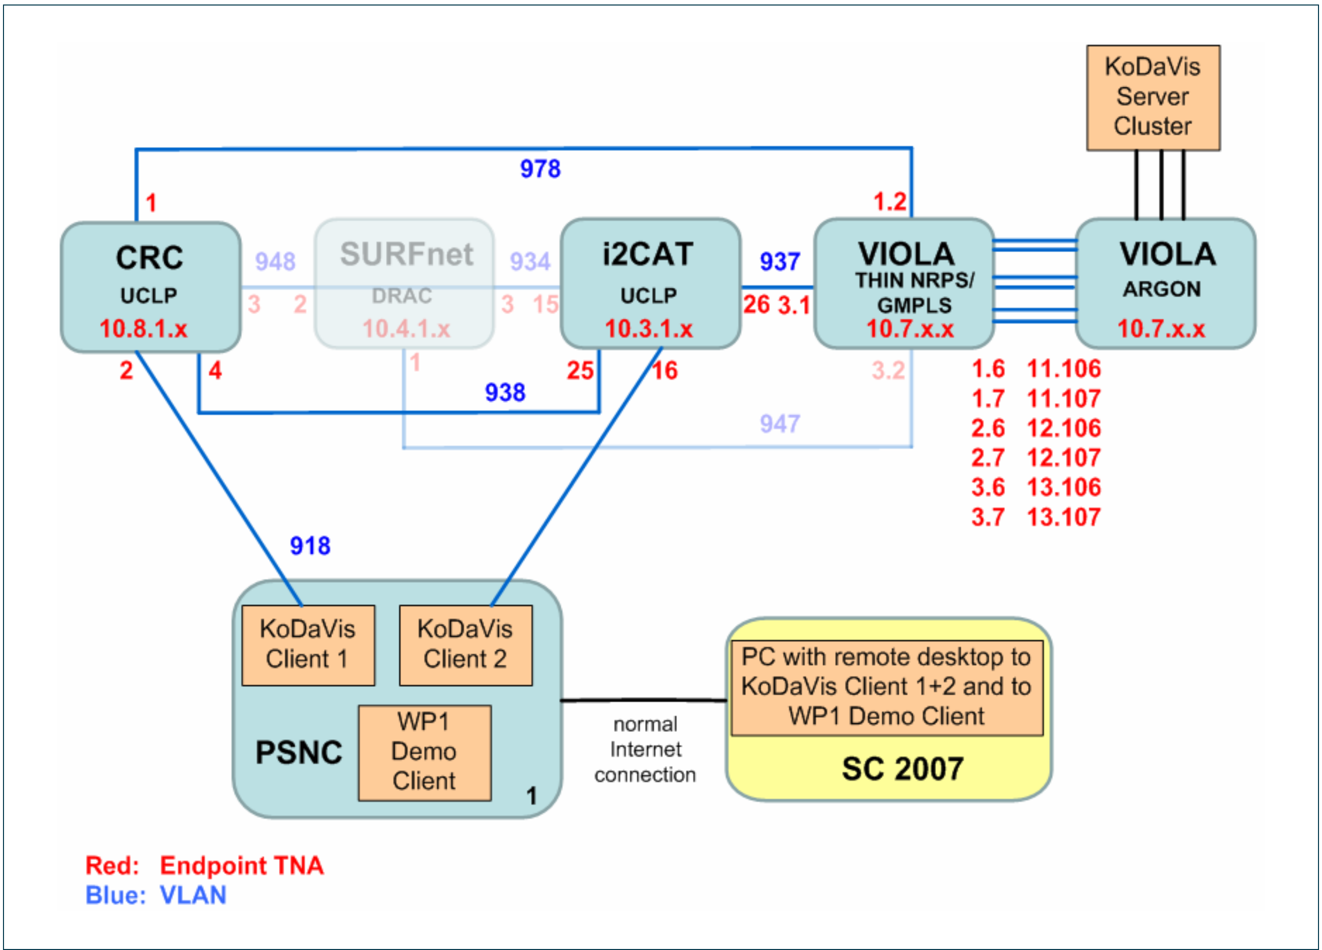

Figure 1-4: Test-bed for SC07 demo

The demonstration consisted in the configuration of inter-domain connections between the tes-beds to establish data paths using the NSP and the NRPSs (ARGON for VIOLA nodes and UCLP for i2CAT and CRC). SURFnet was configured statically since the DRAC NRPS did not participate in the demonstration. Through the establishment of these connections and the use of applications to transmit data through the test bed, the functionalities of the NSP and the NRPSs were shown in the conference venue.

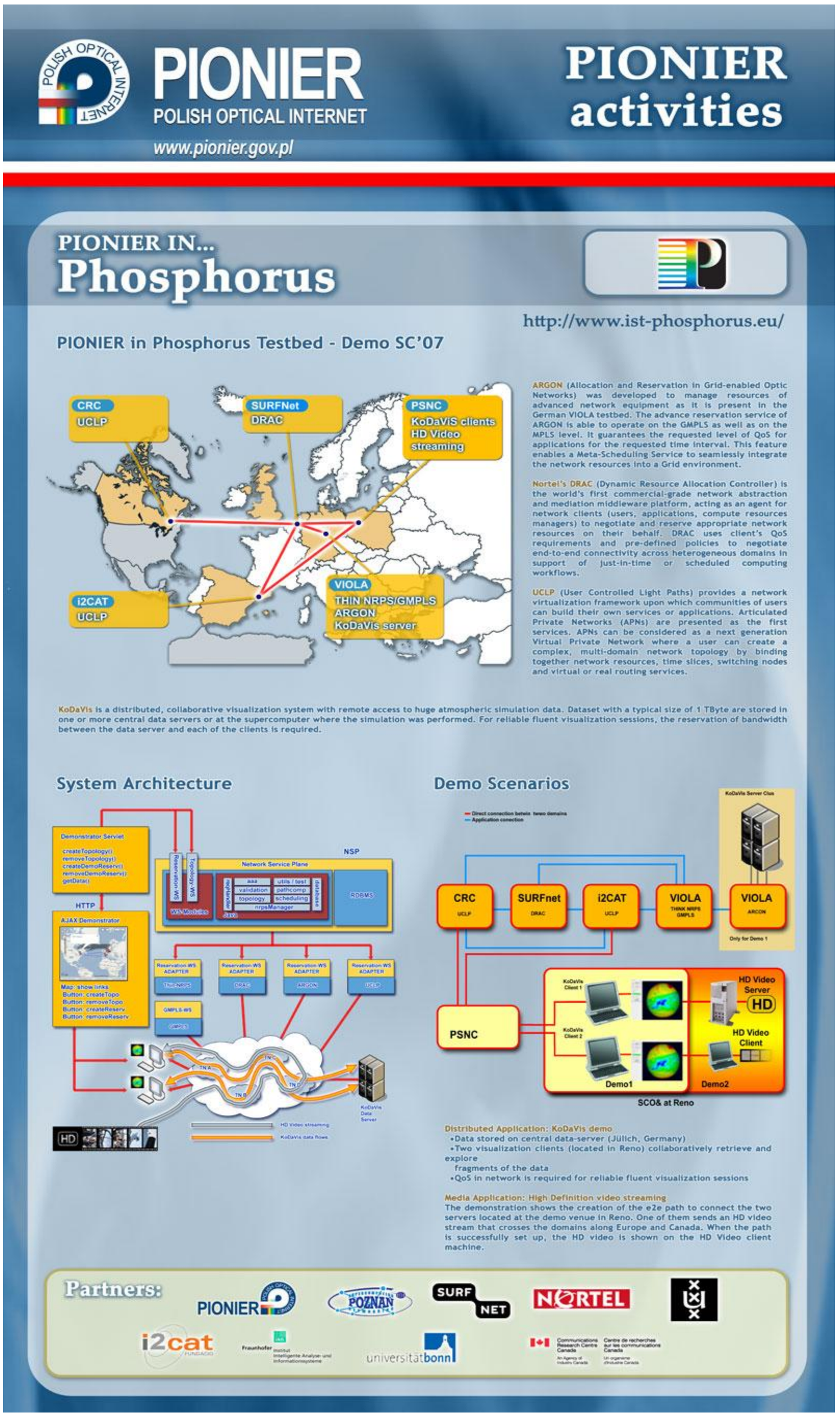

Figure 1-5: Super Computing 2007 poster

Task 1.2 Interoperability of NRPS and GMPLS control plane

After the definition of east/west interfaces for NRPS interoperability in Task 1.1, this task was focused on the southbound interfaces towards the standard GMPLS control plane. Thus, this activity integrated the working version of a GMPLS Control Plane of the VIOLA german test-bed with the NRPSs previously developed interfaces.

This task was divided in two activities, each one focused on the definition of the boundaries between NRPSs and Control Planes and the development of the southbound interfaces for NRPSs respectively. The work done under this task can be briefly summarized by:

- Definition of the requirements for the southbound interfaces to the standard GMPLS control plane
- Analysis and identification of common functionalities and services
- Development of interfaces for interoperability between NRPS and GMPLS under an overlay model
- Definition of use cases to focus on feasible interoperability local demonstrations

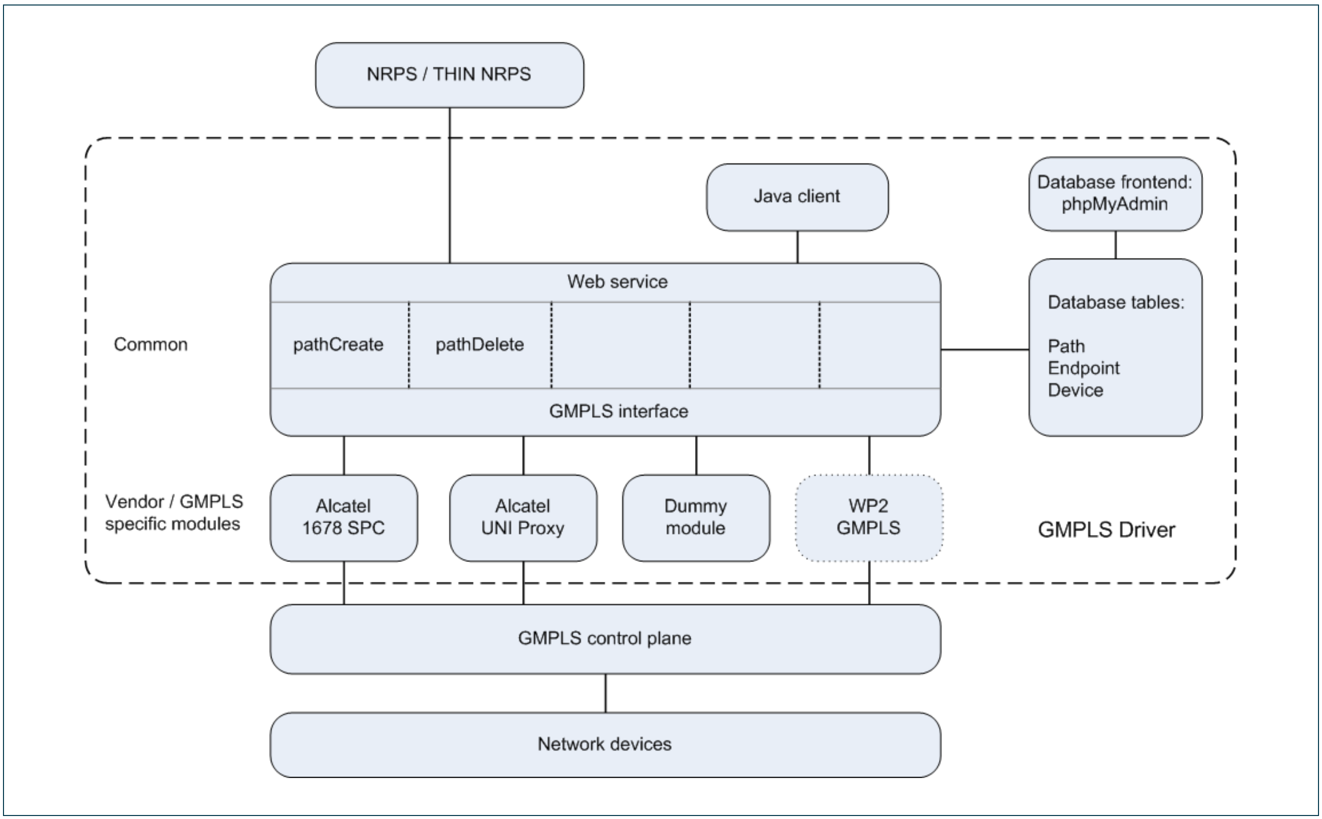

Figure 1-6: Thin NRPS entity implemented for interfacing the GMPLS control plane towards the NSP

Task 1.3 Integration of the Network Service Plane with the Service Layer of the middleware

This task was in charge of the development of the northbound interfaces of the NRPSs within the Network Service Plane with the Service Layer of the Middleware, where Network resources and Grid resources are exposed in a seamless environment. The main goal achieved within this task has been providing the Network Service Plane between heterogeneous NRPS towards the Middleware. This Network Service Plane manages all the NRPSs building a multi-domain scenario and enabling the capability of creating dynamic paths in multi-domain scenarios, one of the main objectives of the project. Moreover, the developments carried out under this task required strong collaboration of WP3 and WP4.

Activities within this task were:

- Activity 1.3.1 – Definition of requirements from the Network Service Plane to the Service Layer of the middleware
- Activity 1.3.2 – Development of northbound interfaces for interoperability of the Network Service Plane to the service Layer
- Activity 1.3.3 – Development of a framework for network resources and Grid resources. The Service Layer and the Meta-Scheduler (MSS)

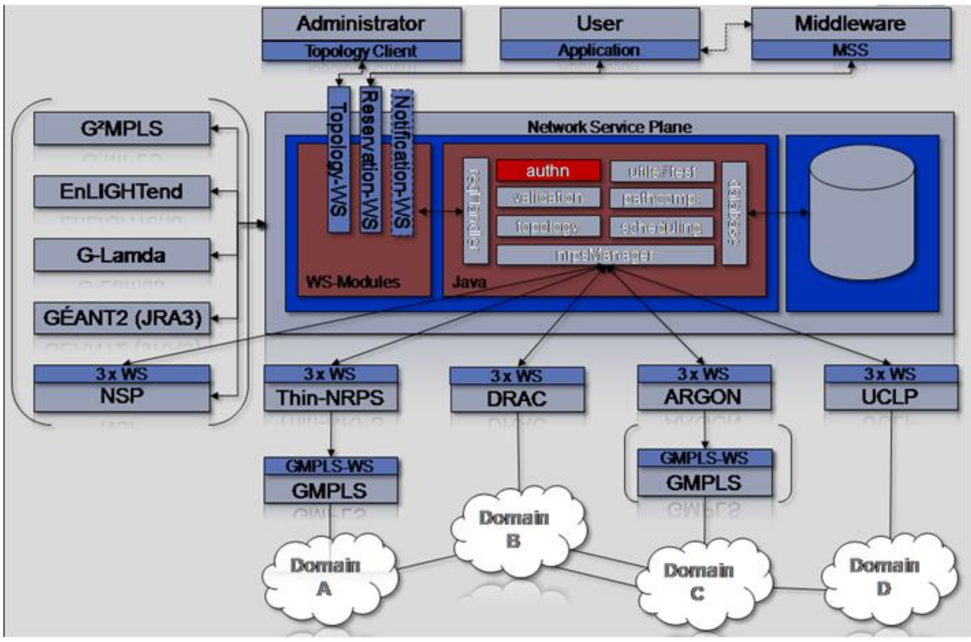

Figure 1-7: Global overview of the internal NSP architecture and interfaces

Task 1.4 Interoperability between NRPS, GMPLS control plane and the Service Layer

The main goal of this task was the integration of the architecture designed, implemented and tested on T1.3 in the whole system in Phosphorus allowing Grid applications to ask the middleware for network resources in advance and let the middleware and its MSS being able to request for the resources to the NSP. This task provided directly input for WP6.

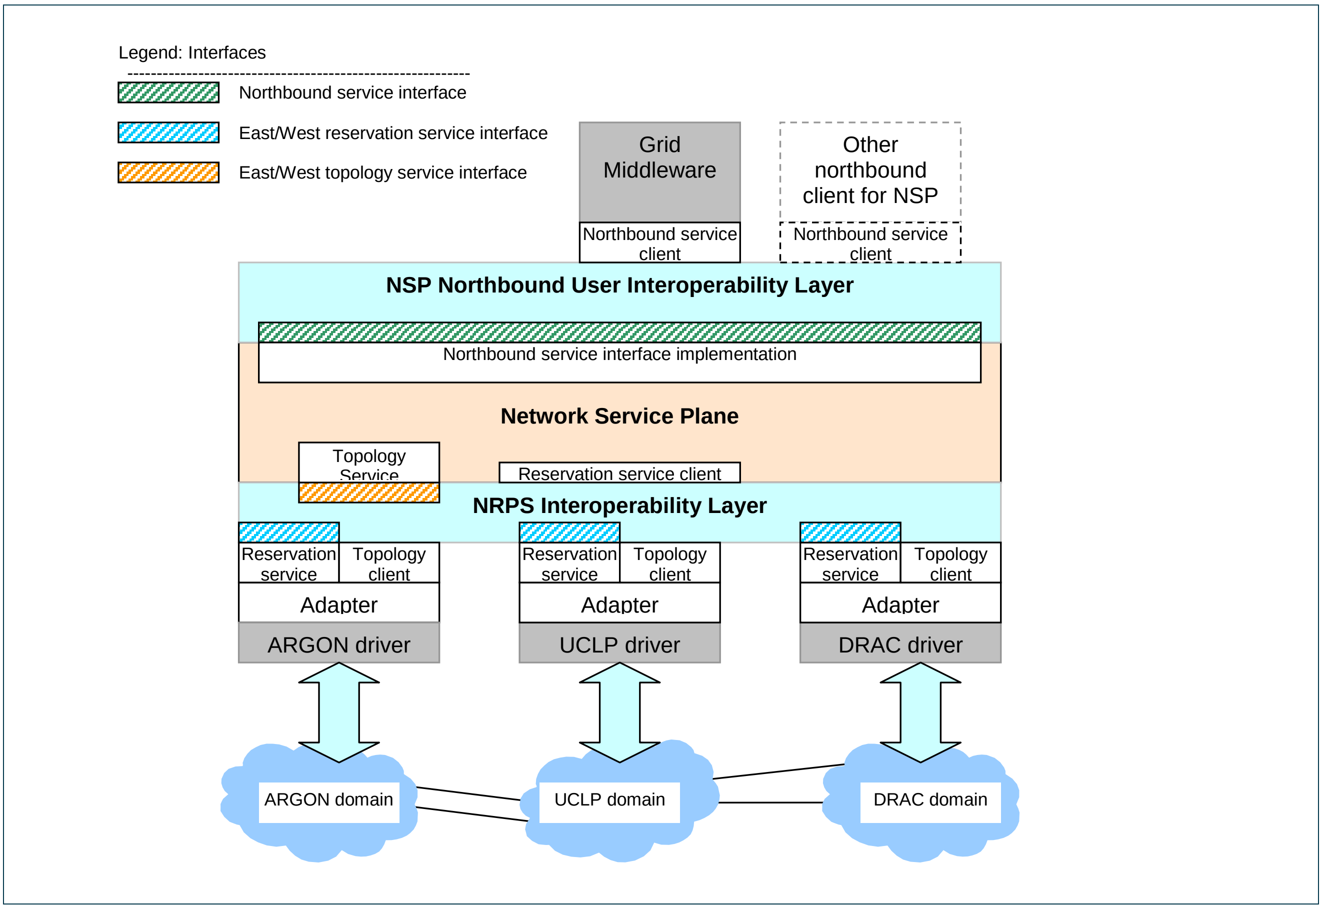

Figure 1-8: Initial design of the architecture for NRPS and GMPLS control plane interoperability.

The task was achieved by means of its two activities. On the one hand, activity 1.4.1, demonstrator interoperation of NRPSs with GMPLS Control Plane. This activity was focused on the demonstrations fo the interoperability. On the other hand, activity 1.4.2, interoperation of the Service Layer, GMPLS and NRPSs. This activity focused on preparing use cases through the platforms provided by the partners in order to demonstrate the functionalities of the Service Layer.

Task 1.5 Interoperability between the NSP and the G^2MPLS Control Plane

This task was focused on the development of the interfaces required to allow interoperability with the G^2MPLS system developed under WP2. Thus, this task has been achieved with strong collaboration of WP2. The interoperability has enabled the system to provision end-to-end paths across G^2MPLS domains and Harmony domains. The task, divided in three activities, has been carried out during two phases. The first phase had as a main outcome the Harmony-to-G^2MPLS-gateway (HG^2-GW); while the second phase mainly consisted of enhancing the HG^2-GW of the first phase.

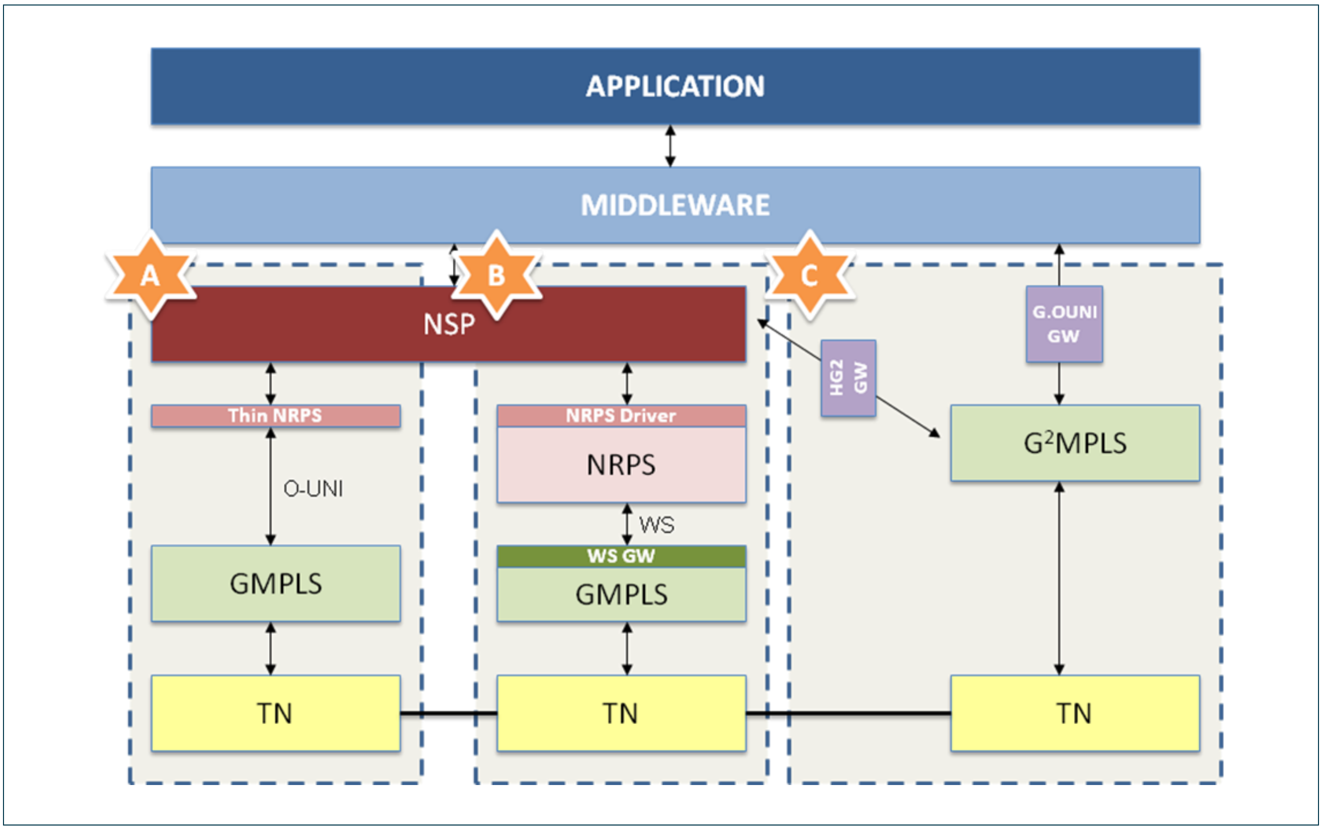

Figure 1-9: GMPLS/G²MPLS interworking with Harmony.

