



034115

PHOSPHORUS

Lambda User Controlled Infrastructure for European Research

Integrated Project

Strategic objective:
Research Networking Testbeds



Deliverable D6.1 Test-bed design

Due date of deliverable: 2007-01-31
Actual submission date: 2007-01-31
Document code: Phosphorus-WP6-D6.1

Start date of project:
October 1, 2006

Duration:
30 Months

Organisation name of lead contractor for this deliverable:
Instytut Chemii Bioorganicznej PAN

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



Authors

Lada Altmannová	CESNET
Artur Binczewski	PSNC
Wojbor Bogacki	PSNC
Maciej Brzeźniak	PSNC
Gino Carrozzo	NXW
Nicola Ciuli	NXW
David Domingo	i2CAT
Thomas Eickermann	FZJ
Sergi Figuerola	i2CAT
Marcin Garstka	PSNC
Leon Gommans	UVA
Paola Grosso	UVA
Ferdinand Hommes	FHG
Gigi Karmous-Edwards	MCNC
Cees de Laat	UvA
John MacAuley	NORTEL
Tomasz Makiela	PSNC
John Moore	MCNC
Reza Nejabati	UESSEX
Bram Peeters	SURFnet
Jan Radil	CESNET
Carsten Rosche	FHG
Angel Sanchez	i2CAT
Michel Savoie	CRC
Jochen Schon	FHG
Stanislav Šíma	CESNET
Dimitra Simeonidou	UESSEX
Maciej Stroiński	PSNC
Berenguer Vilajoliu	i2CAT
Josef Vojtěch	CESNET
Jan Węglarz	PSNC





Abstract

One of the main objectives of Phosphorus is to demonstrate an application development environment that will validate the new services and the infrastructure developed in the project. To achieve this objective a distributed testbed will be built, in which the project developments will be incrementally introduced. The testbed will constitute a real environment in which the project outcome will be demonstrated with a set of real scientific applications in a set of real-life scenarios. The testbed will be constructed from communications equipment (optical switches, TDM switches, Gigabit Ethernet switches, transmission equipment) and advanced GRID resources (like computing nodes, visualisation and storage resources). The communications equipment will be the platform of the project's developments and implementations in order to allow applications to utilise the functionality of an advanced optical network and assure seamless cooperation of various testbed equipment.

The Phosphorus testbed will consist of multiple local testbeds located in several sites in Europe and outside Europe. The local testbeds will be interconnected with other local testbeds using multiple international optical networks (including GÉANT2 and others) to create a global, heterogeneous testbed comprising several technological and administrative domains. The structure of the testbed will be based on requirements of other Phosphorus activities, which will utilise the testbed in order to verify and demonstrate their developments.

This document presents the design of the Phosphorus testbed including the communications equipment, the GRID resources and the topology of interconnections between all testbed components. It is divided into two main sections, one of which describes the local testbeds and the resources present in them, while the other one shows the topology of interconnections between local testbeds. Section three of this document summarises the project developments which are going to be deployed in the testbed. The requirements obtained from other Phosphorus activities are shown in appendixes.



Table of Contents

0	Executive Summary	8
1	Local Testbeds	10
1.1	PSNC testbed	12
1.2	CESNET testbed	20
1.3	i2CAT testbed	27
1.4	SURFnet testbed	32
1.5	UvA testbed	35
1.6	UESSEX testbed	41
1.7	VIOLA testbed	47
1.8	NORTEL testbed	55
1.9	MCNC testbed	58
1.10	CRC testbed	64
1.11	IP addressing in the local testbeds	71
1.12	Project developments in the local testbeds	72
2	Topology of the global testbed	74
2.1	Topology requirements and assumptions	75
2.2	Data plane design	77
2.3	Control/provisioning plane design	81
3	Conclusions	82
4	References	83
5	Acronyms	84
Appendix A	Requirements of Workpackage 1	86
A.1	Introduction	88
A.2	Testbed Requirements	88
Appendix B	Requirements of Workpackage 2	104



Test-bed design

B.1	G ² MPLS Controller requirements	104
B.2	Signalling Control Network (SCN) requirements	105
B.3	Transport Plane requirements	105
B.4	Configuration requirements	106
Appendix C Requirements of Workpackage 3		108
C.1	Network requirements	108
C.2	Local testbed hardware and software requirements	109
Appendix D VPN-solutions for realizing a network for Control Plane traffic in Phosphorus		112
D.1	Requirements for the Control Plane Data Communication Network (DCN)	112
D.2	Tested solutions	113
D.3	Conclusion	114
Disclaimer		115



Table of Figures

Figure 1.1: Topology of the PSNC local testbed	15
Figure 1.2: Testbed topology in CESNET	21
Figure 1.3: CESNET Dark fibre	23
Figure 1.4: Detailed view of i2CAT's local testbed	28
Figure 1.5: i2CAT testbed external connections	30
Figure 1.6: SURFnet local testbed	33
Figure 1.7: Connection to other local testbeds	34
Figure 1.8: The VanGogh / Rembrandt cluster nodes in Netherlight / Lighthouse	36
Figure 1.9: The DAS-3 Cluster Node at UvA Netherlight	36
Figure 1.10: The GMPLS testbed setup	37
Figure 1.11: The UCLP Testbed based on NORTEL OME's	37
Figure 1.14: Multi-domain Generic AAA Testbed at UvA Lighthouse	39
Figure 1.15: The UEssex topology of the local testbed and the testbed resources	44
Figure 1.16: Topology of the VIOLA local testbed	49
Figure 1.17: NORTEL local testbed	56
Figure 1.18: Connection to other local testbeds	57
Figure 1.19: Topology of the Enlightened Resource testbed	59
Figure 1.20: Topology of the Enlightened testbed	61
Figure 1.21: CAnet 4 Infrastructure	65
Figure 1.22: CRC's Local Testbed Connectivity Infrastructure	66
Figure 1.23: Detailed View of CRC's Local Testbed	67
Figure 1.24: Detailed view of the AON Demonstrator at CRC	68
Figure 1.25: CRC Testbed Topology Proposal	70
Figure 2.1: Topology of the PHOSPHORUS testbed	79
Figure A.1: Overall topology for WP1	88
Figure A.2: CRC Testbed Topology Proposal	89
Figure A.3: Detailed View of CRC's Local Testbed	91
Figure A.4: Detailed view of the AON Demonstrator at CRC	91
Figure A.5: Eucalyptus poster	93
Figure A.6: i2CAT test bed topology proposal	94
Figure A.7: detailed view of i2CAT's local testbed	95
Figure A.8: i2CAT grid resources	96



Test-bed design

Figure A.9: Nortel Testbed Topology Proposal	97
Figure A.10: Nortel DRAC testbed (simplified).	98
Figure A.11: Global WP1 testbed connectivity – SURFnet perspective	99
Figure A.12: View SURFnet's local testbed	100
Figure A.13: VIOLA test bed topology proposal	101
Figure A.14: Detailed view of the VIOLA testbed (simplified)	102



0 Executive Summary

Phosphorus will address some of the key technical challenges to enable on-demand e2e network services across multiple domains. The Phosphorus network concept will make applications aware of their complete Grid resources (computational and networking) environment and capabilities, and able to make dynamic, adaptive and optimized use of heterogeneous network infrastructures connecting various high-end resources. Phosphorus will enhance and demonstrate solutions that facilitate vertical and horizontal communication among applications middleware, existing Network Resource Provisioning Systems, and the proposed Grid-GMPLS Control Plane.

One of the main assumptions of Phosphorus is that the project propositions and developments should be validated and demonstrated in a real advanced optical network. To achieve this, the project will build a distributed testbed in which the project outcome will be demonstrated with a set of real scientific applications in a set of real-life scenarios.

The testbed will be constructed from communications equipment (optical switches, TDM switches, Gigabit Ethernet switches, transmission equipment) and advanced GRID resources (like computing nodes, visualisation and storage resources). The communications equipment will be the platform of the project's developments and implementations in order to allow applications to utilise the functionality of an advanced optical network and assure seamless cooperation of various testbed equipment. The different types of resources (communications, computing etc) will correspond to the elements of an advanced GRID system in which GRID resources are distributed over a communications network. The structure of the testbed will also be compatible with the Phosphorus paradigm of integration between traditional GRID and the communications infrastructure. The GRID resources will not only include the hardware but also GRID middleware software which will control the resources and act as an interface between the applications and the infrastructure.

The project will demonstrate an incremental approach to the testbed construction. The first version of the testbed will be based on the solutions (equipment, technology, protocols, functionality) available in the first months of the project lifetime. When the first version of the testbed is ready and operational, it will be incrementally extended to support the developments of Phosphorus. The extension of the testbed will include mostly (but not only) new versions of control plane protocols and resource provisioning systems which will be developed in Phosphorus and new functionality they offer.

The Phosphorus testbed will consist of multiple local testbeds located in several sites in Europe and outside Europe. The local testbeds will be interconnected with other local testbeds using multiple international optical



Test-bed design

networks (including GÉANT2 and others) to create a global, heterogeneous testbed comprising several technological and administrative domains. The structure of the testbed will be based on requirements of other Phosphorus activities, which will utilise the testbed in order to verify and demonstrate their developments.

The management and provisioning layers of the testbed will support on-demand optical paths provisioning between GRID resources located in different local testbeds. This feature will make use of GMPLS and Network Resource Provisioning Systems, in the first stage the testbed will use existing versions of GMPLS and NRPSes while in the second stage it will use their new versions which will be developed by the project. The new versions will allow for further integration between the transport network and GRID applications and middleware.

This document presents the design of the Phosphorus testbed including the communications equipment, the GRID resources and the topology of interconnections between all testbed components. It is divided into two main sections, one of which describes the local testbeds and the resources present in them while the other one shows the topology of interconnections between local testbeds. Section three of this document summarises the project developments which are going to be deployed in the testbed. The requirements obtained from other Phosphorus activities are shown in appendixes.

Project:	Phosphorus
Deliverable Number:	D6.1
Date of Issue:	31/01/07
EC Contract No.:	034115
Document Code:	Phosphorus-WP6-D6.1



1 Local Testbeds

The Phosphorus testbed will consist of several local testbeds in which the switching and GRID resources will actually be located and in which the testbed operations will be executed. Each local testbed will be constructed from different resources and will support different project needs which gives the testbed the heterogeneity needed to achieve the project goals and demonstrate that the project developments are not limited to any single technology or any product family. The distribution of testbed will also create natural technology and administrative domains which allows for demonstration and verification of the project results in a multi-domain environment. The distribution of the testbed makes the work in workpackage 6 more challenging and requires some efforts to ensure seamless integration of all local testbeds on multiple layers: physical interconnection of the local testbeds, control plane and/or NRPS integration, application distribution, resource allocation and administrative issues. Workpackage 6 will take care of the integration.

Each local testbed will be provided and operated by one project participant. As a general rule the project participant will provide all the resources to the local testbed but in a limited number of resources will be provided by a project participant who does not have their own local testbed and will contribute to local testbeds operated by other partners.

Local testbeds will have various equipment and play different roles in the global testbed. There are three types of equipment that local testbed will provide:

- Networking – providing connectivity
- Computing and visualisation – providing processing power
- Storage – providing room for data storage

The choice of resources for each local testbed is based on the requirements from other Phosphorus activities and depends on the task that will be performed in each local-testbed. The computing, visualisation and storage resources will be used by applications run by Workpackage 3 and will generate streams of data which will be transmitted over the testbed. The networking equipment in the local testbeds (mostly pure optical switches, Gigabit Ethernet switches and SDH crossconnects) will not only serve to transmit the data but first of all will be the devices on which G²MPLS and NRPSes will provision connectivity for applications. The equipment will be integrated with the control/provisioning layer of the testbed with use of a set of management interfaces. In fact the switching equipment in local testbeds is the most important element of the whole Phosphorus testbed.

Project:	Phosphorus
Deliverable Number:	D6.1
Date of Issue:	31/01/07
EC Contract No.:	034115
Document Code:	Phosphorus-WP6-D6.1



Test-bed design

The testbed will be built from networking devices from multiple vendors to demonstrate the open approach of Phosphorus which developments will not be limited to any single vendor or any set of vendors.

The PHOSPHORUS testbed will consist of 11 local testbeds. The testbeds will be provided by PSNC, CESNET, i2CAT, SURFnet, UESSEX, UvA, NORTEL, MCNC and CRC. Another local testbed (the VIOLA testbed) will be provided jointly by three Phosphorus partners: FHG, FZJ and UniBonn. A detailed description of the local testbeds is given in the next sections of this document. An additional local testbed may be located at CCT at LSU which is not a member of the PHOSPHORUS consortium but has declared its interest in collaboration with the project. The consortium may also use some resources of SARA for testing the TOPS application. In this case SARA will be connected to SURFnet local testbed using SURFnet infrastructure.

All Phosphorus local testbeds will be interconnected using international optical links in order to create a global, heterogeneous testbed for the purposes of Phosphorus.

Project:	Phosphorus
Deliverable Number:	D6.1
Date of Issue:	31/01/07
EC Contract No.:	034115
Document Code:	Phosphorus-WP6-D6.1



1.1 PSNC testbed

The local PSNC testbed is going to be used for DDSS, KoDaVis and WISDOM applications tests, as defined by Workpackage 3. PSNC testbed will contain switching, computing and storage resources. The switching layer will be built from optical switches and Gigabit Ethernet switches, computing resources will include PC cluster nodes and a high-end storage server and storage resources will encompass disk volumes on Fibre Channel disk matrices and tape pools in an automated tape library.

The PSNC local testbed will be controlled by GMPLS/G²MPLS and integrated with other Phosphorus local testbeds with use of the proper GMPLS/G²MPLS interfaces. The GMPLS/G²MPLS control plane will allow for automatic provisioning of optical paths.

1.1.1 GRID resources in the local testbed

Computing resources in the PSNC local testbed include 4 PC cluster nodes which hardware configuration is suitable for testing of all applications intended for this testbed: DDSS, KoDaVis and WISDOM. The nodes are equipped with 2x Itanium2 1.3GHz CPUs, 4 GB RAM and 2 Gigabit Ethernet adapters. The testbed will also include a storage server equipped with 4x Xeon 2.80GHz CPUs, 2 GB RAM, 4x FC and 2 Gigabit Ethernet interfaces.

Additionally, some resources of three IBM x345 PCs can be used if necessary, i.e. if the processing power of the PC cluster nodes is insufficient for the applications. IBM PCs are equipped with 2x Xeon 2.66 GHz CPUs, 1 GB RAM, 2 Gigabit Ethernet interfaces and 2 Fibre Channel interface cards.

Storage resources assigned by PSNC for an application test include 1TB and 2TB disk volumes on IBM FastT700 disk matrix and NexSan SATABeast matrix. These disk volumes will provide enough capacity to perform multiple test rounds for DDSS, KoDaVis and WISDOM applications, i.e. store the actual test data as well as store monitoring and control data. IBM FastT700 used in tests has the total capacity of 15 TB, based on 160 Fiber Channel disks (FC 2Gbps technology). NexSan SATABeast matrix is an FC-SATA matrix with Fiber Channel host ports and 42 SATA disk drives.

Additionally PSNC will assign some resources of an automated tape system for tests of DDSS backup/archive application. PSNC will use for that purpose up to 4 TB of LTO2 tapes (i.e. 20 tapes) and one LTO2 Ultrium tape drive. The resources will provide enough storage space and read/write performance for DDSS B/A tests. The tape library used in the tests is the ADIC Scalar i2000 library with 60 TB of LTO2 tapes and 6 Fiber Channel drives.

Project:	Phosphorus
Deliverable Number:	D6.1
Date of Issue:	31/01/07
EC Contract No.:	034115
Document Code:	Phosphorus-WP6-D6.1



1.1.2 Switching and transmission resources in the local testbed

Most of the connections inside the local testbed as well as all links to the other Phosphorus local testbeds will use Gigabit Ethernet as the transmission technology. Some internal connections to the storage will use Fibre Channel. This way the local testbed will use the same internal transmission technologies as most modern data centres and clusters and external connectivity technologies that are typical for modern GRID systems. Gigabit Ethernet transmissions will be switched in the testbed in two layers – on layer 1 by optical switches and on layer 2 by Gigabit Ethernet switches.

The core of the local testbed will be built from three optical switches to which the other testbed resources will be connected. The switching infrastructure will also contain three Gigabit Ethernet switches and two Fibre Channel switches for internal connections between GRID resources. Fibre Channel switches will be located at the back-end of the some nodes attached to the testbed. Therefore, they are not direct components of the testbed and are not directly connected to the optical switching infrastructure.

The testbed core will switch lightpaths in the optical domain without conversion to electrical signals. This way the optical core – the part of the local testbeds in which most of the project developments will be done – will be independent of the transmission technology and the results of the work can be applied also to optical switching with other transmission technologies.

The optical switches for the testbed have not been selected yet. According to market research conducted by PSNC, there are two manufacturers of optical switches that can be used for the local testbed: Calient Networks and GlimmerGlass. The two manufacturers produce pure optical switches which offer the necessary functions and they can ensure future developments of their products, possibly accommodating the ideas of Phosphorus. PSNC already talked to the representatives of the two companies and they declared their interest in Phosphorus and commitment to collaboration with PSNC and the Phosphorus consortium.

Each of the optical switches will have at least 16 optical ports (16 inputs and 16 outputs) and will be able to switch each input signal to any output port. This number of ports is sufficient for the needs of the testbed, will allow for connecting all the testbed resources as well as for creating multiple connections between the optical switches.

The optical switches will be accompanied by three Gigabit Ethernet switches which will switch Gigabit Ethernet VLANs.

Both optical switches and Gigabit Ethernet switches will be controlled by GMPLS/ G²MPLS.

1.1.3 Topology of connections inside the local testbed

The core of the testbed will be built from three optical switches which will be interconnected in a triangle. To allow for allocation of multiple lightpaths between any pair of the optical switches, they will be connected by multiple parallel fibre links. Multiple links will allow for simultaneous parallel transmission in the testbed as well



Test-bed design

as for switching between such link and simulating more complex testbed topologies according to the different requirements of the test scenarios. Having the switches connected in the triangle will make it possible to have two different paths between GRID resources connected to different optical switches, one direct and the other one via a third switch. This feature will increase the flexibility of serving different test scenarios, including failures of a single connections and switching lightpaths to backup routes.

The testbed will also include three Gigabit Ethernet switches which will be connected to the optical switches (one Gigabit Ethernet switch to each optical switch). The Gigabit Ethernet switches will aggregate traffic from GRID resources which will be connected to them and will switch Ethernet VLANs expanding the switching capabilities of the testbed and adding an additional layer to the heterogeneous Phosphorus network. Each of the Gigabit Ethernet switches will have multiple connections to one of the optical switches in order to provision multiple lightpaths for a single GRID node or independent lightpaths for different GRID nodes connected to the same Gigabit Ethernet switch.

Having Gigabit Ethernet switches as an aggregation layer between GRID resources and the optical network will make the testbed similar to many distributed GRIDs which are built from several local clusters which are interconnected with use of optical network. The local clusters use Gigabit Ethernet for communication between local resources while the optical network is used only for transmitting data to other local clusters.

The GRID nodes will be connected to Gigabit Ethernet switches in the Gigabit Ethernet technology. Each of them will have two Gigabit Ethernet interfaces which will support Ethernet VLANs. One of the interfaces will be used for connecting the GRID node with the control-plane network in order to allow for communication between the GRID node and the control-plane of the switches. In case the test scenario requires two independent physical links between a GRID node and the data network, the Gigabit Ethernet switches will allow for creating a control VLAN on one of the links and routing it on the switch to the control-plane network instead of the optical switches of the data plane.

As the number of Gigabit Ethernet interfaces of each GRID node is very limited (only two interfaces on each GRID node), VLANs can be used to simulate more optical paths to applications and GRID middleware. Each of the VLANs on a single physical link to a GRID node can be mapped on a Gigabit Ethernet switch to a different interface going to an optical switch and then switched separately in the optical domain.

In case a test-scenario requires direct connection between a GRID node and an optical switch, the connection may be easily provisioned by disconnecting a GRID node from a Gigabit Ethernet switch and connecting it directly to a spare interface of any optical switch in the testbed.

Two other transmission technologies will be used in the testbed to connect the GRID nodes with storage resources: Fiber Channel and SCSI. The connections will be controlled by the storage management software in the testbed and will not be controlled by the Phosphorus control/provisioning plane.

The topology of the testbed as well as all the testbed resources is shown in Figure 1.1.

Project:	Phosphorus
Deliverable Number:	D6.1
Date of Issue:	31/01/07
EC Contract No.:	034115
Document Code:	Phosphorus-WP6-D6.1

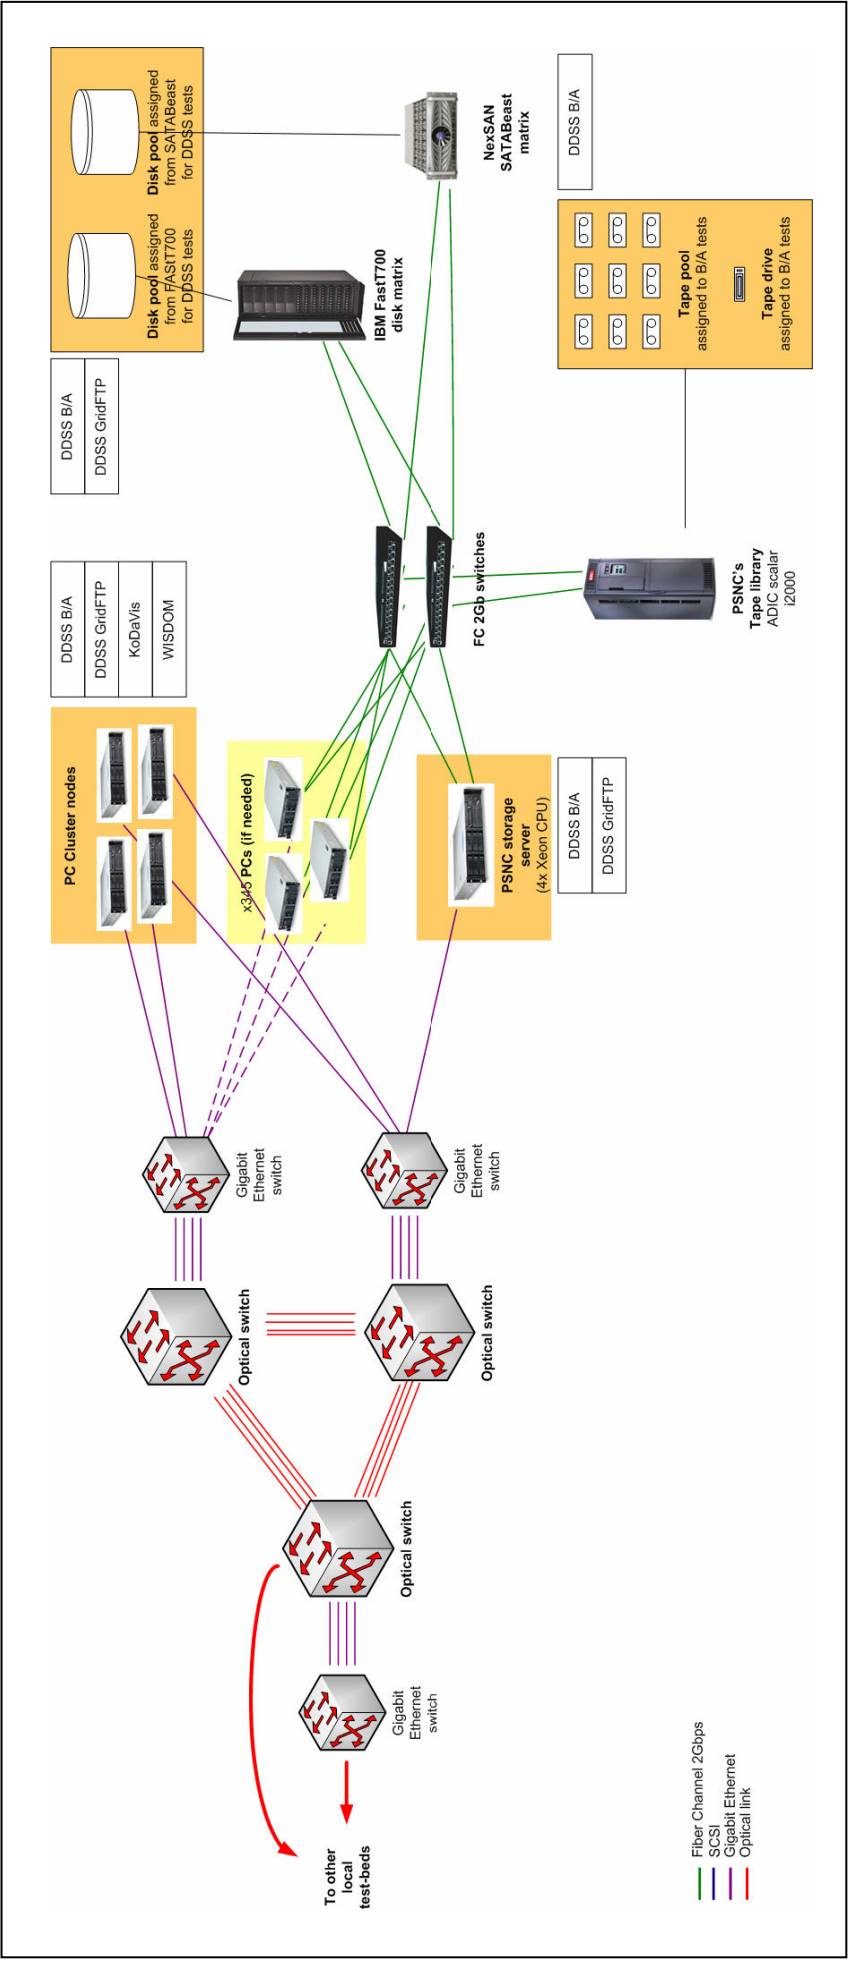


Figure 1.1: Topology of the PSNC local testbed



1.1.4 Control and provisioning plane in the local testbed

The optical switches and the Gigabit Ethernet switches in the will be controlled by GMPLS (in the first phase of the project) and G²MPLS (in the second stage). Each of the switches will have an external GMPLS/G²MPLS controller made of a personal computer with a GMPLS/G²MPLS implementation. The total number of controllers in the local testbed will be six (three controllers for optical switches and three for Gigabit Ethernet switches). The controllers will be connected to the control-plane network and will use the proper GMPLS/G²MPLS interfaces to communicate with other controllers in order to provision a path. The controllers will also use specialised interfaces to communicate with the controlled switches; the interfaces will be defined by Workpackage 2 of Phosphorus and will be based on TL1 and/or SNMP.

The control-plane network will be implemented as a layer 2 segment to which all the controllers will be connected by Fast Ethernet links. Flat structure of the network will allow for communication between each pair of controllers. The control-plane network will also interconnect GRID resources (computing nodes) with the controllers in order to allow for the exchange of control information between the GRID resources and the control plane of the network, e.g. signalling requests from GRID to the network.

The control-plane network will be connected to the control-plane networks of other local testbeds with use of a VPN. There are two possibilities: using software VPN like tinc on a PC or using IPsec tunnels on a dedicated Juniper router. Use of a VPN will protect the control-plane network from unauthorised access from the Internet.

The Phosphorus consortium does not envisage any need of having NRPS in the PSNC local testbed. This testbed is supposed to use only GMPLS/G²MPLS to establish lightpaths and to communicate with the control/provisioning plane of other local testbeds. In case the test scenarios require having NRPS in this local testbed it is possible to add NRPS to this testbed provided that the NRPS has interfaces to the switches installed in this testbed or is able to use GMPLS/G²MPLS to communicate with the switches.

1.1.5 Middleware and applications in the local testbed

Computing servers and the storage server will run the Linux operating system, which is appropriate for the testbed purposes. The details concerning software distribution, kernel version, etc. are going to be discussed with the testbed partners who are responsible for the test applications.

Computing servers (PC cluster nodes) are also going to run Globus Toolkit version 4 and the Grid FTP server and client modules. They will be used by one part of DDSS tests, i.e. Grid FTP transfers. Also the IBM Tivoli Storage Manager clients will be used on PC cluster nodes for DDSS backup/archive application tests.

PC cluster nodes will also run the KoDaVis client and data-server as well as the WISDOM client. Also UNICORE middleware will be installed in these nodes since they are necessary to support KoDaVis and WISDOM applications.

Project:	Phosphorus
Deliverable Number:	D6.1
Date of Issue:	14/11/07
EC Contract No.:	034115
Document Code:	Phosphorus-WP6-D6.1



Test-bed design

The storage server will run the IBM Tivoli Storage Manager server and client software as well as Globus Toolkit 4 and the Grid FTP server and client software that will be used for DDSS B/A and DDSS Grid FTP tests. It will not be used for other application tests.

Besides the test applications that will be provided by PSNC, PSNC cluster nodes will also run the applications provided by the other project partners. These partners' applications include KoDaVis and WISDOM. Therefore KoDaVis (test-only) client and WISDOM client will be installed in PSNC cluster nodes. Running these applications in PSNC testbed requires installation of the UNICORE middleware. Therefore it will be installed in PSNC cluster nodes.

According to WISDOM application tests requirements, Autodock application modules and RMS system will be installed in PSNC in the test cluster nodes. According to Appendix C of this deliverable (Requirements of Workpackage 3) the WISDOM application test requires access to FLEXlm- and Oracle servers. These servers will be located at VIOLA and will be accessed remotely from the PSNC testbed (this requires only the appropriate firewall configuration). RMS system is needed for WISDOM tests but it is also necessary to support KoDaVis tests. PSNC will act as test-only client for KoDaVis application tests, therefore installation of AVS/Express in PSNC is not required.

1.1.6 Interconnections with other local testbeds

As detailed in the next section of this document, the PSNC local testbed will have direct connections to several other Phosphorus local testbeds, namely CESNET, VIOLA, SURFnet, i2CAT, UESSEX and CRC. Each of the connections will be implemented as a lightpath which carry Gigabit Ethernet traffic.

Following the needs of the test scenarios the links will be terminated on one of the optical switches in the testbed or on one of the Gigabit Ethernet switches. Having the links terminated on a Gigabit Ethernet switch will allow for transmitting multiple VLANs on single physical links and this way simulating more lightpaths that are actually available. Each of the VLANs on a single physical link can be mapped on the Gigabit Ethernet switch to a different interface going to an optical switch and then switched separately in the optical domain of the local testbed. Having VLANs will also allow for provisioning transit VLANs – VLANs between two local testbeds which are not connected directly. Each of such VLANs can be seen by the applications and GRID middleware as direct links increasing the flexibility of the topology of the global testbed.

1.1.7 Integration with other local testbeds

To ensure smooth interactions between the PSNC local testbed and other Phosphorus local testbeds, precise integration is necessary. The integration must cover three layers: a control plane to ensure provisioning of lightpaths between local testbeds, the GRID middleware to allow for distributed computing, and use of distributed storage and applications to have the same application using resources in different local testbeds.

