BETWEEN GRIDS AND NETWORKS: GRID-ENABLED NETWORK CONTROL PLANES

Bartosz Belter [†], Artur Binczewski [†], Gino Carrozzo [‡], Nicolla Ciulli [‡], Eduard Escalona [*], George Markidis [∞], Reza Nejabati [*], Dimitra Simeonidou [*], Maciej Stroiński [†], Anna Tzanakaki [∞], Georgios Zervas [*]

[†] Poznan Supercomputing and Networking Center, ul. Noskowskiego 10, 61-704 Poznan, Poland
Email: {bart, artur, stroins} @man.poznan.pl

[‡] Nextworks, Corso Italia 116, 56125 Pisa, Italy
Email: {g.carrozzo, n.ciulli} @nextworks.it

** Department of Electronics Systems Engineering, University of Essex, Colchester, CO4 3SQ, UK*
Email: {eescal, rnejab, dsimeo, zervas} @essex.ac.uk

[∞] Athens Information Technology, P.O. Box 68 Markopoulo Ave., GR - 19002 Peania, Athens, Greece
Email: {gmar, atza} @ait.edu.gr

Abstract

Grid-GMPLS is a Network Control Plane architecture that implements the concept of Grid Network Services. In the PHOSPHORUS framework, GNS is a service that allows the provisioning of network and Grid resources in a single-step through a set of seamlessly integrated procedures. By providing a unified network/Grid infrastructure the control plane can flexibly adapt to the demands of applications having intensive requirements on both computational resources and network resources.

This article introduces the concept of the Grid-GMPLS control plane architecture with particular focus on new services, deployment models, and interoperability issues of Grid-GMPLS and GMPLS control planes.

Keywords

Network Control Plane, Bandwidth on Demand, Grid computing, Grid Network Service, GMPLS, G^2MPLS

Introduction

Today's high capacity optical networks can satisfy even the most demanding application requirements. Determinism (e.g. guaranteed QoS), shared data spaces, large data transfer or low latency are achievable through dedicated optical channels (lambdas). The subject of using the networking infrastructure in an automated way remains in the focus of many research projects. Starting from Layer 3, there is an effort put towards this goal – the GN2 AMPS enables authorized end-users to make a single reservation for Premium IP bandwidth that is effective across a chain of participating domains *(Patil, 2006)*. Going further, bandwidth provisioning for lower layers (from Layer 0 to Layer 2) are in the scope of many projects, including EU-funded (PHOSPHORUS, GÉANT2-JRA3) and non-EU (e.g. DRAGON/OSCARS from Internet2).

Accordingly, there has been a tremendous amount of research and development in the Grid community in terms of Grid services infrastructure and Grid application development. However, there is still a missing gap – a challenging issue in the area of using network as a first-class Grid resource. Currently there is no existing implementation 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.

By providing a unified network/Grid infrastructure, the Grid-enabled control plane *(Markidis,2007)* can flexibly adapt to applications' demands having intensive, combined requirements on CPU, memory and storage resources as well as on the communication network. The main innovation introduced here is a concept where the network (light-paths) and Grid (computational, storage) resources are provisioned in a single-step *(Figuerola, 2007)*: 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 real node-to-node deployment of on-demand Grid services.

Grid-enabled Network Control Plane – G^{2}MPLS

G^{2}MPLS [3] is a Network Control Plane architecture that implements the concept of Grid Network Services (GNS). GNS is a service that allows the provisioning of network and Grid resources at the same time and with the same priority (single-step) through a set of seamlessly integrated procedures. Joint co-allocation of Grid and network resources allows to configure network connections in the same tier of Grid resources by guaranteeing enhanced service availability (advantage for the user) tailored to the user requirements (advantage for the user and the network operator). G^{2}MPLS functionalities comprise discovery and advertisement of Grid resource capabilities and availabilities, GNS setup including resource localization, advance reservation, GNS monitoring, and recovery.

G^{2}MPLS Control Plane models

G^{2}MPLS Control Plane models concern the layering of grid and network resources. Thus, they have a different meaning and scope with respect to the IETF definitions for GMPLS Overlay, Augmented, and Peer/Integrated models *(Ciulli, 2007a)*.

G^{2}MPLS Overlay model

In the G^{2}MPLS Overlay model the Grid layer is supposed to have Grid and network routing knowledge in order to provide Grid resource configuration and monitoring (as in its standard behaviour) plus network resource configuration and monitoring (Figure 1). Most of the computational and service intelligence is maintained on the Grid layer and, upon the occurrence of a Grid job request from a user, it is up to the Grid scheduler to initiate and coordinate the reservation process.