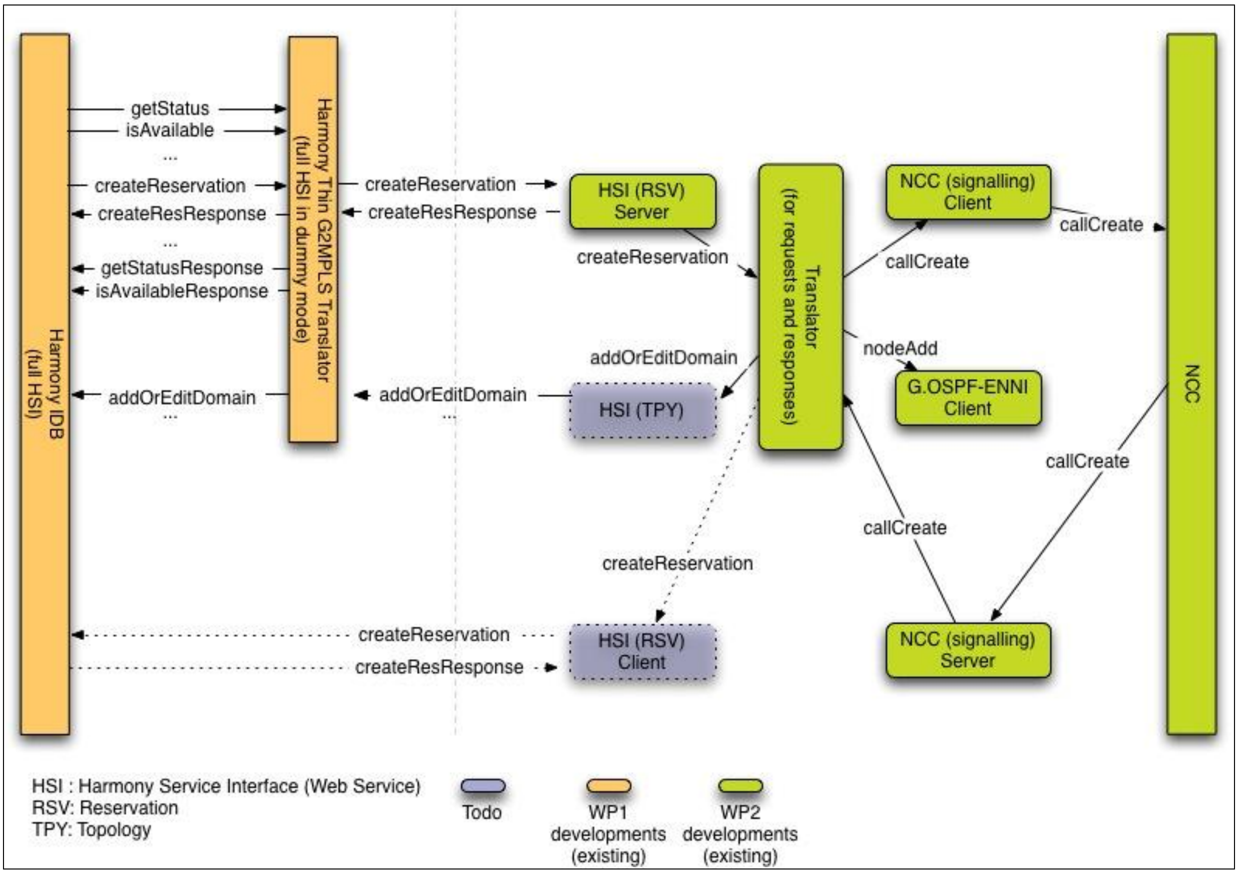

Figure 1-10: Two modules used in the translator architecture

Task 1.6 NSP architecture enhancements

This task was focused on enhancing the early-prototype of the Network Service Plane. The enhancements respond to the need of increasing the NSP flexibility and interoperability with other related BoD projects. The activity was split in five activities, each one focusing in one enhancing line (i.e definition, design and implementation of the distributed service plane operating mode, support for malleable reservations, bandwidth management, or emulation and performance analysis of the service plane). The AAI modules were also included in the Network Service Plane in order to allow the users or applications to communicate with the service plane in a secure way.

Moreover, within this task, WP1 has started one collaboration line with another EU FP7 project: the FEDERICA project. The collaboration started in order to use FEDERICA resources to deploy the Harmony system working over emulated NRPSs and then evaluate the performance and scalability of the Network Service Plane under different loads and scenarios.

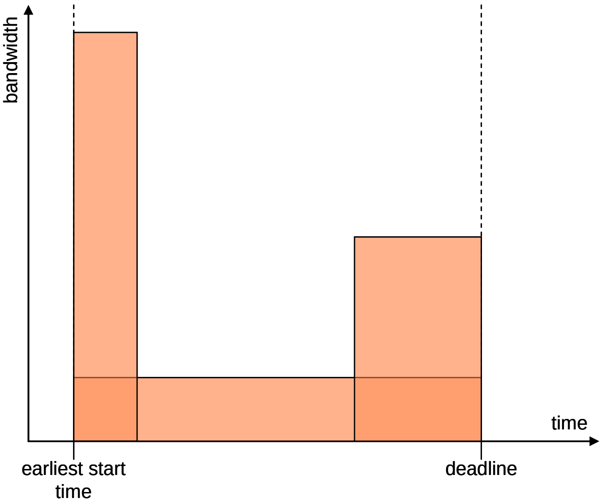

Figure 1-11: Three different malleable reservations

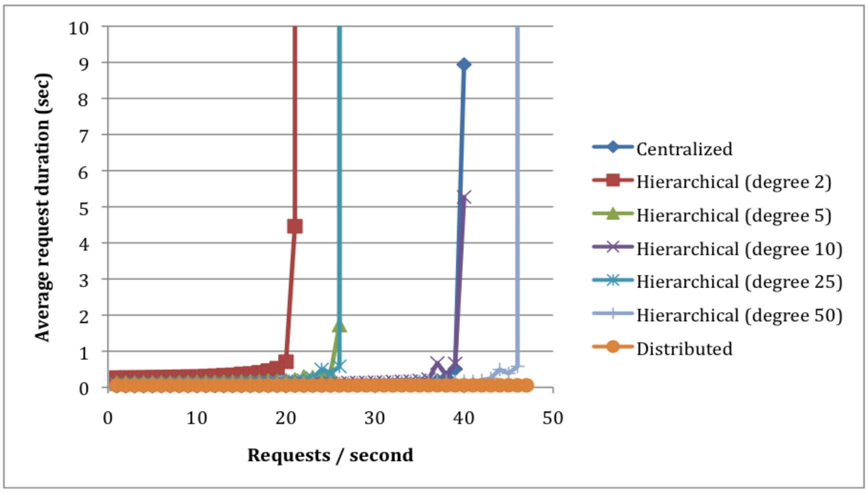

Figure 1-12: Simulation of distinct architecture results: Performance of reservation requests.

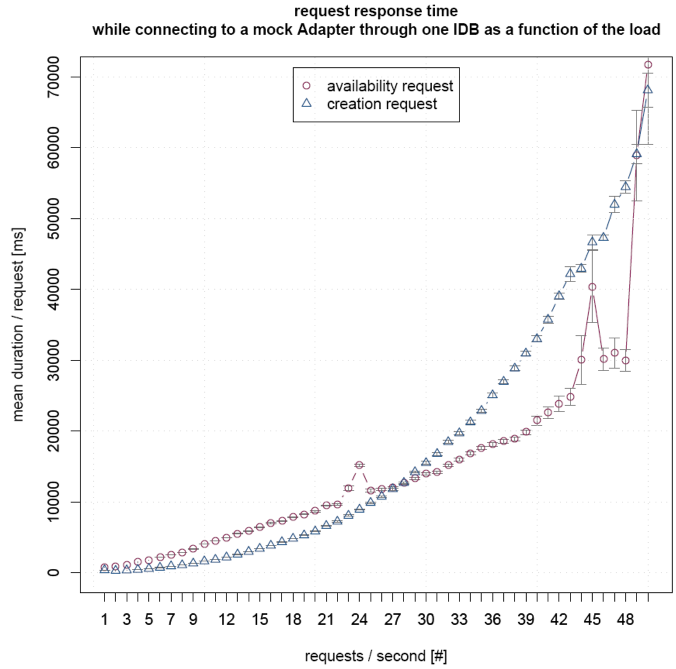

Figure 1-13: Performance tests results. Request response time.

Task 1.7 Interoperability of the NSP with GÉANT2 JRA 3 and other related projects (Internet2, EnLIGHTened, G-Lambda)

The project consortium declared its willingness to co-operate with other research projects. The work done under the umbrella of this task, divided in three activities, consisted of making the Harmony system interoperable with two related projects within the research community: GÉANT2 JRA 3 system and the Internet2 IDC system.

Moreover, within this activity the Harmony test-bed has been extended to Korea, Poland and the United Kingdom. The extension to Korea has been achieved by means of signing a cooperation agreement between the PHOSPHORUS consortium and the Korea Institute of Science and Technology Information (KISTI).

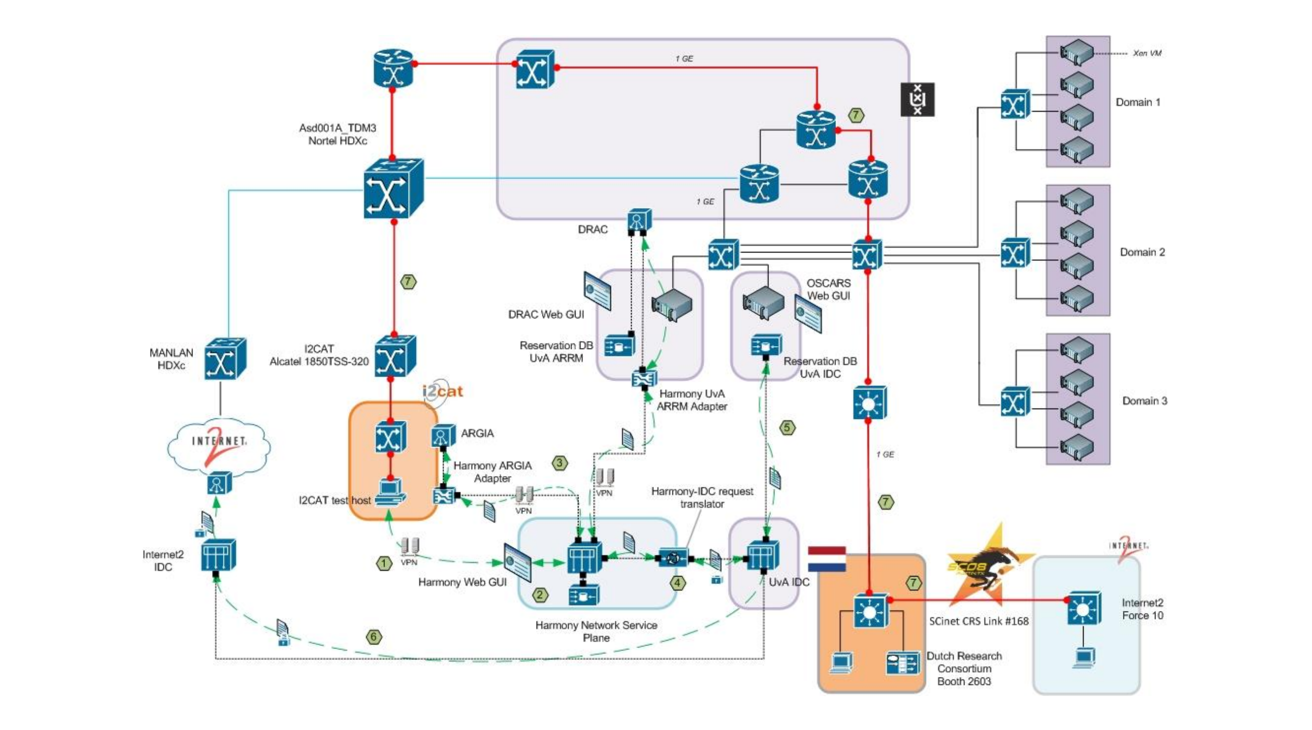

Figure 1-14: Test-bed set up for Internet2 IDC – Harmony SC08 joint demonstration

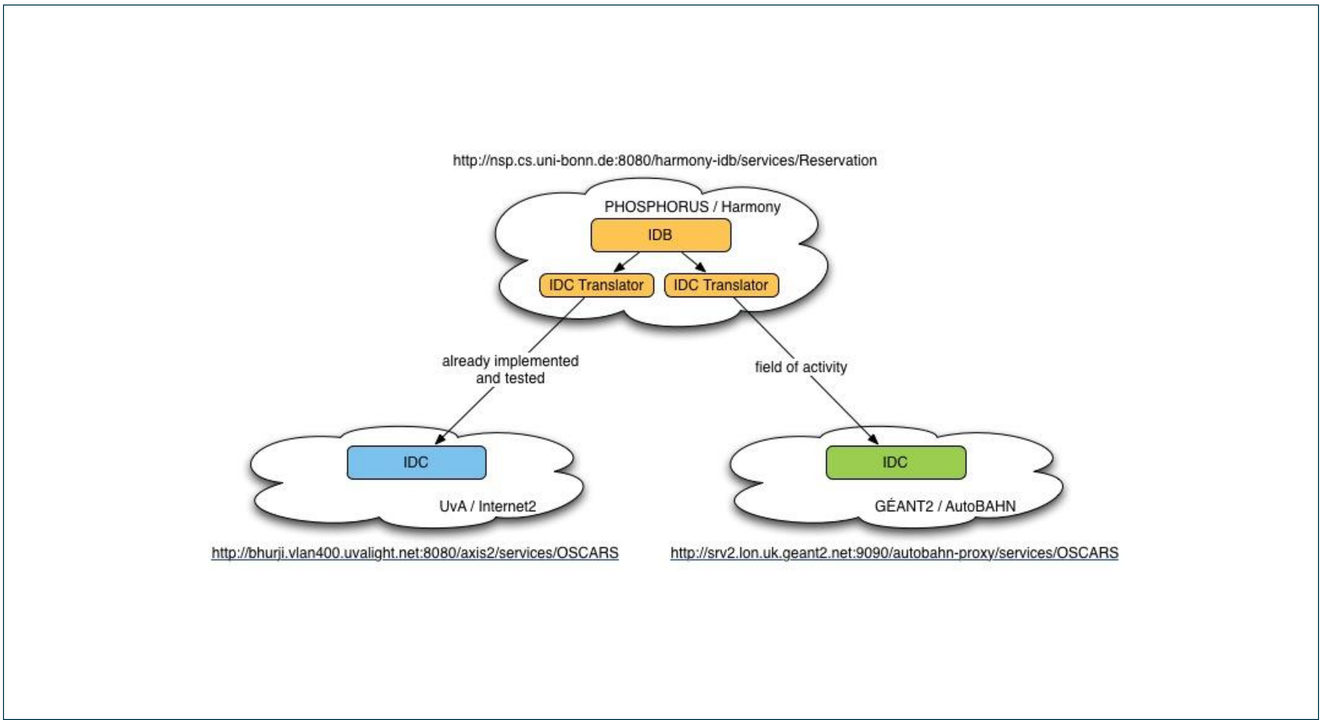

Figure 1-15: Abstract overview of the Harmony/AutoBAHN test-bed.

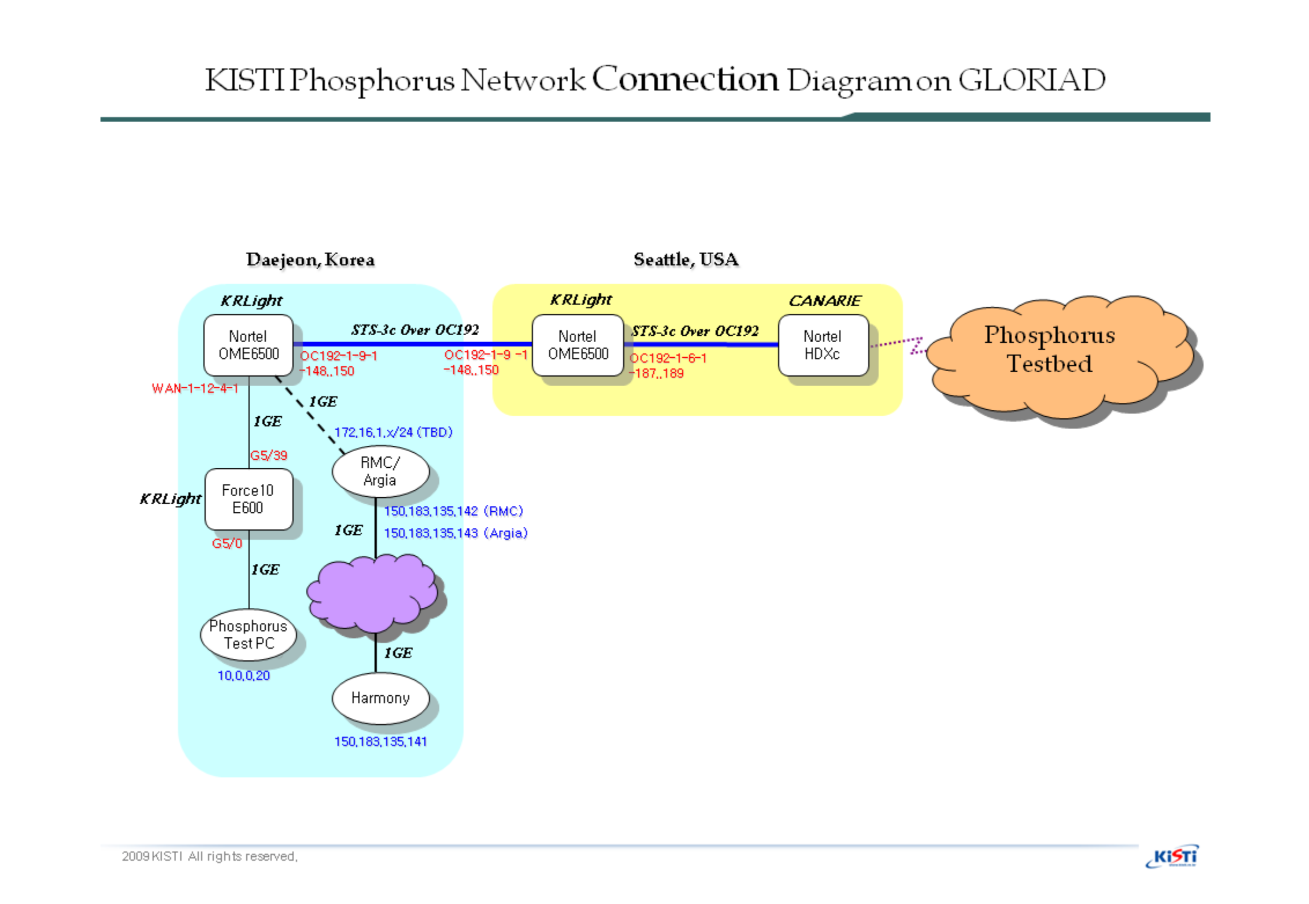

Figure 1-16: KISTI local test-bed topology

1.4.1.4 *Impacts on state-of-the-art in Research Networking*

The HARMONY system developed within the WP1 of the Phosphorus project is an innovative system for the research and the Grid community.

Taking into consideration these initiatives, WP1 has developed a whole new system that has taken advantage of current NRPS (although with generic adapters in order to support new NRPS). Due to the acceptance it has received and the valid tests/demonstrations performed so far, some of the work of the HIS (Harmony Service Interface) is being taken into consideration within the NSI (Network Service Interface) Working Group of the OGF (Open Grid Forum).

As a conclusion, the WP1 has provided to the community an open source tool that allows Grid applications to request on demand and in advance network reservations services, with specific connectivity bandwidth across different heterogeneous domains. The main impact it has had on the state of the art is:

- An open source system to request e2e connections to different heterogeneous NRPS

- A system that offers advance reservation functionalities
- A system integrated with the Grid middleware that allows automatic request of connectivity service from the application
- A Service interface that is being proposed within standardization groups (OGF)
- A real system which is being currently operative on the Phosphorus Testbed

HARMONY is a system that is ready to support the Grid research community and can facilitate their the access to network resources that nowadays is provided in manual basis.

1.4.1.5 *Exploitable results achieved*

According to the WP1 development plan, the interface specification and implementation of the Harmony system has been the main exploitable result achieved. Thus, the main Harmony's features achieved during the development are:

- Network Service Plane implementation with flexible, configurable operating modes (centralised, hierarchical, distributed, among others)
- Inter-Domain Brokers and path computing elements for optical network resources
- Advance reservations for Grid applications.
- Grid Middleware integration with the NRPS or GMPLS.
- Topology abstraction and sharing among (administrative) domains.
- Secure intra-domain topology hiding, using resource abstraction techniques.
- Authentication and Authorisation Infrastructure integration with NRPS and GMPLS.
- WSDL description (OASIS WSRF v1.2 compliant) of the HSI for flexibility, portability and easy adaptation to new systems and/or applications.

Due to all these features, Harmony's test-bed has grown from a first basic scenario, involving five domains of four different countries to the current scenario, where there are ten domains of seven different countries. This growth shows how Harmony can connect independent domains (with distinct technology equipment) which are using any of the above mentioned NRPS and enable domain interoperability to the users or Grid middleware applications.

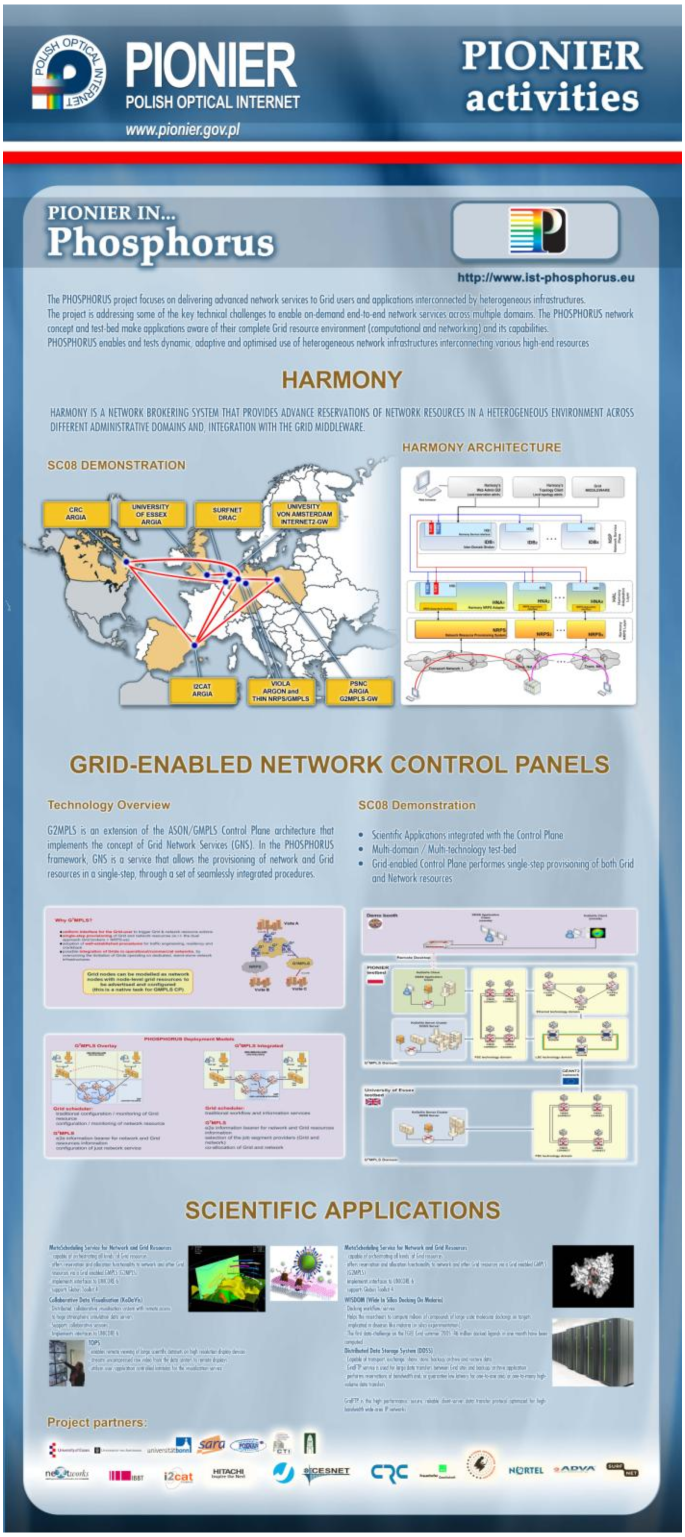

Figure 1-17: SC08 Poster

1.4.1.6 *Future exploitable plans*

National Research and Education Networks (NRENs) represent the most potential group of users of the Harmony system, since NREN are the ones which are using multi-technology, multi-vendor and heterogeneous local test-beds. The future exploitable results should focus on deploying Harmony in NREN's local test-beds, since Harmony fills the gap due to distinct technologies presents in the different test-beds. Moreover, any e-Science application which requires to transmit huge amounts of data during one determined interval of time can use Harmony's services to achieve their purposes.