The GMPLS/G²MPLS control plane of the PSNC local testbed will be integrated with the control planes of other Phosphorus local testbeds which also use GMPLS/G²MPLS. The inter-domain interface used for integration of

Project:	Phosphorus
Deliverable Number:	D6.1
Date of Issue:	31/01/07
EC Contract No.:	034115
Document Code:	Phosphorus-WP6-D6.1



Test-bed design

the control plane will allow for exchanging information about each local testbed and provisioning optical paths between two or more local testbeds. In the first stage of the project the legacy GMPLS NNI interface can be used while the G²MPLS inter-domain interface are defined by Workpackage 2 and then deployed in the local testbed.

Direct integration of the PSNC local testbed and local testbeds using NRPSes will be possible when the NRPSes support interfaces to GMPLS/G²MPLS. Such interfaces will be defined jointly by Workpackage 1 and Workpackage 2.

On the GRID middleware layer the integration with other local testbeds will include:

- For Globus Toolkit:
 - arrangement of version 4 of the Globus Toolkit,
 - configuration of a common Certification Authority unit as well as creation and exchange of host and user certificates,
 - mapping of the distinguished names of certificates to local user accounts.
- For UNICORE:
 - installation and configuration of the UNICORE toolkit,
 - configuration of PSNC UNICORE nodes for collaboration with FHG nodes

Configuration of common Certification Authority for Globus Toolkit as well as installation of compatible versions of the Globus Toolkit and UNICORE are very important to assure the actual compatibility of the GRID middleware in local testbeds. This part of integration will be performed in collaboration with other Phosphorus partners.

On the application layer integration with other local testbeds will require:

- DDSS GridFTP: configuration of PSNC GridFTP servers and clients for collaboration with other partners (FZJ, FHG, UESSEX)
- DDSS B/A: configuration of PSNC Backup/Archive clients for collaboration with FZJ's TSM server
- KoDaVis: configuration of the application data-server and test-only clients
- WISDOM: configuration of the application clients

According to an agreement between Workpackage 3 partners (this workpackage is responsible for GRID middleware and applications in Phosphorus) each application will be configured and integrated by one Phosphorus partner – the partner who provides this application for use in Phosphorus. PSNC is responsible for the integration of GridFTP and B/A applications between different local testbeds while the other applications will be integrated by other Phosphorus partners.

1.1.8 Operations in the local testbed

The local testbed will be used by four applications run by Workpackage 3. When the applications are equipped with the G²MPLS interface, they will be able to communicate with the control plane of the testbed network in

Project:	Phosphorus
Deliverable Number:	D6.1
Date of Issue:	31/01/07
EC Contract No.:	034115
Document Code:	Phosphorus-WP6-D6.1



Test-bed design

order to automatically provision optical paths inside this local testbed and between this local testbed and other Phosphorus local testbeds. Two types of devices will be controlled by the control plane: optical switches and Gigabit Ethernet switches.

There will be four applications running in the local testbed: DDSS GridFTP, DDSS B/A, KoDaVis and WISDOM. Each of them will use part of the GRID resources available in the testbed. Distribution of the applications over the testbed resources is shown in Table 1.1. As multiple applications are supposed to use the same GRID resources and simultaneous use of the same infrastructure by different applications may not be possible, the resources will be assigned to applications based of a test schedule. Each of the applications will use the same switching and transmission infrastructure.

The use cases which will be tested in this local testbed have been described by Workpackage 3 in the WP3 Application Use-Case questionnaire.

Resource	DDSS GridFTP	DDSS B/A	KoDaVis	WISDOM
Hardware resources				
PC nodes (IA-64)	yes	yes	yes	yes
Storage server	yes	yes		
IBM x345 servers	if needed	if needed		
Disk pool on FASTT		yes		
Disk pool on SATA Beast		yes		
Tape drive & tapes on Adic library		yes		
Software resources				
Globus Toolkit middleware	yes			
Unicore middleware			yes	yes
TSM client/server		yes		
Grid FTP client/server	yes			
KoDaVis client/server			yes	
WISDOM client/server				yes

Table 1.1: Application distribution over PSNC local testbed



1.2 CESNET testbed

For a global Phosphorus testbed, CESNET can contribute with metropolitan, intercity and cross-border dark fibres for experimental purposes, pure optical lambdas in CESNET2 network and optical switching and routing equipment. Equipment can be used to switch fibres (optical switch), lambdas (SDH/SONET crossconnects) and to route or switch 1GbE/10GbE signals. The testbed can be connected to the PHOSPHORUS testbed via lambdas provided by GN2 or via cross-border dark fibre provided by CESNET. In addition we can provide connection services on 10Gb lambda from Prague to StarLight in Chicago.

Most of equipment have been already installed, one fibre line will be lighted later (DF Brno-Cieszyn).

1.2.1 Equipment and connectivity to the local testbed

Figure 1.2 shows a map of the local testbed. There are two main PoPs in Praha and Brno. The core network is based on Cisco 15454 MSP (3 boxes in Praha, 1 in Brno). These SONET crossconnects are connected via 10 G DWDM lambdas (ie lambdas with OC-192c framing) and are able to aggregate 1 GbE signals. In addition to SONET equipment, 1Gb/10Gb LAN and WAN PHYs GbE switching and routing platforms from Cisco (Catalyst 6503), Force10 (E300) and Foundry (BigIron, Edgeline, SuperX) are available. All equipment supports DWDM pluggable optics (XENPAK, XFP, GBIC, SFP), Cisco 15454 support full tuneable (32 channels in C band) transponders.

There are no dedicated core/edge routers or switches, but they can be deployed and used as needed.

Dark fibres Praha-Brno-Cieszyn are designed to support 10 Gbps transmission speeds today but may be used for higher bit rates (40 Gbps and higher), as in the optical core cloud testbed described by the UEssex.

A smaller testbed is dedicated for higher speeds (40, 80, 160 Gbps) and all optical signal experiments (2R regenerators, multiwavelength conversions, OTDM). This high-speed testbed can be connected to the local testbed as well.

An emulation of Grid-like resources is possible with powerful PCs with 1 GbE and 10 GbE cards. 10 GbE cards support pluggable DWDM transceivers (XFP) and can be connected to the local testbed either directly (DWDM) or through tuneable transponders. We believe it'd be very interesting and useful for our community to provide the end users with more than just 1 GbE channels over DWDM/MPLS networks.

Just to summarize, equipment from all three categories can be connected to the local CESNET testbed or directly to the Phosphorus testbed.

Project:	Phosphorus
Deliverable Number:	D6.1
Date of Issue:	31/01/07
EC Contract No.:	034115
Document Code:	Phosphorus-WP6-D6.1



Test-bed design

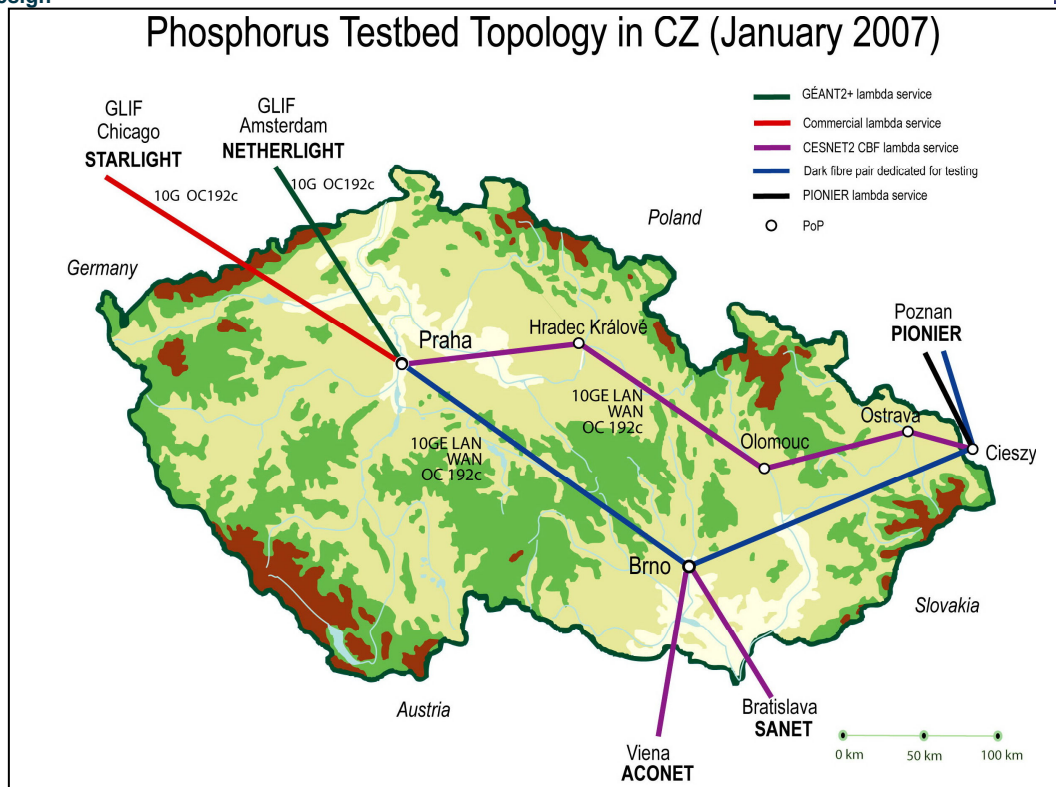


Figure 1.2: Testbed topology in CESNET

1.2.2 Management/control plane

CESNET is using commercial software to manage 15454 (Cisco Transport Controller) via a 'common' IP network and we're discussing with Cisco their plans for support of GMPLS.

We've been informed that Cisco has no plans to support GMPLS on MSPP 15454 this year, so we're interested in development efforts on external PC-based GMPLS controllers. We believe it'll be important also for our intent to deploy semiconductor optical switches within our local testbed, because these optical switches have no embedded GMPLS support. External controllers are perhaps the only solution how to integrate them with the rest of our equipment.

We'd like to add some remarks to the latest discussion about optical impairments. As our colleagues from UoE, we can perform different types of emulating/stressing/measurements (OSNR, CD). What we can append here is a possibility to perform these tasks not only in our lab (ie fibre reels) but on intercity and cross-border dark fibre links (ie Praha-Brno, Brno-Cieszyn). It may be difficult to do similar tests on lambda-based testbeds.

Equipment (ie SONET crossconnects, routers/switches, PCs) has no GMPLS built-in controllers but may be controlled via an external controller (PC with GMPLS or G²MPLS implementations) if such controllers are available.

Project:	Phosphorus
Deliverable Number:	D6.1
Date of Issue:	31/01/07
EC Contract No.:	034115
Document Code:	Phosphorus-WP6-D6.1



Test-bed design

1.2.3 Experimental dark fibre line Praha - Brno (CESNET)

1.2.3.1 Line

The line is composed of 298,5 km of G.652/G.655+/G.655- fibres having 67 dB of attenuation and 440 ps/nm of chromatic dispersion. There's one inline amplifier and no compensators of chromatic dispersion.

1.2.3.2 Deployment

The line is the longest one in the testbed and was planned to be deployed as NIL with utilization of Raman amplification. It all started as an experiment over 302 km/65 dB of mixture G.652/G.655+ on fibre reels because we were trying to simulate the real line. Dispersion compensation fibres were used together with EDFAs and Raman amplifiers. This worked quite well in our lab. The next step was to deploy all equipment on the real line. We found that the combination of G.655+/G.655- is almost balanced and therefore no compensation was needed (440 ps/nm is well within the margins of transceivers). The real attenuation 67 dB was too high (even the attenuation coefficient is 0,225 dB/km which is really excellent) and even with powerful Raman amplifiers the line experienced too bad BER. Therefore Raman amplifiers had to be displaced and one inline EDFA had to be deployed. As a matter of fact, a NIL solution with Ramans was too sensitive to planned future upgrades to more DWDM channels and the final solution with inline amplification was not that bad.

Perhaps one important result we have found is that G.652 fibres are a rather good choice when one need to work with high power equipment (some EDFA boosters and all Raman equipment) because of a larger effective area (i.e. a bigger core diameter). Therefore G.652 fibres are not so sensitive to nonlinear effects like self phase modulation (SPM) and cross phase modulation (CPM) for DWDM systems. Both last miles in Praha and Brno are G.652 fibres and not G.655 (we asked to extend the length of G.652 to approx. 25 km as last miles).

1.2.3.3 Advantages for project

The line connects two major cities Praha and Brno. Praha has connections to NetherLight (via Geant2+) and to StarLight (a direct lambda), Brno has Nx10G connections to Slovakia and Austria, which can be used for the project too. Another Nx10G line Brno – Cieszyn is being prepared – please see the following paragraphs.

Project:	Phosphorus
Deliverable Number:	D6.1
Date of Issue:	31/01/07
EC Contract No.:	034115
Document Code:	Phosphorus-WP6-D6.1

Test-bed design

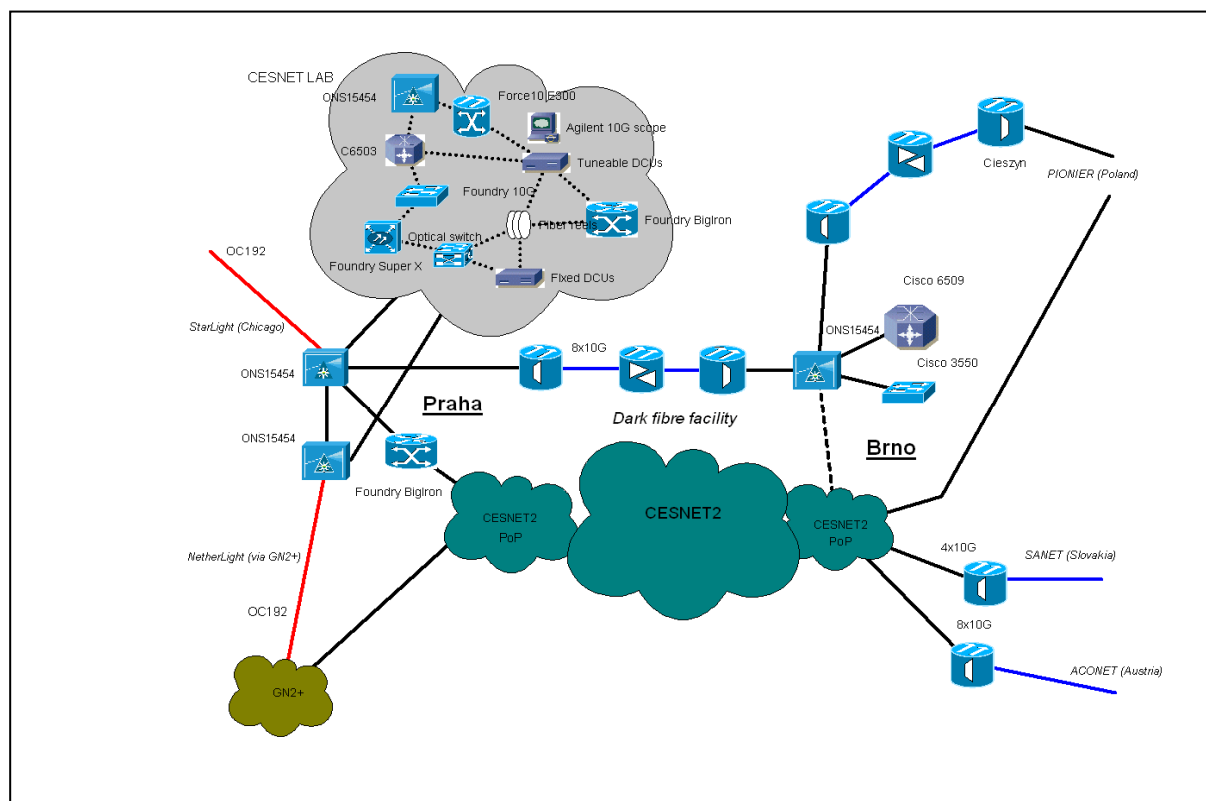



Figure 1.3: CESNET Dark fibre

The line can be used for speeds above 10 Gb/s if needed.



Project: Phosphorus
 Deliverable Number: D6.1
 Date of Issue: 31/01/07
 EC Contract No.: 034115
 Document Code: Phosphorus-WP6-D6.1



Test-bed design

1.2.4 Experimental dark fibre line Brno (CESNET) – Cieszyn (PIONIER)

1.2.4.1 Line

The line is composed of 264 km of G.652 fibre showing totally about 63 dB of attenuation and 4500 ps/nm of chromatic dispersion.

1.2.4.2 Planned optical deployment

A deployment is planned as cost effective static DWDM system based on CLA devices allowing up to 8 avelengths with speeds up to 10 Gbps. A NIL (Nothing-In-Line) solution is possible but would need distributed Raman amplification. To avoid high cost and high powers in fibres a solution with one inline amplification point was proposed. All amplification points will use CLA boxes with remote monitoring and management. Chromatic dispersion will be compensated using FBGs (Fibre-Bragg-Grating).

1.2.4.3 CLA amplifiers

The lack of optical equipment suitable for NREN's needs in past years led CESNET to develop optical -building kit composed from commercially available elements and modules. The result is low-profile and reliable rack mountable device having a plenty of management interfaces. Reliability is increased by e.g. power supply redundancy, important parameters monitoring, amplification and management part independency and usage of non-moving parts (hard disc).

Real life usability was proven in CzechLight testbed and then within NREN CESNET2 network in past. In year 2006 a CBF triangle A-CZ-SK was finished enabling the interconnections of neighbouring NREN CESNET2, SANET and ACOnet at the optical level (Layer-1). Two lines Brno – Bratislava (about 200 km, 45 dB attenuation, G.652 fibre) and Brno – Vienna (224 km, 50 dB, G.652) are operated by CLA amplifiers. More details can be found at <http://www.ces.net/doc/press/2006/pr061106.html>.


1.2.4.4 Advantages for project

Among the most important features belong a fact that device is operated by SW based on Linux and allowing easy remote monitoring, management and possibility of future development. Parameters like optical powers, currents, voltages, temperatures, fan speeds can be monitored.

Project:	Phosphorus
Deliverable Number:	D6.1
Date of Issue:	31/01/07
EC Contract No.:	034115
Document Code:	Phosphorus-WP6-D6.1



Test-bed design



Cross Border Dark Fibres Brno (CZ) – Cieszyn (PL), 263km, 8 x 10 Gb/s

1.2.5 Open dark fibre and photonic facility

CESNET has dark fibre facility and photonics lab, enabling building of testbeds for different projects and purposes. This allows integration of L1/L2 elements and support of designing the Internet of the future at all network levels (from dark fibre lighting to network applications in different research areas). Testbeds based on lambda connections only (without full access to dark fibres and programming of lighting devices) are considered useful but not sufficient for above task. CESNET has experience in deployment of dark fibres since 1999 and in deployment of open source photonics products since 2002.

The facility enables fundamental shift from current networking technologies that are optimized for legacy Internet, and do not offer the software programmability and flexibility required by the research community. Modular optical components (photonics industry products) readily available in the market and based on standards allow building of fully programmable optical network platforms. This extends well-known open approach from software development to hardware and optical devices development and management.

CESNET has 4 cross-border fibres (CBF) to neighbor NRENs now. We suppose our testbeds will be connected by CBF to other dark fibre testbeds (dark fibre islands) in different projects, if feasible from budgetary view (supposing fibres are available, cost for multispan lighting is now about 200 kEUR/1000 km for 8 lambdas using open systems). Interconnection of two or more dark fibre testbeds by lambdas will be seen as possible temporary solution. The main purpose of interconnection is to achieve synergy effect by joint research. Dark fibre testbeds will be open for partners for collaboration on the place or by remote access, if technically feasible.

For example, CzechLight Amplifiers (programmable optical devices with remote control) are now manufactured by two independent companies under license of CESNET and CESNET developed prototypes of CzechLight tuneable CD Compensator and CzechLight semiconductor optical switch. This means development of above or

Project:	Phosphorus
Deliverable Number:	D6.1
Date of Issue:	31/01/07
EC Contract No.:	034115
Document Code:	Phosphorus-WP6-D6.1



Test-bed design

similar devices will not be supported by Phosphorus project, but CESNET will support Phosphorus project by possibility to use such devices in the testbed.

For Phosphorus purposes, the facility will allow testing and integration of above devices as well as L1 agnostic switches, impairment monitoring and lightpaths switching and restoration, fibre lighting by new advanced products of photonics industry, testing of reconfigurable optical transport systems, testing of inter-operability in multi-vendor environment (including lighting products and alien waves), implementation of services by multiple lambdas (for example 10x10G or 5x20G), etc.

This means CESNET testbed for Phosphorus project can be modified step by step as result of intra-Phosphorus discussion about:

- what optical parameters of transmission are interesting for monitoring and optical resource management,
- what improvements of programmable optical devices (amplifiers, compensators, semiconductor switches) are recommendable

Project:	Phosphorus
Deliverable Number:	D6.1
Date of Issue:	31/01/07
EC Contract No.:	034115
Document Code:	Phosphorus-WP6-D6.1



1.3 i2CAT testbed

The Phosphorus testbed located at i2CAT is going to be used to test the interoperability between different NRPSs and between UCLP and GMPLS. It is also going to be used to test TOPS application. The switching layer will be composed by L2 switches and optical equipment from Nortel and Alcatel. Our testbed will also include some computational resources that can be connected to the Phosphorus testbed data plane.

The i2CAT local testbed data plane will be connected to other Phosphorus testbeds data planes using the infrastructure of the Spanish NREN (redIRIS) that is connected to Géant 2. However the i2CAT connection to GLIF facilities in Amsterdam will also be used to expand our Phosphorus testbed.

1.3.1 GRID resources in local testbed

Grid resources in the i2CAT local testbed include the following hardware:

- 9 x Blade Dell PowerEDGbE 1855, Pentium Xeon 2.8 Ghz, 512/1024 MB RAM, 73/140 GB HD and Gigabit Ethernet interface.
- 5 x Pentium IV 2.8 GHz , 512 MB RAM, 120 GB HD with Gigabit Ethernet interface.

All grid nodes have a Debian Linux OS and Globus 4.0 installed.

1.3.2 Optical networking and physical layer resources in the local testbed

At this moment it is not planned by i2CAT to use any optical switch like a Glimmer Glass switch or from any other vendor for its local testbed. Electrical switching will be implemented using Cisco switches. Most of them are Cisco 3750 series switches because of its policy, bandwidth adjustments and jumbo frames support. Gigabit Ethernet switches will not be controlled by UCLP version 2.

All connections inside the local testbed will use Gigabit Ethernet as the transmission technology. In some cases Gigabit Ethernet links will be over a DWDM 200GHz spaced ring while others (short distance) will use Gigabit Ethernet PHY over copper or fibre (850nm). In the same way, the connection to and between i2CAT computing resources will use Gigabit Ethernet links. In order to share the computing resources with other i2CAT projects we will use L2VPN's technology in our IP core network.

The DWDM ring will be implemented using two Nortel OPTera Metro 5200 boxes. Each one has the capability to transmit 2 Gigabit Ethernet plus 2 SDH channels up to 2,5 Gbps. The internal architecture of the Nortel Metro OPTera 5200 will permit the possibility to create a Gigabit Ethernet connection and teardown it. The possibility to connect one of the Gigabit Ethernet channels to the other one also exists, acting as a loopback. This capability will be interesting in order to connect different partners acting only as a transit network. The 2 SDH channels will not be used because i2CAT has not any client equipment supporting this

Project:	Phosphorus
Deliverable Number:	D6.1
Date of Issue:	31/01/07
EC Contract No.:	034115
Document Code:	Phosphorus-WP6-D6.1



Test-bed design

The three connections between i2CAT local testbed to VIOLA, SURFnet and PSNC local testbeds arrive as VLANs to i2CAT Ethernet/IP network, which can be transported to a Nortel Optera 5200 using internal L2VPNs and delivered as GigabitEthernet channels (without any L2 label).

The path between the first Nortel Optera 5200 to the applications (TOPS) and final user machines is composed by another Nortel Optera 5200 and some GigabitEthernet switches which all are UCLP capable.

On the other side, the connection between i2CAT local testbed to CRC local testbed can be a dedicated GigabitEthernet link or a VLAN and it will end in a Cisco ONS 15454E which is connected to the second Nortel Optera 5200, so the path to the application or final user is UCLP capable too.

In summary, any connection starting at i2CAT local testbed can end at VIOLA, SURFnet, PSNC or CRC testbeds using the Nortel Optera 5200 to switch between the links.

1.3.4 Control and provisioning plane in the local testbed

As an NRPS developer i2CAT will deploy UCLP on its local test bed. In particular we will use UCLP to control the Nortel OPTera Metro platform and the CISCO ONS 15454. The server running UCLP will be located on i2CAT premises and will have two GigabitEthernet interfaces. One connected to the Control Plane using a VPN created using software and the other one connected to a port of a L2 switch with a dedicated VLAN for the management of the optical equipment involved on this project. This way all the optical equipment that will be controlled on the i2CAT Phosphorus testbed will have their management ports connected to ports on L2 switches with the management VLAN configured on them.

We have mainly two options regarding the management protocol used to configure the equipment: SNMP and TL1. We have decided to use the TL1 interface to set up connections.

In case we want to interconnect UCLP with other NRPS or GMPLS controlled domains, our UCLP server will be able to communicate with them through the interface connected to internet using a VPN software. Moreover, this interface will also allow a user located on a remote site to control the i2CAT using the UCLP software running in the UCLP server.

1.3.5 Middleware and applications in the local testbed

Our test clients will be composed by computing resources in the local testbed that run a Linux operating system and have a GT4 stack installed. During the first phase of the project we plan to have as an application in our testbed TOPS. In particular we will configure a cluster to act as a visualization end point using a tiled display.

Project:	Phosphorus
Deliverable Number:	D6.1
Date of Issue:	31/01/07
EC Contract No.:	034115
Document Code:	Phosphorus-WP6-D6.1

1.3.6 Interconnection with other local testbed

i2CAT local testbed will be connected to four different partners using two international connections.




Figure 1.5: i2CAT testbed external connections

Firstly, i2CAT local testbed will connect to CRC local testbed through the GLIF infrastructure using a dedicated GigabitEthernet channel. The GigabitEthernet link will end in a Cisco ONS dedicated to the project.

Secondly, i2CAT local testbed will connect to VIOLA, SURFnet and PSNC local testbeds using L2VPNs provided by Géant2 and local NRENs. Each of these VPNs ends in a different VLAN in our link to the Spanish NREN, which also has other projects and production channels. It means that these three connections to VIOLA, SURFnet and PSNC will share the bandwidth (1 Gbps) with other projects and our base production Internet connection.

The L2VPNs are provided by Géant2 in collaboration with NRENs, which is a good way to have direct L2 connections between partners around Europe without having a high cost. But it has their own limitations, as the bandwidth through the Géant2/NREN networks is always shared with other connections and we don't have any control or management of the VPN in the entire path from origin to destination.



Test-bed design

1.3.7 Integration with other local testbeds

At control plane level, i2CAT testbed will be connected with the other two types of NRPS: ARGON and DRAC. All the NRPS-controlled domains previously registered with the Network Service Plane, will be able to intercommunicate between them and create interdomain and intradomain connections through the correspondent NRPS.

At the Grid Middleware layer, for TOPS application test, i2CAT testbed will install and deploy Globus Toolkit version 4.

At the application layer, integration with TOPS is required.

1.3.8 Operation in local testbed

The i2CAT local testbed will participate in the TOPS application test. The local UCLP NRPS will be able to configure optical paths inside the local testbed and will communicate with the other NRPSs to create the connections required between the local testbed and the other Phosphorus testbeds. Inside the local testbed, the NRPS will control 2 Nortel Optera Metro 5200 and in the short term the CISCO ONS15454.



1.4 SURFnet testbed

The SURFnet testbed will be dedicated to developments in Work Package 1. It is built from Nortel OME6500 nodes that are controlled by DRAC, and will have client nodes to do proof-of-principle application testing. It is foreseen that in the next evolution of the testbed where integration with applications becomes possible these tests will also be running over the SURFnet testbed, as local connectivity to the UvA and SARA who participate in other parts of the Phosphorus project makes this relatively straightforward.

1.4.1 GRID resources in the local testbed

No computing grid resources are currently foreseen.

1.4.2 Switching and transmission resources in the local testbed

The SURFnet testbed consists of three SDH MSPPs (multi-service provisioning platforms), these are interconnected between themselves and allow direct connection of application servers (or when needed applications can be connected across the testbed to allow for proof-of-principle testing). The network is connected to external networks through an STM-64 that can be setup to provision connections through the GLIF network or GbEANT2 as needed. Access to the testbed will be limited, although a public interface to the DRAC test server is planned for Q2 2007.

Currently the client interfaces are GbE only. During the time of the project we envisage the possibility to interconnect locally at 10GbE using a 10G EPL card. This however would need to be mirrored at the other side to allow end-to-end connections.

Test-bed design

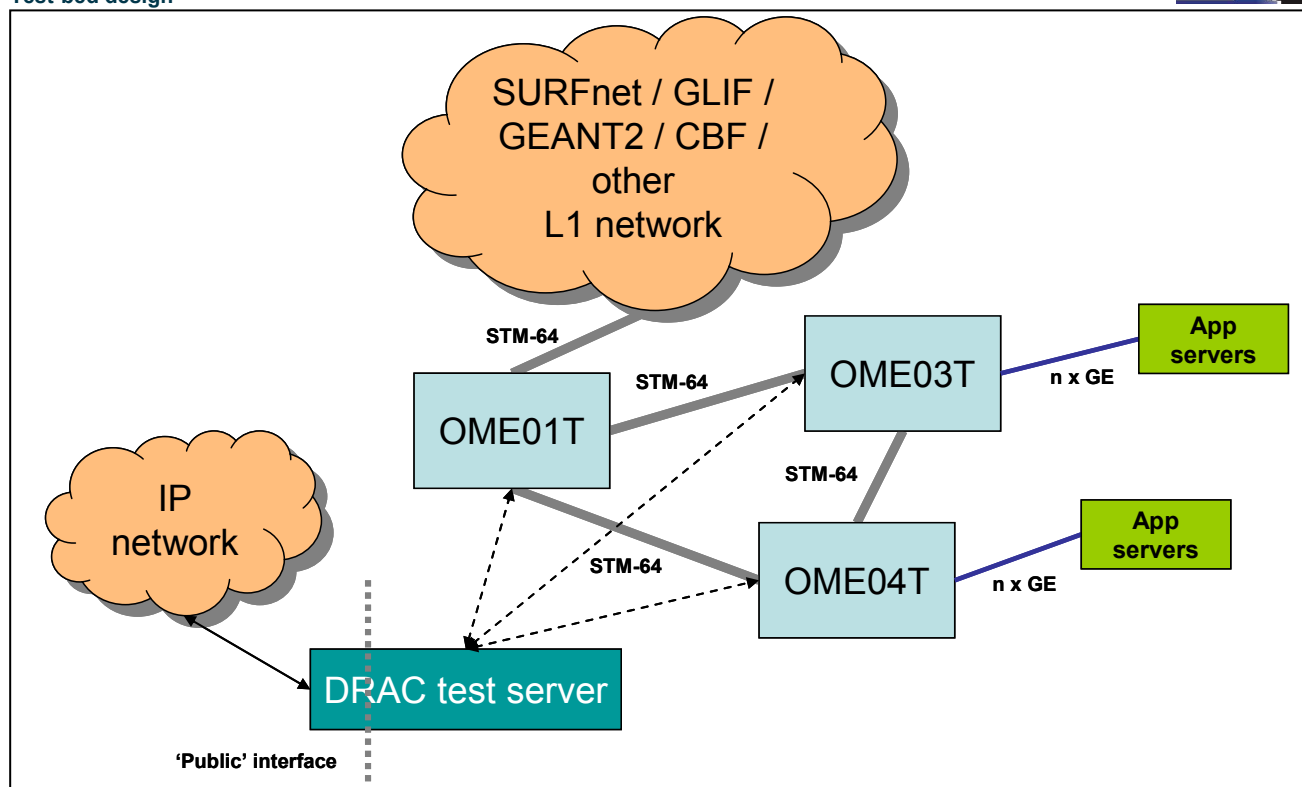


Figure 1.6: SURFnet local testbed

1.4.3 Topology of connections inside the local testbed

The testbed has STM-64s in its backbone, and is connected to the outside network with an STM-64. The three nodes are configured in a triangle, where one OME has an exit towards the external network. This setup allows for testing of basic failure scenarios and tests on constrained routes.