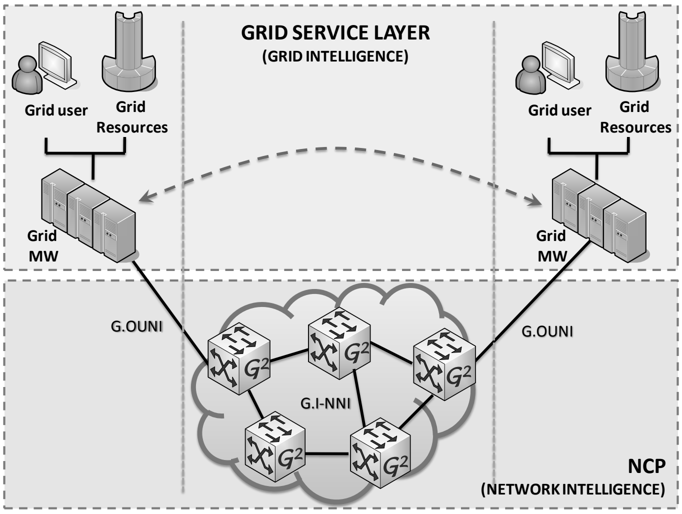

Figure 1 G^{2}MPLS Overlay model

Each involved Grid site is responsible for the scheduling, configuring, and monitoring of the job part to be run on its resources. This functionality is provided by the Local Resource Management System (LRMN).

The G^{2}MPLS is responsible for the scheduling, configuring, and monitoring of the job part related to the network, thus implementing advance reservations, recovery, and connection monitoring mechanisms.

G^{2}MPLS Integrated model

In the G^{2}MPLS Integrated model most of the functionalities for resource advance reservation are moved to the G^{2}MPLS Network Control Plane, and G^{2}MPLS is responsible for selecting the job segment providers and coordinating the co-allocation process (Figure 2).

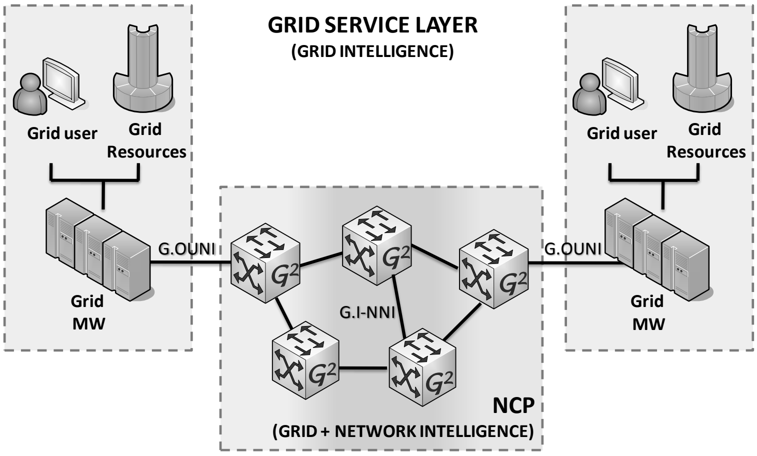

Figure 2 G^2MPLS Integrated model

Each involved Grid site is still responsible for the scheduling, configuring and monitoring of the job part to be run on its local resources. This functionality continues to be provided by the Local Resource Management System (LRMN).

On the contrary, G^2MPLS is responsible for the co-allocation of all the job parts, those related to the Grid sites and those related to the network. Therefore, Grid and network resources are seamlessly allocated in one step, under the coordination of the G^2MPLS NCP.

The Grid and network localization service is included in the G^2MPLS integrated model in order to have knowledge of the capability and availability of all the resources that can be involved in Grid Network Services instantiations. The Grid middleware/application could just declare the identity of the Grid site involved in the job, while it is up to the G^2MPLS to resolve the network attachment end-points for these sites. However, it could be possible for some applications to declare just an anonymous service requirement (e.g. an amount of storage or CPU) despite its location. In this case G^2MPLS is responsible also for localizing the best resource candidate for that service and the best network attachment point.

G^2MPLS network reference points

The deployment of the enhanced Grid-GMPLS Control Plane sets analogous reference points with respect to the ASON/GMPLS ones *(Ciulli, 2007b)*. The resulting network interfaces constitute Grid-aware evolution (Figure 3) of the standard interfaces (UNI, I-NNI, E-NNI) with a set of procedures that maintain the backward compatibility with the original ASON references:

- **G.OUNI**, i.e. the Grid Optical User-Network Interface that supports Grid and network signalling and discovery between the Grid site and the G^2MPLS domain;
- **G.I-NNI**, i.e. the Grid Internal Node-Node Interface (G.I-NNI) that supports the routing and signalling procedures between adjacent nodes;
- **G.E-NNI**, i.e. the Grid External Network-Network Interface that propagates Grid and network topology information across different Control Plane domains and supports the inter-domain signalling mechanisms;
- **SBI**, i.e. the Southbound Interface that retrieves the resource status from the specific Transport Plane and translates Control Plane actions into appropriate configurations of those resources;
- **NBI**, i.e. the Northbound Interface that groups two interfaces towards upper layers: one towards the Grid layer (G.NBI) and one towards the Network Service Plane (including NRPS, N.NBI).