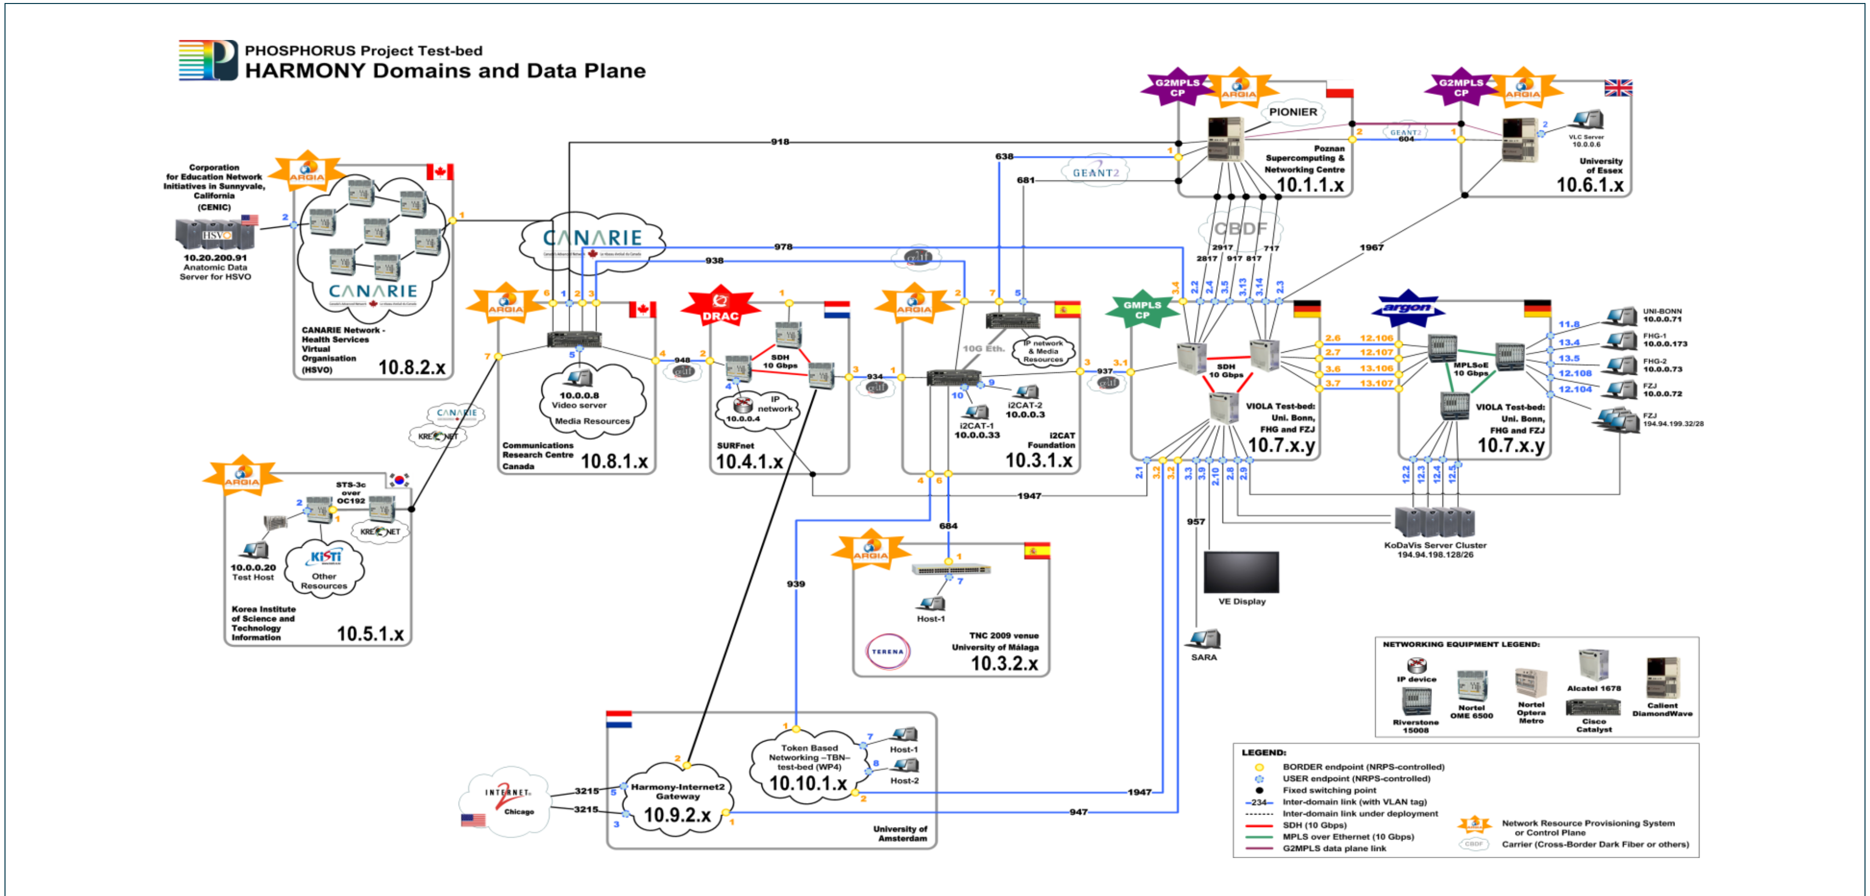

Figure 1-18: Harmony test-bed at the end of the project

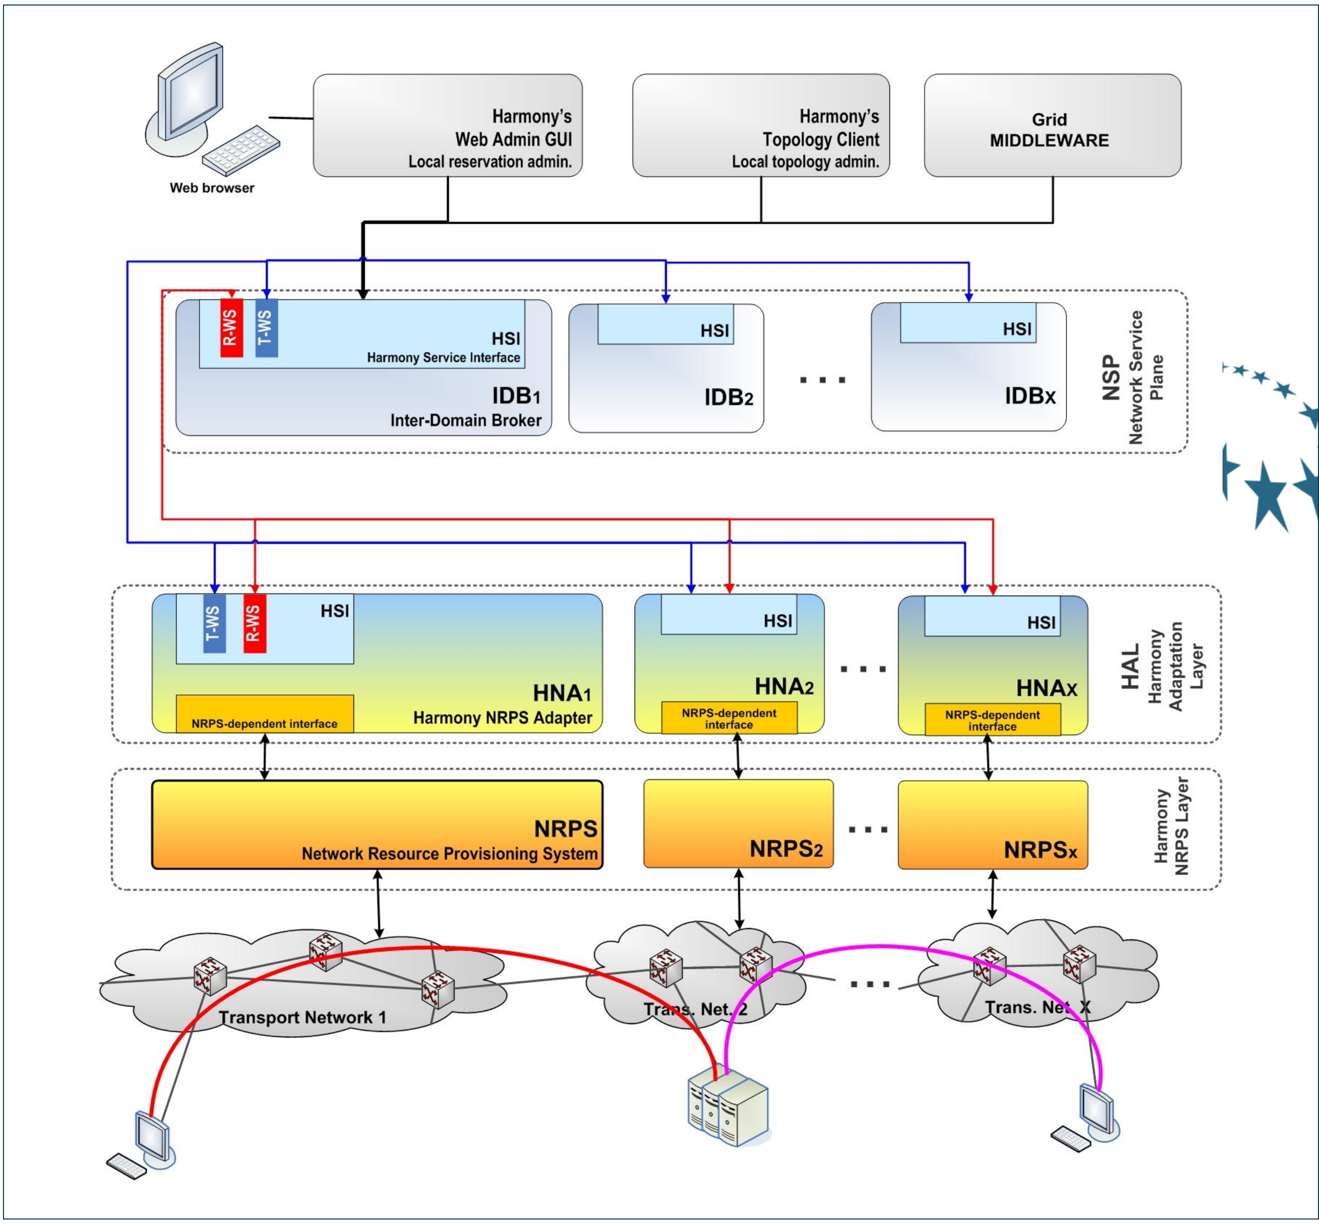

Figure 1-19: Harmony overall architecture

1.4.2 Enhancements to the GMPLS Control Plane for Grid Network Services

Partners involved in this workpackage: NXW (leader), PSNC, ADVA, FHG, i2CAT, AIT, UESSEX, UNIVLEEDS

Work package 2 worked on the specification and development of mechanisms and procedures to implement Grid Network Services (GNS) through an enhanced GMPLS Control Plane, called Grid-GMPLS (G^2MPLS). In the PHOSPHORUS project and in WP2 in particular, the GNS is a service that allows the provisioning of network and Grid resources in a single-step, though a set of seamlessly integrated procedures. G^2MPLS extends the ASON/GMPLS Control Plane architecture to provide part of the functionalities related to the joint selection, co-allocation and maintenance of both Grid and network resources.

The partners involved in this workpackage were: NXW (leader), PSNC, ADVA, FHG, i2CAT, AIT, UESSEX and UNIVLEEDS.

WP2 focused on the G^2MPLS architectural aspects and protocols, the software design and development and the early assessment of functionalities in prototypes interfaced to real optical networking hardware. Complementary WP2 activities were:

- the support to the large-scale G^2MPLS validation activities on the pan-European testbed operated by WP6, including the public demonstrations;
- the interfacing of G^2MPLS Control Plane with the Grid Middleware layer (WP3)
- the interfacing of G^2MPLS Control Plane with the Harmony system (WP1)
- the interfacing of G^2MPLS Control Plane with the Grid-AAA layer (WP4);
- the support to WP5 for simulations/studies on the scalability of some G^2MPLS solutions.

The resulting G^2MPLS Control Plane consists of a set of design documents, prototypes and open source code, all publicly delivered to the R&D community.

1.4.2.1 *WP2 Objectives*

The WP2 work focused at first on the architecture, including protocol extensions and network interfaces specifications, then on the fast-development of the relevant software components, their integration into prototypes and consolidation for testbed activities. As represented in **Figure 1-20**, WP2 directly contributed to most of the Phosphorus macro-objectives and, in particular, to:

- **Objective 1**, ”To demonstrate on demand service delivery across access-independent multi-domain / multi-vendor research network test-beds”
- **Objective 2**, “To develop integration between application middleware and transport networks”.

WP2 casted these Phosphorus macro-objectives into specific research objectives and, subsequently, into a list of measurable outcomes, all described in the following.

WP2 research objective 1

One of the research objectives of WP2 was to define the G²MPLS enhanced services (i.e. GNS) that integrate and complement the standard ASON/GMPLS network services, i.e. automatic setup and resiliency of network connections across the Transport Plane. Therefore, WP2 work did not result in an application-specific architecture and G²MPLS supports any kind of endpoint applications by providing network transport services and procedures that can fall back to the standard ASON/GMPLS ones. This compliance fostered for the possible integration of Grids in real operational networks, by overcoming the current limitation of Grids operating as stand-alone networks with their own administrative ownership and procedures. From a user's perspective, G²MPLS enabled a real node-to-node deployment of on-demand Grid services, while for network operators (particularly the NRENs) it resulted in a means for the integration of Grids and automated network control plane technologies in real operational infrastructures.

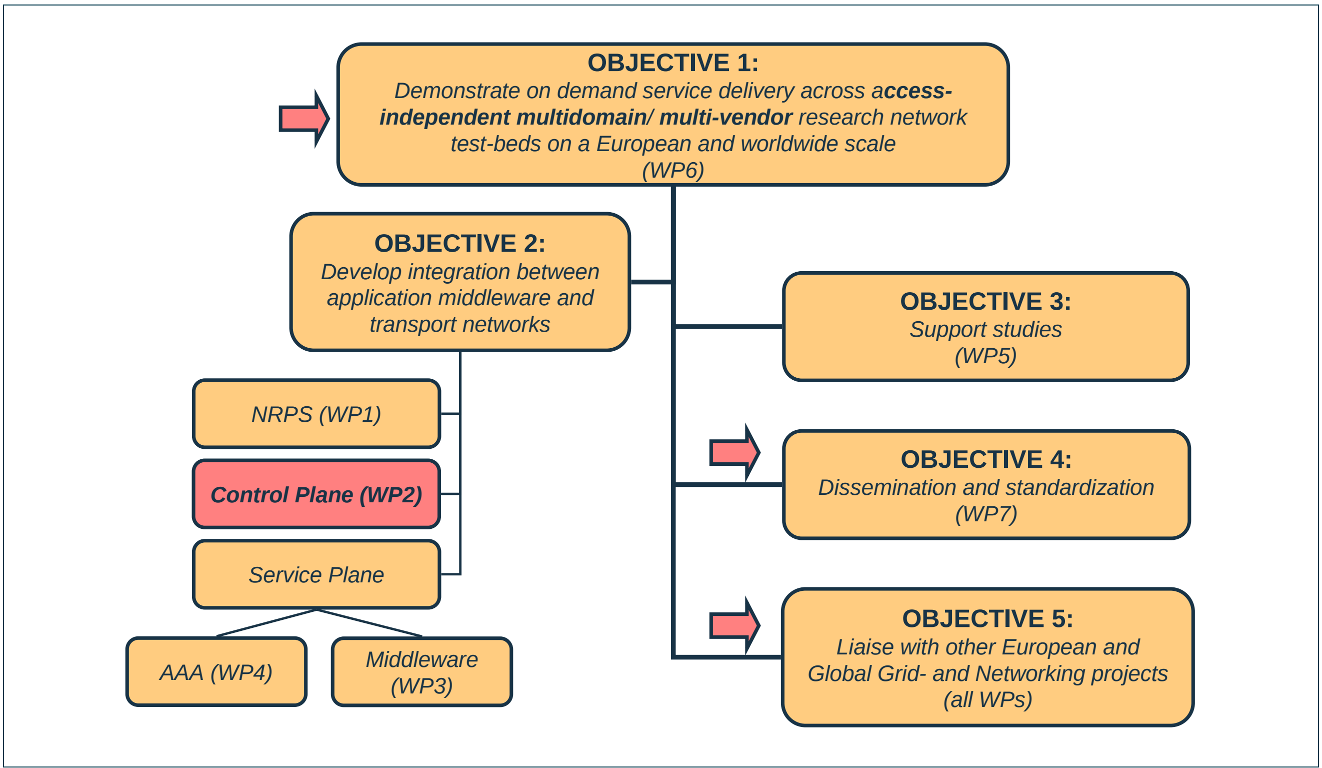

Figure 1-20: Workpackage 2 contribution to the Phosphorus objectives.

WP2 research objective 2

Another objective of WP2 was to provide a flexible Control Plane solution that can adapt to different set of requirements. To this purpose, two Control Plane models were defined for the G²MPLS architecture: the

Phosphorus Overlay and the Phosphorus Integrated. These models concerned the layering of Grid and network resources and had a different meaning and scope with respect to the IETF definitions for Overlay, Augmented and Peer/Integrated models.

WP2 research objective 3

Phosphorus test-bed was an initial complete framework for the deployment and assessment of G²MPLS NCP, because it was structured, technologically heterogeneous and multi-domain. Moreover, the test-bed comprised other non-GMPLS control/management technologies (i.e. Network Resource Provisioning Systems).

A further WP2 objective was to define procedures for and demonstrate the interaction and cooperation of G²MPLS with these technologies. To this purpose, the G²MPLS compliance to the ASON/GMPLS network reference points model (i.e. the inheritance and extensions of the ASON network interfaces, UNI, E-NNI) simplified the work, because it allowed both peer-style interactions (i.e. through E-NNI) and overlay-style relationships (i.e. through UNI). Moreover, the interoperation through well-established network interfaces allowed a generalized approach and simplified also the integration and interoperation.

1.4.2.2 *Starting point of work*

The ASON/GMPLS architecture as defined by ITU-T SG15 was assumed as the main reference and starting point for the WP2 work. The OIF Implementation Agreements defined by the Architecture and Signaling Working Group and the IETF RFCs and IDs by the CCAMP and MPLS Working Groups were the basic references for the specification of the network reference points, both in terms of protocols and procedures (OIF for UNI and E-NNI and IETF for I-NNI, respectively).

In addition to that, the G^2MPLS design was also based on the standardization activities on Grids carried out by Open Grid Forum through its working groups and research groups, such as the Job Submission Description Language WG (jsdl-wg), the Glue Schema Working Group (glue-wg), the Grid High-Performance Networking RG (ghpn-rg) and the Network Service Interface Working Group (nsi-wg).

Concerning the implementation activities, the development of G^2MPLS controllers leveraged on the existence of a number of open source code trees, like

- the OSPFv2 with MPLS-TE extensions from the QUAGGA routing suite (http://www.quagga.net/),
- the KOM-RSVP with MPLS-TE extensions from IST-FP5-M3I project,
- the ATLANTIS RSVP with MPLS-TE and DiffServ-TE extensions from IST-FP5-TEQUILA project,
- the RSVP_AGENT by PSNC for OIF UNI-Client from the IST-FP6-MUPBED project,
- the DRAGON GMPLS stack.
- the omniORB CORBA ORB for C++ and Python (http://omniorb.sourceforge.net/)

Most of these code trees were not natively designed for GMPLS purposes, except for the PSNC RSVP_AGENT and DRAGON GMPLS. However, also these ones represented a limited starting point of work, as detailed below.

Concerning the RSVP_AGENT by PSNC, it implemented just the OIF UNI-Client signalling, which is only a minimal part of a fully-fledged GMPLS RSVP-TE.

The DRAGON GMPLS, instead, provides an Open Source implementation of the IETF GMPLS, comprising mainly the core routing and signalling functionalities for the I-NNI. The DRAGON stack implemented a peer-model for the network connections between equipments of different technologies (e.g. Ethernet switches and optical cross-connects) and spanning different domains. However, DRAGON did not provide any dynamic and automated handling of the dualism between transport network and the signalling network: the availability of this feature could simplify and make flexible the GMPLS controller operation, which instead requires a lot of preliminary static configuration. All these architectural aspects raised major limitations for the adoption of the DRAGON code as the baseline for the WP2 development activities. The choice to modify up to the foundation an IETF GMPLS code (DRAGON) to make it evolve to ASON/GMPLS and G^2MPLS (PHOSPHORUS) was considered more effort-consuming than the choice of a clean non-GMPLS baseline, because the same number of additional modules would need to be added from scratch (e.g. for network interfaces, call and recovery management) but in a context in which their presence was not foreseen.

Finally, all these code trees (except for QUAGGA) originated during other research projects just ended or approaching their ending at the time of WP2 kick-off, and, thus, were not considered a maintained solid reference for the WP2 development activities.

Figure 1-21: OS code trees in the framework of G^2MPLS developments.

1.4.2.3 *Methodology to reach the objectives*

WP2 was structured in two tasks to accomplish its research challenges:

- Task 2.1 – "GMPLS enhancements for Grid Network Services", which was successfully completed with the release of the preliminary G²MPLS prototype stack (D2.5, M22).
- Task 2.2 – "Interfacing the GMPLS control plane with Grid", which completed the interfacing and development work by delivering two further releases of G^{2}MPLS prototypes, the D2.10 (M27) and the D2.11 (M33).

Each task was further split into activities to deliver a coherent outcome. Each activity was coordinated by a leading partner, appointed according to his experience on the specific theme, and overall supervised by the WP leader (NXW). All the work was based on the intense and enthusiastic cooperation between the participating partners, with many meetings and periodic conference calls to align on concepts and design choices, and to track the activity progress. A broader involvement of the seven WP2 partners (NXW, PSNC, UESSEX, ADVA, FHG, AIT, UNIVLEEDS) occurred during the architectural specification phases. On the contrary, the developing activities were carried out by a restricted team of software developers (NXW, PSNC, UESSEX) to keep a narrower focus on the challenging features to be developed. The G^{2}MPLS development team established direct and continuous communications among its members, with a dedicated mailing list and also by interconnecting the three laboratories permanently through the Internet. Software development responsibilities were distributed as follows:

- NXW, responsible for overall stack design and maintenance, the basic GMPLS skeleton modules, the full G^{2}MPLS RSVP-TE implementation (from scratch) at the three network interfaces (G².I-NNI, G².E-NNI and G².OUNI), the Call and Connection Controllers, the G2.PCERA, the TNRC abstract part (equipment-independent)
- PSNC, responsible for G^{2}MPLS OSPF-TE extensions at the three network interfaces (G².I-NNI, G².E-NNI and G².OUNI), the TNRC plugins to Avda, Calient and Allied Telesys equipments, and the maintenance of the Software Configuration Management platform (SVN repository)
- UESSEX, responsible for the G^{2}MPLS gateway modules, one at the G.OUNI and another one towards Harmony.

The inter-workpackage discussions with WP1, WP3, WP4, WP5 and WP6 were numerous and greatly contributed to the work in Task 2.2. The WP leader acted as the main reference point in these discussions, supported in the specific activities with a target WP by those partners working in both WPs. In particular, UESSEX acted as bridging partner for WP1 and WP6, PSNC for WP3 and WP6, FHG for WP3, AIT for WP5.

The major results achieved by WP2 in Phase-I (M1-M18) can be briefly summarized in:

- A set of high-quality and comprehensive documents defining the G^{2}MPLS Control Plane; these documents, as a whole, provided the specifications of:
  - The G^{2}MPLS architecture and functional decomposition (D2.1),

  - The definition of its network reference points (D2.7),
  - The high-level system design and software architecture of the G²MPLS controllers (D2.3)
  - The specific signalling and routing protocol extensions in support of them (D2.2), and
  - The models and scenarios for the coexistence and deployment of GMPLS and G^2MPLS network sections in NREN transport network infrastructures (D2.6)
- Design and development of an initial but fundamental set of G^2MPLS modules:
  - Design and implementation of G^2MPLS stack skeleton modules (e.g. LRM, SCN GW, etc.)
  - Design and implementation of the G^2MPLS RSVP-TE, from scratch, at the three network reference points (G².I-NNI, G².E-NNI and G².OUNI)
  - Design and implementation of G^2MPLS extensions for OSPF-TE at the three network reference points (G².I-NNI, G².E-NNI and G².OUNI)
  - Design and implementation of G.OUNI (-N and -C) modules and utilities, and adaptation of protocols (RSVP and OSPF)
  - The software layer (TNRC) that allowed the abstraction of Data Plane resources and the interoperation with legacy hardware (Adva ROADMs and Calient PXCs)
  - The G²MPLS Call and Recovery layer (G².CCC, G².NCC and G².RC)
  - The G²MPLS Path Computation Routing Algorithm(G².PCERA routing engine)

The major results achieved in Phase-II (M19-M33) can be summarized in:

- The completion of G²MPLS specifications with:
  - The design of the interworking of G²MPLS with the Phosphorus GAAA TK (D2.8)
  - The design of the interoperation of G²MPLS with Harmony (D2.9)
- The completion and consolidation of the G²MPLS software components:
  - The integration and release of the preliminary G²MPLS prototype (D2.5), including the 4 types of controllers (Edge, Core, Border, G.OUNI-C).
  - The AAA logic and G² PEP Gateway to interwork with the GAAA TK delivered by WP4
  - The Harmony-G²MPLS Gateway (HG²-GW)
  - The development of further TNRC plugins for the interoperation of G^2MPLS with Ethernet switches
  - The development of resiliency mechanisms for the Control Plane components (both signalling and routing)
  - Extra modules and auxiliary tools for the Control Plane status visualization and simplified testing.
- Functional tests on a real network test-bed controlled by G²MPLS (planning and results, D2.4)
- The release of three “all-in-bundle” G^2MPLS prototypes (D2.5, D2.10 and D2.11)

Task 2.1 focused on the G^2MPLS core functionalities, including routing and signalling protocol extensions, recovery and path computation issues, flexible interfacing with Transport Plane equipments, both in terms of specification, design and implementation. It was split in seven activities briefly reported in the following.