If links to other testbeds through the international connectivity will be built with VLANs an extra L2 switch will be installed and configured to allow for this. Connectivity will then be at GbE level, through client ports on the OMEs. This approach can also be used to test multiple applications servers.

1.4.4 Control and provisioning plane in the local testbed

The nodes in the testbed are controlled by DRAC. From the DRAC server a direct connection to the DCN of the testbed exists. Setting up services on the testbed will be done through the public interface of the DRAC server. This will be accessible through either a “real” public IP address, or through a VPN server. With access to the DRAC server a web page is reachable that allows for manual provisioning, and web services for automated provisioning. To circumvent the requirement to have an A-Select identity local provisioning of users is possible. For Web Services certificates that allow specific use shall be provided.

1.4.5 Middleware and applications in the local testbed

No middleware and applications are foreseen in the SURFnet testbed. Direct connections to the UvA and SARA application testbeds are envisaged for integration and application testing where applicable.

1.4.6 Interconnections with other local testbeds

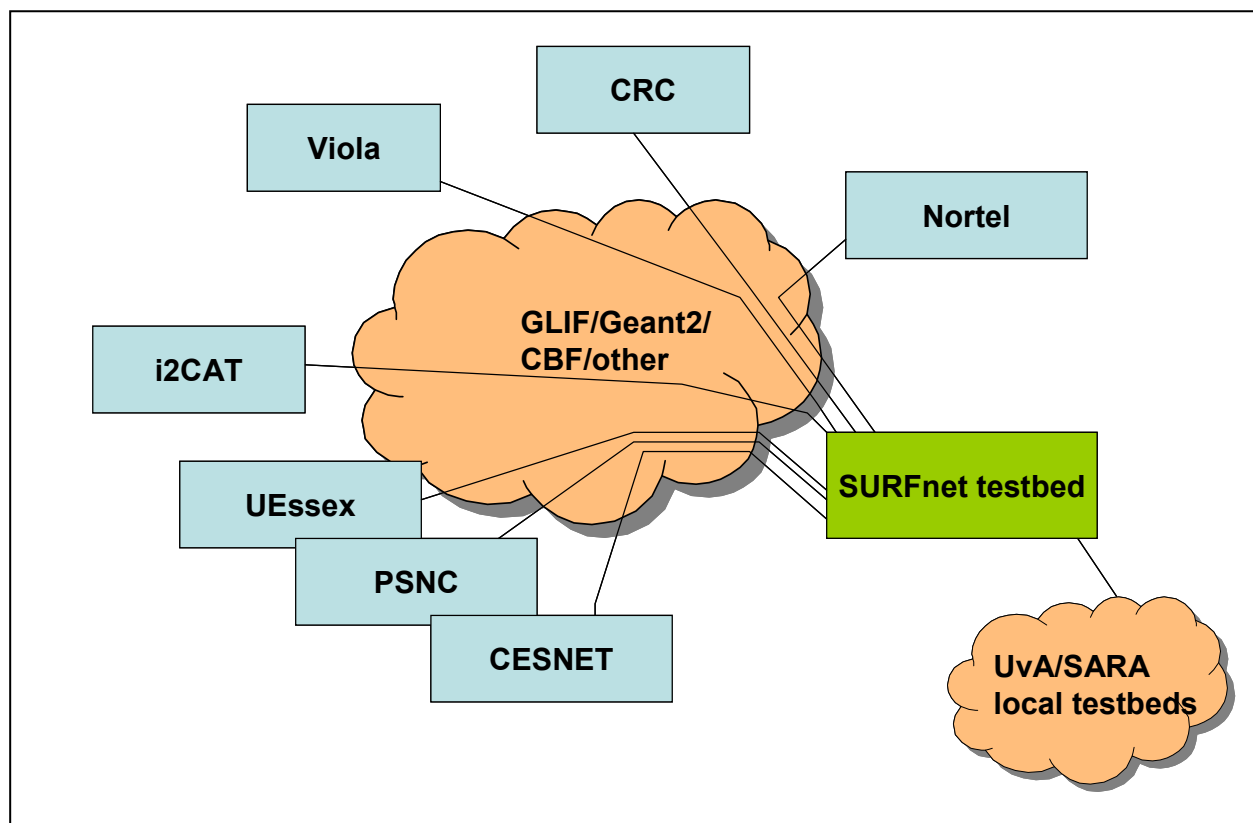


Figure 1.7: Connection to other local testbeds

For testing in WP1 the SURFnet testbed will have direct connections to other local testbeds involved in WP1, namely VIOLA, i2CAT and CRC. Each of the connections will preferably be implemented as a lightpath which carries Gigabit Ethernet traffic.

To allow other parties to connect to each other SURFnet will provide transit when the L2 switch option is chosen. This means that PSNC and UESsex will also be connected, most likely across Geant waves, and that CESNET will have a connection over their 10G connection to NetherLight.

All connections will either locally have access to client ports connecting to a test server, or to client ports connecting to servers running Phosphorus applications running in the UvA or SARA environment.



1.5 UvA testbed

The UvA testbed will be dedicated to developments in Work Package 1 and 4 and may be used by applications if needed and fits in the schedule. The testbed consists of several smaller testbeds, which allows experiment with Authorization models using different types and implementations of control plane. We also intend to investigate use of the Network Description Language (<http://www.science.uva.nl/research/sne/ndl>) as semantic model for topology information exchange between the various UvA testbeds. Currently we deploy GMPLS based on DRAGON, a NPU (IXP2850) based development environment for in band Token Based Authorization which will be developed as a FORCES based GMPLS edge routers, and a UCLP testbed build together with SARA at a co-operated experimental facility LightHouse. We also have several cluster computer setups that can be used to experiment with applications using the various control plane implementations. We expect SURFnet to set up a DRAC based testbed. A new testbed consisting of 3 clusters of control plane nodes and 802.1Q VLAN switches will be integrated to allow interoperability testing with various GMPLS implementations deploying G2MPLS, DRAGON, etc. Also, authorization interoperability with domains using NPRS's such as DRAC, ARGON and UCLP are intended to be supported.

1.5.1 GRID resources in the local testbed

We have three different cluster computer setups that can be connected to the testbed. Those are the Van Gogh cluster, the Rembrandt cluster and DAS-3 connected via StarPlane (www.starplane.org)

Furthermore there are several individual computers that support service, control and data plane functions in the testbed. Here are some pictures of the clusters:

Project:	Phosphorus
Deliverable Number:	D6.1
Date of Issue:	31/01/07
EC Contract No.:	034115
Document Code:	Phosphorus-WP6-D6.1

Test-bed design

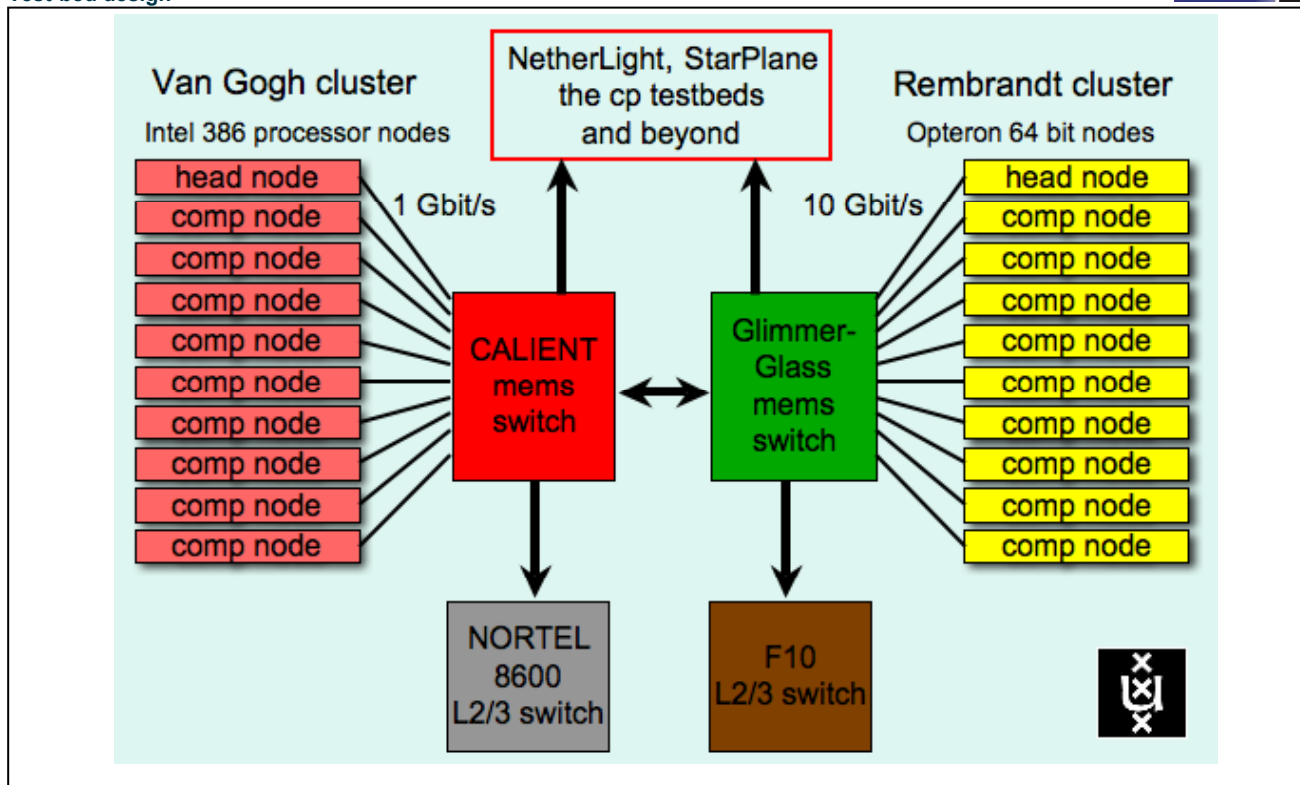


Figure 1.8: The VanGogh / Rembrandt cluster nodes in Netherlight / Lighthouse

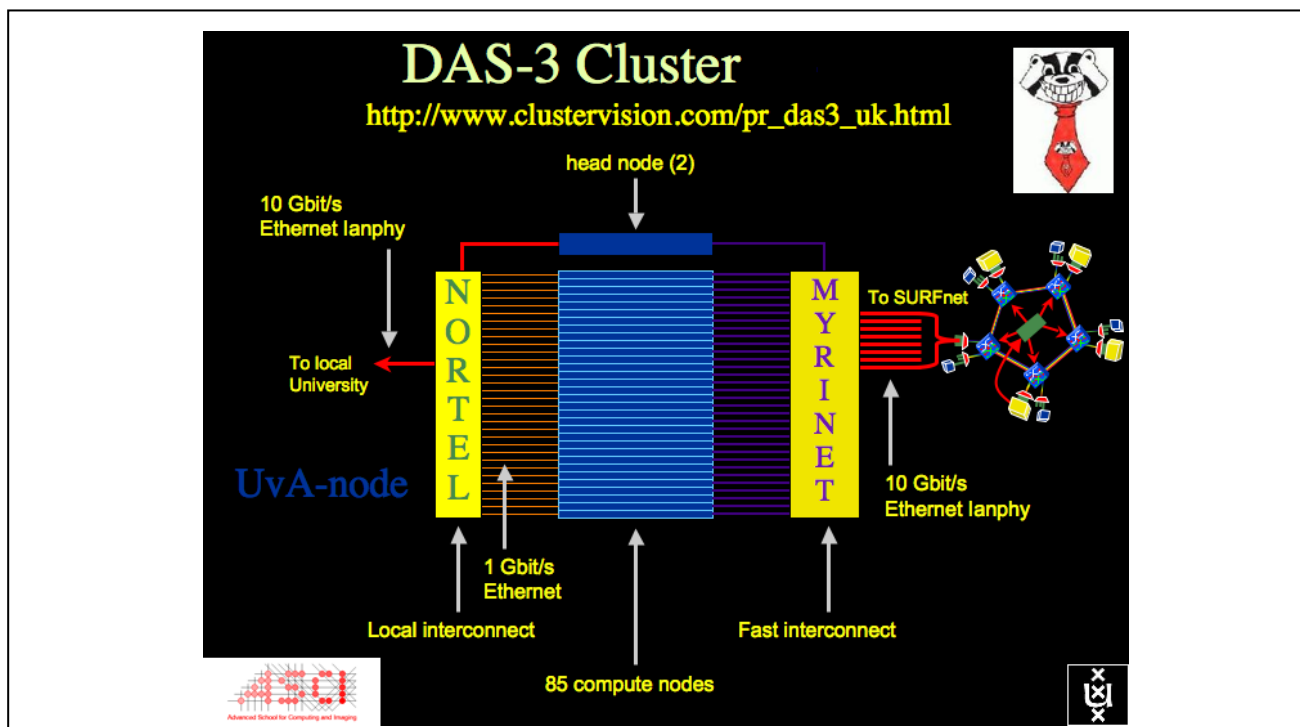


Figure 1.9: The DAS-3 Cluster Node at UvA Netherlight



Test-bed design

1.5.2 Switching and transmission resources in the local testbed

Here we describe only the switching elements in our control plane test setups in our laboratory. Obviously the MEMS devices in the picture above are also included. The StarPlane network is described on <http://www.starplane.org/>.

The first described CP-testbed is the GMPLS setup.

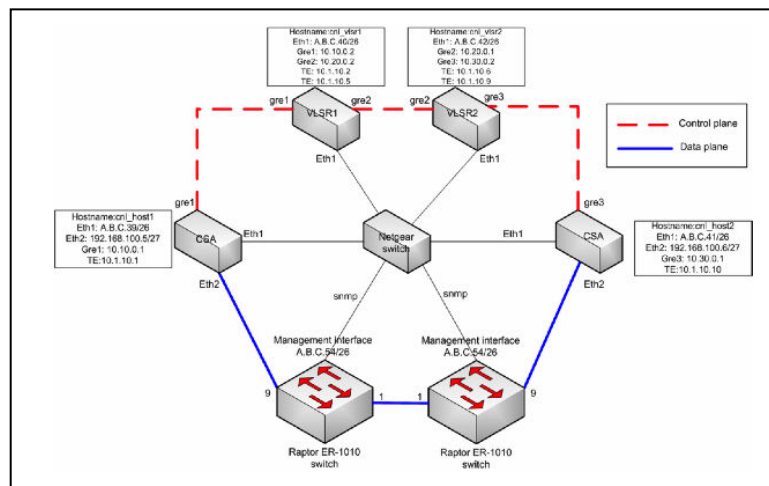


Figure 1.10: The GMPLS testbed setup

The second CP-testbed is the UCLP2 setup:




Figure 1.11: The UCLP Testbed based on NORTEL OME's

The third testbed to note is the Token Based Network testbed: Token Based Switch is a low-level system for traffic routing at high speeds (multi gigabits/sec) based on packet authentication. TBS helps high-performance computing and grid applications that require high bandwidth links between grid nodes to bypass the regular Internet for authorised packets by establishing shortcuts on network links with policy constraints. TBS is fast and safe and uses the latest network processor generation (Intel IXP2850). The associated testbed architecture is depicted in the next pictures. The first figure (1.12) shows an overview of the context where TBS-IP systems will be plugged in. In this first case, we suppose to build an end-to-end connection composed of lightpaths based only on IP packets. The second figure (1.13) shows another context where TBS-IP systems work



Test-bed design

together with TBS-GMPLS. In other words, in this scheme we suppose to have some GMPLS lightpaths within the end-to-end connection that uses normally IP packets.

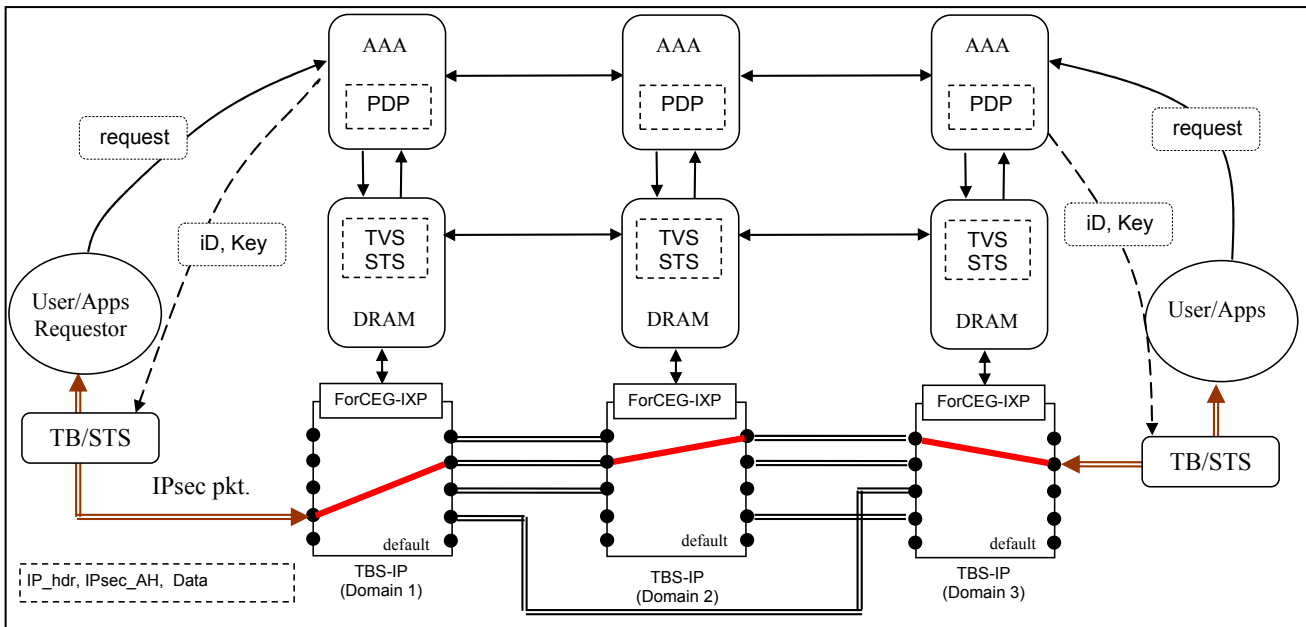


Figure 1.12: The planned Token Based FORCES router setup

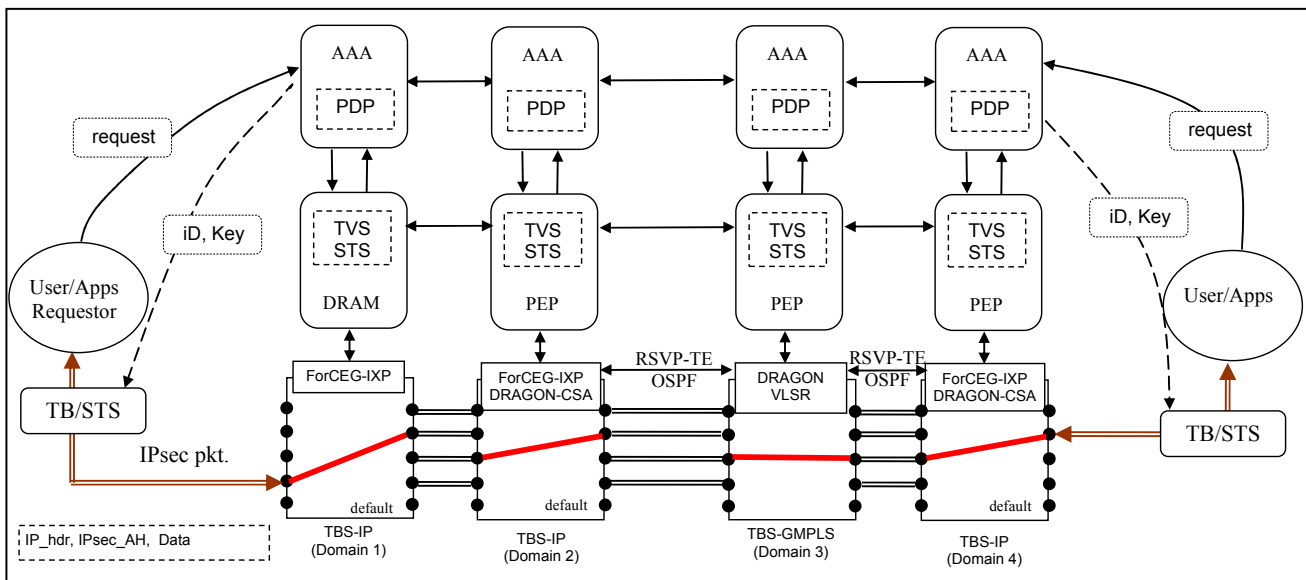


Figure 1.13: A future multi-domain Token Based test setup using GMPLS

1.5.3 Multi domain AAA testbed.

This testbed is being constructed and should be operational end of Q1 2007. It is intended to support various GMPLS implementations such as DRAGON and, when available from WP2, G2MPLS. Also NPRS based domains can be implemented in the testbed. Each domain in shown picture has 3 CPU nodes, 2 nodes are intended to act as a control plane node, driving 802.1Q VLAN switches (or other types if necessary and made available to this testbed) and accepting and forwarding signalling messages via an east-west interface and communicating operation and control messages via a North/South bound interface, which are generated by the 3rd node. The 3rd node will act as a service plane / AAA node that provisions the control plane with service and AAA information. The testbed will be made accessible to both partners from the Phosphorus project and organizations that want to collaborate in the area of AAA, such as Internet2.

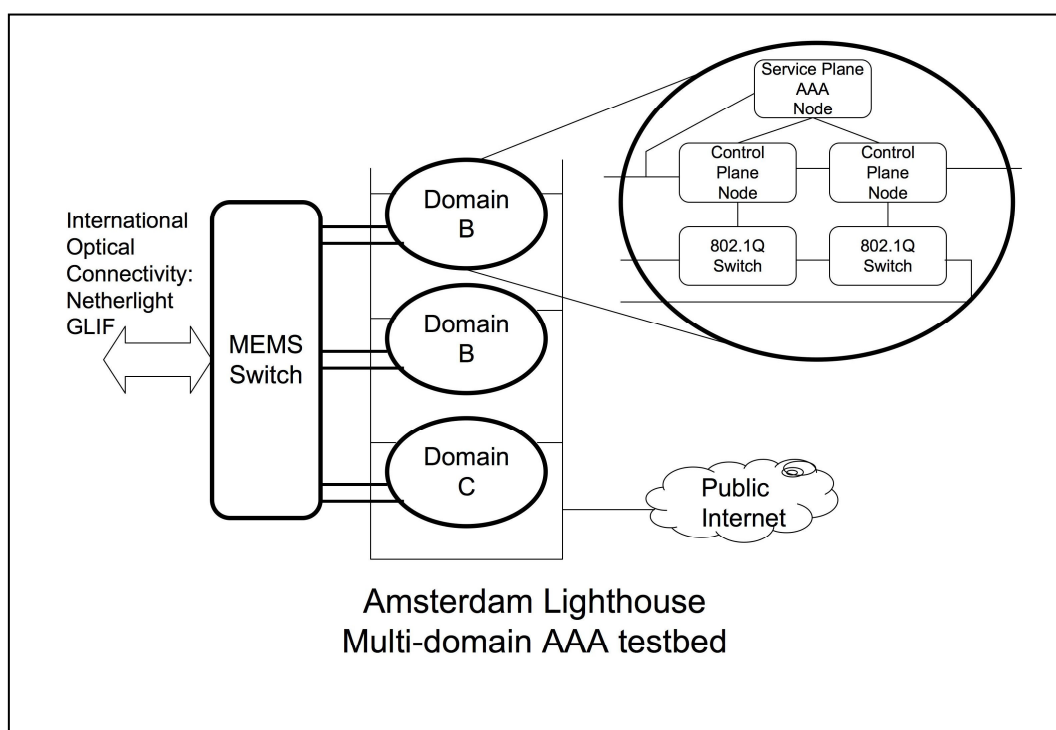


Figure 1.14: Multi-domain Generic AAA Testbed at UvA Lighthouse

1.5.4 Topology of connections inside the local testbed

The different parts of the testbeds are connected with fibre supporting 1 and 10 Gbps, passing through the MEMS devices from Calient and GlimmerGlass. This will enable dynamic reconfiguration of the inter-connectivity.



Test-bed design

1.5.5 Control and provisioning plane in the local testbed

Part of the GMPLS testbed will initially be controlled by the DRAGON project implementation. This implementation has been augmented with our own AAA modules to allow Authorization. This functionality is intended to be incorporated into the Phosphorus GMPLS/G2MPLS implementations when they become available. Other parts of our testbed will be controlled by NPRS style control/service plane implementations using for example webservices based interfacing. Examples are the UCLPv2 implementation.

1.5.6 Middleware and applications in the local testbed

The middleware deployed in this testbed is based on the current releases of UCPL2 and Dragon GMPLS software augmented with our components to enable the topology information framework as developed in WP1 and Authorization modules as developed in WP4.

1.5.7 Interconnections with other local testbeds

Interconnection with other local testbeds will be done via the MEMS devices and our layer 2 switches as depicted in the first picture in this section. Connectivity to other testbeds is handled via NetherLight.

Project:	Phosphorus
Deliverable Number:	D6.1
Date of Issue:	31/01/07
EC Contract No.:	034115
Document Code:	Phosphorus-WP6-D6.1



1.6 UESSEX testbed

In Phosphorus testbed, UEssex will contribute a hybrid optical network testbed involving wavelength switching, sub-wavelength switching (OBS) and GigEthernet/VLAN switching. Additionally University of Essex is under process of building a small Grid computing Cluster which is going to be connected to optical testbed and provide computing resources for PHOSPHORUS test applications. The UEssex can connect to the PHOSPHORUS testbed through dark fibre provided by UKLight and GN2.

Currently, UEssex closely liaises with UKERNA (United Kingdom Education and Research Networking Association) in order to prepare the connectivity between the UEssex optical network testbed and the phosphorus testbed.

UKERNA has agreed to provide the required connectivity through UKLight Dark-Fiber. UKLigth is a national facility to support projects working on developments towards optical networks and the applications that will use them. The UKLigth Dark-Fiber will connect UEssex photonic Network Lab to the international access point of UKLigth, where it will be connected to Géant2 and GLIF and consequently to the Phosphorus testbed. UKERNA has agreed to provide free Gigabit Ethernet connectivity to PHOSPHORUS.

The UEssex local testbed is going to participate in INCA, DSSS and KoDaVis application tests as defined by workpackage 3.

1.6.1 GRID resources in local testbed

It is a small Grid computing cluster which is under development and it is planned to be ready for second phase of the project. It comprises:

- One CPU cluster node with a specific number of processor (vendor and configuration to be decided)
- Storage and file server recourses (vendor to be decided)
- Two PC clients (Dual core Intel Lithium processor with Linux OS) with Gigabit Ethernet interfaces to participate as a test-only client in test applications
- Anritsu Network/data analyzer/generator which can emulate traffic and analyse data traffic in the testbed at 1GbEthernet and 10GbEthernet.

The computing resources in UEssex include one CPU cluster which the vendor, specification and configuration still to be confirmed. The storage resources allocated by UEssex include a disk matrix and/or two file servers which the vendor, specification and configuration still to be confirmed. UEssex Grid resources also include a network monitoring device which is an Antirust network/data analyzer/generator (MD1230) with the capability to emulate traffic and also analyse application data traffic at speed of 1Gigabit Ethernet and 10Gigabit Ethernet and in different networking layers. Furthermore, two PC clients are included in the Grid resources which are equipped with Gigabit Ethernet interfaces and they will participate as a test-only client in testbed applications.

Project:	Phosphorus
Deliverable Number:	D6.1
Date of Issue:	31/01/07
EC Contract No.:	034115
Document Code:	Phosphorus-WP6-D6.1



Test-bed design

The UEssex will participate into the first phase of the project with two PC clients (test-only clients) and the network monitoring device. The full Grid resources will be ready for participation in the second phase of the project.

1.6.2 Optical networking and physical layer resources in the local testbed

The optical networking testbed in the UEssex comprises of three main optical testbeds: the optical core wavelength switched testbed, the optical burst switching testbed or sub-wavelength switched testbed, high-speed all optical signal processing testbed.

A. Optical core or wavelength switched testbed: It comprises one Calient DiamondWave FiberConnect switch. This switch has been partitioned into four virtual independent switches. These switches are connected in a semi-mesh topology and interconnect different parts of the local testbed as shown in figure 1.15. The optical core testbed has the capability for DWDM transmission at Bit rates: 1 Gbps, 10 Gbps, 40 Gbps, and 160 Gbps (to be tested).

B. Optical burst switching (OBS) testbed or sub-wavelength switching testbed: it is a sub-wavelength switching testbed with capability of data burst transmission at 2.5 Gbps. The testbed has been developed in UEssex and comprises two edge routers and one core router. Each edge router can accept Gigabit Ethernet traffic and is able to map them into optical burst according to the application requirements. The core OBS router is an active two section hybrid optical switch. The active part of the router comprises a high-speed network processing engine which is able to process data carried within OBS control header and perform advanced application-aware routing functionalities. The optical switching part of the core OBS router comprises of a high-speed (ns regime) SOA based switch for routing of small data bursts and a slow (20ms) 2D MEM-based switch for routing of large data bursts.

C. High-speed all optical signal processing testbed: this is a high-speed optical transmission testbed with capability to support 160 Gbps transmission rate. With this testbed, it is possible to emulate different types of transmission impairments such as loss and dispersion. It also has multi wavelengths 2R & 3R regeneration facilities to compensate transmission impairments. Furthermore it supports all optical signal processing and OTDM functionalities such as :all optical clock recovery, all optical data recovery and all optical time slot interchanging.

The optical core (wavelength switched testbed) provides connectivity between Phosphorus test bed and UEssex local optical and Grid testbed facilities as shown in Figure 1.15.

The UEssex local testbed can provide high-speed and intelligent sub-wavelength switching within the sub-wavelength switching testbed. This test bed will be participated in PHOSPHRUS testbed if it is necessary and appropriate.

High-speed all optical testbed can emulate different types of transmission impairments such as loss, dispersion and transmission links with different lengths and qualities. This testbed will participate to the PHOSPHORUS testbed to simulate different optical paths with different types of transmission quality and impairments. This will

Project:	Phosphorus
Deliverable Number:	D6.1
Date of Issue:	31/01/07
EC Contract No.:	034115
Document Code:	Phosphorus-WP6-D6.1



Test-bed design

allow the PHOSPHORUS control plane to incorporate and consider constrain based routing with physical layer constraints as described in technical annex.