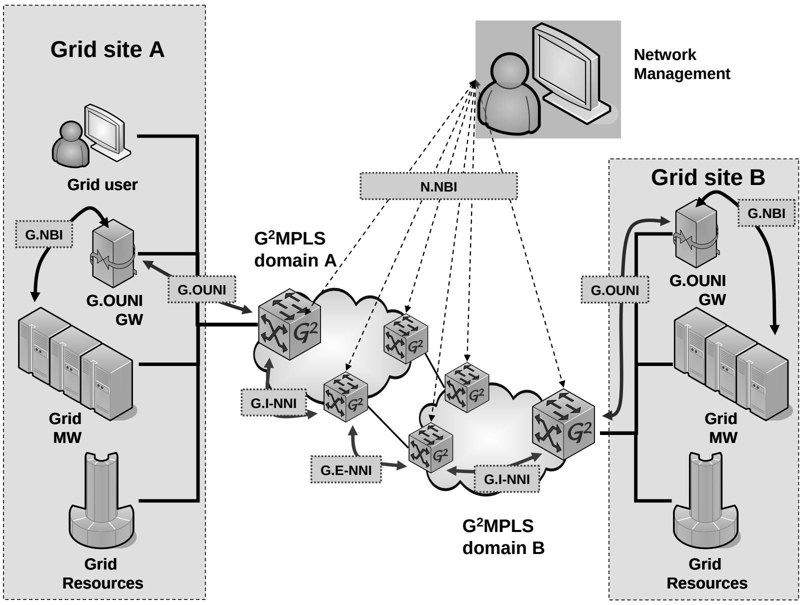

Figure 3 G^{2}MPLS network reference points

G^{2}MPLS Use Cases

In this chapter some use-cases for the G^{2}MPLS architecture are provided with the aim of identifying the roles and requirements for the system components. Those use-cases involving local and remote Grid resources are described under the assumption of just one call in one GNS transaction. This does not compromise the scope and general applicability of the architecture, since more complex use-cases can be obtained by the composition of these basic ones.

Loopback GNS transaction

The first use-case provides the details of the Loopback GNS transaction. In this scenario a Grid user configures and requests the execution of a job that involves Grid resources located only at his/her local grid site – Vsite A (Figure 4).

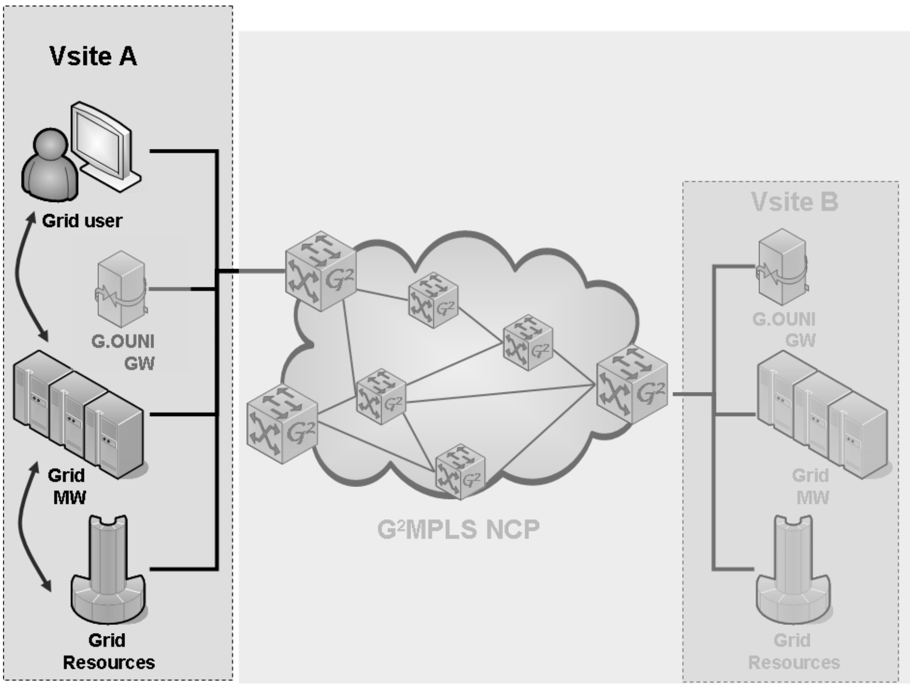

Figure 4 Loopback GNS transaction

As a triggering action, the Grid user defines the job in the Client Application and issues the request. The Client Application sends a request to local Middleware at Vsite A. Next, the Grid broker in the middleware resolves the local availability of Grid resources needed for the job, and schedules and configures the local Grid resources via Local Resource Management System (LRMS). After that the middleware sends job request confirmation to the Client Application. Client Application notifies the Grid User of the job request confirmation. Optionally, when advance reservation is requested, the middleware sends a job execution notification to the Client Application and then the Client Application notifies the Grid User of the job execution.

GNS transaction – direct invocation

A Grid User configures and requests the execution of a job that involves Grid resources located at his/her local Vsite A and Vsite B (Figure 5).