Activity 2.1.1 – System design for the Grid-GMPLS Control Plane

A2.1.1 was the key activity for the whole WP2, because it provided the overall specification of the G^2MPLS architecture. A2.1.1 had the main purpose to produce the deliverable D2.1 ("The Grid-GMPLS Control Plane architecture"), which defined the G^2MPLS reference architecture and functional decomposition of G^2MPLS controllers, and collected all the requirements, design principles and solutions identified for the G^2MPLS implementation. Therefore, D2.1 represented the cornerstone for any further specification of G^2MPLS modules, in terms of both functionalities and software design.

Activity 2.1.2 – Routing protocol extensions

A2.1.2 started the translation of G^2MPLS functionalities and procedures (identified in A2.1.1) into low-level specifications and, then, software implementation. To this purpose, the team used the architectural reference provided by D2.1 and complemented it with the background expertise on grid middleware contributed by some of the involved partners (e.g. FHG-IAIS, PSNC and UESSEX in particular).

The main focus in A2.1.2 was on Grid resources modelling and description (e.g. CPU, storage) and on the specification of the mechanisms for flooding this information throughout the Grid-GMPLS network.

For what concerned Grid resources modelling, WP2 chose the GLUE Schema (Grid Laboratory Uniform Environment), because it provided a hierarchical structure for modelling Grid sites, their supported services and resources. The GLUE information model was evaluated flexible and middleware-independent, further supported by its increasing consensus in the Grid research community (e.g. EGEE, OSG, NorduGrid, OMII-UK, OMII-EU, UNICORE, NGS). The G^2MPLS routing extensions were obtained by mapping the relevant GLUE schema elements and attributes into appropriate objects to be transported in Opaque Link State Advertisements of the OSPF protocol (Grid LSA). The specified routing protocol extensions were part of the milestone M2.1 and were reported in deliverable D2.2.

Activity 2.1.3 – Signalling protocols extensions

A2.1.3 was one of the major activities in the WP2 workplan in terms of low-level design and software implementation. In fact, the signalling protocol is the core enabler of any automatic network service and GNS setup, as well as the basic template for the implementation of the signalling at the network interfaces external to a G^2MPLS domain (G.OUNI, G.E-NNI). This activity continued the fruitful cooperation with grid middleware experts in the PHOSPHORUS project.

The main focus in A2.1.3 was on Grid job requests modelling/description, and on the specification and implementation of the signalling procedures and object extensions that are needed to implement the GNS concept. The OGF Job Submission Description Language working group (JSDL) is a standardized language for the submission of jobs among organizations using a variety of local job management systems. It was selected by the team because of its consolidated position as a core vocabulary for describing job submissions in a number of existing operational systems such as: Condor, Globus Toolkit, LSF (Load Sharing Facility), PBS (Portable Batch System), SGE (Sun GridEngine), UNICORE (Uniform Interface to Computing Resources). The identification of the G^2MPLS signalling extensions was obtained by identifying the JSDL elements relevant for GNS description and then by mapping these attributes into specific objects of the signalling protocol. The new RSVP objects were part of the milestone M2.2 and were reported in deliverable D2.2.

Activity 2.1.4 – Development of a Recovery Controller and a Call Controller

A2.1.4 added a key layer to the G²MPLS stacks: the Call and Recovery control layer. This layer implemented the ASON "intelligence" to the whole stack, and was a fundamental piece of its architecture. It contributed to WP2 with a large development effort, and with design documents incorporated in D2.3.

The Call Controller and the relative positioning of the Recovery Controller were designed fully compliant with the ASON architecture and the ITU-T's rec. G.7713/Y.1704, and also undertook some basic recommendations in IETF RFC 4974. Of course, they incorporated all the planned Grid extensions, as per D2.1 and D2.2.

The G²MPLS Recovery Controller (G².RC) allowed the association of LSPs into Recovery Bundles (a.k.a. "connections" in ASON terminology, or "tunnels" in IETF terminology) for recovery or restoration purposes. It was the bridge between the Call signalling layer and the LSP signalling layer. Its design allowed the easy introduction of new recovery strategies, or of the reworking of the existing ones.

Both the Call Controller and the Recovery Controller basically coordinated the whole G²MPLS signalling actions and implemented a decoupled signalling behaviour that particularly fitted the multi-domain environment of NRENs.

Activity 2.1.5 – Full-optical enhancements to the Path Computation Engines in the GMPLS Control Plane

A2.15 dealt with the design and software development of the GNS routing decision entity for the G^2MPLS stack. The activity got a major input from D2.1and D2.2, and consisted of a large development effort that was summarized with the design notes included in D2.3 and the prototype module integrated in D2.5.

G^2.PCE-RA was the main consumer of the grid and network TE information that is flooded by the G.OSPF-TE protocols at the different network interfaces. As such, it stored the global view of the G^2MPLS network topology, including the multi-domain TE information and the grid sites with their own resources (services, computing elements, sub-clusters, storage elements).

G2.PCE-RA was the producer of all the route computations and/or other routing responses (e.g. TNA resolution) in favour of the signalling entities of the G2MPLS stack (NCC and G2.RSVP-TE). The design choice made by WP2 development was to support in G2.PCE-RA the computation of pairs of disjoint routes in the topology between the ingress and egress TNAs specified, by following the Two Step Algorithm (TSA).

Activity 2.1.6 – Interoperation with legacy GMPLS equipment

This activity was aimed at producing the G^2MPLS Transport Network Resource Controller (TNRC).

The TNRC was one of the most complex software modules in the G²MPLS architecture. It consisted of an Abstract Part (TNRC-AP) to represent any kind of Transport Network equipment with any kind of switching capability and its related operations (e.g. cross-connections or resource alarms notifications), and a Specific Part (TNRC-SP).

The complexity of the TNRC-AP lied in its generic role: it had to translate between the G²MPLS-level resources and the TN-level ones; it had to buffer and bridge between the timings and behaviours of G²MPLS operations and those that needed to be sent to the equipment box in order to control it (e.g. some atomic actions at the G²MPLS level needs to be serialized into multiple commands to the equipment; their completion need to be coordinated or roll-backed in case of failure, etc.)

The TNRC-AP was designed to attach software plugins with a well-defined but flexible API. These plugins were the TNRC Specific Parts (TNRC-SPs), which were specialized interfaces for a specific device model and switching capability (if multiple ones are supported by the same equipment box). The developed TNRC SPs included:

- A generic stub for LSC and FSC switching capabilities, that allows to run a test-bed of G²MPLS controllers without any physical equipment sitting under them
- An LSC plugin for the Adva FSP3000 ROADM, that talk to the equipment box via its TL-1 interface, plus an Adva FSP3000 emulator (which allows to run it detached from the real equipment)
- An FSC plugin for the Calient DiamondWave FiberConnect switch that talk to the equipment box via its TL-1 interface.
- Plugins for L2 Ethernet switches from various vendors.

The results of this activity fed D2.3 in terms of design specifications and D2.5 in terms of software modules composing the preliminary G²MPLS prototype.

Activity 2.1.7 – Functional tests

A2.1.7 produced a testing framework for the integration and validation of the G^2MPLS Control Plane software modules. This activity mainly focused on the assessment of the correct behaviour of the G^2MPLS stack in terms of:

- protocol procedures,
- management/control of the Network and Grid information,
- robustness and performance of the communications with Transport Plane equipments
- support of all the G^2MPLS designed features

The main contributors to this activity were the WP2 developers (NXW; PSNC, UESSEX) and the G^2MPLS testbed owners (PSNC and UESSEX). The testing activity mostly benefited from the background know-how of the cited partners in terms of network operations and from three software integration meetings, in PSNC during June 2008, in UESSEX during August 2008 and in NXW-PSNC during September 2008. Moreover, fundamental contribution to the activity towards the achievement of its milestone (M2.4) derived from the continuous testing in developers' laboratories, remotely accessed by the whole WP2 developers' team. The scope of G²MPLS functional tests was limited to the laboratory environment, i.e. internal to the WP2 and the project, in the perspective of a self-assessment of the produced developments.

A2.1.7 directly contributed to the D2.4 for what concerns reporting and D2.5 for the preliminary G²MPLS Control Plane public prototype.

Task 2.2 focused on the interfacing issues of the G²MPLS Control Plane with two orthogonal approaches:

- in the northbound direction, interfaces were identified and specified towards the Grid Middleware (mediated by the G.OUNI) or any Network Resource Provisioning System, which in this case behaves like a decision entity in the Management Plane and uses the G²MPLS downgraded to standard GMPLS;
- in the east-west direction, interfaces were specified towards external domains of similar (i.e. G²MPLS) or different control technologies (e.g. Phosphorus Network Service Plane or GN2-BoD Grid-enhanced overlay model).

Task 2.2 was split in five activities briefly reported in the following.

Activity 2.2.1 –Integration/co-existence issues between the different GMPLS Control Plane models

A2.2.1 aimed to complement the architectural design produced in A2.1.1 with the analysis of deployment models and coexistence issues for G²MPLS and GMPLS Control Planes.

The main focus of this activity was on network scenarios, procedures and functionalities of GMPLS and G²MPLS that could attract NRENs' interest on these Control Plane architectures. The participation of many relevant EU NRENs in the PHOSPHORUS consortium simplified the work in this activity, which was also enriched with the submission of public questionnaires to European NRENs and HPC centres. The questionnaire for NRENs was circulated to the wider NREN community during the TERENA European Future Networking Initiative Workshop, held on February 22nd 2007 in Amsterdam. The questionnaire for HPCs was directly circulated by the Project Coordinator to the most relevant Grid community channels. Answers were received by CESNET2, PIONER, FCCN, DFN, GARR, HEAnet, Barcelona Supercomputing Center and CINECA. Questionnaires contained positive positioning mainly by NRENs who are currently missing an automatic control/management system for network and Grid resources or have test-beds focused on that (PIONER, CESNET2, GARR, DFN). Furthermore, a preliminary conditional consensus was received from NRENs and HPCs external to PHOSPHORUS consortium.

The results of this activity converged into deliverable D2.6 ("Deployment Models and Solutions of the Grid-GMPLS Control Plane").

Activity 2.2.2 – Grid-GMPLS interfaces and related Grid-based semantics

A2.2.2 dealt with the network interfaces external to a G^2MPLS domain, i.e. the G.OUNI and the G.E-NNI, as well as the G.I-NNI, both in terms of specification and developments. It received major inputs from the overall Task 2.1 in terms of architectural design and software modules/protocols, while it contributes to WP2 (and the project at large) interface specifications and software prototypes.

One main focus of A2.2.2 was the specification and implementation of the Grid Optical User Network Interface (G.OUNI), i.e. the interface between the Grid users and the GMPLS optical transport network. The G.OUNI provided flexible bandwidth allocation, scheduled services (i.e. advance reservations), AAA and security, automatic service discovery, fault detection/protection/restoration, propagation of service and agreement related events, traffic classification/grooming/shaping.

Other focuses were on the design and development of the Grid External Network-to-Network Interface (G.E-NNI) and the high-level formalization of the Grid Internal Node-to-Node Interface (G.I-NNI).

The interfaces' specifications elaborated on the protocol extensions (G.OSPF-TE and G.RSVP-TE) defined in D2.2 in order to obtain a formal and more abstract description converged in deliverable D2.7. This activity also fed D2.3 in terms of design specifications and D2.5 in terms of software modules composing the preliminary G^2MPLS prototype.

Activity 2.2.3 – Integration of Grid-GMPLS Control Plane in Phosphorus Grid AAI

This activity merged the results of all the previous architectural, specification and software design work of WP2 with those related to the GAAA infrastructure in WP4. For this reason, specifying the interworking between G^2MPLS and the GAAA TK required a number of WP2-WP4 inter-workpackage discussions, both remotely and in face-to-face meetings.

These discussions started from defining the correct usage scenario of the GAAA TK by the G^2MPLS domains, including some basic terminology and definitions alignment; e.g. security domains vs. G^2MPLS network domains; authorization processes at the GAAA TK vs. the Call setup verification procedures in G^2MPLS (according to G.7713/Y.1704). The next step was to define the interaction points and methods by which G^2MPLS could access the authorization facilities of the GAAA TK; this step required small adaptations on both sides. The final step was to sort out the software issues, i.e. integrating software modules written by different workpackages, in different languages, with different architectures in mind.

This work procedure concluded successfully with the release of an enhancement to the preliminary G^2MPLS prototypes, with the ability to perform service authorization through the WP4 GAAA infrastructure.

The implementation consisted of a G^2MPLS PEP Gateway (G^2 PEP GW) that performed authorization actions using the GAAA TK Java libraries developed in WP4, and of all the authorization logic at the G^2.CCC and G^2.NCC (inspired by ITU-T's rec. G.7713/Y.1704). The G^2 PEP GW is integrated within the Call Controllers architecture via XML RPC.

The results of this activity fed D2.8 in terms of specifications and design and D2.10/D2.11 for what concerns prototypes.

Activity 2.2.4 – GMPLS integration with Grid NRPS-es (UCLP, DRAC, ARGON)

This activity, similarly to A2.2.3, required deep inter-workpackage interactions, with WP1 (A1.5.2 in particular). The thinking and planning of an interoperability solution between G²MPLS and Harmony began much before the official start date of A2.2.4, when the Harmony and G²MPLS architectures and developments were still evolving to a stable condition. This early start helped to finalize the two architectures in order to taken into account the future needs of interoperability. E.g. in the G²MPLS case, the G².NCC has been equipped with an East-West CORBA interface that allows to easily plug a G.E-NNI gateway module that can serve as a bridge to the Harmony world.

The core discussion and developments in A2.2.4 focused on the most complex problem of East-West interfacing G²MPLS and Harmony (i.e. peering in an inter-domain fashion), involving the exchange of both signalling and routing information.

The peering at signalling level was more challenging. The starting assumption on service granularity is that Harmony service requests needed to be mapped into G² Call signalling, and vice versa. However, Harmony services worked with a 1-tier signalling fashion (i.e. service request blocking until a response is received), whereas G²MPLS (according to ASON and GMPLS) uses 3-tiers. For this reason, the role of the Harmony-G²MPLS Gateway (HG²-GW) was not just to perform semantic translations, but also to perform a stateful buffering of Harmony single signalling tier into G²MPLS multiple ones.

The development of the HG²-GW was a task that jointly undertaken by A2.2.4 and A1.5.2, since it was a common and compatible objective. A careful and commonly achieved software design allowed to split the coding work into pieces that could be executed either in A2.2.4 or A1.5.2.

The results of this activity fed D2.9 in terms of specifications and design and D2.10/D2.11 for the prototype parts.

Activity 2.2.5 – Consolidation of the Grid-GMPLS Control Plane

This activity carried out the final consolidation of the G^{2}MPLS Control Plane based on the validation feedbacks from the experimental activities in the Phosphorus test-bed (WP6) and the recommendations provided by the Phosphorus supporting studies (WP5).

Consolidation was focused on fixing those software components (network nodes, adapters, gateways, etc.) developed during the Task 2.1 and most of the activities in Task 2.2. No major revision of the WP2 design choices was needed, but just some fixings and further improvements of the produced software modules.

The activity was undertook by the same partners previously involved in the WP2 developments (NXW, PSNC, UESSEX) and released two G^{2}MPLS Control Plane prototypes, namely the D2.10 at M27 and the D2.11 at M33.

D2.10 collected and integrated all the G^{2}MPLS developments produced in the first two years of Phosphorus execution. D2.11 fixed and extended D2.10 with new further functionalities and support tools, such as:

- strategies and modules for the G^2MPLS Control Plane resilience, in order to preserve data traffic in case of any Control Plane failure (e.g. G^2MPLS controller fault, manual restart, etc.)
- control and configuration of carrier-class L2 Ethernet switches
- tools for Control Plane visualization, in which the status of network topology (e.g. TE Links) and the network transport services installed on them (i.e. Calls and LSPs) was shown
- tools for the easy access to the G2MPLS G.UNI, in the form of a set of shell tools capable of dialling and terminating G^2MPLS Calls related to GNS.

Among the consolidation activities the team worked also on the production of G^2MPLS operation guidelines, which were included in D2.11 to ease the take-up and use of the G^2MPLS Control Plane by the wider R&D community.

A complementary activity carried out by the WP2 team in the framework of A2.2.5 was also the preparation, setup and operations of the G^2MPLS public demonstrations with WP6. The public G^2MPLS demonstrations were held at SC'08 (Austin – Texas, USA), ICT'08 (Lyon, FR) and TNC'09 (Malaga, ES) with a great success of audience and follow-ups.

1.4.2.4 *Impacts on state-of-the-art in Research Networking*

G^2MPLS as delivered by WP2 is an operational and mostly open Network Control Plane solution for implementing Grid Network Services. Its permanent deployment in the multi-domain Phosphorus testbed represented an innovative research networking framework for trialling advanced Grids integrated with the latest optical technologies (e.g. ROADMs). Extremely challenging concepts have been investigated and validated through the G^2MPLS Control Plane, such as:

- the Transport Network abstraction and unified description (TNRC), with the support of many switching capabilities like fibers (FSC), lambdas (LSC), L2 switching
- the multi-domain bandwidth-on-demand via E-NNI,
- the generalized integration of Grid applications and middlewares with Network Control Plane (GNS),
- the seamless AAA chain composition
- the flexible interfacing with other BoD systems (e.g. Harmony).

Multiple key factors contributed to the achievement of this major result.

At first, we had the planning and delivery of a solid knowledge base on the Grid-enabled GMPLS Control Plane, including coherently both specifications and prototypes. It followed a process of incremental improvements from the high-level (architecture) functionalities down to the low-level details (software design and prototypes). The synoptic graph in **Figure 1-22** shows the relationships between the G^2MPLS deliverables and could be used also as a reading guide for the G^2MPLS Control Plane public documents.

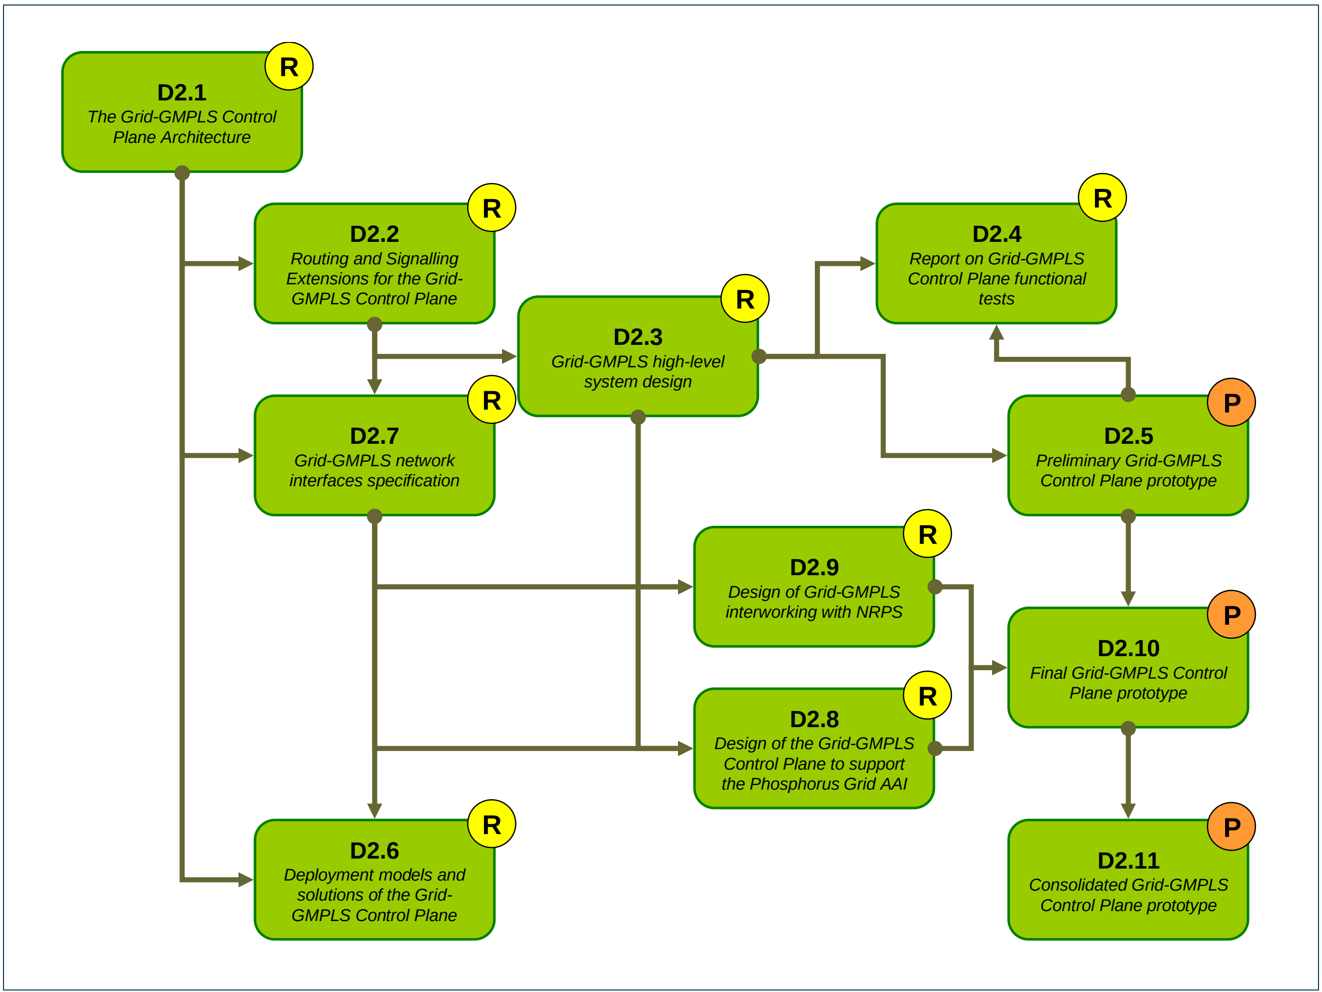

Figure 1-22: Synoptic graph of WP2 deliverables.

Then, we had the design and fast-development of a modular G^{2}MPLS prototype (ref. **Figure 1-23**), capable of implementing the behaviours and functionalities of the four main G^{2}MPLS controllers (G^{2}MPLS Edge Controller, G^{2}MPLS Border Controller, G^{2}MPLS Core Controller and G^{2}MPLS UNI-C Controller) depending on the configuration of the different modules by the user.

Finally, we had the testbed deployment and public demonstrations (ref. **Figure 1-25**, **Figure 1-24**, **Figure 1-26** and **Figure 1-27**). G^{2}MPLS functionalities were demonstrated with DDSS application for remote anycast storage of data contents (at SC'08 and ICT'08) and with KoDaVis/UNICORE for the remote anycast processing and client collaboration.

It is also worth mentioning the high value of background expertise and know-how contributed by the partners, their strong commitment to produce a solid piece of work for the Research Networking area, and the highly cooperative team spirit.

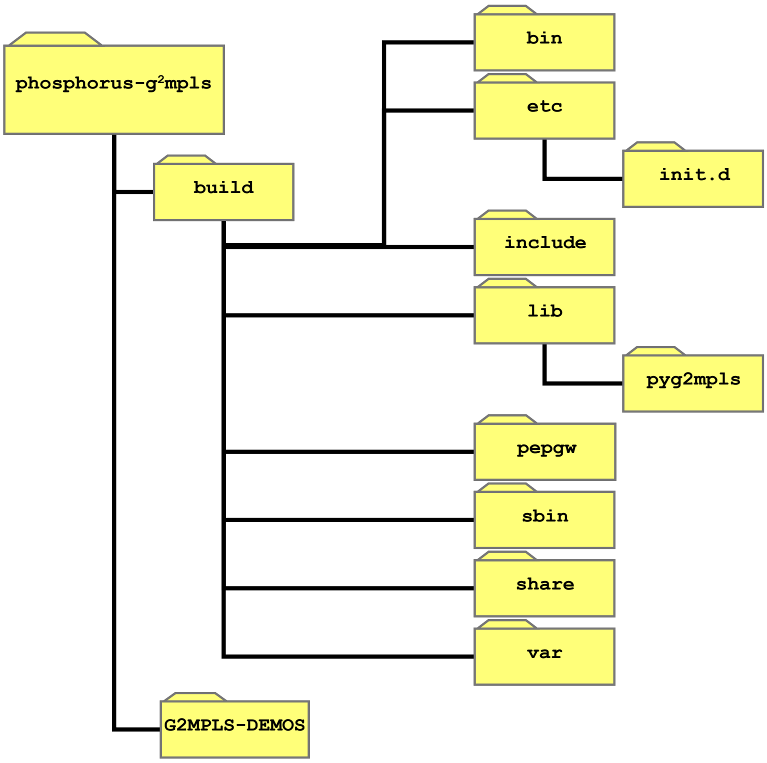

Figure 1-23: Phosphorus G^{2}MPLS prototype structure.

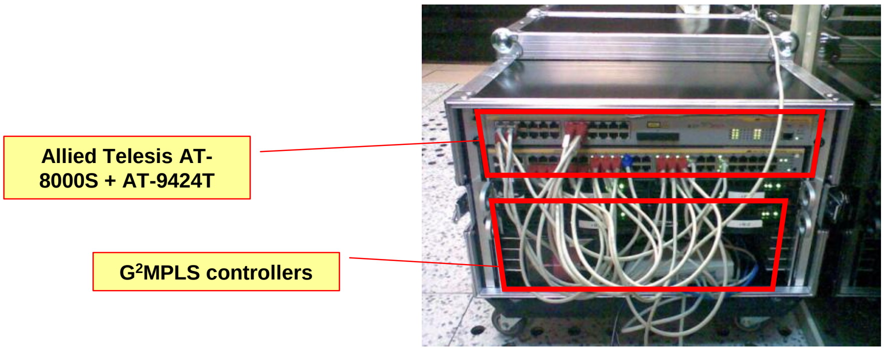

Figure 1-24: Phosphorus G^{2}MPLS testbed in PSNC: L2 Ethernet part.