The local testbed also comprises two VLAN Gigabit Ethernet switches (vendor and type to be decided). One switch connects the local Grid testbed to the optical core wavelength switched testbed. The local Grid resources use Gigabit Ethernet or Fast Ethernet for communication and their traffic is aggregated by the VLAN Gigabit Ethernet switch for transmission over optical network. The second switch interconnects the phosphorus global testbed to the optical core wavelength switched local testbed. It segregates/aggregates VLAN tagged traffic from/to different local testbeds and interconnect them individually to the core of the UEssex local testbed.

1.6.3 Topology of connections inside the local testbed

Figure 1.15 shows the UEssex hybrid optical network and Grid testbed. As It was mentioned earlier it comprises four main clouds: optical core cloud (wavelength switched testbed), optical burst switching testbed (sub-wavelength switched cloud), high-speed all optical signal processing testbed and Grid computing testbed.

Optical core wavelength switched testbed comprises one 96X96 Calient DiamondWave FiberConnect switch. This switch has been partitioned into 4 virtual independent 24X24 switches (VS1,VS2,VS3,VS4). These switches are connected in a semi-mesh topology as shown in figure 1.15. They are connected together with two main fibres (one for each direction) and two protection fibres. These switches interconnect the Grid testbed, the optical sub-wavelength switched testbed and the high-speed all optical testbed to the PHOSPHORUS testbed

One VLAN Gigabit Ethernet switch (VG1) interconnects the UEssex local testbed to the PHOSPHORUS global testbed. It provides an aggregation/segregation point for outgoing/incoming VLAN tagged traffic to/from other local testbeds in PHOSPHORUS. Deployment of the VG1 switch in the testbed allows segregation of VLAN tagged traffic from different sources into individual links for switching through different part of the local testbed.

Another VLAN Gigabit Ethernet switch (VG2) interconnects one of the optical switches (VS1) to the local Grid-testbed. This provides an aggregation/segregation point between phosphorus testbed bandwidth/traffic and the local Grid resources which use Fast or Gigabit Ethernet links.

In summary, in the UEssex local testbed, the incoming wavelengths (Gigabit Ethernet links) from Phosphorus testbed enter to VLAN Gigabit Ethernet switch VG1 and VLAN tagged traffic tributaries are segregated into individual optical links. This links enter to the optical switch VS2 and they are switched though a mesh of four optical switches to a suitable testbed. Each incoming wavelength or each optical link (each VLAN tagged traffic tributary) can be routed independently through VS4 into the all optical testbed where it can experience different optical paths with different qualities and different types of physical layer impairments. After that it can be switched to exit the local testbed through VS4 and VS1 optical switches. Each incoming wavelength or optical link (VLAN tagged traffic tributary) can be also routed through VS1 into the Grid testbed where it can be routed by the Gigabit Ethernet switch VG2 to different resources according to application requirements. In the similar way the traffic from Grid resources will be aggregated into the PHOSPHORUS testbed wavelength by the VG2 switch and will be switched to exit the local test bed through VS2, VS1 and VG1 optical switches.

Project:	Phosphorus
Deliverable Number:	D6.1
Date of Issue:	31/01/07
EC Contract No.:	034115
Document Code:	Phosphorus-WP6-D6.1



Test-bed design

Topology of the testbed and also the testbed resources is shown in Figure 1.15.

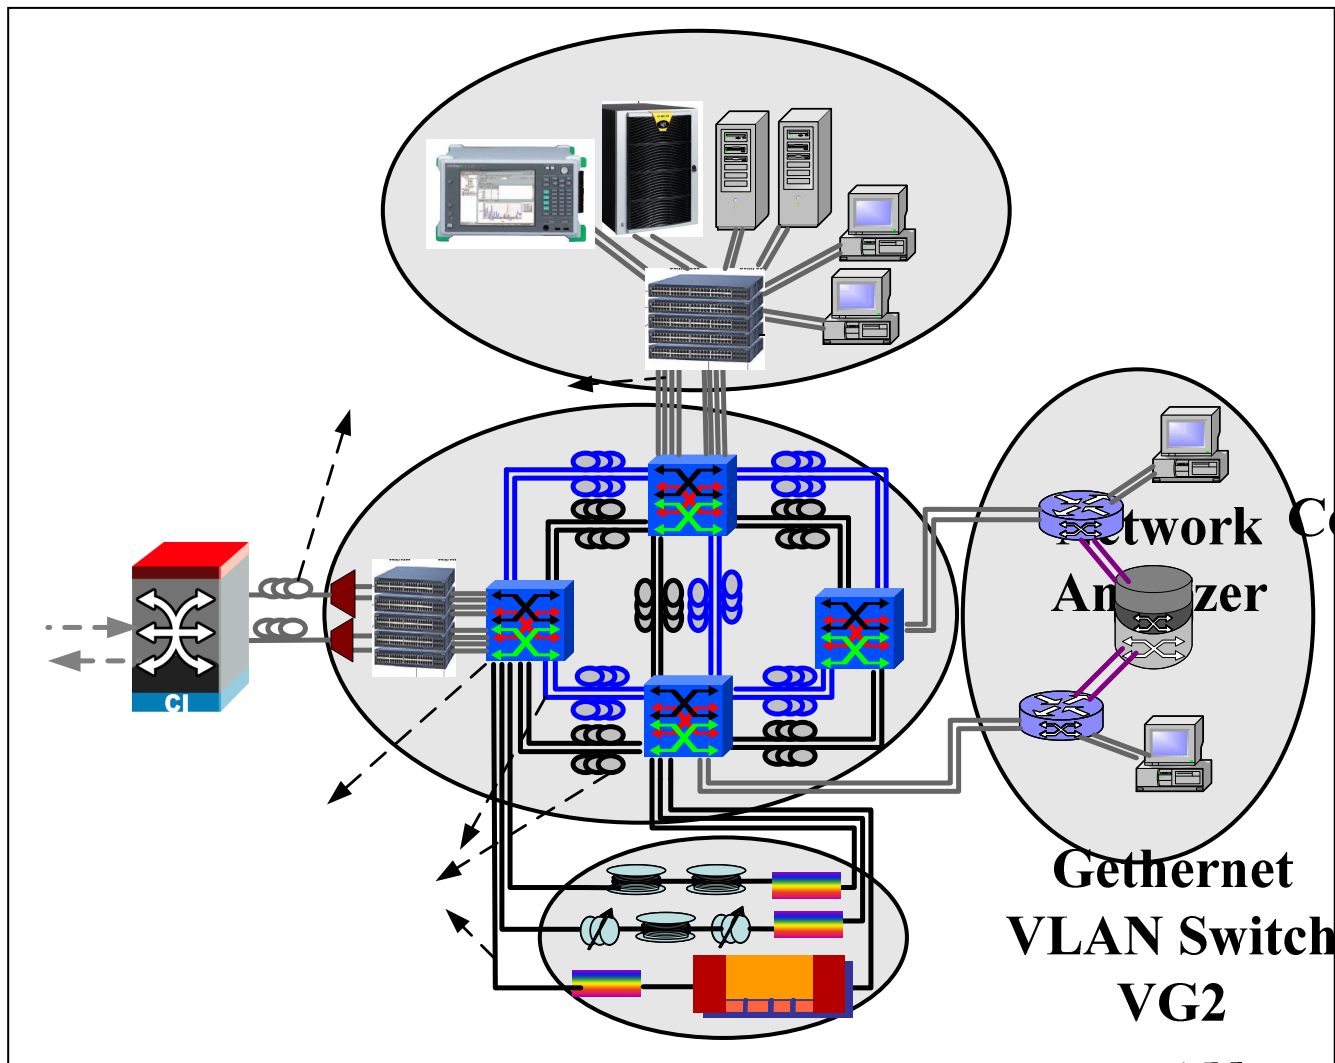


Figure 1.15: The UEssex topology of the local testbed and the testbed resources

1.6.4 Control and monitoring planes in the local testbed

The UEssex local testbed will deploy GMPLS control plane in the first phase of the project and G²MPLS control plane in the second phase. The GMPLS control plane will control the optical switches at the core and two VLAN Gigabit Ethernet switch at the edge of the testbed. Each switch in the testbed will have an external PC based controller with a GMPLS/G²MPLS implementation. These controller uses TL1 or SNMP based interfaces to communicate with the switches. There will be four control nodes for core switches and one control node for each edge Gigabit Ethernet switch. The interconnection between control nodes will be through a control plane network using fast Ethernet link. The communication between core control nodes and also communication between edge VG1 control node and core VS2 control node will be based on a specific

Project: Phosphorus
 Deliverable Number: D6.1
 Date of Issue: 31/01/07
 EC Contract No.: 034115
 Document Code: Phosphorus-WP6-D6.1



Test-bed design

interface (Grid enabled I-NNI). The communication between Edge VG2 and core VS1 will be based on a specific interface (Grid enabled UNI). These interfaces will be developed by WP2.

The control plane network will also deploy a router (vendor, type to be decided) which will connect the control plane network with the internet in order to allow communication with control/provisioning plane of the other local test-beds. These communications will be based on a specific interface (E-NNI) which will be developed and provided by WP2. The router will support IP-SEC for secure communications with other local testbeds. As an alternative solution UEssex can also deploy software based VPN instead of a router with IP-SEC. this solution will be based on PCs running VPN software which allows using a common flat control plane network and also covers different sites.

The PHOSPHORUS testbed doesn't require for UEssex local testbed to have NRPS provisioning plane. However, in case the testbed scenario will require having NRPS in the UEssex local testbed it is possible to add NRPS to this testbed

1.6.5 Middleware and applications in the local testbed

All the test clients, the computing and the storage resources in the local testbed will run Linux operating system. Furthermore, a test client is going to run Globus Toolkit and Grid FTP client modules. This will be used by DDSS application test. For other applications which UEssex is participating, in the first phase of the project only a test client with Linux operating system will be used. Furthermore UEssex will liaise with the application test organisers to install and deploy the required applications and softwares on the clients. In the second phase, when the UEssex Grid testbed is fully available with the computing resources, the UNICORE middle ware will be installed in these nodes since it is necessary for KoDaVis application.

1.6.6 Interconnection with other local testbed

The UEssex local testbed will have direct connection to PSNC local testbed and also Surfnets network. Each of the connection will be a single wavelength which carries VLAN Gigabit Ethernet traffic. In the first phase of the PHOSPHORUS project testbed, UKERNA will terminate the Gigabit Ethernet links from PSNC and Surfnets at a Gigabit Ethernet switch in the international access point of UKLighth and switch the traffic from these testbeds to UEssex through a single Gigabit Ethernet connection. In the second phase of the project, UKERNA will provide two Gigabit Ethernet links connectivity between the international access point of UKLighth and UEssex (i.e one link for PSNC and One for Surfnets) which will be through two wavelengths in a single fibre. At the entrance of the UEssex testbed, these wavelengths will be demultiplexed into two separated optical link (fibre) and also at the exit point of the UEssex testbed these wavelengths will multiplexed into a single fibre (see Figure 1.15).

According to the needs of the testbed scenario the Gigabit Ethernet links will be terminated on VLAN Gigabit Ethernet switch at the ingress/egress point of the UESSEX local testbed and individual VLAN tagged traffic tributaries (incoming/outgoing traffic from different local testbeds) can be routed to a suitable part of the local testbed.

Project:	Phosphorus
Deliverable Number:	D6.1
Date of Issue:	31/01/07
EC Contract No.:	034115
Document Code:	Phosphorus-WP6-D6.1



Test-bed design

The UEssex will need a connection to PSNC testbed for DDSS application and a connection to VIOLA testbed for KoDaVis application. The connection with PSNC testbed will be direct connection and the connection to VIOLA testbed will indirect through PSNC testbed/HUB (see global testbed description).

1.6.7 Integration with other local testbeds

At control plane and bandwidth provisioning layer the UEssex local testbed will be connected to two types of the network domains: GMPLS/ G²MPLS domain (PSNC) and NRPS domain (Surfnet). The GMPLS/ G²MPLS Control plane in the UEssex local testbed will use E-NNI specific interfaces to communicate to other GMPLS/ G²MPLS domain and this interface is going to be developed within WP2. GMPLS/G²MPLS control plane will deploy a specific interface solution for interconnection with NRPS domain. This interface solution will be developed jointly by WP2 and WP1.

At the Grid middleware layer, for DSSS application test Globus Toolkit will be deployed and for KoDaVis application test UNICORE will be deployed (in second phase when the local Grid testbed installation is completed).

At the application layer, integration with DDSS, KoDaVis and INCA will be required. For DSSS application a Grid FTP client will be deployed. For KoDaVis application a test only client will be deployed in first phase and a data server in second phase. For INCA application a test only client will be deployed.

1.6.8 Operation in local testbed

The UEssex local testbed will participate in three application testes:

- In DSSS application as a Grid FTP client;
- In KoDaVis application as a test only client in the first phase and a data server will be added in second phase;
- For INCA application as a test only client.

As it was mentioned earlier the UEssex local Grid testbed comprising a CPU cluster and storage resources will be fully operational by the second phase of the project. Upon availability of this testbed UEssex will liaise with other WP3 partners and application providers to participate in computing and storage intensive test applications such as TOPS and WISDOM

Project:	Phosphorus
Deliverable Number:	D6.1
Date of Issue:	31/01/07
EC Contract No.:	034115
Document Code:	Phosphorus-WP6-D6.1



1.7 VIOLA testbed

The local VIOLA testbed is going to be used for TOPS, DDSS, KoDaVis, and WISDOM applications tests, as defined by Work Package 3. The VIOLA testbed will contain switching, computing, visualization, and storage resources. The switching layer will be built from SDH cross connects and Gigabit Ethernet switches, computing resources will include a Cray X-D1 Opteron cluster and an Intel Xeon cluster. The main visualization resource will be the i-CONE™ display system. It consists of four Barco 909 projectors and provides a resolution of up to 2000x1460 pixels at 96 Hz. The pixels will be sent over the network to a PC with a gigabit network card per projector. An STK tape library will serve as the main storage resource.

The VIOLA local testbed will be controlled by the NRPS ARGON with GMPLS interface and will be integrated with other Phosphorus local testbeds with use of the proper NRPS and G²MPLS interfaces. The GMPLS/G²MPLS control plane will allow for automatic provisioning of optical paths based on SDH connections.

1.7.1 GRID resources in the local testbed

The main compute resources in the VIOLA testbed are a Cray X-D1 system at FZJ (14 dual Opteron compute nodes, each of them attached to the testbed with a 1 GbE adapter) and an Intel Xeon cluster named PACK at FhG-SCAI (17 dual Xeon compute nodes, each of them attached to the testbed with a 1 GbE adapter). These clusters will provide the compute capacity for the WISDOM and KoDaVis applications. They will also serve as client nodes for the DDSS application.

At FZJ, access to a TSM server with an STK tape library will be granted for the DDSS application (backup/archive).

At FhG-IAIS we will use the i-CONE™ display system for TOPS application. The i-CONE™ projects virtual environments on a curved, 240 degree wide-angle screen. The screen is slightly tilted backwards to improve visual and acoustic qualities. This technology eliminates the edges and corners of a CAVE™ – the cubic virtual theatre widely used today. The i-CONE™ is driven in active stereo mode by a sgi Onyx or a PC Cluster (HP xw8200, Quadro FX3000G) with a total resolution of up to 6200x1460 pixels. The i-CONE™ consists of four channels with a resolution of 1440x1320 pixel or 1600x1460 pixel at 94 Hz per BARCO 909 projector. For the initial test a single HP-workstation will be used with one projector. This PC is connected with its Gigabit network card to the VIOLA testbed.

1.7.2 Switching and transmission resources in the local testbed and their connections

Most of the connections inside the local testbed as well as all links to the other Phosphorus local testbeds will use Gigabit Ethernet as the transmission technology. This way the local testbed will use the same transmission

Project:	Phosphorus
Deliverable Number:	D6.1
Date of Issue:	31/01/07
EC Contract No.:	034115
Document Code:	Phosphorus-WP6-D6.1



Test-bed design

technologies as most modern GRIDs. Gigabit Ethernet transmissions will be switched in the testbed in two layers – on layer one by SDH cross connects and on layer two by Gigabit Ethernet switches using MPLS.

The core of the local testbed will be built from three SDH cross connects, which will switch SDH connections by using GMPLS. Within this GMPLS domain G²MPLS developed in WP2 will be evaluated and tested

The SDH cross connects are ALCATEL 1678 MCC systems, which are currently part of the VIOLA testbed. Within GÉANT2 Alcatel 1678 MCC systems are also used for providing Ethernet services. The cross connects are connected via 10 Gbps SDH links, building a triangle. Each of the SDH cross connects will have at least 12 Gigabit Ethernet ports and will be able to switch each GbE port to any other GbE port using SDH connections. On these connections different bandwidths can be used (multiples of 155 Mbps, 1 VC4 container can carry 155 Mbps). The Alcatel cross connects can be connected to OIF UNI 2.0 clients, which can be used for creating user initiated end-to-end connections. Three Proxy UNI 2.0 clients from Alcatel are available within the testbed. They can be used for initiating GbE connections between the Ethernet-Switches using SDH transport.

The SDH cross connects will be accompanied by three Riverstone 15008 Gigabit Ethernet switches between which services based on MPLS tunnels can be defined. The Ethernet switches build the MPLS domain and add an additional layer to the heterogeneous Phosphorus network. The connections between the Ethernet switches themselves are 10 GbE links. The connections between the MPLS domain and the GMPLS domain will be GbE connections, which can be aggregated.

Both SDH cross connects and Gigabit Ethernet switches will be controlled by the NRPS ARGON. The SDH cross connects can additionally be controlled by using the GMPLS/G²MPLS interfaces.

The Riverstone Gigabit Ethernet switches will aggregate traffic from GRID resources which will be connected to them. Having Gigabit Ethernet switches as an aggregation layer between GRID resources and the SDH network will make the testbed similar to many distributed GRIDs, which are built from several local clusters which are interconnected with use of MPLS and GMPLS networks. Having multiple connections within the MPLS and GMPLS domain will allow for simultaneous parallel transmissions in the testbed. This feature will increase the flexibility of serving different test scenarios, including failures of a single connection and switching connections to backup routes.

The Alcatel 7750 Ethernet switches are used for connecting the VIOLA testbed to the other PHOSPHORUS testbeds (see chapter “Interconnections with other local testbeds”).

All Gigabit Ethernet Switches, SDH cross connects and Proxy UNI systems are connected to a common control network (DCN network). This network is used for local management and configuration of the network devices.

The topology of the testbed as well as all the testbed resources is shown in Figure 1.16.

Project:	Phosphorus
Deliverable Number:	D6.1
Date of Issue:	31/01/07
EC Contract No.:	034115
Document Code:	Phosphorus-WP6-D6.1

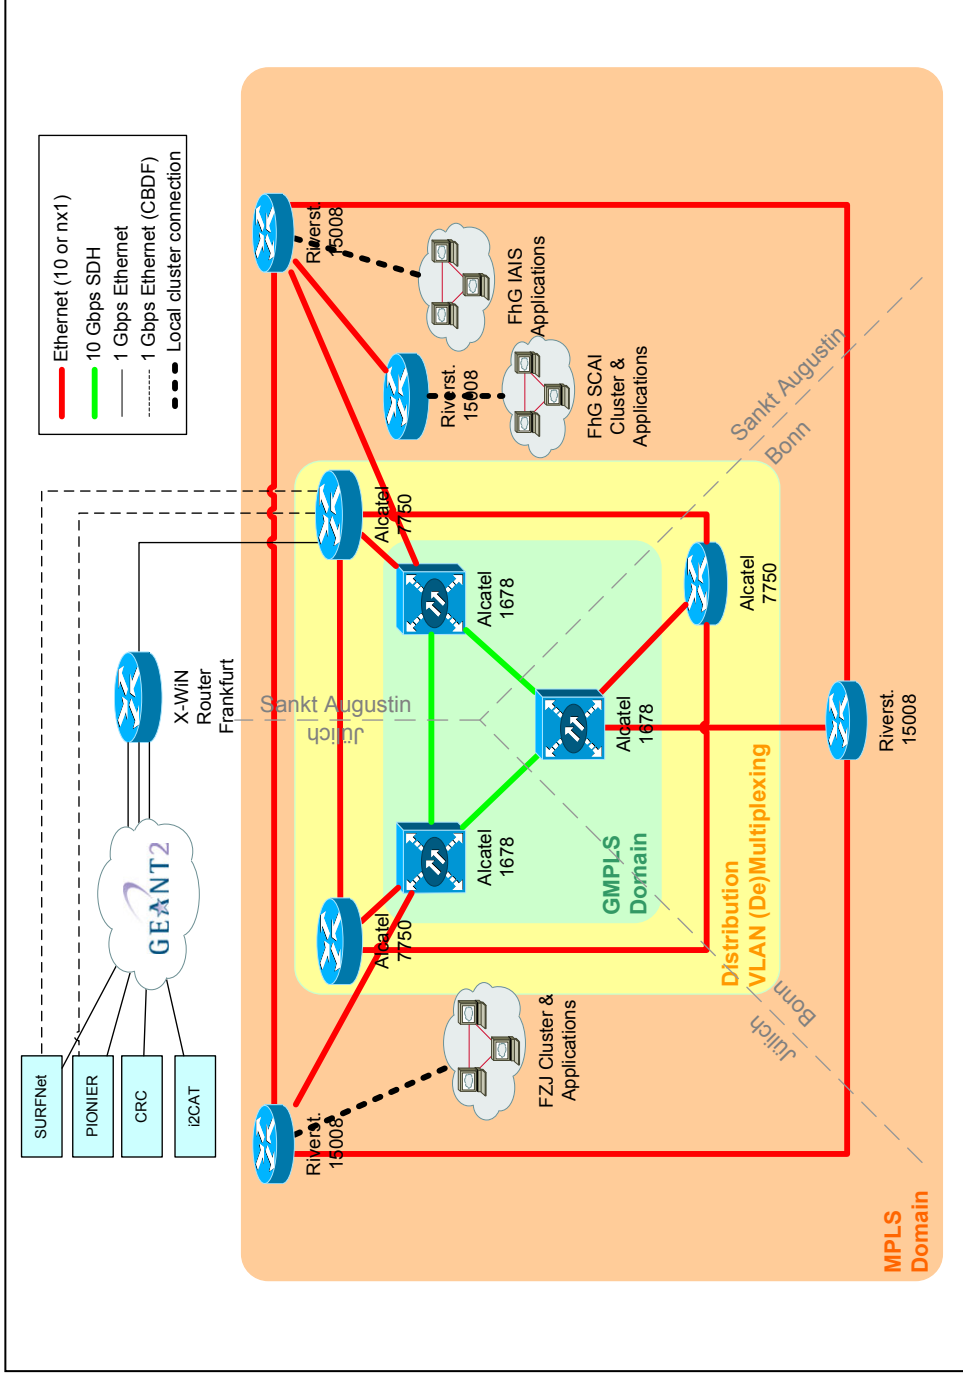


Figure 1.16: Topology of the VIOLA local testbed



1.7.3 Control and provisioning plane in the local testbed

The SDH cross connects will be controlled by GMPLS (in the first phase of the project) and G²MPLS (in the second stage). For GMPLS these switches have a built-in controller. For G²MPLS they will have an external controller made of a personal computer running a G²MPLS implementation. The GMPLS/G²MPLS controllers and the UNI proxies are connected to a common Ethernet network for exchanging control plane traffic (out-of-band signalling). Flat structure of the network will allow for communication between each pair of controllers. The control plane network will also interconnect GRID resources (computing nodes) with the controllers in order to allow for the exchange of control information between the GRID resources and the control plane of the network.

The control-plane network will be implemented as a layer 2 segment to which all the controllers will be connected by Fast or Gigabit Ethernet links. The control plane network will also include a router (Cisco 1712) which will connect the control-plane network with the Internet in order to allow for communication with control/provisioning plane of other local testbeds. The router will support packet filter to protect the control-plane network from unauthorised access from the Internet and IP-SEC for secure communication with other local testbeds. Instead of using a router for separating the control plane at different locations a VPN solution based on the publicly available free software TINC can be used, too. This software also allows to use a common flat control plane network, which covers different sites. Such a common network will allow capturing all control plane packets of the whole testbed at each side.

1.7.4 Middleware and applications in the local testbed

The normal VIOLA middleware stack is installed on the X-D1 and PACK clusters. It contains UNICORE 5 (which will be upgraded to UNICORE 6 during the course of the project), and an RMS with advance reservation capabilities (PBS Pro on the X-D1, and EASY on PACK). Access to network reservation services (ARGON) is provided via the MetaScheduling Service (MSS) located at SCAI. Distributed MPI applications are supported by the MetaMPICH library and a specific MetaMPICH UNICORE client-plugin. This middleware is required by the WISDOM and KoDaVis applications.

The Globus Toolkit will be installed on the X-D1 and PACK for use by the GridFTP part of the DDSS application (GridFTP).

A TSM (IBM Tivoli Storage manager) client is already installed on the X-D1 for use by the backup/archive part of the DSS application.



1.7.5 Interconnections with other local testbeds

As detailed in the next section of this document, the VIOLA local testbed will have direct connections to several other Phosphorus local testbeds, namely SURFnet, i2CAT, CRC and PSNC. Each of the connections will be implemented as a logical light path, which will carry Gigabit Ethernet traffic.

Following the needs of the test scenarios the links will be terminated in one of the Alcatel 7750 Gigabit Ethernet switches. These switches are able to de/multiplex VLANs and they are interconnected to each other to distribute the VLANs to the SDH cross connects which are located at different locations within the VIOLA testbed. The VLANs on the link to GÉANT have to be demultiplexed in Frankfurt where the German access point for the GÉANT network is located. Demultiplexing will be done by one of the X-WiN Routers by using its Ethernet capabilities.

Having the links terminated on a Gigabit Ethernet switch will allow for transmitting multiple VLANs for simulating light paths on a single physical link. Each of the VLANs on a single physical link can be mapped to a port on one of the three Alcatel 7750 Gigabit Ethernet switches going to an SDH cross connect and then be switched separately in the domain of the local testbed. Having VLANs will also allow for provisioning transit VLANs – VLANs between two local testbeds which are not connected directly. Each of such VLANs can be seen by the applications and GRID middleware as a direct link increasing the flexibility of the topology of the global testbed.

Currently the exact topology for the connections to the other testbed has not yet been determined. There are ongoing discussions on this issue with SURFNET and DFN.

1.7.6 Integration with other local testbeds

To ensure smooth interactions between the VIOLA local testbed and other Phosphorus local testbeds a precise integration is necessary. The integration must cover three layers: control plane to ensure provisioning of lightpaths between local testbeds, GRID middleware to allow for distributed computing and use of distributed storage and applications to have the same application using resources in different local testbeds.

The GMPLS/G²MPLS control plane of the VIOLA local testbed will be integrated with the control planes of other Phosphorus local testbeds which also use GMPLS/G²MPLS. The inter-domain interface used for integration of the control plane will allow for exchanging information about each local testbed and provisioning optical paths between two or more local testbeds. In the first stage of the project the legacy GMPLS NNI interface can be used while the G²MPLS inter-domain interface will be defined by work package 2 and then deployed in the local testbed.

On the GRID middleware layer the integration with other local testbeds will include:

- For UNICORE:
 - inclusion of the UNICORE installation at PSNC and potentially other sites into the VIOLA UNICORE Grid



Test-bed design

- For Globus Toolkit:
 - installation of version 4 of the Globus Toolkit,
- For both Grid middlewares:
 - agreements with PSNC and potentially other sites about Certification Authorities (CAs), exchange of certificates and mapping of the distinguished names of certificates to local user accounts. Using EUGridPMA acknowledged CAs would be the natural choice.

On the application layer integration with other local testbeds will require:

- DDSS GridFTP: configuration of GridFTP servers and clients at FZJ and FhG-SCAI.
- DDSS B/A: configuration of FZJ TSM clients for access to PSNC's TSM server
- KoDaVis: configuration of the application data-server and test-only clients
- WISDOM: configuration of the application workflows
- TOPS: configuration of the visualization client. The connection of the video signal of the visualisation client to the i-CONE™ display will be done by local switching procedures, the projectors have to be switched on by hand. While the TOPS application is running the researcher at the display site will view the pictures.

FZJ will provide documentation and support for the installation of the UNICORE middleware at PSNC, UESSEX and potentially other sites that will participate in the WISDOM or KoDaVis application tests.

1.7.7 Operations in the local testbed

All of the four applications defined in Work package 3 will access the VIOLA testbed: WISODM, KoDaVis, TOPS, and DDSS (both GridFTP and TSM). These applications will use the Phosphorus services mainly through the Grid middleware. This middleware is provided by Work package 3 and will access these services either directly (e.g. the compute resources via UNICORE, the MSS and the local RMS) or through the interfaces to Work package 1, which are currently being defined. At least in the first phase of the project, the MSS uses the ARGON system to access the network services provided by the Phosphorus testbed.

The distribution of the applications over the testbed resources is shown in table 1.2. The use cases which will be deployed and tested in this local testbed have been described by Work package 3 in the Deliverable D3.1 "Use-cases, requirements and design of changes and extensions of the applications and middleware". More details can be found in that document.

Project:	Phosphorus
Deliverable Number:	D6.1
Date of Issue:	31/01/07
EC Contract No.:	034115
Document Code:	Phosphorus-WP6-D6.1



Test-bed design

Resource	WISDOM	KoDaVis	TOPS	DDSS GridFTP	DDSS B/A
Hardware resources					
Cray X-D1	yes	yes		yes	yes
PACK Cluster	yes			yes	yes
Visualization system (Display, PC graphics cluster)			yes		
STK tape library					yes
Grid middleware					
Globus Toolkit middleware				yes	
Unicore middleware	yes	yes			
Other software					
FlexX/Autodock FLEXIm, Oracle	yes				
C/C++ compiler, MPI	yes	yes	yes		
Other standard Linux packages			yes		
TSM					yes

Table 1.2: Application distribution over VIOLA local testbed

Forschungszentrum Jülich: Cray X-D1

# Nodes	14 x 2-CPU for VIOLA and Phosphorus
CPU	Opteron, 2.2 GHz
Memory per node	4 GB per Node (2 CPU)
Home-Data	1.1 TB (global via NFS)
Temp-Data	8 GB per node local (/tmp)

Project: Phosphorus
 Deliverable Number: D6.1
 Date of Issue: 31/01/07
 EC Contract No.: 034115
 Document Code: Phosphorus-WP6-D6.1



Test-bed design

Fraunhofer Institute SCAI: PACK Cluster

# Nodes	17 x 2-CPU SMP + Login/Service-nodes (2 x SMP) +File/server-nodes(2 x SMP)
CPU	3,0 Ghz
Memory / node	2 GB ECC-RAM
Home-Data	320 GB
Temp-Data	80 GB

# Nodes	6
CPU	2 x Ultrasparc-IIIli 1,1 Ghz SUN Fire V240
Memory per node	2 GB RAM
Temp-Data	35 GB per node local (/space)

Project:	Phosphorus
Deliverable Number:	D6.1
Date of Issue:	31/01/07
EC Contract No.:	034115
Document Code:	Phosphorus-WP6-D6.1



1.8 NORTEL testbed

The Nortel testbed will be dedicated to developments in Work Package 1. It is built from Nortel OME6500 MSSP (multi-service provisioning platforms) nodes and Metro ESU 1850 Ethernet switches that are controlled by DRAC. This testbed will also have access to client nodes to do proof-of-principle application testing, as well as a network element simulation testbed for modelling large network deployments.