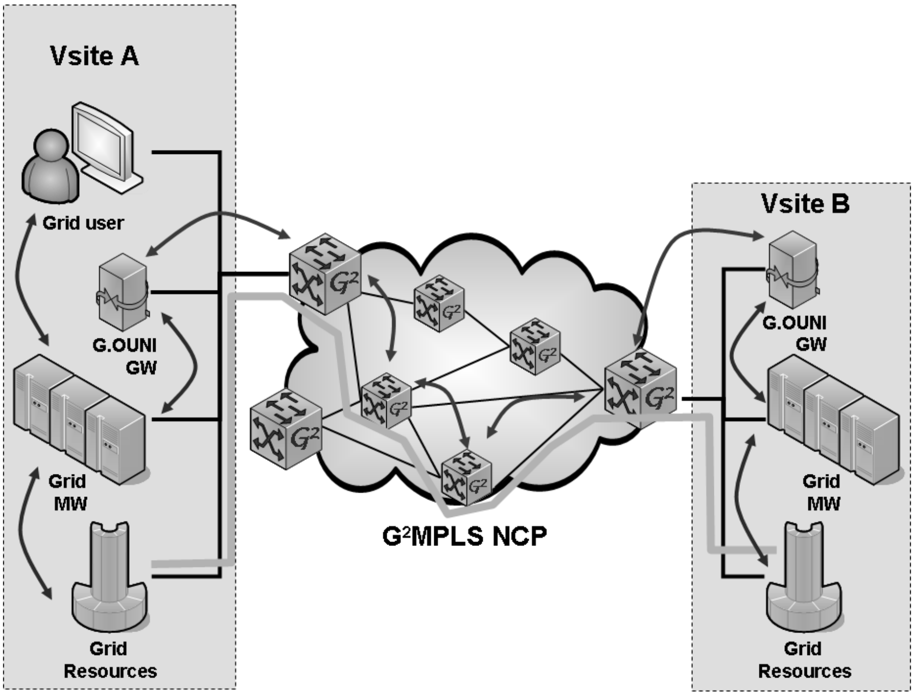

Figure 5 GNS transaction - direct invocation

It is assumed that Grid Middleware at Vsite A and Vsite B (the local Grid scheduler in particular) are aware of the types and number of Grid resources located at their own site (directly) and at remote sites (via G.OUNI service and resource discovery).

As the triggering action, the Grid user defines the job in the Client Application and issues the request. The Client Application sends a request to the local Middleware at Vsite A. The Grid broker in the middleware resolves the Grid resources needed for the job (local and remote) and requests the local service (G.OUNI GW) to build a GNS transaction between the local Vsite A and remote Vsite B. The gateway schedules and configures the local Grid resources via LRMS on Vsite A and afterwards initiates a GNS transaction – the GNS transaction request is sent through the nodes towards the selected egress node. The egress node sends the GNS transaction request to the peering G.OUNI gateway. The G.OUNI gateway on site B requests its LRMS in the middleware to schedule and configure the local Grid resources. When done, the G.OUNI gateway at Vsite B sends a GNS transaction response to the selected ingress edge node. Then the ingress node sends the GNS transaction response to the G.OUNI gateway. The GNS transaction response is sent back by the G.OUNI gateway at Vsite A to the Grid broker. The middleware sends the job request confirmation to the Client Application which notifies the Grid User of the job request confirmation.

GNS transaction – indirect invocation

A Grid User at Vsite C configures and requests the execution of a job that involves Grid resources located remotely at Vsite A and Vsite B (Figure 6).

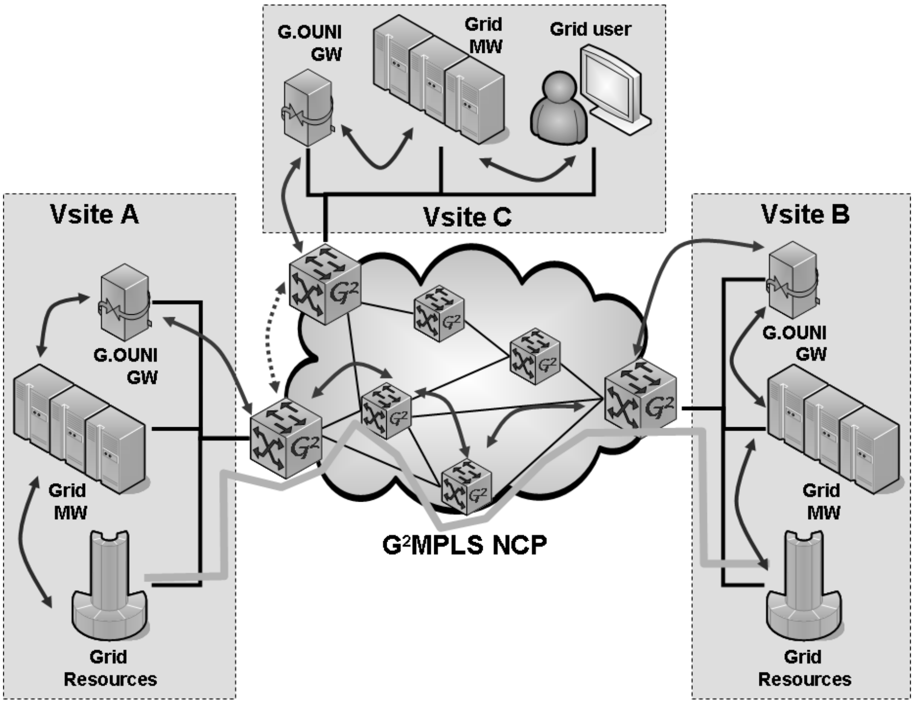

Figure 6 GNS transaction - indirect invocation

It is assumed that Grid Middleware at Vsite A, Vsite B, and Vsite C (local Grid scheduler in particular) are aware of the types and amount of Grid resources located at their own site (directly) and at remote sites (via G.OUNI service and resource discovery).