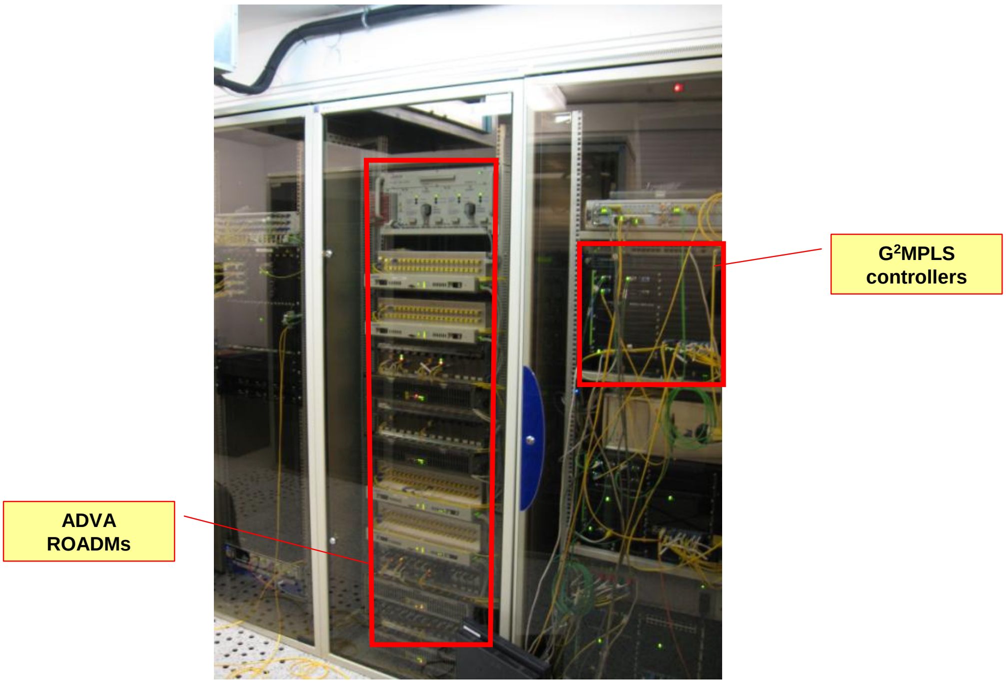

Figure 1-25: Phosphorus G^{2}MPLS testbed in PSNC: LSC part.

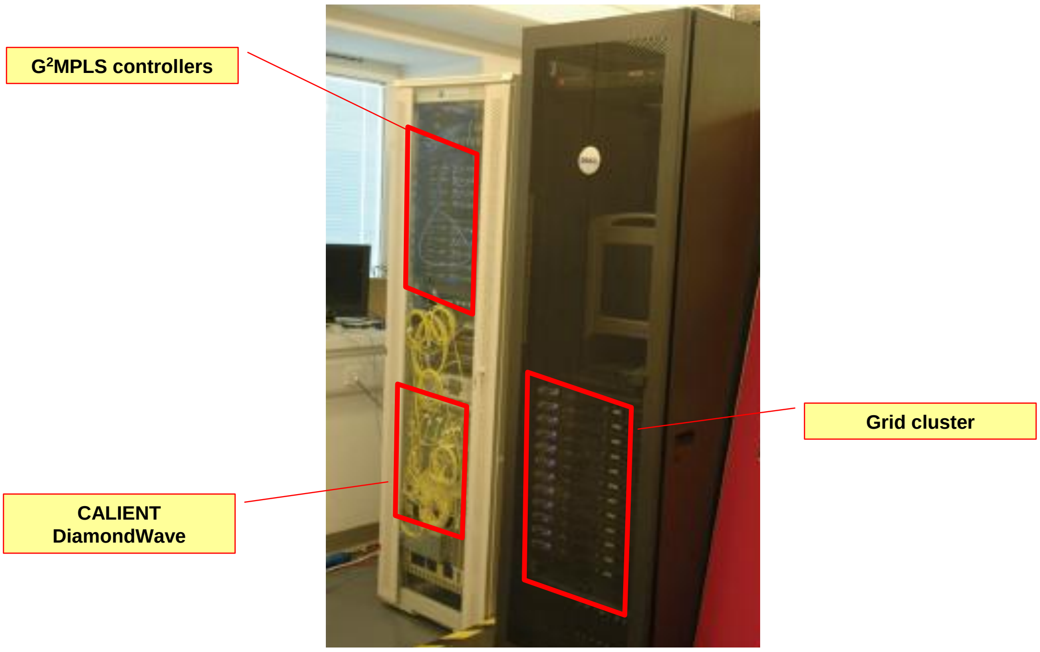

Figure 1-26: Phosphorus G^{2}MPLS testbed in UESSEX: FSC part.

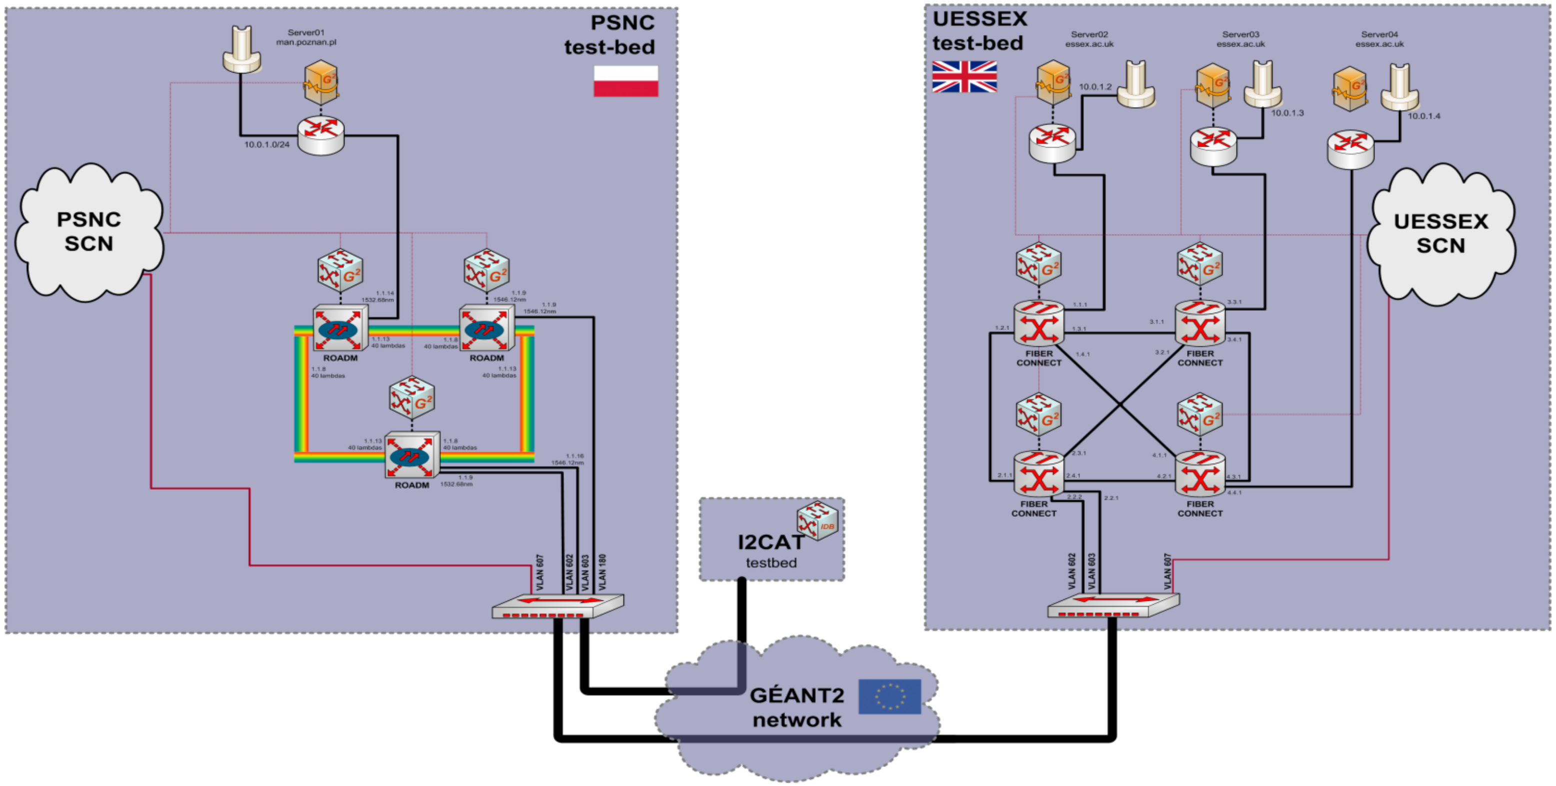

Figure 1-27: Phosphorus G^{2}MPLS multi-domain testbed.

All the challenges investigated in WP2 are still in the worldwide research agenda and Phosphorus highly contributed to them with its G^2MPLS prototypes. The G^2MPLS Control Plane represents a framework on top of which further research initiatives on advanced experimental networking may arise. In fact, the valuable foreground produced by WP2 partners is just spurring further initiatives in research contexts like the Future Internet/Future Networking, in which the core WP2 partners plus some other Phosphorus partners will evolve the G^2MPLS Control Plane.

1.4.3 Middleware and Applications

Partners involved in this workpackage: FHG (leader), PSNC, FZJ, HEL, CTI, SARA, UNIBONN, UESSEX

The mission of WP3 in the PHOSPHORUS project is twofold: (1) development of middleware components for the integration with both the network service plane and the network control plane (as developed in WP1 and WP2) and (2) the adaptation and provisioning of advanced applications to benefit from the advanced network and middleware environment in the PHOSPHORUS testbed.

1.4.3.1 *WP3 Objectives and Starting Point of Work*

Often Grid applications have complex workflows, in which certain parts make extensive use of the network e.g. to transfer huge amounts of data between storage and compute or visualisation resources. In today's production Grids, resources are provided either exclusively for an application or are shared on a best-effort base. Within the German project VIOLA, the UNICORE Grid middleware has been enhanced to make synchronous reservations of compute nodes and network bandwidth between the nodes. This is achieved by interfacing UNICORE to a co-allocation service that provides this feature, relying on resource reservation systems with advance reservation capability. Within Phosphorus (WP3), this system will adapted to the network reservation services provided by NRPS through the HARMONY system and G2MPLS (WP1& WP2) and is extended to orchestrate network, compute and other Grid resources in a more general way, as required by the selected showcase applications. This generalised system should be able to:

- handle network resources and services as any other Grid resource on the middleware layer and orchestrate them jointly with other resources for complex application workflows,
- perform resource selection for computational and other resources connected to the network on the network layer while using middlware services for co-allocation,
- create a chain of trust from the middleware layer to the network layer using the user credentials from the middleware layer

WP3 will mainly contribute in two ways to the overall project mission and vision to develop solutions for some of the key technical challenges that enable on-demand end-to-end services for Grid-based applications.

- The first contribution is to develop the Grid-middleware services that are required to make these network services accessible for applications.

- The second contribution to the PHOSPHORUS objectives is the enhancement of existing applications with high communication requirements regarding bandwidth and/or QoS. These applications were being enabled to utilise the capabilities of the PHOSPHORUS test-bed through adapted workflows. They serve to provide requirements towards the middleware and the network, to evaluate the developments of the other Work Packages in test-bed experiments and finally demonstrated the results.

The PHOSPHORUS middleware developed in WP3 will implement Grid service abstractions exposing Grid resources, including conventional computational resources (CPU, memory, storage) as well as network resources (end-to-end QoS management). Hence, it will be possible for an application to request conventional Grid resources with the additional ability to demand network QoS (e.g. guaranteed bandwidth or latency among computational nodes). Interfaces between services are relying on existing standards like the OGF WS-Agreement or the OASIS WSRF, thus providing open solutions that simplify the adaptation to other middleware than that used within PHOSPHORUS, e.g. Globus Toolkit 4.2.

Major objectives addressed by WP3:

- Formulation of user requirements to contribute to the definition of the functions and services to be implemented by the overall project

- Development and provisioning of the middleware services to orchestrate requested for an application, i.e. network, compute resources, visualisation devices, etc.

- Integration with network layer developments (Control Plane and Grid-NPRS) with Grid middleware to enable high performance applications running efficiently in the test-bed

- Adaptation of selected Applications to showcase the benefit from the integrated test-bed environment developed in the project

- Definition of the properties that network resources (services) should have in order to integrate seamlessly into an environment with other grid services

- Enabling user and applications to automatically get access to the resources that match their requirements e.g. in terms of performance and QoS.

- Enabling users and application to plan resource usage in advance and to dynamically adapt the resource usage to the changing requirements of the application thus avoiding expensive waste of resources.

- Managing all resources required to run an application in an integrated manner with a single service based interface towards to user or the application respectively.

- Testing the provided networking functionality based on application use cases defined in the first project stage and deriving requirements for extensions and enhancements to be carried out in the second 12 month period from the test-bed experiments.

- Integration of the Resource Selection Service into the UNICORE Rich Client, final experiments.

- Evaluate the results of the testbed experiments done in the previous phase and the demonstrations at the SC08 and ICT 2008 and define missing functionality, implementation of missing functionality.

- Providing authentication and authorization information obtained in the middleware layer to the network layer supporting AAI mechanisms implemented there.

- Evaluate and define possible integration of INCA and G2MPLS

1.4.3.2 *Work performed and results*

The research and development in WP3 was organised in four tasks, each one comprising several activities. The following paragraphs describe the work performed and the results achieved.

Task 3.1 – Adaptation of applications

This task defined the scope of the demonstrations to be prepared and carried out during the lifetime of the project and designs the changes and extensions of applications, which enable them to verify the concept and implementation of the Phosphorus services. Based on this design, the applications and the execution environment, e.g. the respective workflows, have been adapted to make use of the resources and services in the test-bed. Key activities in this task have been:

Activity 3.1.1 Use-Cases

This activity defined the use-cases for the applications (WISDOM, KoDaVis, Collaborative CAVES, Distributed Storage), mapping them onto the Phosphorus test-bed.

Activity 3.1.2 Application Enhancement

The applications WISDOM, KoDaVis, Collaborative CAVES, Distributed Storage have been adapted to make use of the enhanced middleware and underlying network services. This included both modifications of the applications, e.g. to interact with the network service plane without middleware, and improvements of the applications' workflows.

Task 3.2 – Middleware

This task is responsible for the adaptation, development and implementation of middleware to seamlessly integrate with new network capabilities and to allow the co-allocation of multiple, heterogeneous resources to the application. Key activities in this task are:

Activity 3.2.1 Definition of middleware-NRPS interface

Existing and evolving standards, protocols and implementations have been reviewed and evaluated. Based on this evaluation the required changes and extensions of the middleware and applications to use the new services have been designed and implemented.

The requirements of applications and middleware towards the resource provisioning system and the common framework under development on WP1 have been analysed designed.

Functionality, protocols and interfaces to resource provisioning system have been design and implement in this activity. The work under this and the previous paragraph was carried out jointly with WP1 (Activity 1.3.1). The results of this activity have been published in the report that is part of D1.6 (Definition and development of the Network Service Plans and northbound interfaces development).

Activity 3.2.2 Implementation of Middleware extensions

The middleware UNICORE and the MetaScheduling Service from the VIOLA project have been extended to allow changes in resource utilisation, either by agreements made in advance or dynamically.

Activity 3.2.3 Implementation and integration of an ontology for service abstractions

An integrated CIM-based ontology of the service abstractions of resources was designed and made available in the test-bed including the work done in WP1 for network services. This activity also provided the reasoning interface to allow automatic selection of services (semantic mediation) according to the users' request or the requirements exposed by the applications.

Task 3.2 interacted with WP1, WP2, WP4, and WP6 to reach a common understanding of requirements and to define common interfaces based on this.

During the last period of the project new activities of task 3.2 have been defined:

Activity 3.2.1 Integration of the Middleware with the AAI for the network layer

The solution implemented by WP1 and WP4 has been analyzed and the relevant information for authentication and authorization available on the middleware layer has been identified. Based on this the interface towards the network layer create a seamless chain of trust from the middleware client to the network layer has been designed and implemented.

Activity 3.2.2 Resource Selection Service

The Resource Selection Service developed and implemented in the previous period of the project has been integrated into the UNICORE Rich Client. The integration allows a resource selection for the MetaScheduling Service based on the requirements of the application selected by the user and the resources available in the testbed.

Activity 3.2.3 Integration of INCA and G^{2}MPLS

A possible integration of INCA and G^{2}MPLS has been evaluated through simulation.

Task 3.3 – Interfacing between G²MPLS and Resource Management System via G-OUNI

This task was responsible for the implementation of the interface of the Grid Resource Management Systems to the G²MPLS system via the G-OUNI protocol as specified jointly with WP2. Successfully completed key activities in this task comprised:

Activity 3.3.1 Middleware – Resource Management System Interface implementation

Implementation of the direct interface between the Co-allocation service and the Resource Management Systems for phase I of the project, used until the G2MPLS/G-OUNI became available.

Activity 3.3.2 G-OUNI specification and implementation

The functionality and protocol of the G-OUNI was designed jointly with WP2 and implemented on the Resource Management System side as an extension of the MetaScheduling System.

During the last period of the project a new activitiy of task 3.3 has been defined:

Activity 3.3.1 Consolidation of the Middleware integration with G^{2}MPLS

The results of the testbed experiments done in the previous phase and the demonstrations at the SC08 and ICT 2008 have been analysed and the missing functionality defined. In a second step the missing functionality. Has been implemented. The activity was concluded by final experiments with the enhanced interface.

Task 3.4 – Integration and evaluation of middleware

This task was responsible for the integration of the middleware services and the network services. The integration resulted in a seamless environment with a single service based interface where the user can define the execution environment for his application. The resulting execution environment comprises both hardware resources like compute nodes, storage systems, visualisation devices, and network resources. This task depends on the services and interfaces provided by WP1 and WP2.

1.4.4 Authentication, Authorization and Accounting

Partners involved in this workpackage: UvA (leader), FHG, HEL, CTI, UNIBONN

1.4.4.1 *WP4 objectives, tasks, and major results*

The Authentication, Authorisation and Accounting (AAA) services constitute an important component of the infrastructure supporting on-demand Optical Network Resource Provisioning across multiple domains and different target consumer applications. A consistent AAA infrastructure requires the interaction among the related AAA services at all networking layers including network/forwarding elements, control plane, reservation and provisioning service, and user/target applications

The Phosporus Work Package 4 activity aimed to develop, design and implement in the project testbed a generic AAA Authorisation architecture for on-demand multidomain Network Resource Provisioning (GAAA-NRP) in heterogeneous Grid and network resource provisioning middleware environment.

This development covered three main areas:

- GAAA-NRP architecture that should allow adding AAA authorisation services at each networking layer and support the whole lifecycle of the provisioned resource from reservation to decommissioning;
- Pluggable GAAA Toolkit library (GAAA-TK) that could be integrated with the used in the project networking and Grid middleware; and
- ForCES based Token Based IP Switch (TBS-IP) that could allow fine-grained access control, and policy enforcement at lower networking layer.

The WP4 produced valuable results in all mentioned above areas that contributed to the general development of on-demand NRP concept and architecture, further development and extension of the generic AAA Authorisation architecture (GAAA-AuthZ) proposed in RFC2904, in particular, definition of the basic security services and mechanisms for on-demand NRP, and provided a use case and implementation scenario for the ForCES architecture currently being developed at IETF.

The WP4 conducted its work in tight cooperation with other WP's, in particular, defining authentication and authorisation services functionality and interfaces and providing pluggable GAAA-TK library that supported basic policy based authorisation scenarios in Harmony/NSP (WP1), G2MPLS system (WP2) and allowed using UNICORE Grid security credentials when integrating WP3 Grid middleware applications with Harmony/NSP and G2MPLS.

WP4 developments to achieve/ensure Phosphorus GAAA-NRP interoperability at larger scale with existing Authentication and Authorisation infrastructure (AAI) used by Grid community and NREN community included defining XACML attribute and policy profile for NRP and Grid and supporting eduGAIN secure credentials correspondingly. This work was done in cooperation with EGEE/OSG Grid oriented projects and in coordination with the GN2 JRA5 activity.

The following tasks provided framework for the WP4 research, development and testbed activity:

Task 4.1 – AAA Architecture selection

This task included detailed study and and analysis of the related security and AAA technologies and standards what allowed to propose a generic AAA Authorisation architecture for on-demand optical network resource provisioning. This work leveraged existing among WP4 partners experience in the application security and networking area and defined directions for further GAAA-NRP development. The proposed GAAA-NRP architecture also included definition of the Token Based Networking to allow access control and policy enforcement at the networking layer for dataflow.

Task 4.2 – AAA scenarios

This activity was conducted in tight cooperation with WP1, WP2, WP3 to define basic authorisation scenarios to support WP1 Harmony/NSP and WP2 G2MPLS use cases and WP3 Grid middleware applications. The specified use cases and scenarios motivated the definiton of the general Network Resource Provisioning model that specified such stage of the provisioned resources lifecycle as reservation, deployment, access/use, and decomissioning. Such NRP model or workflow definition allowed better definition of AAA/Authorisation service and provided a framework for heterogeneous AAA and NRPS systems integration and consistent security context management in cross-domain network resource provisioning.

Task 4.3 – Integration and testing

This activity provided a framework for cooperation with other work packages to implement the proposed GAAA-NRP infrastructure and developed GAAA Toolkit library (GAAA-TK) library with the project testbed and Harmony/NSP and G2MPLS systems in particular. This activity included also UvA cooperation with Internet2 DCN project and was resulted in a few demonstration of interworking between Phosphorus and Internet2 DCN testbeds.

Task 4.4 – AAA components for token based authorisation model

This task was focused on actual GAAA Toolkit library development that provided a reference implementation for the GAAA-NRP infrastructure and its functional components and mechanisms defined in Task 4.1 and Task 4.2. The token-based authorization model creates the notion of a user session for reserved network resources in multidomain scenarios and allows creating a dynamic trust associations of the provisioned resources, it uses authorisation tickets and tokens for inter-domain signalling and access control.

Task 4.5 – ForCES Token Based Switch

This task was focused on the development and implementation of the Token Based Switch that implement the Token Based Networking architecture and uses IETF ForCES standard for adding TBN functionality for a standard IP Switch. This work delivered a high performance, hardware based implementation to generate and recognize tokens that can be added to application specific dataflows to achieve better granularity in access control and policy enforcement at networking Data plane.

Task 4.6 – AAA infrastructure to support Optical Network Resource Provisioning at larger scale

This task provided a framework for summarising WP4 experience in developing and implementing GAAA-NRP authorisation infrastructure and GAAA-TK library in the project testbed and identifying the area of the further developments to support policy based resource provisioning at larger scale. The work was focused on the development of the AAA infrastructure and functional components to support heterogeneous Grid and Optical Network Resource Provisioning based on developed AAA/AuthZ architecture and components.

The following sections provide more detailed information about the major WP4 results that were achieved in the project and can provide a basis for the further development in Network Resource Provisioning and generic AAA Authorisation areas.

1.4.4.2 *Network Resources Provisioning Model*

The typical on-demand Network Resource Provisioning (NRP) process includes four major stages, as follows:

(1) resource reservation;

(2) deployment (or activation);

(3) resource access/consumption;

(4) resource de-commissioning after it was used.

Additional stage (5) re-location is considered to include combination of all 4 basic stages (1) - (4) starting from de-commissioning the resource that should be relocated, e.g. changing lightpath, or moving jobs/experiment to another Data Center. However, in case of relocation the general security context, and reservation ID in particular, should be inherited from the initial reservation.

In its own turn, the reservation stage (1) typically includes three basic steps:

(a) resource lookup;

(b) complex resource composition (including alternatives), and

(c) reservation of individual resources.

The reservation stage may require the execution of complex procedures that may also request individual authorisation for resource access. This process can be controlled by an advance reservation system or a meta-scheduling system; it is driven by the provisioning workflow and may also include Service Level Agreement (SLA) negotiation. At the deployment stage, the reserved resources are bound to a reservation ID, which we refer to as the Global Reservation Identifier (GRI). The decommissioning stage is considered as an important stage in the whole resource provisioning workflow from the provider point of view and should include such important actions as global provisioning/access session termination and user/process logout, log information sealing, accounting and billing.

The rationale behind defining different CRP workflow stages is that they may require and can use different security models for policy enforcement, trust and security context management, but may need to use common dynamic security context. Defining CRP models will also allow simplifying the integration of the NRP provisioning with the higher level scientific workflow.