Additional computational resources and video client/servers from the Nortel NTAC lab can be added to the testbed for limited periods of time as needed.

1.8.1 GRID resources in the local testbed

The following PC resources are currently connected to the testbed for DRAC client testing. Additional resource can be provided if required:

- 2 x Dell PowerEdge 1800, 2 x Xeon 2.8GHz, 2 GB RAM, 2 x 80 GB RAID HDD
- 2 x Pentium 4 2.2GHz, 1 GB RAM, 20 GB HDD

1.8.2 Switching and transmission resources in the local testbed

Four Nortel OME 6500s interconnected via OC-192 links will be used in support of the testbed. A single OME 6500 in the Nortel NTAC lab will act as a gateway into CAnet 4. Edge connectivity in the testbed is achieved via GbE links directly out of the OME 6500, with VLAN assignment and Ethernet grooming performed by two interconnected Metro ESU 1850s. It is possible that later in 2007 a pair of MERS 8600 supporting PBB/PBT will be connected to the testbed to provide advance Ethernet services.

A subset of the bandwidth on this configuration will be dedicated to the Phosphorus testbed as multiple DRAC instances will be using the network. Additional resources can be added to the test bed as needed. Currently the client interfaces are GbE only, however, the EPL card can also support 10/100/1000 BaseT. In addition, SONET/SDH client interface cards could be added to the configuration if needed.

The Nortel testbed also has access to a simulated environment of OME 6500 and HDX/HDX-c network elements for the testing of large scale networks. Complex topologies can be modeled using the same behaviors real network elements provide. Although real traffic cannot be carried across the simulated network, manual topologies can be entered to simulate connections into the real testbed. Currently, the DRAC testbed is simulating a 60 node SURFnet network.

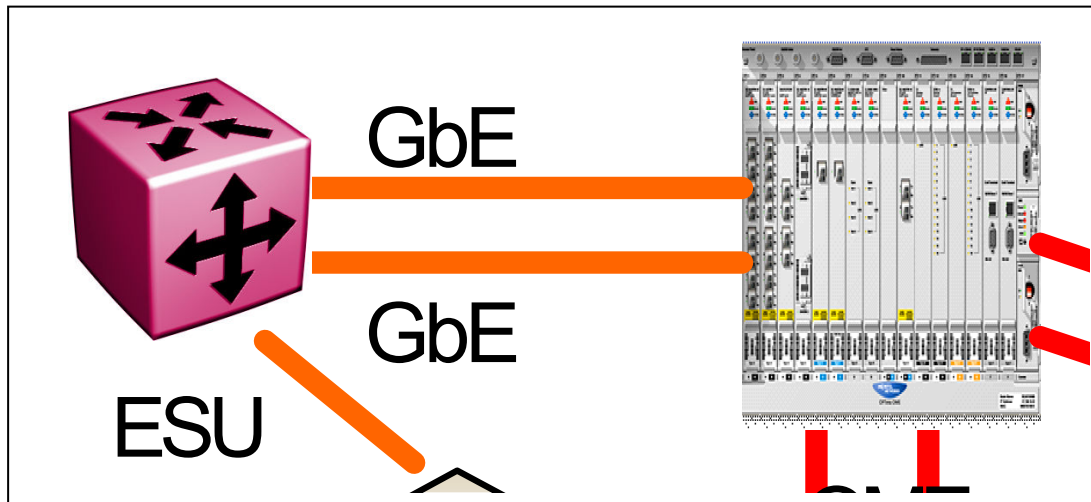


Figure 1.17: NORTEL local testbed

1.8.3 Topology of connections inside the local testbed

A single OME 6500 in the Nortel NTAC lab will act as a gateway into CANet 4. This gateway network element will then be connected via an OC-192 link into a three node OME 6500 network. Each of the three OME 6500s are interconnected via two OC-192 links, and each OME is populated with a single four port EPL card holding two pluggable GbE ports (of the four available). These GbE interfaces can be connected into other equipment if needed (test sets, PC, etc). Two Metro ESU 1850 Ethernet switches will be interconnected to the OME 6500 by GbE interfaces to provide VLAN tagging of data flows into the optical network.

Connectivity between the OME 6500 nodes and the simulated nodes is achieved via manually provisioned virtual links. DRAC support the building of arbitrarily large and complex networks using the simulated OME and HDX/HDX-c network elements.

1.8.4 Control and provisioning plane in the local testbed

The nodes in the testbed are controlled by DRAC. From the DRAC server a direct connection to the DCN of the testbed exists. Setting up services on the testbed will be done through the public interface of the DRAC server. This will be accessible through either a “real” public IP address, or through a VPN server. With access to the DRAC server a web page is reachable that allows for manual provisioning, and web services for automated provisioning. The Nortel testbed will support authentication through either A-Select or local account provisioning of users. For Web Services certificates that allow specific use shall be provided. It is assumed a generic Phosphorous user account will be created for use but the Phosphorous testbed. Permitted resources will be provisioned against this account.



Test-bed design

1.8.5 Middleware and applications in the local testbed

No middleware and applications are foreseen in the Nortel testbed. Requests for deployment of specific applications will be considered.

1.8.6 Interconnections with other local testbeds

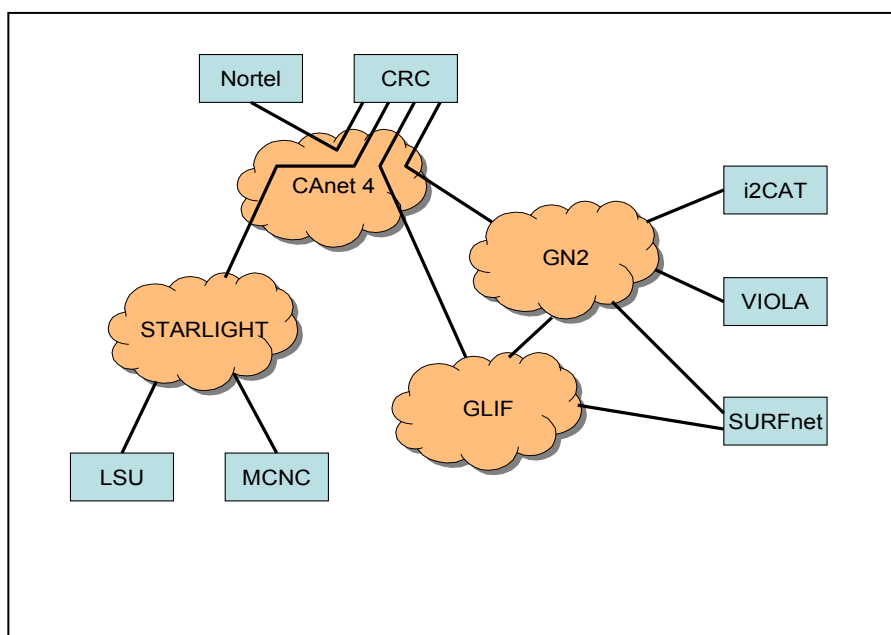


Figure 1.18: Connection to other local testbeds

Nortel will use one 10 Gbps (OC-192) link from the Nortel NTAC lab to the CAnet 4 Ottawa POP, where an STS-24c worth of bandwidth will be mapped through to a Gigabit Ethernet (GbE) service into the Broadband Applications and Demonstration Laboratory (BADLABTM) at the CRC Carling campus.



1.9 MCNC testbed

The existing Enlightened Computing testbed is a national-footprint optical network that has been deployed to facilitate the middleware and application goals set forth by the project, as well as to support investigations into new network and control plane architectural models. The core of the testbed is built using Calient Networks Diamond Wave optical cross-connect (OXC) switches interconnected by Ten Gigabit Ethernet (10 GbE) circuits provided by Cisco Systems and NLR.

GMPLS is used as the control plane protocol to allow dynamic instantiation of end-to-end paths across the testbed, in order to interconnect computing and storage resources as required by applications. Each of the four original sites that host an OXC also provides 10 GbE-attached switches and/or end hosts as well as the potential to extend connectivity to users and resources via local Regional Optical Networks (RONs). Individual institutions within the team administer their local nodes and help with planning and day-to-day maintenance of the complete cooperatively managed domain.

1.9.1 GRID resources in the local testbed

Below is a figure of the Enlightened Resource testbed, to include compute resources off of each OXC. Access to those resources can be discussed with the individual owners of the compute resources.

Test-bed design

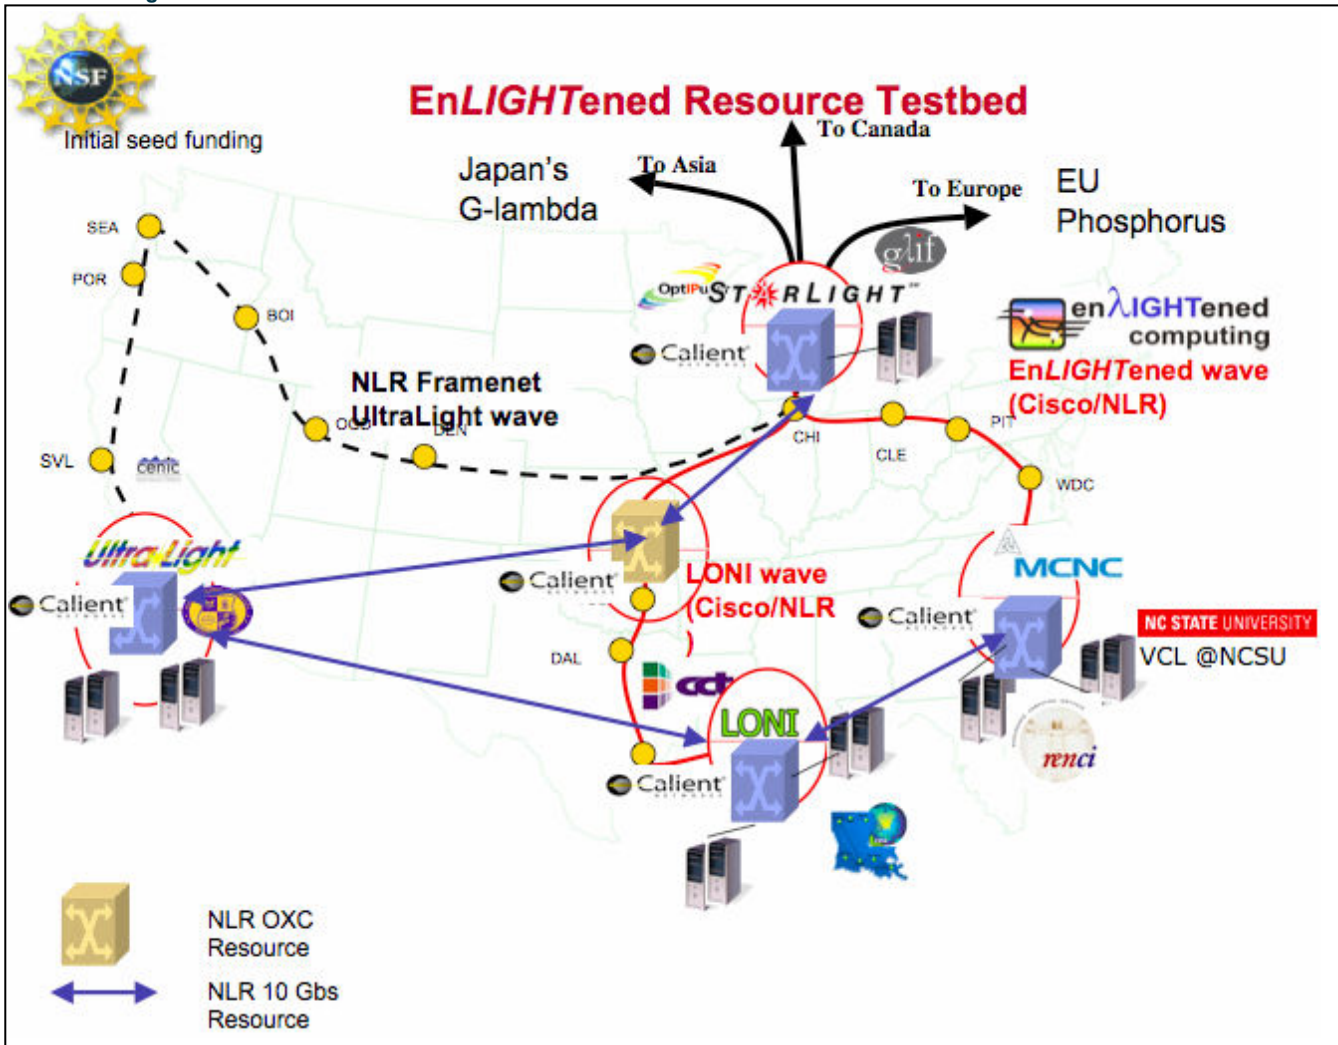


Figure 1.19: Topology of the Enlightened Resource testbed

1.9.2 Switching and transmission resources in the Enlightened testbed

All internal connections within the local testbed are realized with 10 Gigabit Ethernet circuits from National Lambda Rail (NLR). Two circuit types are currently deployed: layer 1 circuits (NLR WaveNet service) and layer 2 circuits (NLR FrameNet service). The layer 1 handoff to the Enlightened testbed is directly from the NLR DWDM equipment at each site. The layer 2 handoff is from NLR Ethernet switches at each site.

The core of the Enlightened testbed is built with optical cross-connect switches (OXC), interconnected by the 10 Gigabit Ethernet circuits described above. There are currently four OXCs deployed in the following locations by Enlightened partners: Chicago, Illinois (Starlight), Los Angeles, California (Ultralight), Baton Rouge, Louisiana (Louisiana Optical Network Initiative - LONI) and Raleigh, North Carolina (MCNC).

Project:	Phosphorus
Deliverable Number:	D6.1
Date of Issue:	31/01/07
EC Contract No.:	034115
Document Code:	Phosphorus-WP6-D6.1



Test-bed design

Since all direct connections to the OXCs are using 10 Gigabit Ethernet interfaces, the switching granularity of the testbed is 10 Gigabit Ethernet. At the Starlight location, there is an Ethernet switch managed by LONI that can be used to aggregate 10 Gigabit Ethernet circuits as needed.

In most cases, Ethernet switches are attached to tributary ports on the OXCs forming the edge of the testbed. Computing and storage resources are then connected to the edge switch via Gigabit Ethernet of 10 Gigabit Ethernet connections. In some cases, computing and storage resources are attached directly to tributary ports on the OXC.

Currently all Ethernet ports that connect directly to the OXC are configured for untagged framing. This allows for any to any connectivity between tributary ports, since trunk configuration don't have to be matched. It also eliminates the need to negotiate VLAN tag assignments between sites. The limitation of this method is that trunks between sites are not supported.

End-to-end provisioning between sites is accomplished via a GMPLS control plane initiated from the Grid middleware. The GMPLS control plane only establishes lightpaths across the Calient optical switches.

1.9.3 Topology of connections inside the Enlightened testbed

The current topology of the testbed is a hub and spoke, with the hub in Chicago at Starlight. One circuit connects each site to the hub OXC.

The network team is currently working on extending the core topology of the testbed through a close collaboration with NLR. A new OXC owned by NLR is being deployed in Kansas City, along with additional circuits between key sites. The OXC and circuits will be full integrated into the GMPLS-controlled domain and managed cooperatively with the other OXCs as an equal partner. Figure 1.20 illustrates the new extended topology.

NLR is intending to use the combined testbed to pilot a resilient wave service that it would like to deploy as an alternative to its' existing unprotected wave service. Combining the testbeds in this way allows both projects to benefit from additional capacity and complexity of the combined facility. For example, the Enlightened team can design new features in the middleware that allow the use of alternate paths to a single destination, as a means of providing path recovery from that layer. Additionally, the team could investigate how the middleware should react to a network layer recovery scenario, much like the resiliency function that NLR is planning to investigate.

Project:	Phosphorus
Deliverable Number:	D6.1
Date of Issue:	31/01/07
EC Contract No.:	034115
Document Code:	Phosphorus-WP6-D6.1

Test-bed design

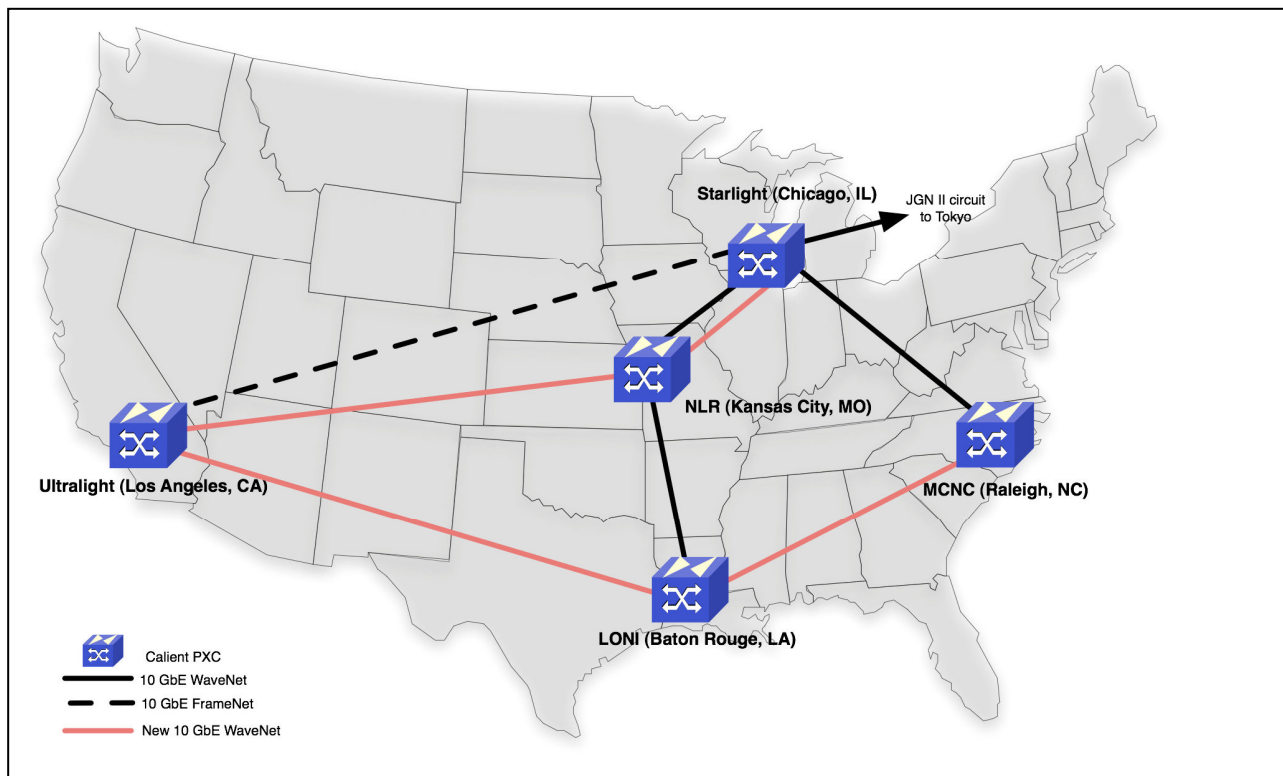


Figure 1.20: Topology of the Enlightened testbed

1.9.4 Control and provisioning plane in the Enlightened testbed

The optical switches in the local testbed use GMPLS for the control plane. The Ethernet switches at the network edge are currently configured statically. Using PC-based GMPLS controllers (Dragon's VLSRs) to instantiate dynamic VLANs on the Ethernet switches has been discussed, but there are no firm plans for implementation at this time.

The control plane transport between OXCs on the Enlightened testbed is currently provided by normal NREN IP service paths. To better protect and observe the traffic on this vital connection, MCNC has proposed a two-phase project to use an NLR-provided FrameNet (Layer 2) circuit to build the control plane transport. The first phase entails connecting the OXCs together via the multi-point FrameNet connection, and has been approved by the NLR Network Research Council. Phase one is being deployed as part of the testbed expansion described above.

The second phase of the proposed project will add the capability to support multiple control plane projects (via VLAN segmentation) and capture control plane traffic to aid in debugging and protocol development. The phase 2 proposal is currently in process, and MCNC is hopeful that sufficient funding to move forward will be found in the next few months.



Test-bed design

Middleware provisions end-to-end paths by attaching to the management interface of the OXCs and using TL-1. Manual provisioning can also be accomplished via the Calient Web-based GUI. We currently manually provision the Ethernet switches attached to the OXCs. We are now investigating dynamic layer two provisioning via control plane or middleware.

Investigation into using G2MPLS on the Enlightened testbed is just beginning.

1.9.5 Middleware and applications in the Enlightened testbed

The Enlightened middleware is currently capable of instantiating a GMPLS request to setup dynamic lightpaths across the OXCs and coordinate with requested compute resources via HARC middleware. This is based currently on a static topology, however, work has begun on implementation of topology discovery via OSPF information captured by the OXC control plane.

1.9.6 Interconnections with other Enlightened testbeds

Based on the footprint of the Enlightened testbed, a natural interconnection point with other local testbeds is the Starlight facility in Chicago. Discussions about which trans-continental circuits might be used and details of the interconnections have just begun.

Currently there is an interconnection at Starlight between the Enlightened testbed and the JGNII GMPLS testbed in Japan, using a JGNII-owned 10 Gbps circuit that terminates in Tokyo. Both data plane and control plane connections use that circuit, segmented in different VLANs. The data plane connection terminates on the Enlightened OXC in Chicago with a 10 Gigabit Ethernet interface. The control interfaces of the OXCs are connected via Fast Ethernet connections to a small JGNII switch at Starlight.

The control plane connection between domains uses GMPLS E-NNI. We have successfully demonstrated interdomain provisioning across this interface.

A similar strategy for interconnection could be used with other Phosphorus local testbeds. Data plane connections via 10 Gigabit Ethernet interfaces could be easily accommodated directly on the OXC. Gigabit Ethernet interconnections could possibly be accommodated using the Ethernet aggregation switch mentioned above. Fast Ethernet control plane interconnections could possibly be hosted on the Ethernet switch owned by JGNII.

1.9.7 Integration with other local testbeds

To ensure smooth interactions between the Enlightened testbed and some of the Phosphorus local testbeds, precise integration is necessary. The integration must cover three layers: a control plane to ensure provisioning

Project:	Phosphorus
Deliverable Number:	D6.1
Date of Issue:	31/01/07
EC Contract No.:	034115
Document Code:	Phosphorus-WP6-D6.1



Test-bed design

of lightpaths between local testbeds, the GRID middleware to allow for distributed computing, and use of distributed storage and applications to have the same application using resources in different local testbeds.

The Enlightened testbed currently supports interdomain provisioning using GMPLS E-NNI, as stated above. Thus, interoperability with OXCs on other local testbeds using GMPLS E-NNI will require some detail coordination as a first step. Supporting G2MPLS directly on the Enlightened testbed is still under investigation.

Middleware and resource integration: requires further discussion between Phosphorus partners.

1.9.8 Operations in the local testbed

This section also requires further discussion between Phosphorus partners.

Project:	Phosphorus
Deliverable Number:	D6.1
Date of Issue:	31/01/07
EC Contract No.:	034115
Document Code:	Phosphorus-WP6-D6.1



1.10 CRC testbed

The CRC testbed comprises equipment located in the Broadband Applications and Demonstration Laboratory (BADLAB™) and the Optical Networking Laboratory (ONL) facilities and extends across CANet 4. The CANet 4 network elements assigned to the Phosphorus project will be managed and controlled via the UCLP NRPS.

The LightPaths (LPs) established across CANet 4 will terminate at the GÉANT2 POP in New York, the GLIF POP in New York and the STARLIGHT POP in Chicago enabling connections to other Phosphorus partner testbeds.

1.10.1 Networking, computing, storage and visualization resources in local testbed

The Table 1.3 shows the equipment available at CRC for the present project.

Networking	Computing & Storage	Visualization
<ul style="list-style-type: none"> • Cisco Catalyst 6509 (GbE, 10 GbE) • Cisco Catalyst 3750 • 8 Channel CWDM MUX/DEMUX • Cisco 12008 router (GbE, ATM) • Nortel OPTera Metro 5200 (GbE, DWDM) • 8 x 8 Photonic Switches • 16 Channel ROADM • 16 Channel DWDM MUX/DEMUX • EDFAs • VOAs • OCPM • BERT Tester • OSA 	<ul style="list-style-type: none"> • 3 x Pentium IV Dual Xeon @ 2.8GHz, 2.5 GB RAM, 2 x 40 GB HDD • 3 x Pentium IV Dual Xenon @ 1.7 GHz, 512 MB RAM, 1 or 2 40 – 120 GB HDD • 9 x Pentium IV 3.2 GHz, 512 MB – 1 GB RAM, 100 – 160 GB HDD 	<ul style="list-style-type: none"> • UltraGrid Uncompressed HD (Xena HD and DVC Centaurus capture cards) • Pleora Technology, Uncompressed SD • Pleora Technology, DVB-ASI • Pleora Technology, HDV • ISABEL (Clients (26), flow servers, license servers (24 users)) • H.323 (Polycom VSX 7000 with MCU (x3), Polycom PVX (x10)) • VideoLAN (HDV, DV) • Skype • Sony HVR-Z1U HD camcorder • JVC JY-HD10U HD camcorder • JVC GR-HD1U HD camcorder • JVC HM-DH5U recorder (x2)

Table 1.3: Equipment available at CRC for the project.



Test-bed design

1.10.2 Optical networking and physical layer resources in the local testbed

The BADLAB facility is connected to the CANet 4 Ottawa POPs via an 8-channel CWDM infrastructure. The core switch/router in the BADLAB is a Cisco Catalyst 6509. A Nortel OPTera Metro 5200 is also available as part of the Optical Networking Laboratory (ONL) adjacent to the BADLAB and supports GbE interfaces which can be connected to the Catalyst 6509 and/or PC's.

GbE is the common interface used on the CANet 4 network elements shown in Figure 1.21. The GbE ports are mapped onto STS-24c circuits (or lesser depending on availability and the bandwidth required by the applications). The networking equipment in the BADLAB is currently being upgraded to support 10 GbE with the addition of two switches from Force10 Networks as depicted in Figure 1.22. Once the upgrade has been completed, we will be in the position to support LPs based on circuits up to STS-192c. Within the BADLAB, the Cisco Catalyst 6509 provides standard switching and routing functionalities.

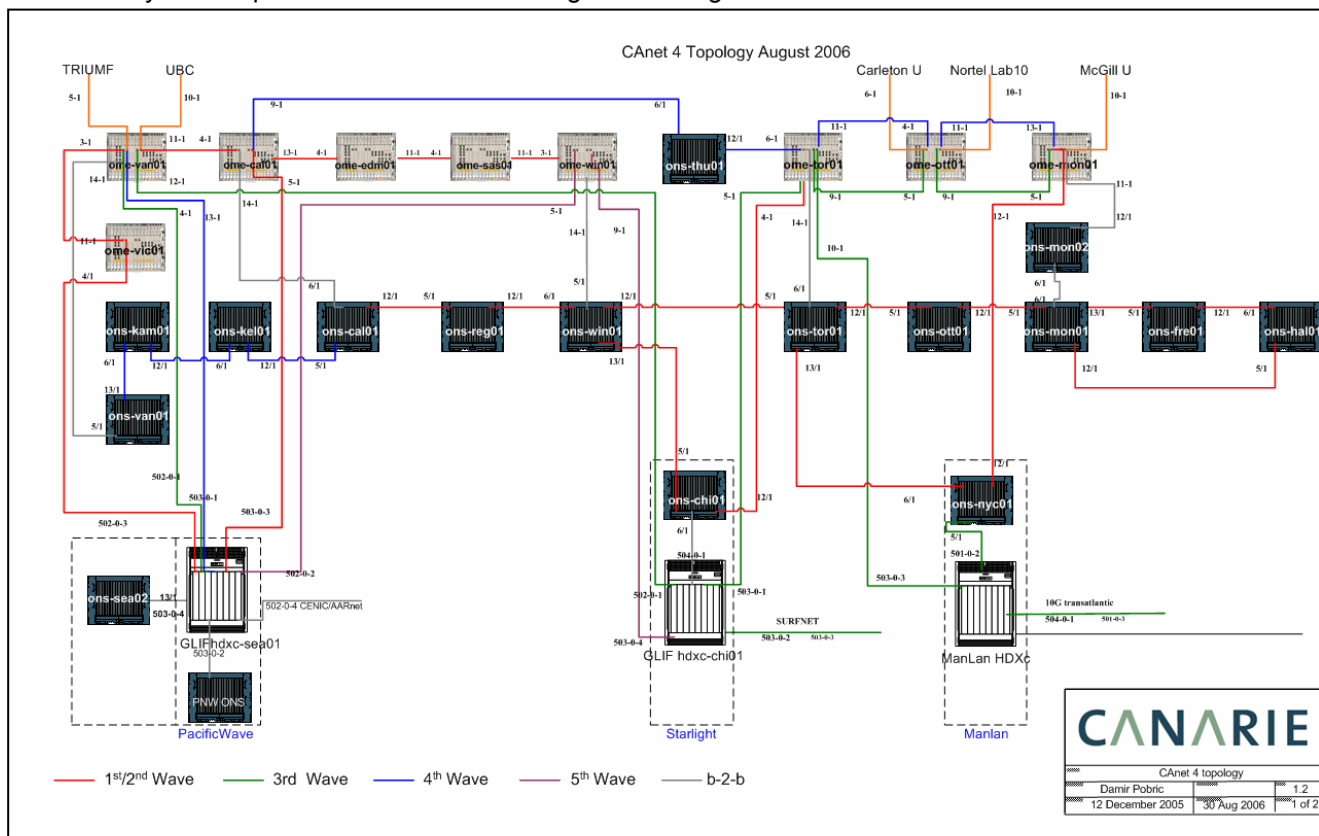


Figure 1.21: CANet 4 Infrastructure

Project: Phosphorus
 Deliverable Number: D6.1
 Date of Issue: 31/01/07
 EC Contract No.: 034115
 Document Code: Phosphorus-WP6-D6.1

Test-bed design

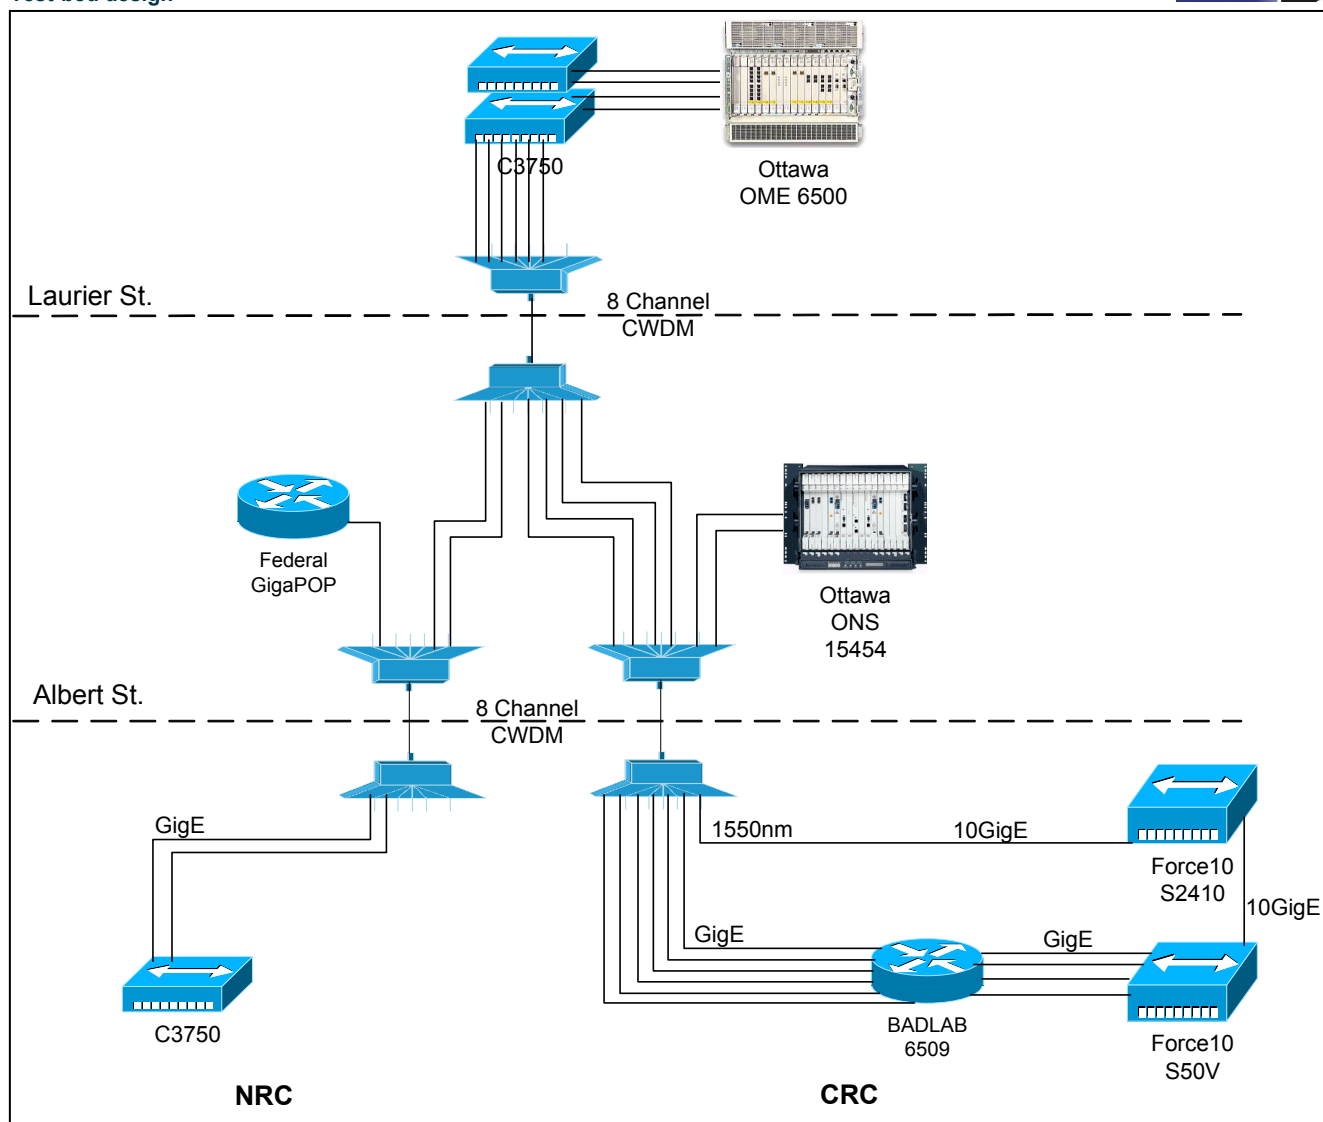


Figure 1.22: CRC's Local Testbed Connectivity Infrastructure

The Cisco Catalyst 6509 supported ports are: copper at 10/100/1000 Mbps Ethernet, 100BaseFX, 1000BaseT/SX/LX/ZX; the Force10 Networks S50V and S2410 supported ports are: copper at 10/100/1000 Mbps Ethernet, 1000BaseT/SX/LX/ZX, 10000BaseSR/LX/ZX; the Nortel Optera Metro 5200 supported ports: protocol and bit rate transparent with 1310 nm optics.