As the triggering action, the Grid user at Vsite C defines the job in the Client Application and issues the request. The Grid broker in the middleware resolves the Grid resources needed for the job. It requests the local G.OUNI gateway to build a GNS transaction between Vsite A and Vsite B.

Next, the G.OUNI gateway initiates a GNS transaction and the GNS transaction request is sent through the nodes towards the selected edge node peering with Vsite A. The edge node sends the GNS transaction request to the peering G.OUNI gateway. The G.OUNI gateway at Vsite A requests the LRMS in the middleware to schedule and configure the local Grid resources for Vsite A. When done, LRMS responds back with a confirmation to the G.OUNI gateway. The GNS transaction request is then sent through the network nodes towards the selected egress node, peering with Vsite B. The egress node sends the GNS transaction request to the peering G.OUNI gateway. The G.OUNI gateway requests the LRMS in the middleware to schedule and configure the local Grid resources for Vsite B. The LRMS at Vsite B schedules and configures the local Grid resources. The LRMS at Vsite B sends a job request confirmation to the G.OUNI gateway and the GNS transaction response is sent back through the network nodes towards the selected ingress node (the one peering with Vsite A). Afterwards, the node peering with Vsite A sends a remote GNS transaction response to the edge node peering with Vsite C. This is a G^{2}MPLS ISI (Internal Signalling Interface). The GNS transaction response is sent back by the G.OUNI gateway at Vsite C to the local Grid scheduler. The middleware sends a job request confirmation to the Client Application which notifies the Grid User of the job request confirmation.

G^{2}MPLS control plane deployment

The deployment of the G^{2}MPLS control plane in research networks implies a set of requirements in terms of management of the resources, partitioning and virtualization. This chapter provides the results of analysis of the deployment models of G^{2}MPLS control plane with particular focus put on the abovementioned requirements. The proposed framework for resource partitioning, integration with the existing Network Management Systems (NMS), and existing Data Communication Networks (DCN) are carefully examined, aiming at providing the specification of the requirements for the deployment of the G^{2}MPLS control plane.

Resource partitioning and virtualization

G^2MPLS, as an enhancement to the ASON/GMPLS control plane architecture, is designed to autonomously operate on the Transport Plane resources under its ownership. Therefore, it is not possible for the Control Plane to work on resources that are under the authority of the other managing entity (e.g. Management Plane).

The atomic functions provided by the Transport Plane equipment towards the Management Plane are at the same hierarchical level of the GMPLS/G^2MPLS Control Plane ones and compete with them. On the contrary, the Management Plane can modify the resources exported by the Control Plane and manage them as virtual resources, but according to the G^2MPLS information model. Thus, it is important to define what are the expected actions that can be exported by the Control Plane to the Management Plane in order to make both interoperable in the same environment. The actions should cover at least the following items:

- Configuration of protocols, behaviours and Control Plane entities (e.g. identifiers for TE-links, interfaces, metrics, etc.);
- Request for calls/connections setup/tear-down in the Control Plane (Soft Permanent Connection in the ASON/GMPLS terminology);
- Some OAM functionalities on the installed connections.

These functionalities provide virtualization of the Transport Network resources towards the Management Plane which could treat them as unitary Transport Plane resources. For example, if we consider the pure case of GMPLS, the Soft Permanent Connection (SPC) is the combination of two Permanent Connection (PC) segments at the source and destination user-to-network interfaces, plus a Switched Connection (SC) segment within the core network. The PC parts are solely owned by the Management Plane, while the SC part is directly owned by the Control Plane but exported to the Management Plane that can treat it as a direct forwarding adjacency (link) between the ingress and egress nodes (virtual resource).

In the work carried by the PHOSPHORUS project it is highlighted that it would be valuable to integrate the G^2MPLS control plane on a disjoint partition of the available network resources, by moving some of them under the authority of the G^2MPLS control plane, and maintaining some others under the management plane (e.g. Network Resource Provisioning Systems). Moreover, because G^2MPLS includes also advance reservation services, there is a need to take into consideration another aspect of partitioning – dealing with the temporal dimension of Transport Plane resources. There are two approaches proposed by PHOSPHORUS. The first one implements the selection of the resources to be used for bookings from a disjoint set with respect to those dedicated for immediate reservations. This approach simplifies the routing and signalling operations, guaranteeing the creation of complete and coherent resources and uniform semantic for resource availability specification in time. A more complex approach might include resource re-partitioning in time, e.g. triggered by the notification of a complete use of the resources in one partition or by other performance benchmarks. The exact definition of the schema (i.e. the amount of the resources dedicated to support immediate requests) relies on the operator of the network offering connection services and it strongly depends on the type of expected connection requests. In case of high deployment of Grid services in the operator's network it is expected to have a significant amount of advance reservations, thus it is up to the operator to allocate the relevant level of resources for them. On the contrary, in case of traditional end-to-end permanent connection services it is expected to have more immediate reservations in the system.