1.4.4.3 *Generic AAA Authorisation infrastructure for NRP*

Figure 1 illustrates major interacting components in the multi-domain CRP using an example of provisioning multidomain network connectivity between a User and a Destination resource or application. Each networking domain is presented as

- Network Elements (NE) (related to the network Data plane);
- Network Resource Provisioning Systems (NRPS) acting as a Domain Controller (DC) (typically related to the Control plane);
- Inter-Domain Controller (IDC) managing cross-domain infrastructure operation, often referred to as Network Service Plane (NSP).
- Access to the resource or service is controlled by the DC or NRPS and protected by the generic Authentication, Authorisation, Accounting (AAA) service that enforces a resource access control policy. The following functional elements comprise the proposed authorisation infrastructure for CRP which we will refer to as GAAA-CRP:
  - Policy Enforcement Point (PEP), Policy Decision Point (PDP), and Policy Authority Point (PAP) as major functional components of the Generic AAA AuthZ infrastructure (GAAA-AuthZ) [RFC2904].
- Token Validation Services (TVS) that allow efficient authorisation decision enforcement when accessing reserved resources.

Depending on the basic GAAA-AuthZ sequence (push, pull or agent) [RFC2904], the requestor can send a resource access request to the resource (which in our case is represented by NRPS) or an AuthZ decision request to the designated AAA server which in this case will act as a PDP. The PDP identifies the applicable policy or policy set and retrieves them from the PAP, collects the required context information, evaluates the request against the policy, and makes the decision whether to grant access or not.

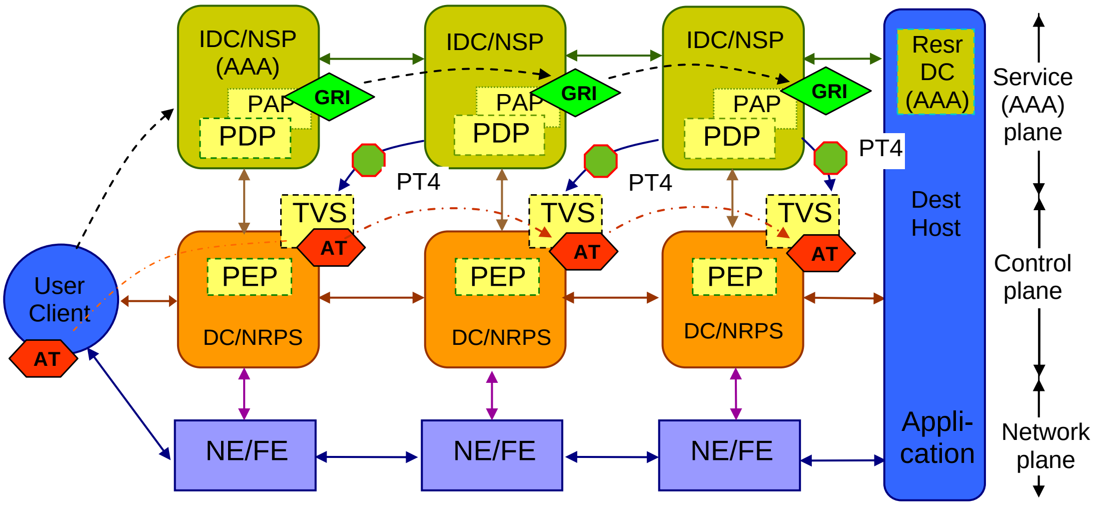

Figure 2.2. Components involved in multidomain network resource provisioning.
CRP stages reservation, deployment and access are presented by the flows corresponding to GRI (forward from the user to the resource), pilot tokens PT4 (backward), and access tokens AT (forward).

The following provisioning sequences can be executed when composing a cross-domain lightpath:

Chain reservation sequence (also referred to as a provider sequence). The user contacts only the local network domain/provider that provides the destination address. Each consecutive domain provides a path to the next domain.

Polling sequence. The user client polls all resources or network domains, builds the path and makes the reservation.

Agent (or tree) sequence. The user delegates the network provisioning negotiation to an agent that will take care of all necessary negotiations to provide the required network path to the user. A benefit of “outsourcing” the resource provisioning is that the agents can maintain their own reservation and trust infrastructure. This can be considered as a basic provisioning sequence for currently used Grid resource management and advance reservation systems.

Depending on the used authorisation and attribute management models, some attributes for the policy evaluation can be either provided in the request or collected by the PDP itself. It is essential in the Grid/Web service oriented environment that AuthN credentials or assertions are presented as a security context in the AuthZ decision request and are evaluated before sending request to PDP.

Based on a positive AuthZ decision (in one domain) the AuthZ ticket (AuthzTicket), containing AuthZ decision and context, can be generated by the PDP or PEP and communicated to the next domain where it can be processed as a security context for the policy evaluation in that domain.

Based on a positive AuthZ decision (in one domain) the AuthZ ticket (AuthzTicket), containing AuthZ decision and context, can be generated by the PDP or PEP and communicated to the next domain where it can be processed as a security context for the policy evaluation in that domain.

In order to get access to the reserved resources (at the access stage) the requestor needs to present the reservation credentials that can be in form of an AuthZ ticket (AuthzTicket) or an AuthZ token (AuthzToken) which will be evaluated by the PEP with support of TVS for ticket or token evaluation, to grant access to the reserved network elements or the resource. In more complex provisioning scenarios the TVS infrastructure can additionally support an interdomain trust management infrastructure for off-band token and token key distribution between domains that typically takes place at the deployment stage when access credentials or tokens are bound to the confirmed GRI by means of shared or dynamically created interdomain trust infrastructure. Token and token key generation and validation model can use either shared secret or PKI based trust models.

Using AuthZ tickets during the reservation stage to communicate the interdomain AuthZ context is essential to ensure effective decision making. At the service access/consumption stage the reserved resource may be simply identified by the assigned GRI created/confirmed as a result of the successful reservation process.

It is an important convention for the consistent NRP operation that GRI is created at the beginning and sent to all polled/requested domains when running the (advance) reservation process. Then in case of a confirmed reservation, the DC/NRPS will store the GRI and bind it to the committed resources. In addition, a domain can also associate internally the GRI with the Local Reservation Identifier (LRI). The proposed TVS and token management model allows for hierarchical and chained GRI-LRI generation and validation.

Correspondingly, we define the following sessions in the overall NRP process (discussed in details below): provisioning session that includes all stages; reservation session, and access session. All of them should share the same GRI and AuthZ context.

In the discussed NRP model we suggest that the resources are organised in domains that are defined (as associations of entities) by a common policy or a single administration, with common namespaces and semantics, shared trust, etc. In this case, the domain related security context may include:

- static security context such as domain based policy authority reference, trust anchors, all bound by the domain ID and/or domain trust anchor;
- dynamic or session related security context bound to the GRI and optionally to a Local Reservation Identifier (LRI).

In general, domains can be hierarchical, flat or have irregular topology, but all these cases require the same basic functionality from the access control infrastructure to manage domain and session related security context. In the remainder of the paper we will refer to the typical use case of the network domains that are connected as chain (sequentially) providing connectivity between a user and an application.

The following main issues should be addressed when developing and implementing authorisation infrastructure in heterogeneous multidomain environment:

- topology aware policy definition, support for different logical organisation of resources, including possible constraints on resource combination and interoperation expressed as policy rules and policy obligations.
- multiple policies processing and combination related to domain and resources;
- support of different format of attributes describing both network resources and identity presentation and credentials,
- ability to be configured or to recognise multiple attribute authorities,
- identity and attributes mapping,

- multiple trust and administrative domains that should be automatically recognised or manually pre-configured,
- need for consistent security context management when provisioning cross-domain paths.
- additionally, multidomain scenarios may require the support of Service Level Agreement (SLA) negotiation and creating dynamic security associations bound to provisioning sessions or SLA and user jobs.

The proposed GAAA-CRP infrastructure includes the following access control mechanisms and functionalities that extend the generic GAAA-AuthZ model described in [2] with the specific functionality for on-demand CRP, in particular:

- AuthZ session management to support complex AuthZ decision and access to multiple resources, including multiple resources belonging to different administrative and security domains.
- AuthZ tickets with extended functionality to support AuthZ session management, delegation and obligated policy decisions.
- Access and pilot tokens used for managing authorisation context in interdomain reservation process as part of the cross-domain policy enforcement that can be used in the control plane and in-band.
- Policy obligations to support conditional policy decisions in multidomain environment, usable/accountable resources, and additionally global and local user account mapping widely used in Grid based applications and supercomputing.

The solutions proposed in the GAAA-CRP framework are based on using such structural components and solutions as the Token Validation Service, the Obligation Handling Reference Model (OHRM), and the XACML attributes and policy profile for multidomain CRP that can combine earlier defined XACML-Grid and XACML-NRP profiles, all described in the deliverables D4.3.1 and D4.5 and implemented in the pluggable GAAA Toolkit library.

1.4.4.4 *Authorisation session management in NRP*

a) Session types in GAAA-NRP/CRP

The session management functionality in GAAA-NRP and GAAA-TK is based on the general NRP model. The overall network provisioning process initiates the provisioning session inside of which we can also distinguish two other types of sessions: reservation session and access session. Although they may require different security context, all of them are based on the (positive) AuthZ decision, may have a similar AuthZ context and will require a similar functionality when considering distributed multi-domain scenarios.

Figure 2 illustrates the relationship between all sessions which are bound by a common GRI. The diagram also indicates what types of policies or protocols are used at each stage. The access control is done at each stage, it may be related to different services but can use the same AuthZ service with different policies.

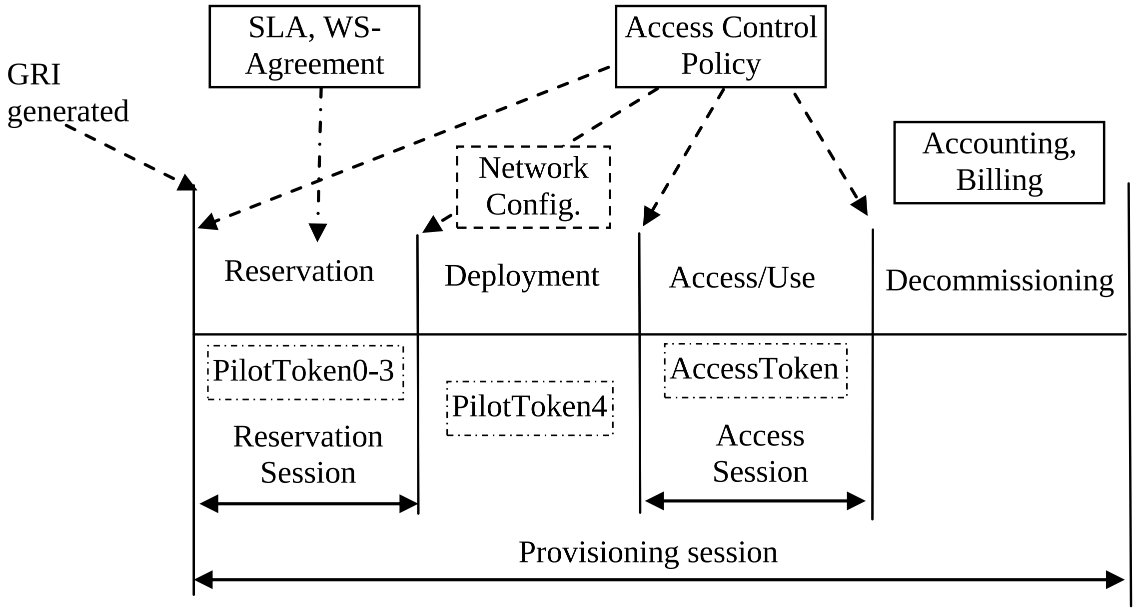

Figure 2. NRP stages and session types.

In multidomain NRP authorisation, tickets and tokens are used to transfer necessary security/authorisation context information between domains and serve as a session or access credentials. Using these mechanisms ensures the integrity and consistency of the provisioning process. When used together, AuthzTicket and AuthzToken share the SessionId attribute which can be either a global or a local reservation/session ID.

b) Using XML Tokens for Signalling and Access Control

In the proposed GAAA-NRP architecture tokens are used for both access control when accessing the reserved resources and for signalling during reservation and deployment stages. Correspondingly, we distinguish the two major types of token in the GAAA-NRP architecture: access tokens and pilot tokens. Access tokens are used in rather traditional manner and described in details in [4]. Pilot tokens functionality and format was proposed and defined as a result of the current development of the AuthZ infrastructure as an integral component of the NRP.

After initial implementation in the GAAA-TK library released in the D4.3.1 deliverable (M22) both the access token and pilot token concepts have been integrated with and tested in WP1 Harmony/NSP and WP2 G^2MPLS testbeds. This motivated changes both in extending the token data-model and adding methods to support new AuthZ scenarios consequently implemented in the updated GAAA-TK library and described in the D4.5 deliverable (M30). Detailed discussion how pilot tokens can be used for building dynamic trust associations is described in chapter 3.

The updated XML Token data model that combines functionality of both access tokens and pilot tokens is presented in Appendix B. Although the tokens share a common data-model, they are different in the operational model and in the way they are generated and processed. When processed by AuthZ service components they can be distinguished by the presence or value of the token type attribute.

Access tokens used in GAAA-NRP have a simple format and contain three mandatory elements: the *SessionId* attribute that holds the GRI, the *TokenId* attribute that holds unique token ID attribute and is used for token identification and authentication, and the *TokenValue* element, and two optional elements: the *Condition* element that may contain two time validity attributes *notBefore* and *notOnOrAfter*, and the *Decision* element that holds two attributes *ResourceId* and *Result*, and an optional element *Obligations* that may hold policy obligations returned by the PDP.

The GAAA-NRP architecture defines two types of access tokens:

AType1 – this pilot token type is used as access credentials and cryptographically binds SessionId/GRI, domainId and TokenId.

AType2 – extends AType1 with the Obligations element that allows communicating policy obligations between domains.

The GAAA-NRP architecture defines four types of pilot tokens that have different profiles of the common data model and different processing/handling procedures:

PType1 – this pilot token type is used just as a container for communicating the GRI during the reservation stage. It contains the mandatory SessionId attribute and an optional Condition element (it does not contain a TokenValue element).

PType2 – this pilot token type is the origin/requestor authenticating token. Its TokenValue element contains a value that can be used as the authentication value for the token origin. The token value is calculated on the GRI by applying e.g. an HMAC function to the GRI together with the requestor symmetric secret or private key.

PType3 – this pilot token type extends the PType2 with a Domains element that allows to collect domains security context information (in the Domains/Domain element) when passing multiple domains during the reservation process. Such information includes the previous token and the domain's trust anchor or public key.

PType4 – this pilot token type is used at the deployment stage and can communicate between domains security context information about all participating in the provisioned lightpath or network infrastructure resources. This token type can be used for programming/setting up a TVS infrastructure for consistent access control tokens processing at the resource access stage.

Pilot token PType3 and PType4 can be used for communicating AuthZ ticket containing extended authorisation session context. When used together with an AuthzTicket the ticket and token identification elements `TokenID`, `SessionID`, and `Issuer` can be shared.

1.4.4.5 *XACML policy and attributes profile for NRP*

a) Main use cases for policy definition in NRP

In general, the access control policy comprises of rules and conditions that specify what user with what attributes may access or execute what action on the resource with what attributes.

Two basic use cases for access control in Network Resource Provisioning (NRP) can be expressed in a simple narrative form:

Use case 1: "User A is only allowed to use user endpoints X, Y and Z", or

Use case 2: "User A is only allowed to use endpoints in domain N and M".

Two other use cases are related to the more complex access control scenarios that take place during the multidomain/multiple resources reservation and require simple delegation functionality:

Use Case 3: User/Group A is only allowed to invoke method X, Y, and Z

Use Case 4: User/Group A is only allowed to invoke method X, Y, and Z based on session delegation

Supporting use cases 3 and 4 requires introducing simple group or session based delegation with full or limited delegation profile that should be supported by both new PEP and related policy definition.

b) Network or resource related attributes

Network related attributes allow building policy depending on the network topology or other network characteristics.

Topology format should provide necessary information about the network resource to allow consistent policy evaluation, and vise versa the policy format may be defined by the network topology to which the policy is applied. Topology semantics will define the resource attributes semantics, and vice versa.

Network related attributes are considered as a part of the XACML Resource definition. The following resource/network related attributes can be specified and used for authorisation:

- Domain ID (network domain)
- Subdomain (or relationship)
- VLAN
- Node or TNA and TNA prefix, or
- Interface ID
- Device or resource-type
- Link ID
- Link parameters: average delay and maximum bandwidth
- ReservationEPR that may directly or indirectly define the resource federation or administrative domain
- Federation that defines a number of domains or nodes sharing common policy and attributes

As it was mentioned before, some more advanced scenarios may require that particular network route or path is provided, in this case the policy definition should allow decision on the specific features of path or on the path in total.

c) Subject related attributes

Subject related attributes allow building policy depending on the properties of the request Subject or user. The following subject related attributes can be specified:

- Subject ID
- Subject confirmation that contains AuthN assertion/token or other attribute confirming subject's ID by trusted AuthN authority
- Subject Role
- Subject Group
- Subject Federation (e.g., Virtual Organisation, or Shibboleth AAI federation) or domain

Subject context that can provide additional information about the Subject other than Subject federation e.g. such as Session ID, or project/experiment name

Typically Subject attributes are provided as Subject credentials which depending on user client implementation and middleware may take a form of X.509 public key and attribute certificates (PKC, AC), SAML Authentication and Attribute assertions, proprietary AuthN system credentials.

d) Action and Environment related attributes

Action related attributes represent a limited number of the specific actions that requesting party can ask to initiate network resource reservation, access or management.

Environment related attributes allow providing additional information for policy definition and evaluation. There is no specific Environment attributes identified for the XACML-NRP profile but this may be a place to put security context related information from the previous domain.

e) Policy Obligations used in NRP and conditional policy definition

It was identified that some resource reservation and access scenarios may imply conditions that are not known at the time of making policy decision, or require some actions that must be performed at the time of the resource use or after it was used. Policy obligation is one of the authorisation policy enforcement mechanisms that allow adding AuthZ decision enforcement components that can not be defined in the policy at the moment of making policy decision by the PDP, or may not be known to the PDP or policy administrator/writer. The obligations can be also included in the extended access token context (see Deliverable D4.4).

The suggested functionality that can be achieved by using obligations includes but not limited to:

- Intradomain network/VLAN mapping for cross-domain connections, that can be used to map external/interdomain border links/TNA's to internal VLAN and sub-network
- Account mapping
- Type of service (or QoS) assigned to a specific request or policy decision
- Quota assignment

- Service combination with implied conditions (e.g., computing and storage resources)
- Usable resources/quota

The need of account mapping may exist in cases when domain based Network Resource Provisioning Systems (NRPS) have pre-installed/built-in pool accounts to which are different types or quality of service are assigned. In such situation authorised user need to use one of such accounts, e.g. “silver”, “gold”, “platinum”. A number of different individual accounts of the same type may be limited, consequently a dynamically assigned account should be selected from the pool of available or free accounts. Such dynamic account assignment can not be specified in the typically stateless policy and cannot be done by PDP. However, the access control policy may contain instruction to PEP to do such mapping, in particular with achieved with the proposed obligations handling model.

The proposed in WP4 (Task 4.4 and Deliverable D4.3.1) the Obligations Handling Reference Model (OHRM) specifies the obligations processing stages in the general case of the distributed authorisation infrastructure that uses the Domain-Central AuthZ service (DCAS) that can be part of the NRP system. The DCAS means that all domain located resources and services use a central AuthZ service that maintains a common set of policies for this domain. The proposed obligations processing model is compliant to the model used in XACML (refer to XACML2.0 standard) and provides a basis for defining obligations handling application progragramming interface (API) which was implemented in the GAAA-TK library.

1.4.4.6 *Pluggable GAAA Toolkit library implementation*

The GAAA-TK library provides all the base GAAA-NRP functionalities and interfaces for the operation heterogeneous NRPS environment using agreed in the Phosphorus set of user credentials and access control or reservation policies. This allows adding AAA/AuthZ services at all functional network planes: the Network service/provisioning plane, the Control plane and Data-forwarding plane.

Figure 3 illustrates the functional diagram of the pluggable GAAA Toolkit library that implements the GAAA-NRP infrastructure and major functional modules. The diagram alos illustrates their interaction when when evaluating a service request.

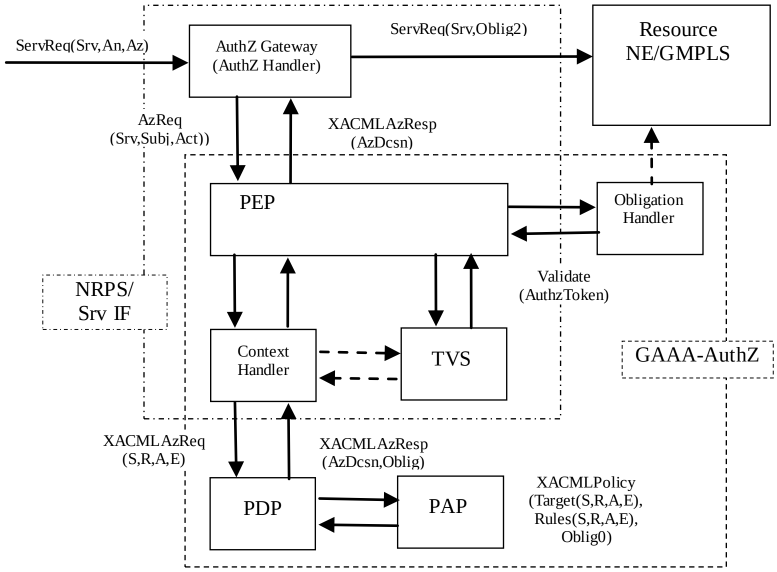

Figure 3. GAAA Toolket functional components

The authorisation service is called from the service/application interface via the AuthZ gateway (that can be just an interceptor process called from the service or application) that intercepts a service request ServiceRequset (ServiceId, AuthN, AuthZ) that contains a service name (and variables if necessary) and AuthN/AuthZ attributes. The AuthZ Gateway extracts necessary information and sends an AuthZ request AuthzRequest (ServiceId, Subject, Action), that contains a service name ServiceId, the requestor's identification and credentials, and the requested Action(s), to the PEP.

The major PEP's task is to convert AuthZ request's semantics into the PDP request which semantics is actually defined by the used policy. When using an XACML policy and correspondingly an XACML PDP, the PEP will send an XACML AuthZ request to the PDP in the format (subject, resource, attributes, (environment)). If in general case the XACML policy contains obligations, they are returned in the XACMLAzResponse (AuthzDecision, Obligations). The PEP calls the Obligation Handler to process obligations which are defined as actions to be taken on the policy decision or in conjunctions with the service access (like account mapping, quota enforcing, logging, or accounting).

If the service request contains an AuthZ token that may reference a local or global reservation ID, or just identifies an AuthZ session in which context the request is sent, the token validation is performed by the Token Validation Service (TVS). The TVS is typically called from the PEP and returns a confirmation if the token is valid.

The basic TVS functionality allows checking if a service/resource requesting subject or other entity, that posses/presents current token, has right/permission to access/use a resource based on advance reservation to which this token refers. During its operation the TVS checks if a presented token has reference to a previously reserved resource and a request resource/service confirms to a reservation condition. When using pilot tokens for signalling during interdomain path building, TVS can combine token validation from the previous domain and generation of the new token with local domain attributes and credentials.

The token generation and handling model can use both shared secret cryptography and public key cryptography and uses HMAC-SHA1 algorithm or digital signature for calculating token value correspondingly. Current implementation uses shared secret, which for the sake of simplicity of testbed implementation is provided as a part of the TVS/GAAA-TK library distribution.

1.4.4.7 *Token Based Networking and In-band Policy Enforcement*

The proposed GAAA-NRP architecture is easily integrated with the Token Based Networking (TBN) technology developed as a part of the WP4 activities to achieve in-band policy enforcement at dataflow layer. The TBN allows binding dataflows to users or applications by labeling application specific traffic, in particularly, our IPv4 implementation uses IPoption field to add a binary token to each IP packet. The token value is calculated similar to the XML token value by applying HMAC-SHA1 transformation to concatenated binary strings of the masked IP packet payload and GRI.

The Token Based Networking (TBN) infrastructure allows handling packets according to a specific tag, built-in the packet, called token. TBN uses tokens built from small encrypted pieces of data (e.g., a few bytes). By its mean, a token allows for applying different access/enforcement/filtering rules/criteria when processing packets that carry the tokens.

Tokens are a simpler way to authorise resource usage than certificates due to their concept: a policy enforcement point authorise the resource usage (e.g., a packet takes a certain path) only if the built-in token matches a cryptographic result applied to the packet using a set of keys provided beforehand by an authority at the moment of the path reservation. Tokens can bind to different semantics (in particular, it can be associated with the group of users and share its use when accessing designed resources), and they decouple the time of authorisation decision from the time of use. Tokens are a simple way of controlling access to authorised/reserved network resources that separates reservation/policy decision and access stages.

The TBN architecture offers to the user applications an authenticated access control mechanism to the reserved high-speed links (lightpaths) across multi-domain hybrid networks. The procedure to establish a lightpath consists of two phases that are decoupled in time: (1) a high-level set-up phase (obtaining tokens from an AAA web-service), and (2) a reserved path access including authenticity and authorisation checks (per-packet token checks at network edges within a multi-domain end-to-end connection).