1.10.3 Topology of connections inside the local testbed

Computing resources connected to the local testbed are depicted in Figure 1.23. The port types have been described in the previous section.

Test-bed design

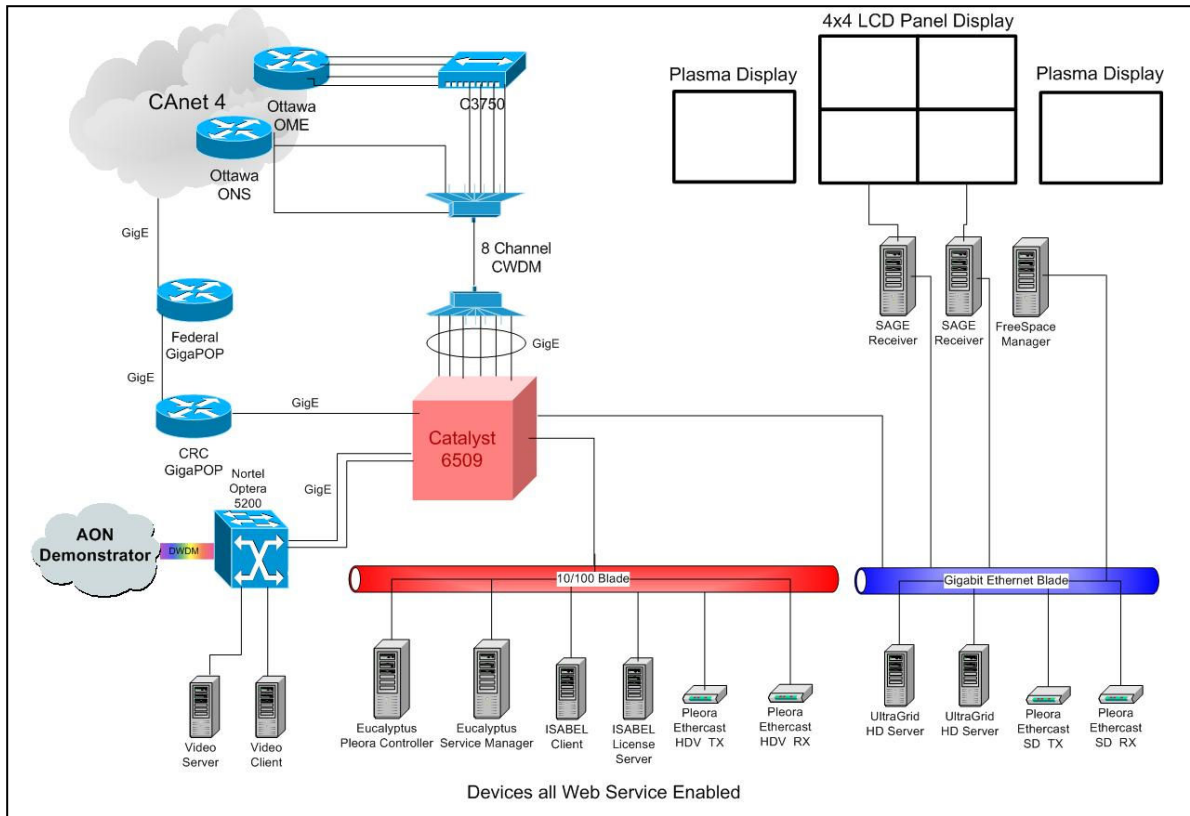


Figure 1.23: Detailed View of CRC's Local Testbed

The equipment associated with the AON Demonstrator shown in Figure 1.24 can be used to extend the local testbed to include an optical infrastructure with lambdas as the end-to-end connections or lightpaths.

Test-bed design

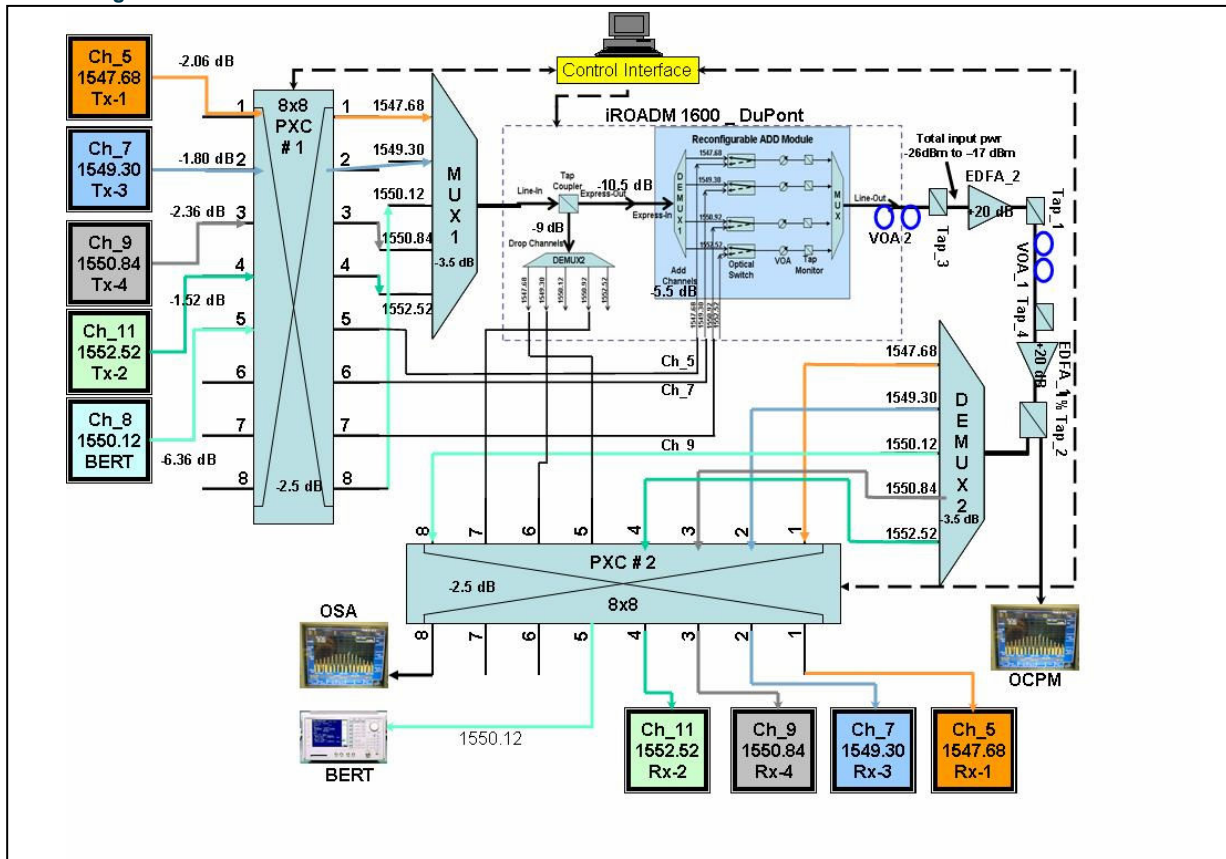


Figure 1.24: Detailed view of the AON Demonstrator at CRC

1.10.4 Control and provisioning plane in the local testbed

As an NRPS developer, CRC will deploy UCLP on the portion of the local testbed which extends across CANet 4. The network elements consist of the following platforms: Cisco ONS 15454, Nortel OME 6500 and Nortel HDXc. There is also the possibility of including the Nortel OPTera Metro 5200 platform in the ONL. A server running UCLP will be located in the ONL. This equipment will be managed via their TL1 interface.

Software based VPNs can be deployed over the Layer 3 service offered by CANet 4, GN2 and other NRENs for control plane connectivity between the local testbeds.

It is expected that software will be made available through WP-2 to implement a GMPLS (and later G²MPLS) based control plane running on a number of PCs in the BADLAB. It is assumed that this control plane will be used to control a subset of the network elements comprising the local testbed which extends across CANet 4. Initially, this would be done making use of networking resources assigned to the UCLP Virtual Lab on CANet 4 in order to perform the first round of tests and achieve some level of confidence with the control plane software.



1.10.5 Middleware and applications in the local testbed

Some of our computing resources that run the Linux operating system will be configured with the GT4 stack which will enable these resources to be accessed by Grid applications. Other resources will be virtualized as Web Services and made available for use by the Eucalyptus solution based on SOA and UCLP technologies which supports collaborative multi-site design for advanced network communities and will be tested with i2CAT. The Web Services include: the ISABEL Computer Supported Collaborative Work application, Pleora's Ethercast uncompressed SD video streaming, the UltraGrid uncompressed HD video streaming, rendering, visualization, etc. Since we have already deployed the Scalable Adaptive Graphics Environment (SAGbE) to drive a Tiled Display, the TOPS application developed by SARA could also be tested.

1.10.6 Interconnection with other local testbeds

The CRC local testbed will be interconnected to various partner testbeds using three international connections and a local connection as depicted in Figure 1.25. CRC will use 4 Gigabit Ethernet (GbE) connections from the Broadband Applications and Demonstration Laboratory (BADLAB™) to the CANet 4 Ottawa POP. These terminating points will be connected via LightPaths (LPs) to: the GN2 New York POP, the GLIF New York POP, the STARLIGHT in Chicago and Nortel in Ottawa.

There is an option to use VLANs over the LPs to provide richer connectivity and therefore a more complex demonstration.




Figure 1.25: CRC Testbed Topology Proposal

1.10.7 Integration with other local testbeds

At the control plane level, the CRC local testbed will be connected with all three types of NRPS: ARGON, DRAC and UCLP. All the NRPS-controlled domains previously registered with the Network Service Plane will be able to intercommunicate between them and create interdomain and intradomain connections through the corresponding NRPS.

At the Grid Middleware layer, for the GridON application test with i2CAT, the Globus Toolkit version 4 will be installed on a number of computing resources. At the application layer, integration with GridON is required.

For the Eucalyptus application test with i2CAT, equipment and software will have to be configured within the i2CAT local testbed in order to implement a collaborative multi-site design capability using SOA and UCLP technologies.



1.10.8 Operation in local testbed

CRC will participate in the GridON and Eucalyptus application tests with i2CAT. The local UCLP NRPS will be able to configure the LPs on the CANet 4 portion as well as interfacing with the software that implements the GMPLS (and later G²MPLS) based control plane running on a number of PCs in the BADLAB to control local networking equipment. There is also the possibility of using the Virtual Lab feature offered on CANet 4 in order to implement more complex network topologies.

1.11 IP addressing in the local testbeds

As in most cases no IP routing will take place between a pair of GRID resources (located either in one local testbed or in different local testbeds), the GRID resources (computing nodes, storage servers etc.) must have IP addresses from the same IP subnet. An IP subnet must be assigned to each of the connections.

The addressing scheme may be different in each test and will depend on the test scenario and the resources and network topology used in this test. Test leaders will be responsible for allocating addresses for the resources and connections which are used in the tests.



1.12 Project developments in the local testbeds

1.12.1 Workpackage 1 developments

Work package 1 developments outcome are: NRPS interoperability, which will allow different local testbed (using any of the 3 NRPS that are involved in the project: ARGON, DRAC and UCLPv2, or a GMPLS Control Plane during phase 1) to interconnect their experimental local testbeds and publish their available network resources to upper layers (as described in Deliverable D1.1). Each one of the NRPS systems running in a local testbed will be enhanced with the interfaces developed under the scope of WP1, which will allow a seamless network services interoperability through the connectivity provided by GN2 and other providers. Each NRPS system will control all the equipment located in its local testbeds under their domain. Once the Grid Middleware and the applications will be deployed in the selected local testbeds, the Service Layer being developed within WP1 will be also deployed in the testbed in order to integrate networking functionalities into the middleware, and therefore provide the capability to request end to end connections with advance reservation services to the selected applications.

WP1 will study and analyze interoperability requirements with JRA3 from GN2, however decisions on system adoption or interfaces development will be carried out depending on the status development of the IDM and DM from JRA3.

Moreover, the interfaces developed under the WP1 will allow the implementation of a set of use cases defined in the Deliverable D1.1, where a seamless interoperability, from the middleware point of view, will be demonstrated between the NRPSs and G²MPLS.

1.12.2 Workpackage 2 developments

One of the Phosphorus goals is to design, implement and test some extensions to the GMPLS Control Plane. For this purpose PSNC will build a testbed where G²MPLS prototype controllers will be deployed in three local testbeds: PSNS, UESSEX and VIOLA. Each controller will be a PC computer with G²MPLS implementation and management interface towards a data plane device (an optical switch or Gigabit Ethernet switch in the PSNC and UESSEX testbeds and Alcatel 1687 SDH cross connect in VIOLA). The main parts of the G²MPLS controller are G.RSVP-TE, G.OSPF-TE and G.LMP protocols. In the first phase the testbed will use the existing versions of these protocols, taken from the standard GMPLS. New versions of these protocols, proposed by Phosphorus with enhancements for GRIDs, will be deployed in the testbed in the second. The first of the three protocols will be used for the advance or on-time path creation, modification and termination. The Phosphorus extensions to this protocol will allow for requesting GRID resources. The second protocol will propagate information about GRID and transmission resources, network topology and GRID capabilities. The third protocol functionalities are related to service and neighbour discoveries between the Vsite (a grid resource) and the accessed network.

Project:	Phosphorus
Deliverable Number:	D6.1
Date of Issue:	31/01/07
EC Contract No.:	034115
Document Code:	Phosphorus-WP6-D6.1



Test-bed design

The G²MPLS will use three interfaces to communicate with other entities in the testbed:

- the Grid-capable Optical User-Network Interface (G.OUNI)
- the Grid-capable External Network-Network Interface (G.E-NNI)
- the Grid-capable Internal Network-Network (or Node-Node) Interface (G.I-NNI)

All the three interfaces will be developed by Phosphorus from existing GMPLS interfaces. The new features of the interfaces are related to GRIDs. The first interface will be used by the GRID middleware and GRID resources to communicate with the control plane of the PSNC local-testbed. The second one will be used at the edge of PSNC testbed towards others Phosphorus local testbeds. The third interface will be used between the G²MPLS controllers inside the local testbed.

There will be also a new interface used towards the control/provisioning plane of the local testbeds which use an NRPS. This interface is in a very preliminary phase of definition. It will be probably based on the JRA3 IDM module

1.12.3 Workpackage 3 developments

Workpackage 3 developments take place in two areas: Grid middleware (UNICORE and to some extent also Globus Toolkit 4) will be enhanced to make use of the network reservation capabilities developed in work packages 1 and 2 and thus provide seamless unified access to Grid and network resources for the applications. The second area of work is the enhancement and enabling of a set of four applications (WISDOM, KoDaVis, TOPS, DDSS) to make use of these new services. In the first phase of the project, the current versions of the Grid middleware and the applications will be deployed in the local testbeds. Section C.2 in Appendix C “Requirements of Workpackage 3” provides a detailed description, which middleware and applications are being deployed in which local testbeds. During the course of the project, the developments of Workpackage 3 in the areas of middleware and applications (which are described in detail in Deliverable D3.1 “Use-cases, requirements and design of changes and extensions of the applications and middleware”) will be deployed in the testbed for testing and demonstration purposes. Obviously, these deployments will cover the same local testbeds, partner sites, and resources as the original deployments in the first phase.

All of the Workpackage 3 developments affect the local testbeds that form the Phosphorus testbed, so there are no deployments specific to the global testbed.

Project:	Phosphorus
Deliverable Number:	D6.1
Date of Issue:	31/01/07
EC Contract No.:	034115
Document Code:	Phosphorus-WP6-D6.1



2 Topology of the global testbed

The global testbed will consist of multiple local testbeds located in several places in Europe, United States and Canada. For the integration of the whole Phosphorus testbed all local testbeds must be interconnected on data plane as well as on the control/provisioning plane. The data plane connections will be used to transit user data between GRID resources located in different local testbeds while the control/provisioning plane connections will be used for integration of the control planes (GMPLS, G²MPLS) 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 will be based on dedicated lightpaths capable of transmitting huge amounts of data – the amounts which will be generated by Phosphorus applications. As the Phosphorus project is developing optical technologies and the testbed has to allow for demonstration of the project developments, it was decided that the data plane would be build as an optical network with switching capabilities in local testbeds and transparent lightpaths between local testbeds.

The topology of interconnections between local testbeds results from requirements from other Phosphorus activities – the activities which will utilise the testbed to demonstrate their developments. To make the testbed as cost-effective as possible, the lightpaths will be 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 use the GÉANT+ infrastructure which allows for provisioning lightpaths between European NRENs and between Europe and New York. The GÉANT+ lightpaths are complemented by lightpaths from other sources (like CANET4 in Canada, Cross Border Dark Fibre provided by NRENs, GLIF and others).

The control/provisioning network is not intended to transport huge amounts of data but to provide reliable connectivity suitable to transmit only the necessary signalling information between local testbeds. To achieve reliable control/provisioning communication even when the experiments conducted in the testbed may disrupt the data plane connectivity, it was decided to separate the control/provisioning plane from the data plane and create dedicated connections for signalling. The control/provisioning plane will use transmission over the Internet which is reliable enough for the testing purposes and very easy to provision.

One of the tasks of Workpackage 6 is to make new types of transmission equipment available to the project. The new resources will include an optical transmission system made available by CESNET. The system will be installed between PSNC (Poznan) and CESNET (Prague) where the two project participants have dark fibre

Project:	Phosphorus
Deliverable Number:	D6.1
Date of Issue:	31/01/07
EC Contract No.:	034115
Document Code:	Phosphorus-WP6-D6.1



Test-bed design

which can be used for Phosphorus. Unlike the lightpaths obtained from external providers, this system will be under full control of Phosphorus participants. This will allow the control/provisioning layer of the testbed to get transmission-related information from the transmission system like, for example, information about error rates in the transmitted signals.

2.1 Topology requirements and assumptions

The testbed topology is based on the requirements which were submitted by other Phosphorus activities – the activities which will utilise the testbed to demonstrate their developments, especially Workpackages 1, 2 and 3. The requirements refer to the physical topology (which local testbeds are directly connected), the transmission technology, link capacity as well as control/provisioning interconnections. In most cases the data plane links connect local testbeds in which similar tasks will be executed. To give some examples, Workpackage 1 requested interconnections between all local testbeds in which NRPSes will be deployed while Workpackage 3 would like to interconnect all sites where a given application will run.

The main assumption for the data plane topology was that the topology should correspond to the needs of the whole Phosphorus consortium and all Phosphorus activities in the best possible way. It was also assumed that each local testbed should have at least two connections to two other local testbeds to make switching between such connections possible at each site. On the other hand, the number of physical connections should be kept limited to make the testbed manageable and cost-effective. The topology of interconnections should offer the necessary flexibility to facilitate connections between local testbeds which are not directly connected, provide multiple links (not necessarily direct) between a given pair of local testbeds when suitable for a test scenario or provide paths passing multiple technology and administrative domains in order to demonstrate multi-domain path provisioning. To increase the flexibility of interconnecting any pair of local testbeds accordingly to the test scenarios, it is planned to use multiple Ethernet VLANs on a single physical link. Using VLANs the testbed can simulate multiple logical connections which share the same physical link. In this case the upper layers (applications and GRID middleware) will see each of the logical connections as a separate link between local testbeds. Using VLANs will require that the physical links are terminated on Gigabit Ethernet switches in which the VLANs can be demultiplexed to different physical ports. If the optical switches which make the switching resources of the testbed are connected with the Gigabit Ethernet switches with multiple links and each VLAN are demultiplexed to a different link between the Gigabit Ethernet switch and the optical switch, the optical switch will be able to switch VLANs as if they were physical links. Using VLANs will also allow for creating logical links between local testbeds which are not connected directly by a physical link. They can be connected by a VLAN which will be seen by the rest of the testbed as a direct link.

To minimise the costs of the testbed the consortium decided to use transmission services which are available to research community free of cost or whose services are included in subscription fees already paid by the community. For this reason most of the connections use GÉANT+ infrastructure which allows for provisioning lightpaths between European NRENs and between Europe and New York. The GÉANT+ lightpaths are complemented by lightpaths from other sources (like CANET4 in Canada, Cross Border Dark Fibre provided by NRENs, GLIF and others). Some of the local testbeds (PSNC, SURFnet, CESNET) will be directly connected to the international networks used for Phosphorus while in some cases (VIOLA, UESSEX, i2CAT) local NRENs

Project:	Phosphorus
Deliverable Number:	D6.1
Date of Issue:	31/01/07
EC Contract No.:	034115
Document Code:	Phosphorus-WP6-D6.1



Test-bed design

will provide connectivity between the local testbeds and the nearest GÉANT+ nodes. As the number of lightpaths and the available bandwidth in some NRENs is very limited and going beyond the existing limits would incur extra costs to the project partners, in some cases several connections going out of a given local testbed will share the same bandwidth. This will limit the amount of data which can be transmitted over the testbed and must be considered in the test scenarios but will not jeopardise the project which main area of interest is in switching, not in data transmission.

The bandwidth requirements have been defined by Workpackage 3 which is going to run a set of demanding scientific applications in the testbed. The applications will act as the main source of data which passes the testbed. One of the applications (Distributed Data Storage) is able to consume up to 10 Gbps of bandwidth. This application will run in PSNC and VIOLA with a possible extension of other sites. The other applications require connections with capacity between 320 and 700Mbps.

As the testbed is supposed to carry GRID traffic and virtually all GRID resources use a technology from the Ethernet family to connect to the external world, it was decided that the testbed must be able to transmit Ethernet frames, natively or encapsulated in any other transmission technology (SDH containers in most cases) with the condition that such encapsulation is transparent to the GRID resources, applications and middleware. To prevent any network element which is outside the control of Phosphorus from affecting the flow of Phosphorus data, the testbed will use only transparent lightpaths in providers' networks, without any routing or Ethernet switching outside Phosphorus local testbeds. This way the transmission parameters of all interconnections (e.g. transmission time or bandwidth) should be virtually constant (changing only in case of a failure in a provider's network).

The Phosphorus consortium has investigated the transmission services available in GÉANT+ and other networks and found that the most suitable and also the easiest service which fulfils the needs of Phosphorus is a Gigabit Ethernet lightpath with 1 Gbps bandwidth. Such bandwidth is sufficient for almost all Phosphorus experiments and applications; only one application can utilise more than this. The cost of this service in Europe is already included in the subscription fee which NRENs already pay to DANTE. This way using this service does not incur any additional costs.

The Gigabit Ethernet lightpaths carry Ethernet frames over an optical backbone network. The encapsulation of Ethernet frames in SDH containers is transparent to the applications and the Phosphorus control/provisioning layer and is executed by GÉANT+ access devices or NREN's transmission equipment (when the NREN is connected to GÉANT+ service in SDH technology).

Besides the data plane on which data between GRID resources will be transmitted, the testbed will include control/provisioning plane which will be used for communication between the network control instances (GMPLS, G²MPLS), between the instances of the Network Resource Provisioning Systems and – when a suitable interface is ready – between GMPLS/G²MPLS and NRPSes. The control/provisioning communications must be seen from two perspectives: proper interfaces and interconnections. The interfaces will define the way in which network control and resource provisioning instances can exchange information in order to collaboratively provision an optical path on a request from a user. They will be proposed by other Phosphorus workpackages and supported by the testbed when the appropriate software is available from other

Project:	Phosphorus
Deliverable Number:	D6.1
Date of Issue:	31/01/07
EC Contract No.:	034115
Document Code:	Phosphorus-WP6-D6.1



Test-bed design

workpackages. This document covers only the physical interconnections of control/provisioning instances over the testbed which allows for transmitting packets between them.

The control/management plane is not supposed to transport large amounts of data. It will transport only the information which is exchanged between the different control and provisioning entities in order to build a global view of the network and provision optical path as requested by users and applications. The control/management plane should offer reliable connectivity for the control and provisioning entities even when some data plane connections fail due to a failure in the transport network, a potentially disruptive testbed experiment or a test scenario in which a link failure is simulated. To achieve reliable control/provisioning communication the control/provisioning plane will be separated from the data plane and will not share the same physical links. As the bandwidth of control/provisioning plane connectivity may be very small and the transmission times are not crucial, it was decided to use Internet to interconnect control and management entities. The necessary security of the control traffic can be ensured by use of a VPN in which the data transmitted over Internet is encrypted and the authentication of the VPN sites is managed by the VPN software.

2.2 Data plane design

The data plane will be responsible for transporting data between GRID resources in the Phosphorus testbed. It will comprise of switching resources in local testbeds (described in section 1 of this document) and a set of transmission links between local testbeds. The topology of the interconnections between local testbeds will comply with the requirements which are described in section 2.1.

All links between local testbeds will be able to carry flows of Ethernet frames with the capacity of 1 Gbps. Some exceptions are due to external constraints and are described below. To prevent any network element which is outside the control of Phosphorus from affecting the flow of Phosphorus data, the project aims at using only transparent lightpaths in providers' networks, without any routing or Ethernet switching outside Phosphorus local testbeds. Most links will be fully transparent but some of them will pass a Gigabit Ethernet switch or router which is controlled by an external entity. The transmission parameters (e.g. transmission time or bandwidth) of all fully transparent interconnections should be virtually constant (changing only in case of a failure in a provider's network).

The testbed will use links from different providers: GN2 (the GÉANT+ service), Cross Border Fibre (CBF), CANet, GLIF and others. In some cases the international links must be extended to the premises of the Phosphorus partners by local NREs.

The testbed will comprise the following links:

- PSNC – CESNET over CBF (with capacity up to 10 Gbps)
- PSNC – i2CAT over GÉANT2
- PSNC – SURFNET over GÉANT2
- PSNC – UESSEX over GÉANT2
- PSNC – VIOLA over GÉANT2

Project:	Phosphorus
Deliverable Number:	D6.1
Date of Issue:	31/01/07
EC Contract No.:	034115
Document Code:	Phosphorus-WP6-D6.1



Test-bed design

- PSNC – VIOLA over CBF
- CESNET – SURFNET over GLIF (with capacity up to 10 Gbps)
- i2CAT – SURFNET over GÉANT2
- i2CAT – SURFNET over GLIF
- i2CAT – VIOLA over GÉANT2
- SURFNET – VIOLA over GÉANT2
- SURFNET – VIOLA over CBF
- SURFNET – UESSEX over GÉANT2
- SURFNET – UvA over SURFNET infrastructure
- PSNC – CRC over GÉANT2 (between PSNC and GÉANT2 PoP in New York) and CANet 4 (between GÉANT2 PoP in New York and CRC)
- CESNET – STARLIGHT over GLIF (this link will be available temporarily, it may have capacity up to 10 Gbps)
- SURFNET – CRC over GLIF and CANet 4
- SURFNET – STARLIGHT over GLIF (this link can be used if necessary for test scenarios)
- MCNC – STARLIGHT over EnLIGHTened wave
- NORTEL – CRC over CANet 4
- CRC – STARLIGHT over CANet 4
- CCT at LSU – STARLIGHT over LONI wave (CCT at LSU is not a member of the PHOSPHORUS consortium; this link will be used if CCT and LSU participates in the PHOSPHORUS tests)

The topology of interconnections between local testbeds is shown in Fig. 2.1. In the three local testbeds in Europe in which the highest number of links terminate (PSNC, SURFnet and VIOLA), the interconnections will be terminated on Gigabit Ethernet switches which will make it possible to provision VLANs between local testbeds which do not have direct connections. Those three local testbeds will act as ‘hubs’ for other local testbeds. Each of the European local testbeds has a direct connection to at least one of the free hubs, which ensures the possibility to create a VLAN between any pair of Phosphorus local testbeds in Europe. There are two options for transmitting VLANs over the lightpath. The first option is based on using Q-in-Q for encapsulation of multiple user’s VLAN into a single VLAN which will be transmitted over the lightpath. To achieve this all the switches on the edge of local testbeds must support this technology. The advantage of this option is that the distribution of bandwidth over the VLANs will depend only on the PHOSPHORUS testbed and the providers of the links will not be engaged when the distribution of bandwidth must be changed. In the other option the user’s VLANs can be transmitted natively over the lightpaths but this may require that the providers of the links assign bandwidth to the VLANs and each reconfiguration of the bandwidth distribution must be done by the providers (which may be much more time-consuming than in the first option).