Integration of G^2MPLS management in the existing Network Management Systems

G^2MPLS aims to inherit the rights for full usage of the Transport Plane resources configured under its ownership from the GMPLS Network Control Plane, and in some cases maintains some of these rights under the authority of the management plane. As the network operators may want to continue the usage of the existing Management Plane (e.g. composed of NMS), it is important to understand how to plan the deployment of G^2MPLS in the networks managed by the existing systems.

In order to ensure the Management and Control Plane coexistence and proper interaction, a set of requirements has been identified. The NMS must be capable of receiving notifications from the G^2MPLS control plane. This ensures that the NMS is kept updated about recent changes in the monitored environment (e.g. addition or removal of the Grid site). Additionally, the NMS should be capable of receiving notifications from the control plane for Grid and network resources allocation or release. In order to have the knowledge about the Grid resources status, the NMS must support the capability of querying the control plane for Grid resource localization, capability and availability. Since NMS is not directly involved in all connection setup processes (e.g. SC), monitoring of the existing connections and newly requested calls is very important to ensure efficient and reliable resource management and fault detection. The NMS must be able to query the information about all calls which are using a particular Grid or network resource.

When all these requirements are fulfilled (NMS supports requested features), the network/service operator should be able to view the grid and network resources existence and state. This knowledge will be very helpful during the failure localization and repair processes.

Integration of the G^2MPLS control plane with the existing DCN

G^2MPLS inherits native GMPLS with the separation of Transport Plane and Data Plane. As per ITU-T G.7712, it also requires a communication network, the Signalling Control Network, to transport signalling messages between its components. Since the specific extensions to support Grid functionalities in the control plane do not enforce any additional requirements on the SCN, it is possible to re-use native GMPLS SCN, respecting well known GMPLS needs (e.g. GMPLS SCN must support IP at Layer 3, etc.).

Interoperability of G^2MPLS and GMPLS control planes

The ASON architecture is built around the identification and description of network reference points (namely, UNI, I-NNI, E-NNI). As mentioned in the chapter 'G2MPLS network reference points', the G^2MPLS Control Plane architecture derives from ASON/GMPLS architecture and extends the reference points in order to provide support for Grid applications. Those extensions form a new set of interfaces able to handle combined requirements of the application on both computational resources and the underlying network (see Figure 3).

Starting the analysis with coexistence of both planes at the intra-domain level, it is required to consider the compatibility and possible interworking of Internal Node-to-Node Interfaces (I-NNI) between the control plane performers. G^2MPLS requires a complex set of extensions to the base GMPLS protocols for Grid purposes, both in terms of routing and signalling. While routing extensions rely on OSPF Opaque LSAs and could be just forwarded – though neglected – by standard GMPLS controllers, part of the signalling extensions cannot be ignored by pure GMPLS nodes and could lead into unsupported operational conditions resulting in a final block of signalling. For this reason, the PHOSPHORUS position is to analyze the interoperation/coexistence issues between G^2MPLS and GMPLS just at the domain granularity (i.e. inter-domain co-existence) in order to adapt and progress signalling and routing messages according to the declared neighbouring domain capabilities. At the same time, considering the use of G^2MPLS controllers to handle pure GMPLS functionality (i.e. to allow standard connection services, as defined per ASON and GMPLS architecture) should not raise any internetworking problems. The decision made at the design level to make G^2MPLS backward compatible allows to disable Grid-specific services and make the interoperability of the planes possible.

The inter-domain coexistence of G^2MPLS and ASON/GMPLS architecture implies the analysis of two separable deployment issues. The former involves the deployment of the PHOSPHORUS G^2MPLS/GMPLS and non-PHOSPHORUS standard GMPLS control planes in the networks. The latter covers interworking with the components lying in the Network Service Plane, such as Network Resource Management Systems (NRPSes) and the GN2 IDM (the inter-domain component of GN2 Autobahn). In the second case, G^2MPLS preserves the consistency of its network reference point in the East-West direction (i.e. the G.E-NNI) by moving in an additional network element – the G.E-NNI Gateway – the bridging logic between the Management Plane and the Control Plane. Interworking across G^2MPLS and ASON/GMPLS domains (the first case) can be established according to the same design principle, but without the need of any gateway functionality, being the two parties natively peers. However, in a G^2MPLS domain both Grid and network resources are under the same control, while in an ASON/GMPLS domain only network resources are controlled. This implies a reduction of the information set carried out through the G.E-NNI, both for routing purposes and, above all, for signalling.