The first phase allows individual users, or group of users (e.g., a research institution), or even user applications, to request privileged end-to-end connection across multi-domain networks by contacting, at service plane, an inter-domain or local ISP network provisioning service. Next, the network provisioning service would check whether the requested end-to-end path is available and return a positive result to the requestor containing the “unique” global reservation identifier (GRI) and authorisation key (Key).

The second phase determines how TBN authenticates network traffic (TCP connections, UDP transmissions, or other protocols) and how it checks the traffic for authorisation on behalf of their applications. The second phase is also responsible for preventing non-authorised use of lightpaths in a multi-domain network. Two network components are involved in the data plane: the token builder (TB) and the token based switch (TBS-IP).

Although there are many components involved in the high-level processing (service and control planes) such as NSP/AAA, we designed and implemented the Token Validation Service (TVS) component that validates the TokenKey/GRI requests from higher level (NSP/AAA) and pushes the provisioning requests into the low-level hardware (PEP).

The TBN infrastructure consists of Token Based IP Switches (TBS-IP) that are controlled by inter-domain controllers in each domain. The TBS includes such major components as Token Builder (TB) and TVS that provides a similar functionality as defined in the GAAA-CRP framework. The applications' traffic is first tokenised by the TB of a local domain (e.g., a campus network), after which it is enforced by the TBS-IP at each domain along the end-to-end path.

Tokens are used to label dataflows and can be made independent of upper layer protocols. In this way the token can be regarded as an aggregation identifier to a network service. The following four types of aggregation identifiers that can be combined are defined:

- identifier to link a service to the NE (e.g., a multi-cast, or transcoding);
- identifier that defines the service consumer (e.g., the grid application);
- identifier that defines the serviced object (e.g., the network stream);
- identifier that defines the QoS (security, authorisation, deterministic property, etc.).

The semantics that is referred to by a token (e.g., a certain routing behaviour) can be hard-coded into a TBS or dynamically programmed via TVS. Hence, a token provides a generic way to match/link applications to their associated network services. Tokens can be either embedded in the application generated traffic or encapsulated in protocols where embedding is not supported, such as in public networks.

To provide necessary performance for multi-Gigabit networks, TBS-IP is implemented using Intel IXDP2850 network processor that has a number of built-in hardware cryptographic cores to perform basic cryptographic functions such as required for TBN operation HMAC, SHA1, digital signature and encryption [TBN2007, D4.3.2].

It is important to mention that the TBN functionality can support Multi-Level Security (MLS) model by labelling and encrypting dataflows between security critical applications at data and control planes while GAAA-CRP model allows flexible policy based reservations and access control at service-plane.

TBS-IP control plane relies on a master-slave communication using ForCES protocol described below.

1.4.4.8 *Using ForCES for network management at control and data planes*

ForCES stands for Forwarding and Control Element Separation and is an upcoming IETF standard [RFC3746-ForCES, ForCES-Protocol]. ForCES defines a framework and associated protocol to standardize information exchange between the control and forwarding plane that comprise of Forwarding Elements (FE) and Control Elements (CE) correspondingly.

The basic building blocks of the ForCES model are the Logical Function Blocks (LFBs) described in an XML format. The ForCES protocol [ForCES-Protocol] works in a master-slave mode in which FEs are slaves and CEs are masters. The protocol includes commands for transport of LFB configuration information, association setup, status, and event notifications, etc. The protocol provides an open API for configuring and monitoring the Forwarding Plane in a standard manner. Grouping a number of LFBs, can create a higher layer service like TBS-IP in our case, or a firewall. Similarly any security method at networking layer can be described using the ForCES model.

The ForCES standard framework defines the transport mapping layer (TML) to transfer the ForCES messages from the CE to the FE and vice versa.

We consider the ForCES network management model as a way to integrate networking Control plane and Data plane into the general CRP process that requires heterogeneous networks configuration at least at the deployment and decommissioning stages. Recent works to define Web Services interfaces to ForCES devices makes such integration even simpler [IWAN2005-ForCES]. In our GAAA-CRP implementation we use ForCES protocol for transferring TBS-IP configuration information from the inter-domain controller to TB and TVS.

The ForCES framework provides a standard way of adding security services to both CE and FE. When used in the CRP/NRP Grid/Networking infrastructure the ForCES security framework [16] can benefit from using the common AuthN/AuthZ infrastructure. In this case the standard GAAA-AuthZ components can be added and related policies defined for the basic ForCES security functions such as endpoints and messages authentication.

1.4.5 Supporting Studies

Partners involved in this workpackage: IBBT (leader), i2CAT, CTI, AIT, UNIBONN, UNIVLEEDS

Partners involved in this workpackage: IBBT (leader), i2CAT, CTI, AIT, UNIBONN, UNIVLEEDS

This workpackage aims to provide supporting studies to the other WPs in the Phosphorus project. This has been achieved by presenting novel research on Grid routing and scheduling algorithms, and analysis of the transformation of these algorithms into practical, deployable Grid components. As such, WP5 (Supporting Studies) has developed a simulation environment for evaluation of optical Grid networks, and has performed a number of studies to support the experimental activities of the project. The objectives proposed before the start of the project are:

1. To study and design resource management and job scheduling algorithms. Network-awareness, constraint based routing and advance reservation techniques will be addressed

2. To document and analyze recommendations for the design of an optical grid control plane

3. To develop a simulation environment, supporting all previously discussed objectives. Extensive simulation will serve as the basis for most studies and analysis.

4. To study and propose algorithms for the design and planning of optical Grid infrastructures. Also, optical switching paradigms will be analyzed.

5. To study and propose algorithms and strategies for providing reliable Grid operation. Based on a general analysis to provide resiliency in optical Grid environments, we will define protection and restoration strategies, and present specific case studies.

6. To further develop the simulation environment, supporting both dimensioning algorithms and resiliency features.

A number of key research goals have been achieved (the relevant deliverable and objective are indicated between brackets):

- Accurate job demand models for optical Grid research (D5.1 – necessary for most objectives).

- Providing and enforcing QoS parameters to Grid end users by appropriate resource management (D5.2 – Objective 1).

- Advanced Grid routing algorithms, including anycast services and physical impairment-aware optimizations (D5.3 – Objective 1).

- Provisioning of advance reservations, to allow time-based reservations and increase network performance (D5.4 – Objective 1).

- Recommendations for control plane design contains a requirements analysis to support Grid service delivery and the effect on the supporting control plane (G2MPLS). The deliverable's addendum contains scalability studies and performance analyses of the G2MPLS control plane and Harmony software stacks (D5.5 – Objective 2).

- Simulation environment for analysis of all previously mentioned Grid functionality (D5.6 – Objective 3).

- Algorithms for the planning and design of optical Grid networks (D5.7 – Objective 4)

- Strategies and algorithms for realizing resiliency in optical Grid networks (D5.8 – Objective 5)

- An extended simulation environment which supports dimensioning, resiliency and hybrid optical switching paradigms (D5.9 – Objective 6)

Most results obtained have a large impact on the research sector, as most research pushed the state-of-the-art of currently available scientific studies. With regard to the industry sector, two major achievements can have a potentially large impact:

- The simulation environment (Figure) allows highly detailed simulations of various scenarios for optical Grid infrastructure. As such, it can be highly useful for network operators and/or consultants, seeking to better understand their existing operations or future deployments. The simulations allow rapid evaluation of various networking strategies while offering high quality results corresponding to actual networking infrastructure.

- A number of novel resource management and job routing algorithms have been proposed in WP5 that improve both resource utilization and blocking performance. As such, network operators may increase their operational efficiency and obtain competitive advantages compared to operators of traditional algorithms.

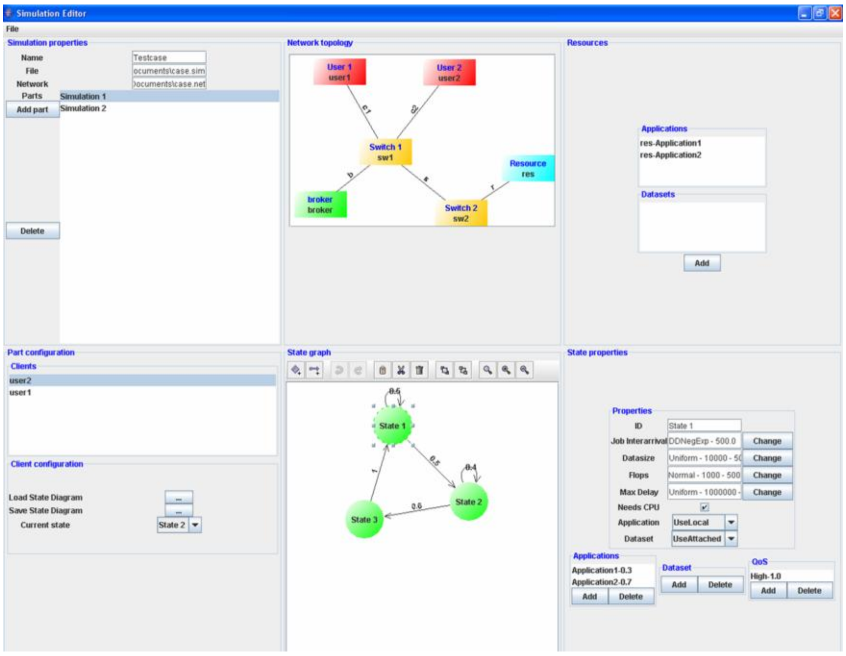

Figure 1-28: Graphical User Interface of simulation environment

1.4.6 Test-bed and demonstration activities

Partners involved in this workpackage: PSNC (leader), ADVA, CESNET, NXW, FHG, i2CAT, FZJ, SURFnet, UNIBONN, UESSEX, CRC

One of the main assumptions of PHOSPHORUS was that the project propositions and developments should be validated and demonstrated in a real advanced optical network. To achieve this, the project built a distributed testbed in which the project outcome was verified with a set of real scientific applications in a set of real-life scenarios. This way the testbed emulated a modern GRID environment in which demanding applications running on computational nodes used a transmission network to exchange data between the nodes, and access external devices.

The PHOSPHORUS testbed consisted of multiple local testbeds located in several sites in Europe and Canada, in which the switching and GRID resources were actually located and the testbed operations were executed. Each local testbed was constructed from different resources and supported different project needs which gave the testbed the heterogeneity needed to achieve the project goals, and demonstrate that the project developments were not limited to any single technology or any product family. The distribution of testbed also created natural technology and administrative domains, which allowed for demonstration and verification of the project results in a multi-domain environment.

The local testbeds were constructed from communications equipment (optical switches, TDM switches, Gigabit Ethernet switches, and optical transport systems) and advanced GRID resources (e.g. computing nodes, visualisation and storage resources). The communications equipment was the main platform of the project's developments and implementations. The developlements of other workpackages were focused on different aspects of control of the equipment and provisioning of services on the equipment. The GRID resources were used by a set of scientific applications which utilised the testbed and served to assess the benefits of the project developments for real applications.

The communication equipment was controlled by the control plane proposed by the project as an extension to GMPLS (G^2MPLS) and three Network Resource Provisiong Systems (UCLP version 2, DRAC and ARGON) as well as the Harmony system invented by the project.

There were eight local testbeds used for the project:

- The PSNC local testbed consisted mainly of three types of switching (fiber switching capable – Calient DiamondWave FiberConnect, lambda switching capable – ADVA and Layer 2 switching capable – Brocade, former Foundry Networks, and Allied Telesyn) as well as GRID nodes and storage used by applications. The local testbed was used for testing the G^2MPLS control plane and enhancements to applications. The switching infrastructure of the testbed represented all three types of switches addressed by Workpackage 2 in the development of G^2MPLS which allowed multidomain tests in which two or three technological domains were controlled by the same control plane – G^2MPLS.

- The CESNET local testbed, which was equipped with fiber switching capable switches developed by CESNET (CzechLight Switches) was used for testing the G^2MPLS software with focus on the Transport Network Resource Controller for the switches.

- The i2CAT local testbed was equipped with a DWDM ring (NORTEL OPTera Metro 5200) with a capability to create a Gigabit Ethernet connection and teardown it. The DWDM system was controlled by UCLP version 2. The local testbed was also equipped with a set of computers used by the TOPS application. The testbed was used mainly for testing NRPS interoperability and Harmony.

- The SURFnet local testbed contained three SDH cross connects – Nortell 6500 MSPPs. The three nodes were connected via 10 Gbps links, building a triangle. Each of the nodes had several GE ports and was able to switch each GE port to any other GE port using SDH interconnections. The cross connects were controlled by DRAC. The testbed was used mainly for testing NRPS interoperability and Harmony.

- The UvA local testbed was focused on testing Authentication, Authorization, and Accounting mechanisms developed by workpackage 4. The testbed was built of three computer clusters and switching resources (lambda capable switches – Calient and GlimerGlass, Layer 2 capable switches – Nortel 8600 and Force 10). The testbed was controlled by DRAGON and GMPLS.

- The UESSEX local testbed comprised of one fiber switching capable switch (Calient DiamondWave FiberConnect) partitioned into four independent virtual switches and several GRID computers. The Calient switch was controlled by G^2MPLS. The testbed was used to test the G^2MPLS control plane.

- The VIOLA testbed contributed jointly by three partners (FHG, FJZ and UniBonn) and distributed over Germany. The testbed contained three SDH cross connects (Alcatel 1678 MCC), which switch SDH connections by using GMPLS. The cross connects were able to switch Gigabit Ethernet lines between their access ports. The testbed also contained three Layer 2 switching capable switches (Riverstone 15008) and several GRID nodes. The switching infrastructure in the local testbed was controlled ARGON. The testbed was used for testing NRPS interoperability, Harmony, and enhancements to applications and GRID middleware.

- The CRC local testbed comprised a DWDM system (NORTEL OPTera Metro 5200), Layer 2 switching capable switches (Cisco Catalyst, Force 10), computational nodes and visualisation equipment. The switching infrastructure was controlled by UCLP version 2. The testbed was used mainly for testing NRPS interoperability and Harmony.

There were several modern scientific applications deployed in the testbed. The applications were selected and delivered by Workpackage 3. The selected applications required much bandwidth so they were able to utilise the transmission infrastructure of the testbed. The applications were used in order to validate the benefits which real applications can have of the modifications to optical networks which the project proposed.

For the integration of the whole PHOSPHORUS testbed, all local testbeds had to be interconnected on the data plane as well as on the control/provisioning plane. The data plane connections were used to transit user data between GRID resources located in different local testbeds while the control/provisioning plane connections were used for integration of the control planes (GMPLS, G^2MPLS) of local testbeds as well as integration of NRPSes to allow for signalling between them and multi-domain processing of users' requests.

The data plane connectivity was on dedicated lightpaths capable of transmitting huge amounts of data – the amounts which will be generated by the PHOSPHORUS applications. As the PHOSPHORUS project was developing optical technologies and the testbed had to allow for demonstration of the project developments, it was decided that the data plane would be built as an optical network with switching capabilities in local testbeds and transparent lightpaths between local testbeds. The lightpaths were capable of transmitting Gigabit Ethernet frames. This transmission technology allowed for simple integration with the GRID infrastructure which uses Gigabit Ethenet as the major transmission technology.

The topology of interconnections between local testbeds resulted from requirements from other PHOSPHORUS activities – the activities which utilised the testbed to demonstrate their developments. To make the testbed as cost-effective as possible, the lightpaths were obtained from several providers, especially those who provide transmission services to the research community free of cost or whose services are included in subscription fees already paid by the community. For this reason most of the lightpaths used the GÉANT+ infrastructure which allows for provisioning lightpaths between European NRENs and between Europe and the United States. The GÉANT+ lightpaths were complemented by lightpaths from other sources (e.g. CANET4 in Canada, Cross Border Dark Fibre provided by NRENs, GLIF, and others).

The topology of connections is depicted in Figure 1-29.

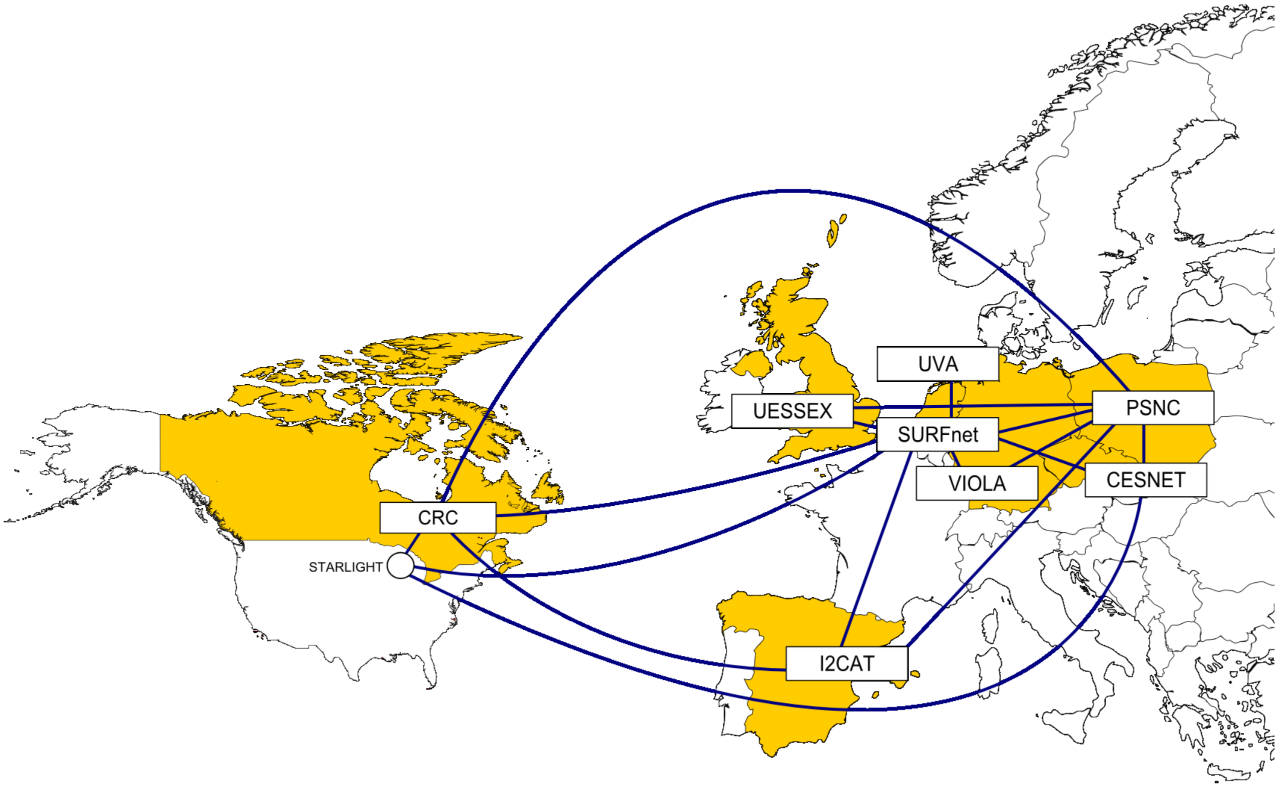

Figure 1-29: Topology of the global testbed

The project demonstrated an incremental approach to the testbed construction. The developments of other PHOSPHORUS activities were gradually deployed in the testbed when they were available from their developers. The most important project developments deployed in the testbed were the enhanced Control Plane for optical networks – G^2MPLS, the Harmony system which integrates several Network Resource Reservation Systems, the interface between G^2MPLS, and Harmony which allows for processing requests in G^2MPLS-controlled domains and NRPS-controlled domains, client interfaces for applications and GRID middleware which allowed them to request services from the network and security features which allowed for authorisation of requests. The developments are described in the previous chapters of this report.

The developments of the project were tested in the testbed in order to verify their applicability in optical networks as well as assess the benefits that modern applications can have when such developments are used in a network. The tests proved that the ideas of the project can be implemented in a network and can be beneficial for modern applications. The tests showed that the project developments met the expectations which formed the basis for the project.

The tests were accompanied by a set of live demonstrations in which the project developments were demonstrated to external audience. Such demonstrations were organised on the most important conferences which are attended by networking researchers and professionals, like TERENA Networking Conference in Europe and Supercomputing conference in the USA. The aim of the demonstrations was to gather interest of the society towards the ideas of the project as well as encourage potential users to deploy and use the project developments in their networks. The live demonstrations used the testbed in order to show the project developments in a real optical network with some real scientific applications. This way the PHOSPHORUS consortium can prove that the ideas of the project may be implemented in real networks and can be beneficial for real applications.

The workpackage was concluded with identification of new types of resources in optical networks which are currently not controlled by the Network Service Plane and Control Plane and which may be controlled with benefits for GRIDs and applications.

1.4.7 Dissemination, Contributions to standards, Liaisons

Partners involved: UESSEX (leader), PSNC, NXW, FHG, i2CAT, CTI, AIT SARA, UNIBONN

The objective of work package 7 is to propagate and publicize the results of the project within and beyond the consortium. This objective includes the promotion and exploitation of the knowledge derived from the various activities of the project. This objective was divided into three main tasks: Dissemination Activities, Contribution to Standards and Liaisons with other organisations and projects.

1.4.7.1 *Dissemination Activities*

The dissemination of the results achieved during the Phosphorus project was acknowledged as one of the most important tasks of the consortium. The strategy taken by the consortium to disseminate the results of the project was two-fold: (I) External Dissemination to the scientific community, and the general public; (II) Internal Dissemination of the results achieved within each work-package to other work-packages within the project. This strategy ensures smooth running and wide visibility of the achievements of the project.

A. Internal Dissemination

Internal dissemination was acknowledged as important to ensure the smooth progress of the project. As such, a number of resources were setup to aid in internal dissemination. These included a project wiki, frequent project and work-package meetings, including management meetings, technical board meeting and general assembly meetings. A private section of the website was also setup which allowed partners to register for, plan and organise these meetings. In addition, a mailing list was setup, one for each work-package and one for each inter-work-package as required.

B. External Dissemination

A structured dissemination plan was carried out during the Phosphorus project to ensure effective visibility and promotion of the results of the project to stakeholders, potential beneficiary network operators as well as potential users of the Phosphorus network (National Research and Education Network (NREN) and Grid Users). The dissemination of the activities, knowledge and technology produced from the project beyond the consortium was achieved via standardisation drafts, international demonstration and test-bed activities, publications, conferences, workshops, public deliverables and the project website.

International Demonstration and Test bed Activities

During the duration of the project, PHOSPHORUS partners demonstrated the results of the project in real network environments with a set of real scientific applications. The demonstrations made is possible for the PHOSPHORUS partners to not only disseminate the knowledge and achievements of the project to a wide external audience at major conferences, but also to verify the results of the projects. This activity was done in collaboration with work-package 6, which is responsible for building and utilising test bed demonstrations. The demonstrations consisted of a set of local test-beds located at the premises of several project partners. The local test-beds were interconnected via dedicated channels in GEANT2, to create the global PHOSPHORUS testbed. This allowed for global and local demonstrations.

At the end of the project PHOSPHORUS successfully carried out over 20 separate demonstrations at 15 major conferences such as Super-Computing Conference, TERENA Networking Conference, ECOC and ICT. [D7.1.3, D6.7 D7.3.1]. The main project demonstrations of the developments of the project include:

- Harmony system demonstrations:
  - ONDM'08, TNC'08, OGF 23, ICT'08: a high bandwidth provisioning for HD video streaming via UCLP, DRAC and ARGON domains,
  - TNC'09: anatomical data transmitted from HSVO to demo booth over provisioned light-path,
  - TNC'09: a high bandwidth provisioning for HD video streaming from KISTI domains to demo booth,
- G^2MPLS Control Plane demonstrations:
  - ICT'08 and TNC'09: multi-domain path provisioning for DDSS backup file application,
  - TNC'09: Ethernet anycast path provisioning for KodaviS application,
- Generalised Token Based Networking demonstration:
  - TNC'09: multiple applications share common network resources in token authorized environment.

Publication

The outcomes and progress of the project were published and document in a wide variety of media. These included publications in conference and workshop proceedings, journal articles, references in the press, as well as PHOSPHORUS public documents such as the deliverables. During the projects, PHOSPHORUS partners attended more than 109 conferences, workshops and meetings.

- Journal Articles, Conferences and Workshop Proceedings
  - At the end of the project, PHOSPHORUS partners had successfully published over 140 articles, papers and posters (50 workshop and meeting proceedings, 68 conference proceedings, and 24 journal articles) in journals and conference/workshop proceedings. The publications cover all aspects of PHOSPHORUS project in the areas of GMPLS and G2MPLS control plane and architecture, Harmony Provisioning System, Grid job routing algorithms, QoS-aware resource scheduling, optical networks and lightpath monitoring and provisioning [D7.1.3, D7.1.2].

- Public Deliverables
  - A requirement of the PHOSPHORUS project was to publish public deliverables that documented the progress of the project. These deliverables also ensured that the PHOSPHORUS partners were on track in achieving their goals. The public deliverables are available on the PHOSPHORUS website (www.ist-phosphorus.eu).

- Magazine Articles
  - During the project, PHOSPHORUS consortium members published two magazine articles. These articles were published in reputable magazines with a wide audience:
    - "Enabling Grid-Network Services via Control Plane: the PHOSPHORUS G2MPLS way to the e-Infrastructures" published in the April 2009 Issue of the Zero-In Magazine
    - "Empowering Grid users with improved services" published in the April 2008 Issue of the eStrategies magazine on Advanced Grid Technologies.