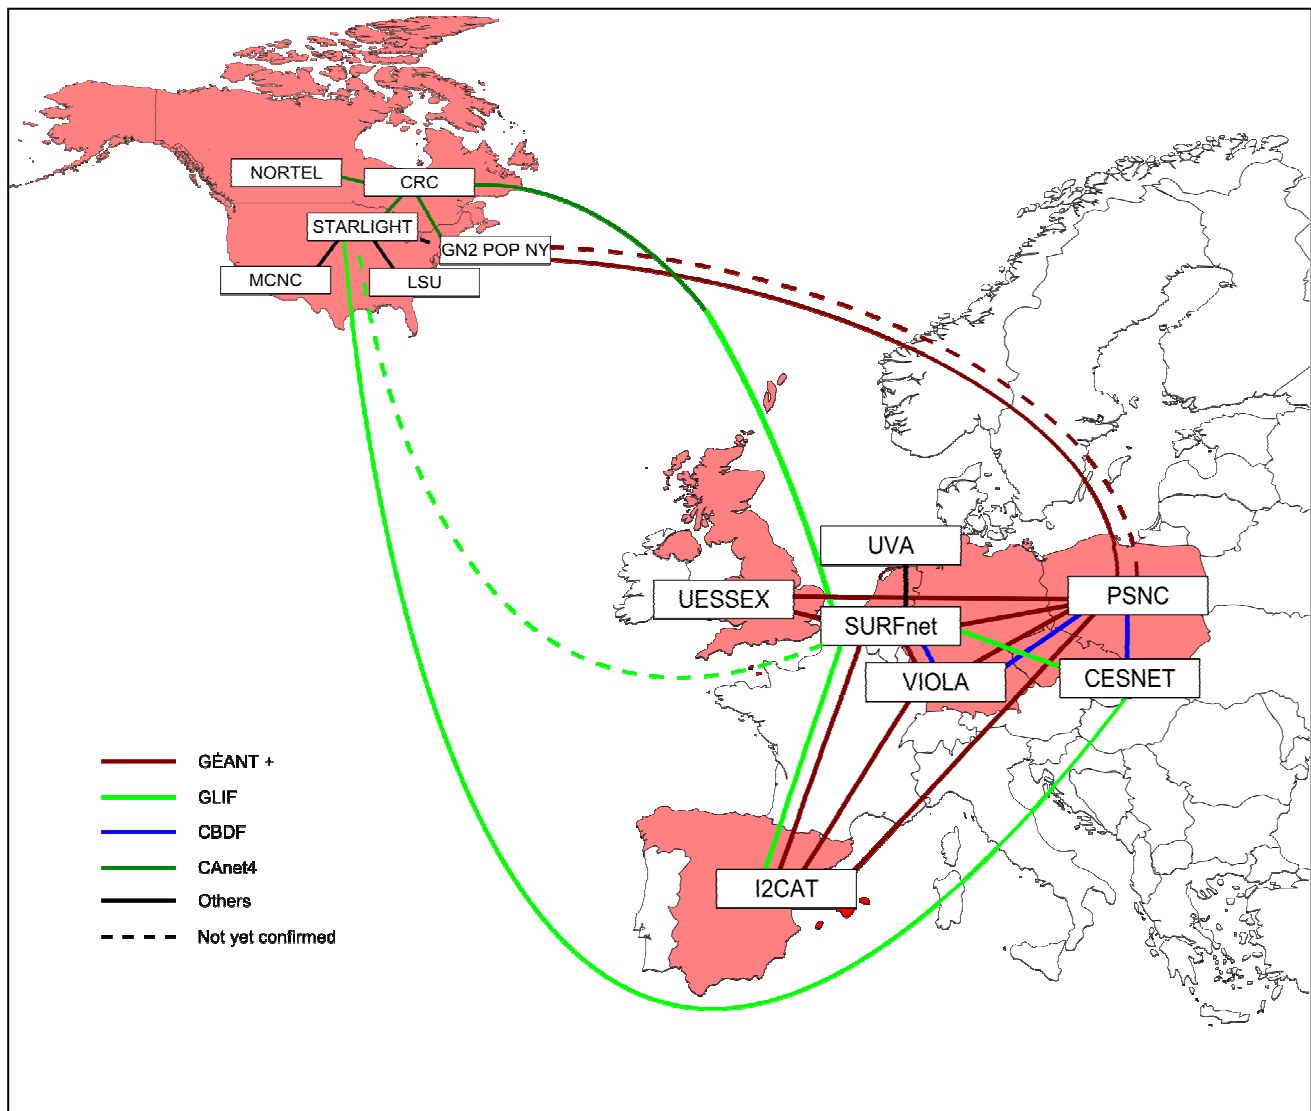


Figure 2.1: Topology of the PHOSPHORUS testbed

STARLIGHT in Chicago serves as a hub for connectivity to the local testbeds located in Canada and the USA. The testbed will use three connections between Europe and America: from PSNC to CRC (over GÉANT2 and CAnet 4 via the GN2 PoP in New York), from CESNET to STARLIGHT (available temporarily over GLIF) and from SURFNET to CRC. The link between SURFNET and CRC, as well as the GLIF connection between i2CAT and SURFNET, will be used mostly for interconnecting i2CAT and CRC (both partners participate in the development of UCLP in Workpackage 1 but can also be used by other PHOSPHORUS partners, according to the project needs and test scenarios).



Test-bed design

CRC will operate as a hub in order to provide connectivity between PSNC and STARLIGHT to which MCNC and CCT at LSU will be connected. This way all three American and Canadian local testbeds can use a single lightpath to PSNC provided by GÉANT2. PSNC will make the lightpath available also for other European local testbeds.

If needed for the test scenarios, it is also possible for the consortium to lease a connection between PSNC and STARLIGHT (provided partly by GÉANT2 and partly by a third party) to increase the number of parallel connections between European and American parts of the PHOSPHORUS testbed.

Although all the connections will have the capacity of 1 Gbps (and some of them even higher), there are some constraints that must be considered in test scenarios:

- The two GÉANT+ lightpaths to UESSEX will share a single 1 Gbps link between UESSEX and UKLight (the lightpaths will be terminated on a switch in UKLight and extended to UESSEX as VLANs on a Gigabit Ethernet trunk); this is going to be changed during the Phosphorus timeframe to two independent 1 Gbps links between UESSEX and UKLight during the project lifetime; the exact date of this upgrade depends on lightpath availability in UKERNA (the provider of connectivity between UESSEX and UKLight)
- The GÉANT+ lightpaths to i2CAT will share a single 1 Gbps link with production i2CAT – Internet traffic so it will not be able to utilise the whole 1 Gbps

The topology of interconnections to VIOLA is still under discussion with DFN and may be changed.



2.3 Control/provisioning plane design

Most of the PHOSPHORUS developments will be implemented in the control and provisioning planes of the network. The developments in the control plane aims in transition from GMPLS to G²MPLS which will offer more specialised services for GRIDS while in the provisioning plane new interfaces for some existing NRPSes will be developed in order to allow for integration between them and between NRPS and G²MPLS control plane in a GRID-enabled network.

To properly test the project developments, the tests must involve multiple local testbeds. There are two main reasons for this:

- Most of the local testbeds will have only one NRPS or only GMPLS/G²MPLS control plane so in order to demonstrate the interoperability between NRPSes as well as between an NRPS and G²MPLS multiple local testbeds must be used in a single test.
- As the choice of hardware for each local testbed is generally limited to the hardware owned by a single PHOSPHORUS partner, which usually means a very limited number of vendors present in a local testbed, having a really multi-vendor demonstration requires using hardware which is available in different local testbeds.

In order to allow for provisioning network services in multiple domains, the control plane entities and/or provisioning systems of different domains must be able to communicate. Secure and reliable communication will be ensured by the use of Virtual Private Network (VPN) technology over the Internet. In the first phase of the project two VPNs will be needed – one for GMPLS signalling and one for NRPS-to-NRPS communication. In the second phase – when the interface between NRPSes and G²MPLS must be tested and demonstrated – the two VPNs will be integrated to allow for communication between NRPSes and the G²MPLS control plane.

Several VPN implementations have been examined by the Phosphorus consortium (see appendix D for further information). After the examination one of the implementations – tinc – has been selected for use in the testbed. The tinc software is available for many hardware platforms and operating systems. It can be obtained and used free of cost. Tinc offers secure, reliable, efficient and scalable communication between several sites over the Internet. Security of tinc ensures private communication with use of encryption and OpenSSL. Scalability means that it is possible to build VPNs comprising many sites and each site is able to communicate with all other sites.

Tinc version 1.0.6 is most likely to be used as it supports many operating systems including Linux, FreeBSD and Windows.

To participate in a VPN each local testbed will need a PC with one of the operating systems supported by tinc. The PC must be connected to the control-plane network of the local testbed or to the NRPS host. It must also be connected to the Internet using another interface.

Project:	Phosphorus
Deliverable Number:	D6.1
Date of Issue:	31/01/07
EC Contract No.:	034115
Document Code:	Phosphorus-WP6-D6.1



3 Conclusions

The purpose of this deliverable was to propose the design of the Phosphorus testbed. The design is based on requirements obtained from the Phosphorus activities which will use the testbed to verify and demonstrate their developments. On the other hand, the design must comply with some external limitations and must be cost effective. To achieve the cost effectiveness the testbed mostly uses the equipment which is already owned by Phosphorus partners and connectivity in research networks which is available to the research community without additional costs.

The Phosphorus testbed comprises of several local testbeds provided by Phosphorus partners. Each local testbed creates a natural administrative and technological domain. To allow for inter-domain testing the local testbeds are interconnected and form a global Phosphorus testbed.

The document is divided into two main sections. The first section describes the local testbeds and the resources which will be used for testing the Phosphorus developments. The other section proposes the topology of interconnections between local testbeds.

This deliverable will be used by Workpackage 6 of Phosphorus as guidelines for the preparation of the testbed. The testbed will be built incrementally over the Phosphorus lifetime. The first version of the testbed will use off-the-shelf protocols and Network Resource Provisioning Systems, which will be then replaced with new versions proposed and developed by the project.

As the design satisfies the requirements of the main users of the testbed, the testbed is believed to offer the opportunities needed by the project to successfully verify and demonstrate the project developments.



4 References

- | | |
|-----------|---|
| tinc | www.tinc-vpn.org |
| OpenSSL | www.openssl.org |
| CAnet 4 | www.canarie.ca |
| STARLIGHT | http://www.startap.net/starlight/ |
| MetaMPICH | Bierbaum, B.*; Clauss, C.*; Eickermann, T.; Kirtchakova, L.; Krechel, A.*; Springstubbe, S.*; Wäldrich, O.*; Ziegler, W.*
Reliable Orchestration of Distributed MPI-Applications in a UNICORE-Based Grid with MetaMPICH and MetaScheduling 3-540-39110-X
Recent Advances in Parallel Virtual Machine and Message Passing Interface : 13th European PVM/MPI User' s Group Meeting, Bonn, Germany, September 17-20, 2006 ; proceedings / ed.: B. Mohr, J. Larsson Träff, J. Worringen, J. Dongarra. - Berlin, Springer, 2006. - (Lecture Notes in Computer Science ; 4192). - 3-540-39110-X, 978-3-540-39110-4. - S. 174 – 183. |
| i-CONE | Simon, A.; Göbel, M.: The i-CONE TM - a panoramic display system for virtual environments
10th Pacific Conference on Computer Graphics and Applications. October 9 - 11, 2002, Tsinghua University, Beijing, Proceedings /ed.: Coquillart, S.:
Los Alamitos, Calif.: IEEE Computer Society, 2002, ISBN: 0-7695-1784-6, pp.3-7 |
| CAVE: | The CAVE: audio visual experience automatic virtual environment
Communications of the ACM, Volume 35, Issue 6 (June 1992) Pages: 64 - 72, Year of Publication: 1992
ISSN:0001-0782, C. Cruz-Neita, D.J. Sandin, T.A. DeFanti, R.V. Kenyon, J.C. Hart
ACM Press New York, NY, USA |



5 Acronyms

AON	All Optical Network
B/A	backup/archive
BADLAB	Broadband Applications and Demonstration Laboratory
BER	Bit Error Rate
CANET 4	The Canadian NREN
CAVE	Cave Automatic Virtual Environment
CD	Chromatic Dispersion
CLA	CzechLight Amplifier
CPM	Cross phase modulation
DDSS	Distributed Data Storage Systems
DWDM	Dense Wavelength Division Multiplexing
EDFA	Erbium Doped Fibre Amplifier
FBG	Fibre Bragg Grating
FC	Fibre Channel
FC-SATA	Fibre Channel to SATA technology (mixed technology used in disk matrices: disk matrix have Fibre Channel ports for hosts connectivity, but contains SATA disk drives)
GÉANT2	Pan-European Gigabit Research Network
GÉANT+	the point-to-point service in GÉANT2
GMPLS	Generalized MPLS (MultiProtocol Label Switching)
G ² MPLS	Grid-GMPLS (enhancements to GMPLS for Grid support)
GN2	GÉANT2
GT4	Globus Toolkit Version 4.0
IP	Internet Protocol
IPSec	IP Security
MEMS	Microelectromechanical Systems
MPI	Message Passing Interface, see http://www-unix.mcs.anl.gov/mpi
MPLS	MultiProtocol Label Switching
MSS	MetaScheduling Service
NIL	Nothing In Line
NNI	Network to Network Interface
NREN	National Research and Education Network
NRPS	Network Resource Provisioning System



Test-bed design

ONL	Optical Networking Laboratory
OSNR	Optical Signal to Noise Ratio
PoP	Point of Presence
Raman	this isn't an acronym, amplification here is simply based on Raman's scattering
RMS	Resource Management System
SAGbE	Scalable Adaptive Graphics Environment
SNMP	Simple Network Management Protocol
SOA	Service Oriented Architecture
SPM	Self Phase Modulation
STK	StorageTek, a company producing tape libraries
TBS	Token Based Switch
TL1	Transaction Language 1
TOPS	Technology for Optical Pixel Streaming
TSM	Tivoli Storage Manager, a commercial backup/archive software by Tivoli, formerly known as IBM's ADSM)
UCLP	User Controlled Light Paths
UNI	User to Network Interface
VLAN	Virtual LAN (as specified in IEEE 802.1p)
VPN	Virtual Private Network

Project:	Phosphorus
Deliverable Number:	D6.1
Date of Issue:	31/01/07
EC Contract No.:	034115
Document Code:	Phosphorus-WP6-D6.1



Appendix A Requirements of Workpackage 1

Compiled by

Carsten Rosche (Fraunhofer Institute IAIS)

`Carsten.rosche@iais.fraunhofer.de`

With contributions from:

- CRC
- i2CAT
- Nortel
- SURFNet
- Fraunhofer IAIS (VIOLA)

Project:	Phosphorus
Deliverable Number:	D6.1
Date of Issue:	31/01/07
EC Contract No.:	034115
Document Code:	Phosphorus-WP6-D6.1



Test-bed design - Appendix
Document Management

History of changes

Version	Status	Date	Person Resp.	Reason for Change
0.1	Draft	28.11.2006	Carsten Rosche	First proposal Including CRC, i2CAT, VIOLA
0.2	Draft	30.11.2006	Albert López	Updated global topology diagram
0.4	Draft		Carsten Rosche	Added SURFNet, Nortel
0.5	Draft	13.12.2006	Carsten Rosche	Proposed final version
0.6	Draft	18.12.2006	Albert López	Version including some corrections regarding i2CAT's test bed
1.0	Final	20.12.2006	Carsten Rosche	Final version. Based on version 0.6 without tracked changes, with minor fixes and embellishment.

Project: Phosphorus
Deliverable Number: D6.1
Date of Issue: 31/01/07
EC Contract No.: 034115
Document Code: Phosphorus-WP6-D6.1

A.1 Introduction

A.1.1 Purpose of the document

- Within the Phosphorus project several testbeds shall be interconnected. This document specifies the requirements of the partners of WP1 towards these interconnections. It shall be presented as a proposal for testbed design towards WP6.
- This document is describing the Data Plane and compute/storage/visualization resources. Design of the Control Plane will be addressed separately. This testbed setup is a proposal intended for development purposes. Therefore the final testbed may differ.

A.2 Testbed Requirements

- In this chapter, the testbed requirements for each local testbed are specified. Figure A.1 shows the overall topology desired by WP1.

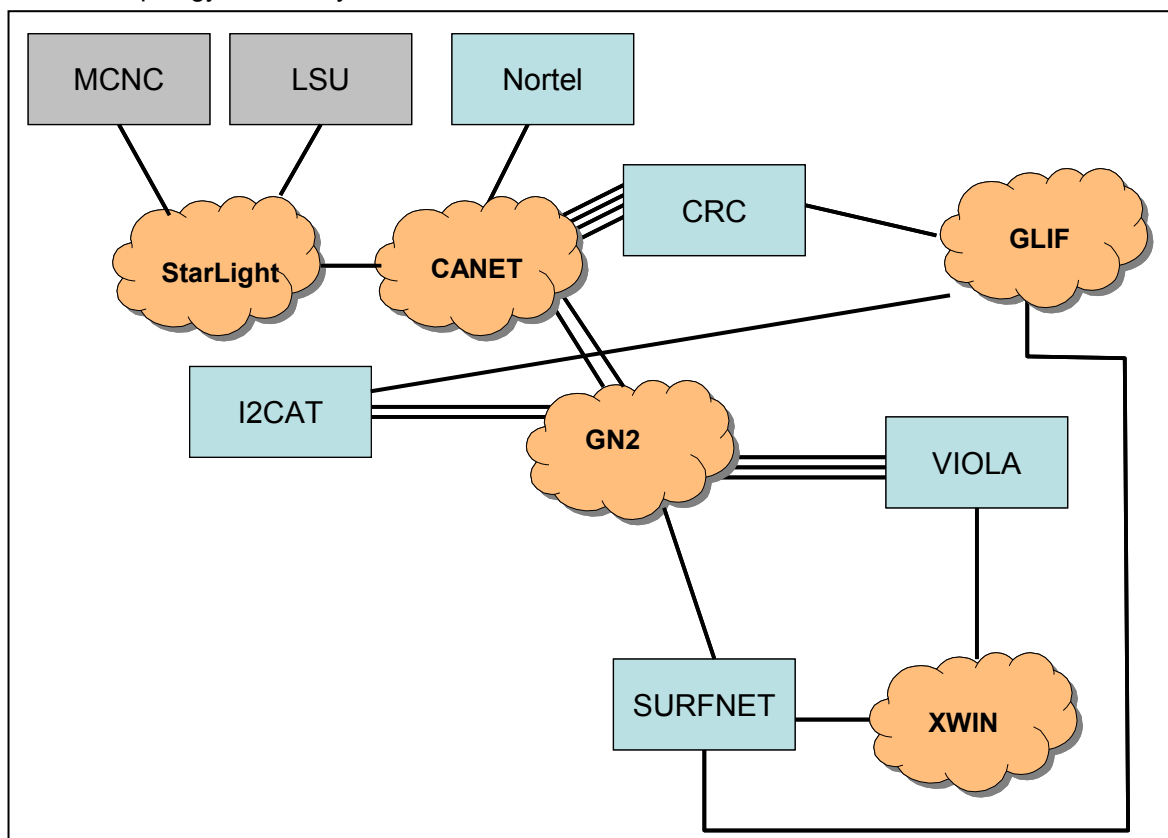


Figure A.1: Overall topology for WP1

A.2.1 CRC Requirements

A.2.1.1 Data Plane Connectivity between Local Testbeds

- Topology of connections: CRC will use 4 Gigabit Ethernet (GbE) connections from the Broadband Applications and Demonstration Laboratory (BADLAB™) to the CANet 4 Ottawa POP. These terminating points will be connected via LightPaths (LPs) to: the GN2 New York POP, the GLIF, the STARLIGHT in Chicago and Nortel in Ottawa. CRC has submitted three individual LP request forms to CANARIE for the first three while Nortel has submitted (or will be submitting) another LP request form to CANARIE for the fourth connection.
- Technology: Gigabit Ethernet over LPs (L2VPN).
- Bandwidth: The LPs across CANet 4 were requested with STS-24c capacities (i.e. 1.25 Gbps to carry GbE traffic). Due to resource constraints, the capacity maybe reduced or only be made available on a shared (i.e. scheduled) basis.
- Time for which they will be needed: LPs were requested from CANARIE for the period from Nov. 20, 2006 to March 31, 2007 (the latest date supported by the form) but included a note stating that the LPs were needed for the duration of the PHOSPHORUS project.

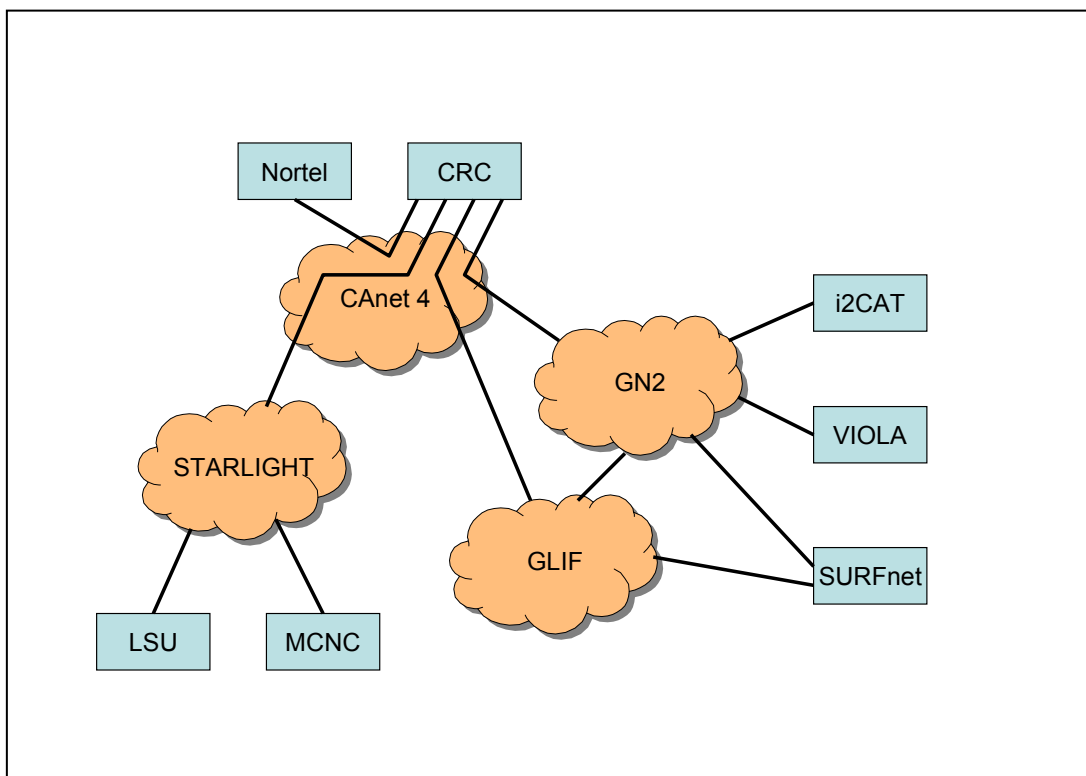


Figure A.2: CRC Testbed Topology Proposal



Test-bed design - Appendix

NOTE: We understand that the links across GN2 that will terminate at their New York POP will be the responsibility of PSNC.

A.2.1.2 Control Plane Connectivity between Local Testbeds

- We can use IPsec tunnels over the Layer 3 service offered by CANet 4, GN2 and other NRENs.

A.2.1.3 Transmission/Switching in Local Testbed

- The topology of the internal connections (inside the local testbed): The CRC testbed will be extended across CANet 4. The LPs across CANet 4 will be managed and controlled via UCLP. The network elements on CANet 4 which are supported by UCLP are: Cisco ONS 15454, Nortel OME 6500 and Nortel HDXc. At the Nortel campus, 4 Nortel OME 6500s will be configured in a gateway plus 3-node arrangement. The BADLAB facility is connected to the CANet 4 Ottawa POPs via an 8-channel CWDM infrastructure. The core switch/router in the BADLAB is a Cisco Catalyst 6509. A Nortel OPTera Metro 5200 is also available as part of the Optical Networking Laboratory (ONL) adjacent to the BADLAB and supports GbE interfaces which can be connected to the Catalyst 6509 and/or PC's.
- Transmission technology inside the local testbed: GbE is the common interface used on the CANet 4 network elements. The GbE ports are mapped onto STS-24c circuits (or lesser depending on availability and the bandwidth required by the applications). The networking equipment in the BADLAB is currently being upgraded to support 10 GbE. Once the upgrade has been completed, we will be in the position to support LPs based on circuits up to STS-192c. Within the BADLAB, the Cisco Catalyst 6509 provides standard switching and routing functionalities.
- Interfaces: Cisco Catalyst 6509 ports supported are: copper at 10/100/1000 Mbps Ethernet, 100BaseFX, 1000BaseT/SX/LX/ZX; Nortel Optera Metro 5200 ports supported: protocol and bit rate transparent with 1310 nm optics.
- Switching (number of switches/interfaces, topology of connections, control plane requirements): Please refer to Figure A.3 and Figure A.4. It is expected that software will be made available through WP-2 to implement a GMPLS (and later G²MPLS) based control plane running on a number of PCs in the BADLAB.



Test-bed design - Appendix

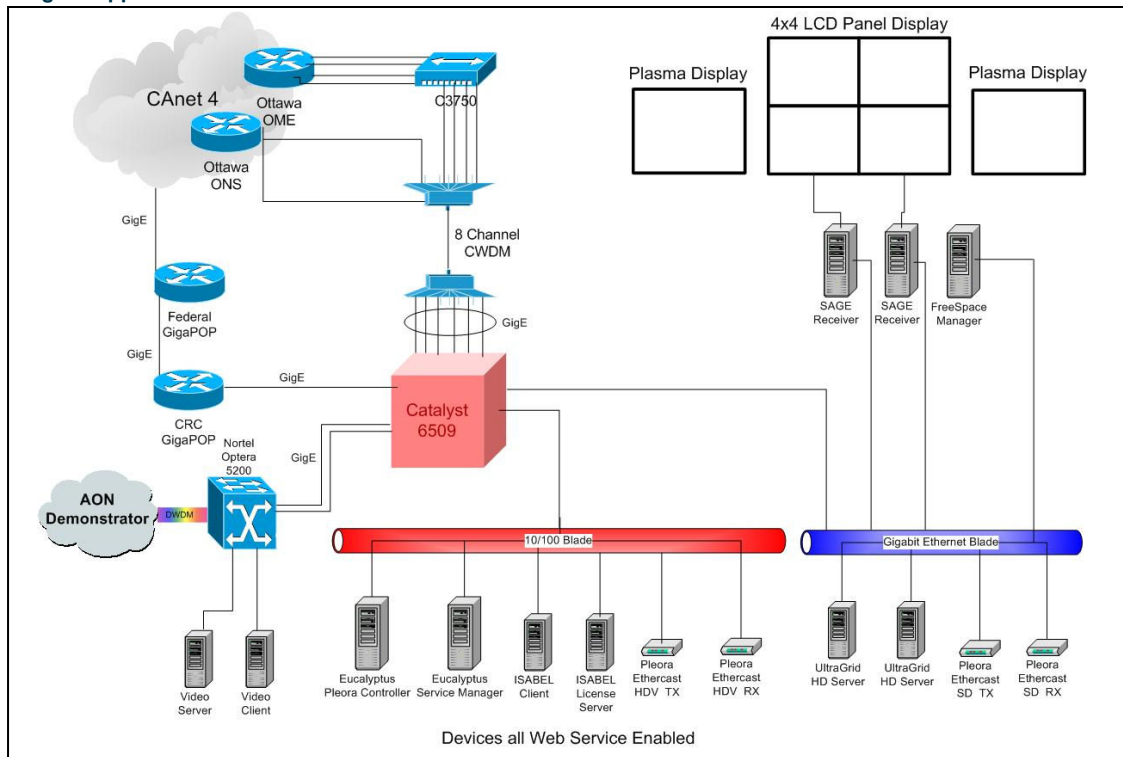


Figure A.3: Detailed View of CRC's Local Testbed

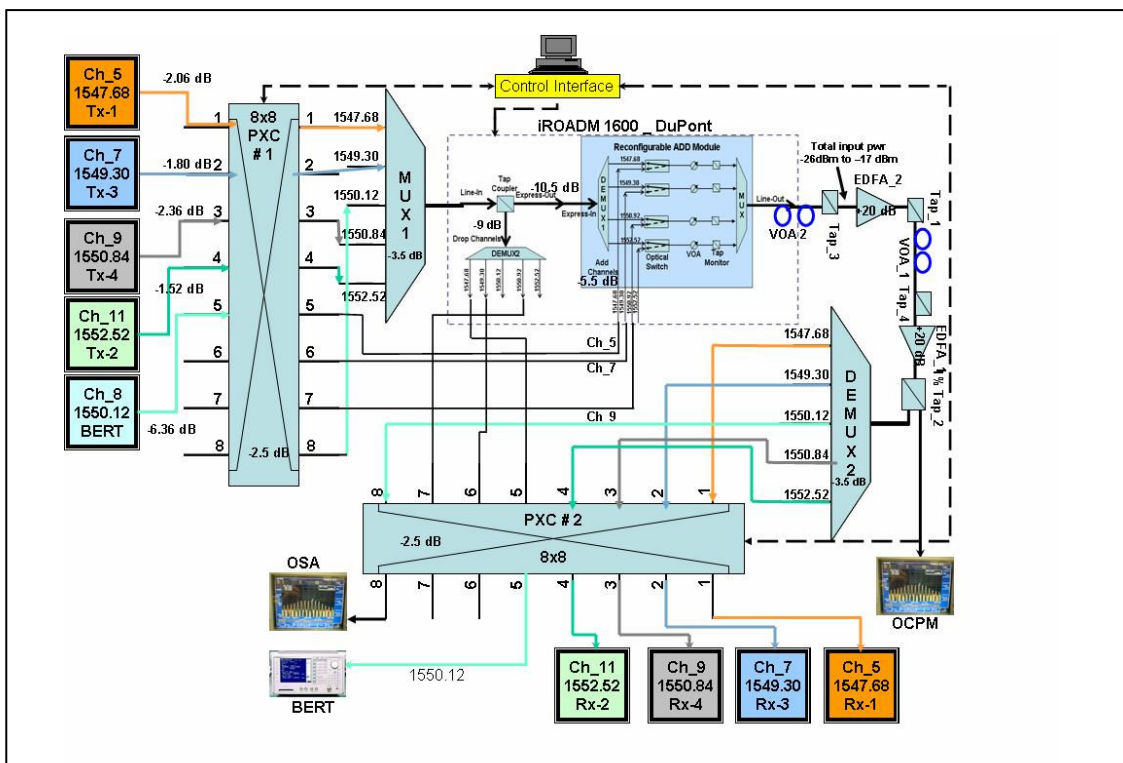


Figure A.4: Detailed view of the AON Demonstrator at CRC

Project: Phosphorus
 Deliverable Number: D6.1
 Date of Issue: 31/01/07
 EC Contract No.: 034115
 Document Code: Phosphorus-WP6-D6.1



A.2.1.4 Networking, Computing, Storage and Visualization Resources in Local Testbed

The table below shows the equipment available at CRC for the present project.