G^2MPLS and its possible impact on the development of research networks

From the NREN perspective, G^2MPLS aims to integrate Grids and automated network control plane technologies in real operational networks. On the one hand, this allows to use well established solutions for setup, recovery and crankback of network connections spanning multiple domains and matching traffic engineering objectives; on the other hand, G^2MPLS allows to overcome the current limitation of Grids that operate as stand-alone overlaid infrastructures upon research networks with different administrative ownership and uncorrelated procedures.

End users of the NRENs may benefit from the fact that G^2MPLS enables a real node-to-node deployment of on-demand Grid services and exposes specific interfaces towards the Grid layer. Part of the middleware functionalities related to selection, co-allocation, and maintenance of both Grid and network resources are provided through these interfaces. Thus, real on-demand provisioning of resources across multiple administrative and network technology domains is possible, enabling a new class of Grid services.

Acknowledgements

This work is supported by the PHOSPHORUS project, which is partly funded by the European Commission under the Sixth Framework Programme, contract number 034115. We would like to express the acknowledgements to all PHOSPHORUS partners for their support and cooperation.

References

Ciulli N. et al. 2007, "The Grid-GMPLS Control Plane architecture", PHOSPHORUS Project, Deliverable 2.1. Available at http://ist-phosphorus.eu/files/publications/Phosphorus-deliverable-D2.1.pdf

Ciulli N. et al. 2007, "The Deployment Models and Solutions of the Grid-GMPLS Control Plane", PHOSPHORUS Project, Deliverable 2.6. Available at http://ist-phosphorus.eu/files/publications/Phosphorus-deliverable-D2.6.pdf

Figuerola, S., Ciulli, N., De Leenheer, M., Demchenko, Y., Ziegler, W., Binczewski, A. 2007, "PHOSPHORUS: Single-step on-demand services across multi-domain networks for e-science", APOC 2007

Markidis, G., Tzanakaki, A., Ciulli, N., Carrozzo, G., Simeonidou, D., Nejabati, R., Zervas, G. 2007, "EU Integrated Project PHOSPHORUS: Grid-GMPLS Control Plane for the Support of Grid Network Services", ICTON 2007

Patil, A., Belter, B., Polyrakis, A., Rodwell, T., Przybylski, M. and Grammatikou, M. 2006, "GÉANT2 Advance multi-domain provisioning system", Terena Networking Conference 2006. Available at http://www.terena.nl/events/tnc2006/core/getfile.php?file_id=760

Biographies of the Authors

***BARTOSZ BELTER** received the M.Sc. degree in Computer Science from Poznan University of Technology in 2002. He works in Poznan Supercomputing and Networking Center in Network Research and Development Department as a Network Applications Engineer. His main interests include network measurements, performance troubleshooting and enhancements, as well as network control planes and bandwidth on demand services. His research involvement also covers the area of software management, design and development. He participated in several FP5/FP6 IST projects (6NET, GÉANT2, PHOSPHORUS) and programmes funded by Polish Ministry of Science and Higher Education (CLUSTERIX, LDAP PROJECT).*

***ARTUR BINCZEWSKI** received the M.Sc. degree in Computer Science from the Poznan University of Technology in 1993. His research interests concern computer networks, routing, multicasting and management. He is the Manager of the Network Division at the Poznan Supercomputing and Networking Center. He was involved in several EC projects: SEQUIN (IST-1999-20841), ATRIUM (IST-1999-20675), 6NET (IST-2001-32063) and others. His research interests concern computer networks protocols and management. He is author or co-author of papers in major professional journals and conference proceedings. He takes an active part in the development of the PIONIER fiber network.*

***GINO CARROZZO**, received the Italian Laurea degree "cum laude" in Electronic Engineering from the University of Pisa, Italy, in 2000, and the PhD degree in Telecommunications Networks from the University of Pisa in 2004. During 1999-2000 he spent one year at Agilent Technologies Laboratories-TTC in Turin, Italy, working on adaptive clock recovery for SDH transport networks. Currently he is R&D system engineer at Nextworks and CPR-DITEL. His research activities concern the design, development and integration of GMPLS Control Plane on carrier-grade SDH and WDM devices, with main focus on constraint-based routing, traffic engineering and traffic recovery procedures. Other research interests include the evolution of Control Plane technologies towards Transport MPLS and carrier-grade development of ITU-T EPL services. Past activities focused on design and implementation of IP-QoS/MPLS architectures (IntServ, DiffServ and MPLS/DiffServ). He participates in many research projects funded by the Italian Research Ministry (NEBULA, TANGO, ARCHIMEDE projects) and in the IST FP5 and FP6 (MOICANE, MEDIANET, PHOSPHORUS).*