- References in the Press:
  - In addition to the magazine article, the work and effort of PHOSPHORUS partners were recognised by parties outside the consortium. This recognition resulted in three references in press releases [D7.1.4].

Project Website and Other Informational Material

Developing and maintaining a project Web site addressed a main dissemination objective. Through the web site, the public documents produced within the PHOSPHORUS project was made accessible, along with a brief summary of the PHOSPHORUS project, all events that took place during the duration of the project, a news section and a comprehensive contact list of all partners involved in the project. In addition links to specific information about the PHOSPHORUS technological issues are also detailed on the website. (www.ist-phosphorus.eu)

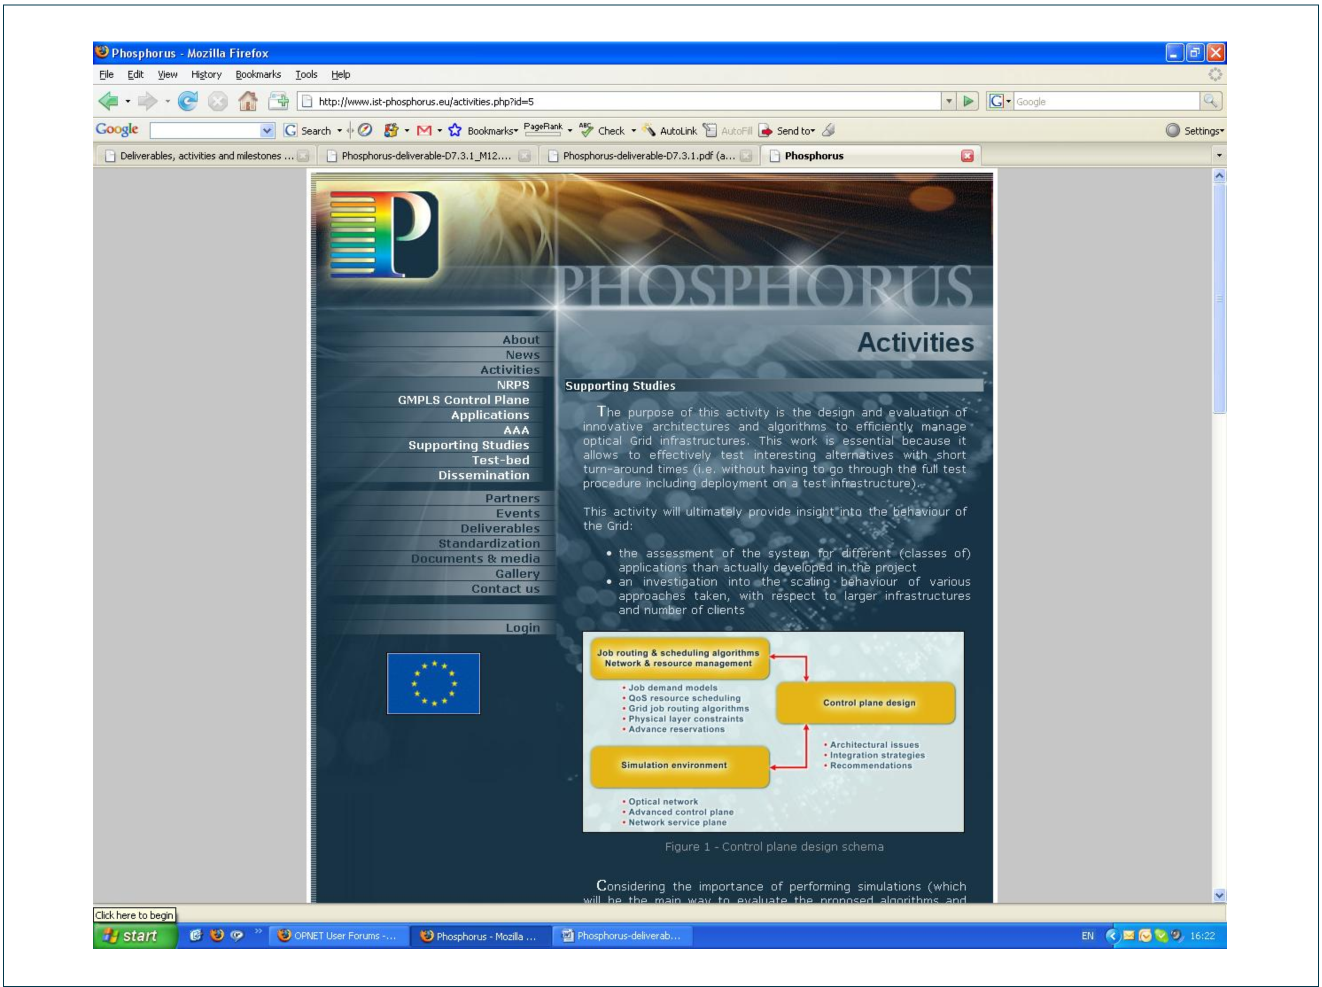

Figure 1-30: Homepage of the updated PHOSPHORUS website

In addition to the website, a number of informational material were produced during the lifetime of the project. These included an official PHOSPHORUS project wall poster that outlines the objectives, architecture, test-bed, partners and the projects contact details. The posters were distributed to all members of the PHOSPHORUS consortium for display in their companies and universities. The posters were also used in PHOSPHORUS booths at conferences and workshops.

A brochure with a one-page description of the objectives of each work-package was prepared. The brochures were widely distributed at conferences and workshops as well as to visitors that attended tours of the laboratories of PHOSPHORUS partners.

Two videos (one short, one long) detailing the ideas and objectives of the PHOSPHORUS project were produced. The long version includes interviews with the work-package leaders. Other information materials prepared by PHOSPHORUS include white T-shirts, labels and stickers and sweets with the PHOSPHORUS logo on it.

Raising Public Awareness and Participation

The PHOSPHORUS consortium recognised the need for disseminating the activities of the project to the wider public [D7.1.4]. To achieve this, the PHOSPHORUS consortium organised workshops targeted at the wider scientific community. Other means to raise public participation involved promoting engineering, optical networking and the PHOSPHORUS project to the engineers of the future from secondary school and universities. These activities included organising seminars, invited speakers, tutorial and lab tours and organisation and participation in summer schools. The activities drew a wide audience with diverse backgrounds. In addition, PHOSPHORUS partners actively participated in activities to raise gender issue awareness and promote engineering as a gender-neutral field.

- Wider Scientific Community - PHOSPHORUS-Organised Workshops
  - To raise public awareness and participation, PHOSPHORUS partners organised and co-organised 6 workshops and 2 “Birds of a Feather”(BoF) workshops. The workshops were widely advertised and gave the wider community of research student and experienced researchers the opportunity to submit papers to these workshops. The BoFs facilitated collaborations with projects and individuals with similar interests to those of PHOSPHORUS.

    i. On-Demand Network Services for the Scientific Community (co-located with TERENA 09)
    ii. Workshop on Grid vs Cloud Computing and Why This Should Concern the Optical Networking Community (co-located with OFC/NFOEC 09)
    iii. Workshop on High Performance Grid Networks (co-located with CCGRID 08)
    iv. PHOSPHORUS-Federica Tutorial and Workshop (co-located with TERENA 08)
    v. PHOSPHORUS-Carriocas Workshop
    vi. Networks for IT: A new opportunity for optical network technologies (co-located with ECOC2007)
    vii. BoF session on "Service Level Agreements” at ISC07

viii. BoF session on "Delivery of Network Services across Heterogeneous Optical Domains" at SC06

- Future Engineers & Youth Activities
  - PHOSPHORUS partners offered opportunities for students to experience what engineers do, and to see and visit the laboratories and super computing and networking rooms. During these visits, students and visitors were given the opportunity to meet experienced engineers and research students, and were introduced to the PHOSPHORUS project. In some cases, this was done within the context of festive and 'fun' activities with the aim of highlighting the appeal of pursuing research careers in optical networking and Grid.

- Gender Action Activities
  - To address the issue of under-representation of women in science, engineering and technology, PHOSPHORUS members participated in a number of collaborations and discussions to promote engineering, science and technology as a gender-neutral field. During the project, PHOSPHORUS members successfully co-organised and supported a Gender Issues Workshop. The aim of the workshop was to present a platform to discuss the development and support of female engineers and scientists. PHOSPHORUS members continually support and promote gender action activities by mentoring female engineers and organising visiting research collaborations with special emphasis on gender awareness. The PHOSPHORUS consortium advocates maximising diversity as a means to enhance the quality of research [D0.2].

1.4.7.2 *Contribution to Standards Activities*

The second of work-package 7's task is to contribute to standardisation activities relevant to the work and results achieved within the PHOSPHORUS project. To this end, PHOSPHORUS partners have been actively involved in contributing to standardisation activities of the Open Grid Forum (OGF) and the Internet Engineering Task Force (IEFT).

A. OGF

PHOSPHORUS participation to the OGF standardisation activities includes active contribution in the following groups:

FI-RG	Firewall Issues Research Group
FVGA-WG	Firewall Virtualization for Grid Applications Working Group
GHPN-RG	Grid High Performance Networking Research Group
GRAAP-WG	Grid Resource Allocation Agreement Protocol Working Group
GSA-RG	Grid Scheduling Architecture Research Group
JSDL-WG	Job Submission Description Language Working Group
NML-WG	Network Markup Language Working Group
NM-WG	Network Measurements Working Group
NSI-WG	Network Service Interface Working Group
OGSA-AUTHZ-WG	OGSA Authorisation Working Group

OGSA-RSS-WG	Open Grid Service Architecture Resource Selection Working Group
WFM-RG	Workflow Management Research Group

Involvement in the above groups resulted in 14 standardisation drafts and draft proposals:

1. Grid Optical Burst Switched Networks – GOBS
2. Grid Optical User Network Interface (G.OUNI)
3. Web Services Agreement Specification
4. Grid Scheduling Architecture – Requirements
5. Grid Scheduler Interaction
6. NDL – Network Description Language
7. Network Topology Descriptions in Optical Hybrid Networks
8. Firewall Issues Overview
9. Requirements on Operating Grids in Firewalled Environments
10. Specification of the Multi-step Negotiation Protocol
11. The Grid Scheduler Interaction
12. XACML Request Context to Obtain an Authorization Decision
13. Use of WS-Trust and SAML to access a CVS
14. Firewall Transversal Protocol (in preparation)

B. IETF

The project also contributed to the Internet Engineering Task Force (IETF) where few project members are active in the following group:

IETF ForCES-WG	Internet Engineering Task Force Work Group on Forwarding and Control Elements Separation

Three IEFT standardisation drafts and drafts proposals incorporated direct outputs from the PHOSPHORUS project:

1. ForCES Protocol Specification [IETF Draft]
2. ForCES Protocol Specification [IETF Draft]
3. ForCES Forwarding Element Model [IETF Draft]

Working groups are mainly focused on the development of standards while research groups focus on presentation and discussion of on-going research. Two of the above research groups are co-chaired by PHOSPHORUS partners. The NSI-WG was formed as a direct result of the work of members of the PHOSPHORUS project. The results of standardisation efforts within the OGF and IETF result in spin-off working groups that target standards. Full details of the standardisation efforts within the PHOSPHORUS project are available in the PHOSPHORUS deliverable [D7.2.2].

1.4.7.3 *Liaisons Activities*

The success of PHOSPHORUS is enhanced through collaborative work between the PHOSPHORUS project and other project within the European Union (EU) and non-EU projects, which ranges from collaborative discussions to the deployment of the results of the PHOSPHORUS projects on other projects test-beds and vice-versa. To this end, one of PHOSPHORUS' main dissemination strategies was to build a strong collaborative framework for participation to test-bed activities from within and external to the EU. To achieve this objective, the PHOSPHORUS consortium included non-EU partners such as MCNC (USA), Nortel (USA), CRC (Canada).

During the project, PHOSPHORUS partners successfully built such collaborative relationships with a number of projects. These projects all carry out research on dynamic provisioning of services that require network and computing resources. These collaborations provided the opportunity to share expertise and technical knowledge, including practical collaboration such as the deployment of PHOSPHORUS results on a collaborating project's test-bed. The main liaison activities are summarised below. Detailed descriptions can be found in the PHOSPHORUS public deliverable [D7.3.1, D7.3.3].

A. EU Collaborations

GEANT2

The PHOSPHORUS project maintains a close collaboration with GÉANT2 project and activities in order to achieve a synergy effect and improve future network services. The JRA3 Bandwidth on Demand (BoD) activity within the GEANT2 project was identified and selected as the liaison point with a high probability of visible common benefits expected based on the limitation of the duration of the PHOSPHORUS project. The AutoBaHN of the GEANT2 project was interoperated with the Harmony and G2MPLS systems of the PHOSPHORUS project. Reservation of requests going between Harmony test-bed and AutoBaHN test-bed was successfully tested. Meanwhile, interoperability between AutoBaHN and G2MPLS domains is promising and has been sketched for future implementation.

Alcatel-Lucent, France Telecom and the CARIOCAS project

The collaboration between PHOSPHORUS, Alcatel-Lucent and France Telecom was initiated to demonstrate the viability of the end-to-end on demand service delivery in multi-domain environment, involving different provisioning technology and heterogeneous network equipment integrated with Grid middleware. Through the collaboration with France Telecom, PHOSPHORUS partners initiated a cooperation with the CARIOCAS project, in which both France Telecom and Alcatel-Lucent, play a major role. Together with the CARIOCAS project, PHOSPHORUS partner co-organised the PHOSPHORUS-CARIOCAS two-day workshop in Paris, France. The participants constituted a balanced mix of experts from Telecom industries, universities, and national laboratories, all representing various Grid technology areas as well as scientific applications.

Federica Project

The FEDERICA project aims to create a European wide "technology agnostic" virtualized infrastructure made of Gigabit circuits, transmission equipment and computing nodes to support and host experimental activities on new Internet architectures and protocols [D7.3.3]. An interesting collaboration between PHOSPHORUS project and the FEDERICA project was initiated to combine the computing facilities of the PHOSPHORUS project and a slice of FEDERICA's virtualized network. The collaboration effort aims to exploit this infrastructure to stress test the results of the PHOSPHORUS project over a wide scenario. The FEDERICA project gains a real user utilizing its virtualized infrastructure. One of the outcomes of this collaboration was a joint organized workshop, the PHOSPHORUS-FEDERICA workshop.

RiNGrid Project

Remote Instrumentation in Next-generation Grids is a 18 months project co-funded by the EC under FP6. The project One of the objectives of the RinGrid project is the "careful analysis of the synergy between Remote Instrumentation and next-generation high speed communications networks and grid infrastructure as a basis for the definition of recommendations for designing next-generation Remote Instrumentation Services" [RinGrid]. The RiNGrid project consortium has expressed great interest in the G2MPLS developed within the PHOSPHORUS project. Collaboration between the two projects resulted in a RinGrid workshop.

EGEE Project

The European Commission-funded Enabling Grids for E-sciencE (EGEE) project aims to build on recent advances in grid technology and develop a service grid infrastructure that is available to scientists 24 hours-a-day. Dialogue between PHOSPHORUS and EGEE projects has been initiated to explores a possible collaboration in the field of Grid middleware and network interoperation.

B. Non-EU Collaborations

CANARIE

Collaboration between CANARIE and PHOSPHORUS has been an active one. CANARIE have provided connectivity between the PHOSPHORUS global testbed and Canada. In addition, CANARIE and PHOSPHORUS share mutual partners, CRC and i2CAT. This interesting collaboration aims to create seamless interoperability between UCLP, DRAC, ARGON and GMPLS Control Plane under Harmony system umbrella. This resulted in the installation of the UCLPv2 in some of the PHOSPHORUS partners' local test-beds (CRC, UESSEX, I2CAT). Finally, one of CANARIE's applications, the HSVO was used to show how external applications, such as HSVO, can achieve the inter-operability with Harmony. The collaboration between CANARIE and PHOSPHORUS have been tested and disseminated via demonstrations at large conferences. [D7.1.3, D7.3.3].

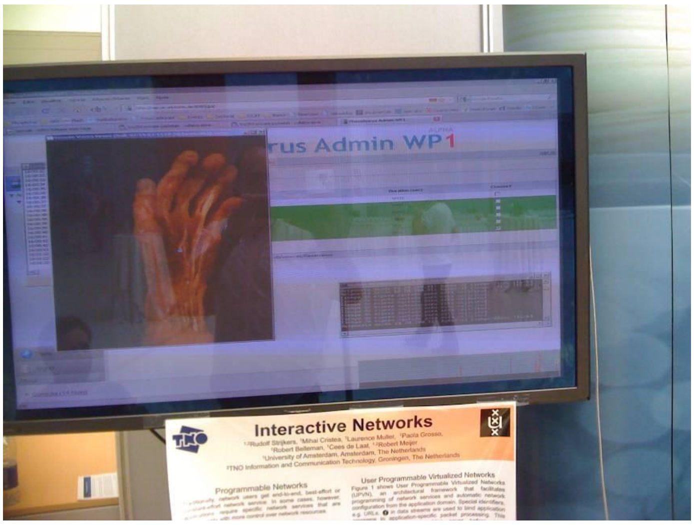

Figure 1-31: HSVO working with Harmony in TNC'09 conference

National LambdaRail/Enlightened Computing/Japan Gigabit Network/G-Lambda

The National LambdaRail/Enlightened Computing, Japan Gigabit Network/G-Lambda and the PHOSPHORUS project share a common goal for collaboration: provisioning services that require optical networks and grid middleware across three continents. All three projects are currently working on establishing standard and open interfaces which can be used for connecting different Bandwidth on Demand system components especially preparing a connection between G-Lambda, Enlightened Computing and PHOSPHORUS test-beds and demonstrating together an advance network and computing service delivery demonstrations. The collaboration between these projects resulted in, among others, a standardisation draft.

Internet2/DRAGON ESNET/OSCARS

Cooperation between Internet2/ESnet and University of Amsterdam (UvA) in the framework of the PHOSPHORUS project was two fold: First, to jointly develop an authorisation service for interdomain lightpath provisioning capable of integrating network resource provisioning systems developed both in the Internet2/ESnet and in the PHOSPHORUS project: OSCARS and NSP/Harmony. Second, to achieve interoperability between OSCARS/DCN and Harmony/NSP systems to provide cross-domain lightpaths. This cooperation resulted in the development of the token based interdomain-signalling concept and the definition of Token Validation Service (TVS) functionality. It was jointly tested via a demo at Supercomputing 2007 in Reno. At the SuperComputing conference 2008, the PHOSPHORUS-DCN interoperability was demonstrated. At this demo, connections through the Internet2 and PHOSPHORUS domains were setup by sending requests to the Harmony NSP.

Figure 1-32: Dutch booth at SC 08 (Austin, Texas) where the Harmony-IDC inter-operability demonstrations were performed

Korea Institute of Science and Technology Information (KISTI)

In 2008, a dialogue was established between Korea Institute of Science and Technology Information (KISTI) and PHOSPHORUS. The collaboration was to investigate the deployment of the results of PHOSPHORUS project (Argia and Harmony) as a Network Resource Provisioning System at KISTI's dedicated network, the KoreLight network. KISTI signed the collaboration agreement with the PHOSPHORUS project. In the beginning of 2009, the test-beds were interconnected and Argia was successfully installed. To allow KISTI join the Network Service Plane of Harmony in the PHOSPHORUS project, the Harmony NRPS Adapter for Argia were set up at KISTI. During TERENA conference 2009 in Malaga (Spain), the collaboration effort between the two parties was successfully demonstrated. The demonstration presented the setup of a network path for the purpose of transmitting an on-demanding HD video streaming from Korea towards the conference venue.

2 Dissemination and use of the knowledge

For the exhaustive description of exploitable results, list of workshops, conferences and Phosphorus publications please refer to Deliverable D7.1.2 'Plan for using and disseminating the knowledge' [D712] available at: http://www.ist-phosphorus.eu/files/deliverables/Phosphorus-deliverable-D7.1.2_M33.pdf

3 References

[D0.2] Phosphorus WP0, “Report on the Gender Issue Plan”, Deliverable D0.2, December 2006

[D7.1.2] Phosphorus WP7, “Plan for using and disseminating the knowledge”, Deliverable D7.1.2, June 2009

[D7.1.3] Phosphorus WP7, “Annual report on dissemination activities”, Deliverable D7.1.3, June 2009

[D7.1.4] Phosphorus WP7, “Plan for raising public participation and awareness” Deliverable D7.1.4, June 2009

[D7.2.2] Phosphorus WP7, “Annual report on standardisation activities” Deliverable D7.2.2, June 2009

[D7.3.1] Phosphorus WP7, “Cooperation agreements with EU NRENs and projects”, Deliverable D7.3.1, June 2009.

[D7.3.3] Phosphorus WP7, “Annual report on EU and non-EU collaboration and technical liaison activities”, Deliverable D7.3.3, June 2009

[D6.7] Phosphorus WP6, “Plan of demonstrations of the project results”, Deliverable D6.7 March 2008

[RFC2904] Vollbrecht, J., P. Calhoun, S. Farrell, L. Gommans, G. Gross, B. de Bruijn, C. de Laat, M. Holdrege, D. Spence, "AAA Authorization Framework," Informational RFC 2904, Internet Engineering Task Force, August 2000. ftp://ftp.isi.edu/in-notes/rfc2904.txt

[TBN2009] Gommans, L., L. Xu, Y. Demchenko, A. Wan, M. Cristea, R. Meijer, C. de Laat, Multi-Domain Lightpath Authorization using Tokens, Future Generations Computer Systems, Vol 25, issue 2, February 2009, pages 153-160

[TBN2007]	"The Token Based Switch: Per-Packet Access Authorisation to Optical Shortcuts", by Mihai-Lucian Cristea, Leon Gommans, Li Xu, and Herbert Bos, in Proceedings of IFIP Networking, Atlanta, GA, USA, May 2007
[D4.3.2]	"ForCES Token Based Switch Design and Implementation ", Phosphorus Project Deliverable D4.3.2, September 30, 2008. [Online]. Available: http://www.ist-phosphorus.eu/files/deliverables/Phosphorus-deliverable-D4.3.2.pdf
[RFC3746-ForCES]	Yang, L., Dantu, R., Anderson, T., and R. Gopal, "Forwarding and Control Element Separation (ForCES) Framework", RFC 3746, April 2004.
[ForCES-Protocol]	Dong, L., Doria, A., Gopal, R., HAAS, R., Salim, J., Khosravi, H., and W. Wang, "ForCES Protocol Specification", http://www.ietf.org/id/draft-ietf-forces-protocol-22.txt (work in progress), March 2009.
[IWAN2005-ForCES]	Haleplidis, E., Haas, R., Denazis, S., Koufopavlou, O., "A Web Service- and ForCES-based Programmable Router Architecture", IWAN2005, France.
[Grid2008]	Demchenko Y, A. Wan, M. Cristea, C. de Laat, Authorisation Infrastructure for On-Demand Network Resource Provisioning, Proceedings The 9th IEEE/ACM International Conference on Grid Computing (Grid 2008), Tsukuba, Japan, Sept. 29 - Oct. 1, 2008. Pp. 95-103. IEEE Catalog Number CFP08GRI-CDR, ISBN 978-1-4244-2579-2.
[XACML-Grid]	"An XACML Attribute and Obligation Profile for Authorization Interoperability in Grids," Joint EGEE, OSG, and Globus document. [Online]. https://edms.cern.ch/document/929867/1
[XACML-NRP]	Demchenko, Y., C. M. Cristea, de Laat, XACML Policy profile for multidomain Network Resource Provisioning and supporting Authorisation Infrastructure, IEEE International Symposium on Policies for Distributed Systems and Networks (POLICY 2009), July 20-22, 2009, London, UK. Accepted paper.
[D4.3.1]	"GAAA Toolkit pluggable components and XACML policy profile for ONRP", Phosphorus Project Deliverable D4.3.1, September 30, 2008. [Online]. Available: http://www.ist-phosphorus.eu/files/deliverables/Phosphorus-deliverable-D4.3.1.pdf
[D4.5]	"Updated GAAA Toolkit library for ONRP (final project release)", Phosphorus Project Deliverable D4.5, March 30, 2009. [Online]. Available: http://www.ist-phosphorus.eu/files/deliverables/Phosphorus-deliverable-D4.5.pdf

4 Acronyms

AAA	Authentication, Authorisation, and Accounting
ARGON	Allocation and Reservation in Grid-enabled Optical Networks
CCGRID	Cluster Computing and the Grid
DRAC	Dynamic Resource Allocation Controller
ECOC	European Conference and Exhibition on Optical Communication
EGEE	Enabling Grids for E-sciencE
EU	European Union
GLIF	Global Lambdagrid
GUNI	Grid User Network Interface
HSVO	Health Services Virtual Organization
ICT	Information and Communication Technology Conference
IDB	Inter-Domain Broker
IETF	Internet Engineering Task Force
ISC	International Supercomputing Conference
KISTI	Korean Institute of Science and Technology Information
NREN	National Research and Education Network
NRPS	Network Resource Provisioning System
OFC/NFOEC	Optical Fibre Communication Conference and Exposition and the National Fibre Optic Engineers Conference
OGF	Open Grid Forum
SC	SuperComputing Conference
TERENA	TransEuropean Research and Education Network Association
TOPS	Technology for Optical Pixel Streaming
UCLP	User Controlled LightPath