NICOLLA CIULLI, *responsible for the Networking Area of Nextworks and CPR - DITEL, received the Laurea degree "cum laude" in Telecommunication Engineering in 1997 at the University of Pisa, and the Degree Diploma at SSSUP S. Anna. His main research activities concern the architectural aspects of Control and Management Planes and of Traffic Control in next generation packet networks: Quality of Service (QoS) provision over the Internet and its interworking with Asynchronous Transfer Mode (ATM), Differentiated Services (DiffServ), MPLS and their integration. A special focus of his research and industrial activity is on GMPLS control and management planes architectures for SDH networks. He participated in several ACTS and FP5/FP6 IST projects (MAESTRO, PETERPAN, BASS, MOICANE, WEIRD, PHOSPHORUS), in research activities funded by the Italian Ministry for Research and Education (MIUR), and in industrial projects (e.g. Marconi ASTN/GMPLS project, where he coordinated the activities of Nextworks, Alcatel projects on 4G mixed packet/TDM switches).*

EDUARD ESCALONA *obtained his M.Sc. degree in Telecommunications Engineering (2003) at the Universitat Politècnica de Catalunya (UPC), Spain after a one year stage in Telecom Italia Lab, Italy, as an undergraduate researcher participating in the IST European Commission project LION. He is working towards completion of his PhD at UPC's Department of Signal Theory and Communications in the Optical Communications Group. He has participated in several research projects like IST Nobel 2, E-Photon/One+ and Phosphorus IST FP6. He is currently working as a research officer at the University of Essex, UK, on Phosphorus IST project. His main research areas are related with control of optical networks, particularly GMPLS protocols, and optical grids.*

GEORGIOS MARKIDIS *obtained his Electrical and Computer Engineering degree from the National Technical University of Athens in 2004 and his M.Sc. degree in Information Networking from Carnegie Mellon University, Pittsburgh, in 2006.He joined AIT as a researcher in February 2006 participating in the European research projects e-Photon/One+ , COST 291 and IST-PHOSPHORUS. His current research interests focus in the field of optical networks for telecommunication and grid computing applications including network dimensioning, impairment aware routing and wavelength assignment algorithms, resiliency and traffic engineering mechanisms.*

REZA NEJABATI *has over 6 years experience in the filed of high-speed electronic and programmable router design. He is currently a senior research officer in the Photonic Network Research Group at the Essex University. The main current areas of his interest are design and control issues for high-speed electronic and optoelectronic interfaces in photonic packet-based networks as well as architectural considerations for photonic Grid networks. Reza Nejabati has published over 25 papers in the area of optical communications.*

DIMITRA SIMEONIDOU *has over 15 years experience in the field of optical transmission and optical network design. In 1987 and 1989 she received the BSc and MSc degree from the Physics Department, University of Thessaloniki and in 1994 the PhD degree from the University of Essex. From 1992 to 1994 she was employed as a Senior Research Officer at the University of Essex in association with the MWTN RACE project. In 1994 she joined Alcatel Submarine Networks to work on the introduction of WDM technologies in submerged photonic networks. . In 1996 became Senior Principal Engineer responsible for the design of wavelength routing nodes for Submarine WDM networks. Prof. Simeonidou is currently leading the optical networking initiative in the global Grid forum. The main current areas of interest are design and control issues for photonic packet-based networks and architectural considerations for photonic Grid networks. Professor Simeonidou is the author and co-author of over 100 papers relating to optical transmission/networking and of 11 international patents on submerged network technologies.*

MACIEJ STROIŃSKI *received the Ph. D. degree in Computer Science from the Technical University of Gdańsk in 1987. Currently he is Technical Director of Poznań Supercomputing and Networking Center. He is also lecturer in the Institute of Computing Science of Poznań University of Technology. His research interests concern computer network protocols and management. He is author or co-author of over 100 papers in major professional journals and conference proceedings.*

ANNA TZANAKAKI *is an Associate Professor at the "Athens Information Technology" center and an adjunct faculty member of Carnegie Mellon University, USA. She has obtained a BSc degree from the University of Crete, Greece, an MSc and a PhD both from the University of Essex, UK. She was a Senior Research Officer and a visiting lecturer at the University of Essex and became a co-founder and a senior engineer of ilotron ltd, a spin-off from the University of Essex. Following ilotron, Dr Tzanakaki joined Altamar Networks, a subsidiary of Ditech Communications, as a principal engineer. She is a co-author of over 80 publications in international journals and conferences and a co-inventor of 1 granted and 11 published patents. She is a senior member of the IEEE and a member of several Technical Program*

Committees. Her research interests include wavelength routed as well as optical burst and packet switched networks & data and control plane issues for telecommunications and grid computing optical network infrastructures.

***GEORGIOS ZERVAS** was awarded the MEng degree of Electronic and Telecommunication Systems Engineering with distinction from the University of Essex. (UK) in 2003. He is currently working towards the PhD degree on Optical Networks and End-Systems for Grid Application and Services at the University of Essex. He is working as Senior Research Officer in an EC funded project Lambda User Controlled Infrastructure for European Research (Phosphorus). He is also involved in EC funded IST e-Photon/One+ project and a number of research activities within the University of Essex. His research interests include high-speed optoelectronic router design, optical burst switched networks, GMPLS networks, user network interface and Grid Networks.*