Networking	Computing & Storage	Visualization
<ul style="list-style-type: none"> • Cisco Catalyst 6509 (GbE, 10 GbE) • Cisco Catalyst 3750 • 8 Channel CWDM MUX/DEMUX • Cisco 12008 router (GbE, ATM) • Nortel OPTera Metro 5200 (GbE, DWDM) • 8 x 8 Photonic Switches • 16 Channel ROADM • 16 Channel DWDM MUX/DEMUX • EDFAs • VOAs • OCPM • BERT Tester • OSA 	<ul style="list-style-type: none"> • 3 x Pentium IV Dual Xeon @ 2.8GHz, 2.5 GB RAM, 2 x 40 GB HDD • 3 x Pentium IV Dual Xenon @ 1.7 GHz, 512 MB RAM, 1 or 2 40 – 120 GB HDD • 9 x Pentium IV 3.2 GHz, 512 MB – 1 GB RAM, 100 – 160 GB HDD 	<ul style="list-style-type: none"> • UltraGrid Uncompressed HD (Xena HD and DVC Centaurus capture cards) • Pleora Technology, Uncompressed SD • Pleora Technology, DVB-ASI • Pleora Technology, HDV • ISABEL (Clients (26), flow servers, license servers (24 users)) • H.323 (Polycom VSX 7000 with MCU (x3), Polycom PVX (x10)) • VideoLAN (HDV, DV) • Skype • Sony HVR-Z1U HD camcorder • JVC JY-HD10U HD camcorder • JVC GR-HD1U HD camcorder • JVC HM-DH5U recorder (x2)

Please refer to the Eucalyptus: a Service Oriented Participatory Design Studio poster (Figure A.5) for a description of a solution based on SOA and UCLP technologies to support collaborative multi-site design for advanced network user communities.



EUCALYPTUS

CANARIE

CRC

ARC-CARC

A Service-oriented Participatory Design Studio

Sandy Liu, Yong Liang, Bo Xu, Bruce Spencer, Libo Zhang – Joint work with: CIMS at Carleton University, CRC UCLP team, John Spence, and Martin Brooks

WHAT IS EUCALYPTUS?

A service-oriented solution to manage and configure the resources needed by users engaging in a participatory design session.

MOTIVATIONS

- Many high-end resources are not available in most labs – remote access needed
- Most end users are not IT experts
- Need to hide the tools' logistical and provisioning complexities
- Design teams are distributed in different geographical locations
- Require on-demand configuration of resources depending on the task at hand, available resources, particularly the available network bandwidth to support participatory design sessions.

PARTICIPATORY DESIGN STUDIO PROJECT

FOR WHOM?

- Architects of buildings
- Can be extended to the Architecture, Engineering, Construction (AEC), industrial design, and the automotive industries.

KEY FEATURES

- Easy to configure
- Business process management
- Programming language and platform independent
- Shared high-performance visualization tools
- Shared high-definition videoconference
- Service-oriented: **Web Services (WS)**
- Supported by 1to10 Gb/s user-controlled Articulated Private Networks (APNs)
- Integration of the SOA with UCLPv2

SUPPORTED BY USER-CONTROLLED LIGHTPATHS (UCLP) SERVICES

- Treats network resources as software objects
- End-users can provision and reconfigure lightpaths on-demand to form their own private end-to-end optical networks, called **Articulated Private Networks (APNs)**.
- High speed, low latency APNs can accommodate rich media traffic required by user communities.

Distance collaboration facilitated by shared access to Alias Systems MAYA with the support of Shared Desktop Web Service and DCV Web Service, allowing visualization resolution up to 9.2 million pixels.

SYSTEM ARCHITECTURE DESIGN

DEEP COMPUTING VISUALIZATION (DCV) WEB SERVICE

Provides a scalable, collaborative, middleware infrastructure to support and enhance the graphics functions of OpenGL software in two visualization modes:

Remote Visual Networking (RVN)

Scalable Visual Networking (SVN)

EUCALYPTUS DASHBOARD - AN INTEGRATED WEB SERVICES CLIENT, DEPLOYED THROUGH JAVA WEB START

COMPOSING WEB SERVICES INTO WORKFLOWS: AN EXAMPLE

A BPEL Workflow for Rendering:

The user sends a rendering request

- 1) Invokes the File Management Web Service
- 2) Selects the set of files for rendering
- 3) Invokes the rendering Web Service with the selected files
- 4) Invokes the visualization Web Service for each visualization client

EXPECTED IMPACT

The emergence of SOA and UCLP herald the beginning of a new age where fully collaborative multi-site design may become the norm.

The advanced network user community will benefit from the structure and function of PDS.

Figure A.5: Eucalyptus poster

Project: Phosphorus
 Deliverable Number: D6.1
 Date of Issue: 31/01/07
 EC Contract No.: 034115
 Document Code: Phosphorus-WP6-D6.1

A.2.2 i2CAT Requirements

A.2.2.1 Data-plane connectivity between local test beds:

- **Topology of connections:** i2CAT will use 2 new VLANs provided by RedIRIS (Spanish NREN) using the same link used for the production IP VLAN provided by CESCO (Local NREN connected to RedIris). These new connections are not available for now. Another connection going to CRC using GLIF facilities will be used.
- **Technology:** GigabitEthernet over lighthouse (L2VPN) in order to reach other partners.
- **Bandwidth:** 1000Mbps shared between the 2 PHOSPHORUS VLANs coming from Géant 2 and the production IP traffic. In a best case scenario, all the capacity can be assigned to only one of those VLANs. Regarding CRC's connection, it will have a dedicated GigabitEthernet through GLIF.
- **Time for which they will be needed:** can be used along the entire project.

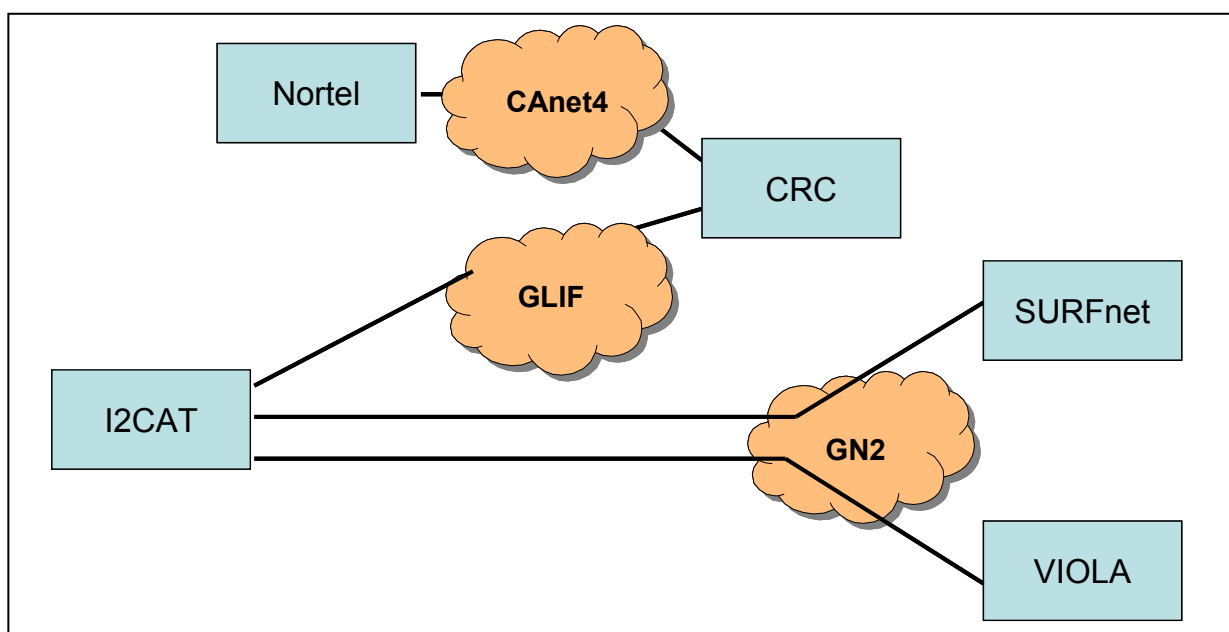


Figure A.6: i2CAT test bed topology proposal

A.2.2.2 Control-plane connectivity between local testbeds

- We can use IPsec tunnels over Internet/GbEANT.

A.2.2.3 Transmission/switching in local test bed

- Topology of the internal connections (inside the local test bed): P2P links (UCLP capable) from i2CAT PoP to application (TOPS and/or EUCALYPTUS) users/PC's. Any connection starting at i2CAT can end at VIOLA, SURFnet or CRC PoP using the Nortel OPTera (also UCLP capable) to switch between the

Test-bed design - Appendix

links. For links that need i2CAT's computer resources we can provide point to point links from one of L2VPNs provided by RedIris or the CRC connection to the Grid farm (using L2VPN technology inside i2CAT network). We will also have the capability to connect vlan-A , vlan-B and vlan-C, creating a Pass-Through connection between any of our three neighbors.

- Transmission technology inside the local testbed: DWDM and GigabitEthernet PHY.
- Interfaces: We will use 2 Nortel Optera 5200. Each Optera is equipped with 2 OCL's and 2 OCLD's slots based on GigabitEthernet technology. All the optics use LX format (1310nm). We will also use Cisco ONS 15454 for the GLIF connection.
- Switching (number of switches/interfaces, topology of connections, control plane requirements): please see Figure A.7.

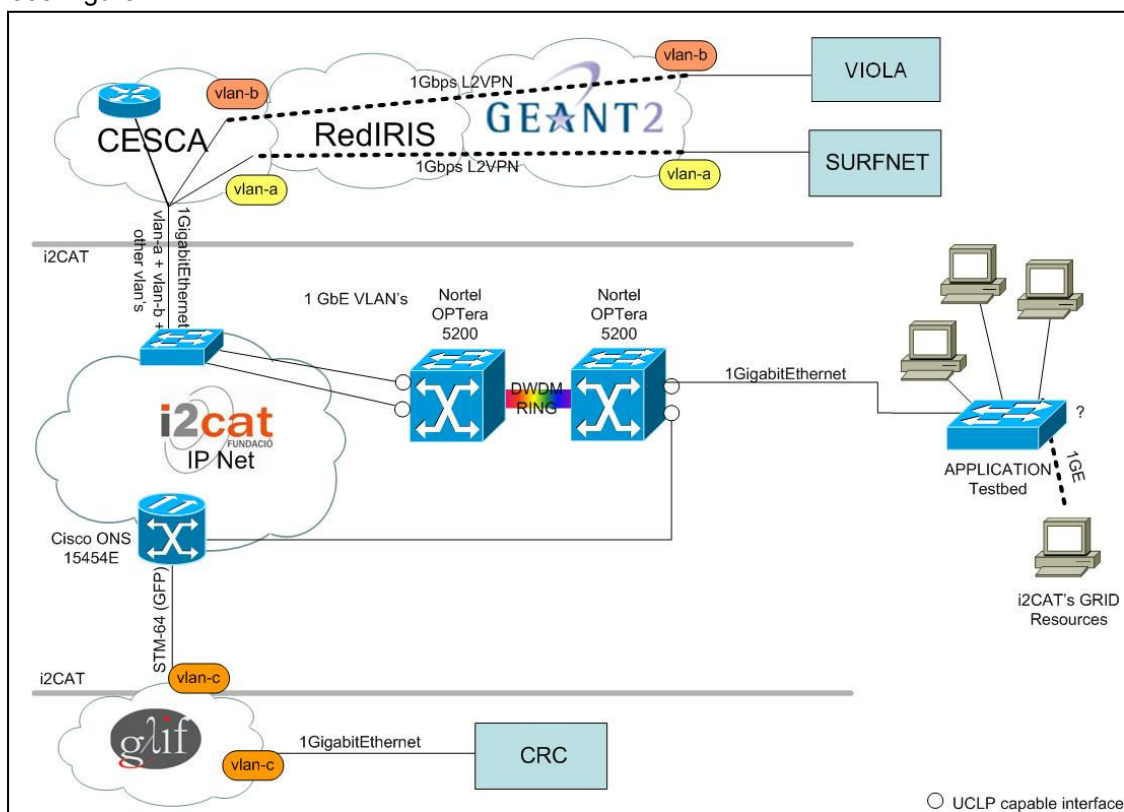


Figure A.7: detailed view of i2CAT's local testbed

A.2.2.4 Computing and storage resources in local test bed

5 x Pentium IV 2.8GHz, 512MB RAM, 120 GB HDD

2 x Pentium IV 3 GHz, 512 MB RAM, 120 GB HDD

4 x Pentium IV 2.8 GHz, 512 MB RAM, 170 GB HDD

9 x Blade DELL PowerEDGbE 1855 Pentium Xeon 2.8 GHz, 1 GB RAM, 140 GB HDD

Project:	Phosphorus
Deliverable Number:	D6.1
Date of Issue:	31/01/07
EC Contract No.:	034115
Document Code:	Phosphorus-WP6-D6.1

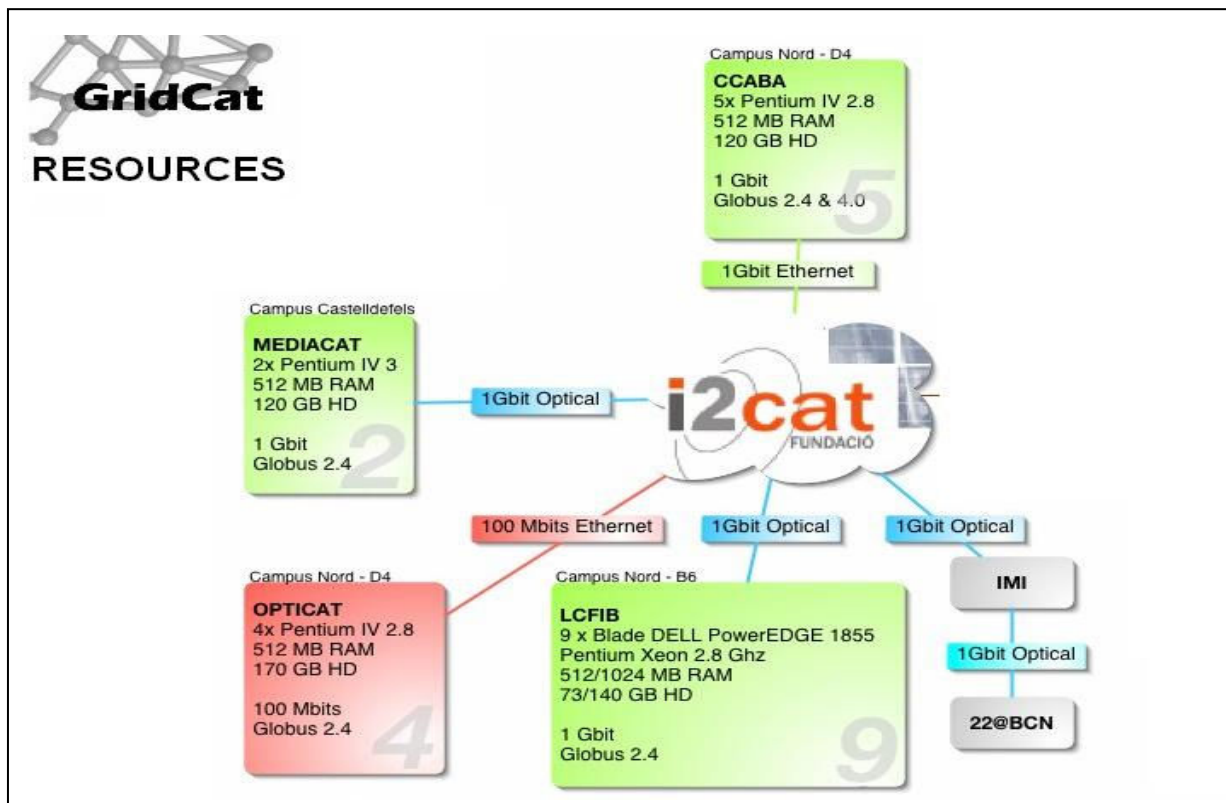


Figure A.8: i2CAT grid resources

A.2.3 Nortel Requirements

A.2.3.1 Data Plane Connectivity between Local Testbeds

- **Topology of connections:** Nortel will use one 10 Gbps (OC-192) link from the Nortel NTAC lab to the CANet 4 Ottawa POP, where an STS-24c worth of bandwidth will be mapped through to a Gigabit Ethernet (GbE) service into the Broadband Applications and Demonstration Laboratory (BADLAB™) at the CRC Carling campus. Nortel is submitting a LP request form to CANARIE for this interconnectivity.
- **Technology:** OC-192 from Nortel into CANet 4 and a Gigabit Ethernet over LPs (L2VPN) from CANet 4 into the CRC.
- **Bandwidth:** The LPs across CANet 4 were requested with STS-24c capacities (i.e. 1.25 Gbps to carry GbE traffic) to connect into the CRC over a single GbE LP.
- **Time for which they will be needed:** LP will be requested from CANARIE for the period from Nov. 20, 2006 to March 31, 2007 (the latest date supported by the form) but included a note stating that the LP were needed for the duration of the PHOSPHORUS project.

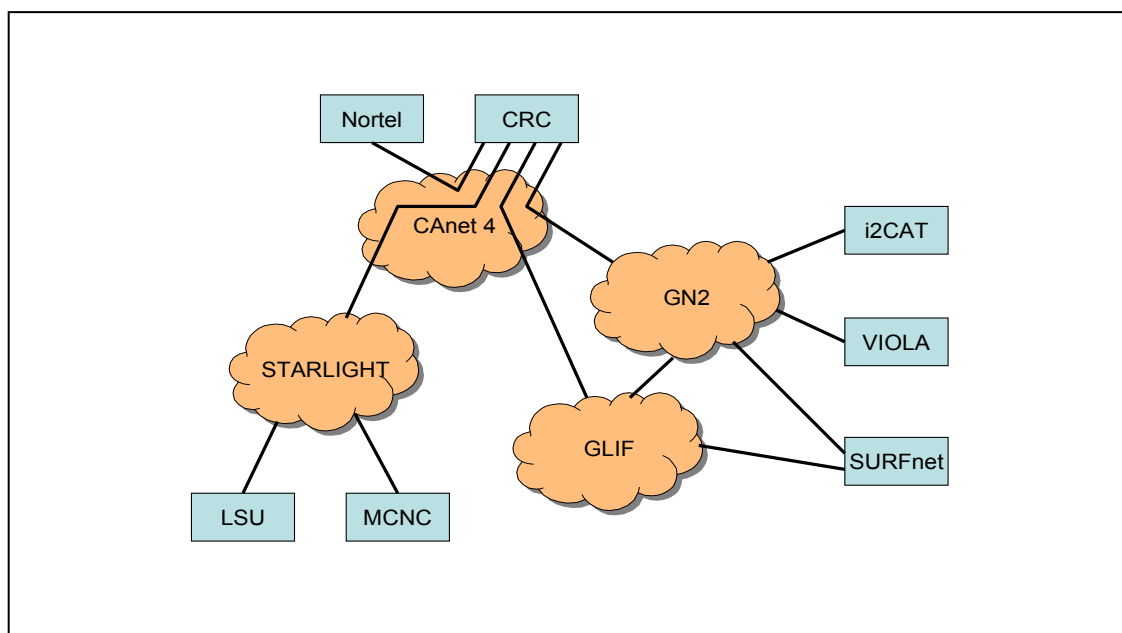


Figure A.9: Nortel Testbed Topology Proposal

NOTE: All Nortel connectivity to the testbed will be through the CRC site to reduce international connectivity costs.

A.2.3.2 Control Plane Connectivity between Local Testbeds

- Will need to be determined.

A.2.3.3 Transmission/Switching in Local Testbed

- **The topology of the internal connections (inside the local test bed):** Four Nortel OME 6500s will be used in support of the testbed. A single OME 6500 in the Nortel NTAC lab will act as a gateway into CAnet 4. This gateway network element will then be connected via an OC-192 link into a three node OME 6500 network. Each of the three OME 6500s are interconnected via two OC-192 links, and each OME is populated with a single four port EPL card holding two pluggable GbE ports (of the four available). The GbE interfaces can be connected into other equipment if needed (test sets, PC, etc). A subset of this configuration will be dedicated to the Phosphorus test bed as multiple DRAC instances will be using the network. Additional resources can be added to the test bed as needed.

Test-bed design - Appendix

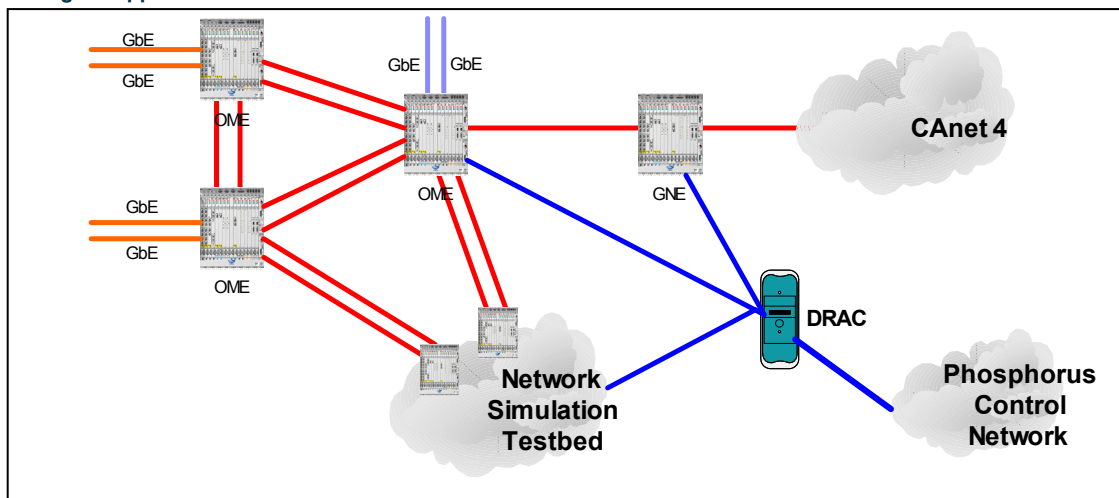


Figure A.10: Nortel DRAC testbed (simplified).

- **Simulation Network:** The Nortel test bed also has access to a simulated environment of OME 6500 and HDX/HDX-c network elements for the testing of large scale networks. Complex topologies can be modeled using the same behaviors real network elements provide. Although real traffic cannot be carried across the simulated network, manual topologies can be entered to simulate connections into the real test bed. Currently, the DRAC test bed is simulating a 60 node SURFnet network.
- **Interfaces:** Each OME 6500 is populated with a four port EPL card holding two GbE ports (of the four available – more can be added if needed). This EPL card can support 10/100/1000 BaseT, and Gigabit Ethernet (GbE) pluggable SFPs, A new version of the card will be available next year supporting 10GbE point-to-point.

A.2.3.4 Computing and storage resources in local test bed

The following resources are currently connected to the test bed. Additional resource can be provided if required:

2 x Dell PowerEdge 1800, 2 x Xeon 2.8GHz, 2 GB RAM, 2 x 80 GB RAID HDD

2 x Pentium 4 2.2GHz, 1 GB RAM, 20 GB HDD

Project:	Phosphorus
Deliverable Number:	D6.1
Date of Issue:	31/01/07
EC Contract No.:	034115
Document Code:	Phosphorus-WP6-D6.1

A.2.4 SURFNet Requirements

The SURFnet testbed will be directly connected to the GLIF network as well as the Geant+ service. This means direct GbEs can be provisioned to other testbeds when needed. As the SARA applications are running from the same facility it is very likely that applications will be able to test directly across the testbed with other testbeds.

A.2.4.1 Data-plane connectivity between local test beds:

- **Topology of connections:** The SURFnet testbed will connect directly to the Geant+ service and the GLIF network. For the WP1 work direct connections to all other networks in the work package is desirable.
- **Technology:** VC4-7v or VC4-8c point-to-point connections to the other testbeds. This will allow us to provision GbEs as well as L1 E-NNI connections.
- **Bandwidth:** the connections should allow for the usage of a full GbE end-to-end
- **Time for which they will be needed:** will be used during the full project.




Figure A.11: Global WP1 testbed connectivity – SURFnet perspective

A.2.4.2 Control-plane connectivity between local testbeds

- We can use secure tunnels over Internet/GbEANT.

A.2.4.3 Transmission/switching in local test bed

Overview: The SURFnet testbed consists of three SDH MSPPs, these are interconnected between themselves and allow direct connection of application servers (or when needed applications can be connected across the



Test-bed design - Appendix

testbed to allow for proof-of-principle testing). The network will be connected to the rest of the world with an STM-64 that can be setup to provision connections to the GLIF network or GbEANT2 as needed. Access to the testbed will be limited, although a public interface to the DRAC test server is planned for early 2007.

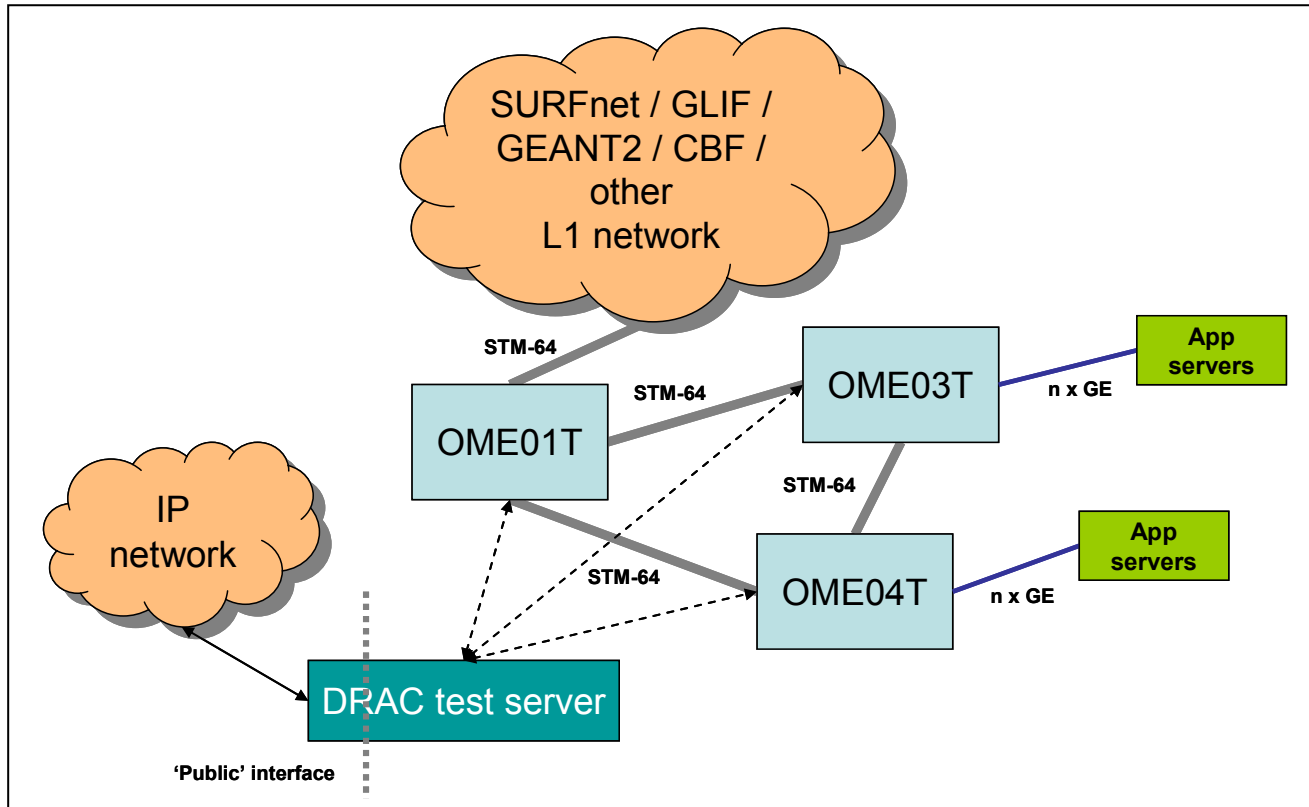


Figure A.12: View SURFnet's local testbed

- **Topology of the internal connections (inside the local test bed):** the testbed has STM-64s in its backbone, and will be connected to the outside network with an STM-64.
- **Transmission technology inside the local testbed:** SDH.
- **Interfaces:** Currently the client interfaces are GbE only. During the time of the project we envisage the possibility to interconnect locally at 10GbE using a 10G EPL card. This however would need to be mirrored at the other side to allow connections

A.2.4.4 Computing and storage resources in local test bed

We will provide a server to allow proof-of-principle testing. The assumption is that applications participating in the project will connect as clients to the testbed. However, these applications are outside immediate control of SURFnet.

A.2.5 VIOLA Requirements

A.2.5.1 Data-plane connectivity between local test beds:

- Topology of connections: VIOLA will use 1 GigabitEthernet connection from VIOLA (at Frankfurt/Main) to GbEANT. Local connection in Frankfurt is not available for now. Optional connections might be established using CBDF to SURFNet and/or PSNC.
- Technology: GigabitEthernet over lighpath (L2VPN)
- Bandwidth: 1000Mbps for each GigabitEthernet connection or 1000Mbps shared for the 4 links. This information is not available for now.
- Time for which they will be needed: used along the entire project.




Figure A.13: VIOLA test bed topology proposal

NOTE: Fraunhofer IAIS is currently negotiating with DFN-Verein (German NREN) and project partners about connectivity. Final topology depends on outcome of these negotiations.

A.2.5.2 Control-plane connectivity between local testbeds

- We can use IPsec tunnels over Internet/GbEANT.

A.2.5.3 Transmission/switching in local testbeds

- Transmission technology inside the local testbed: SDH/GMPLS, GigabitEthernet with GFP/LCAS, L2 Services (VPLS/VLL).
- Interfaces: Riverstone interfaces are all 850nm.
- Switching (number of switches/interfaces, topology of connections, control plane requirements): please see Figure A.14

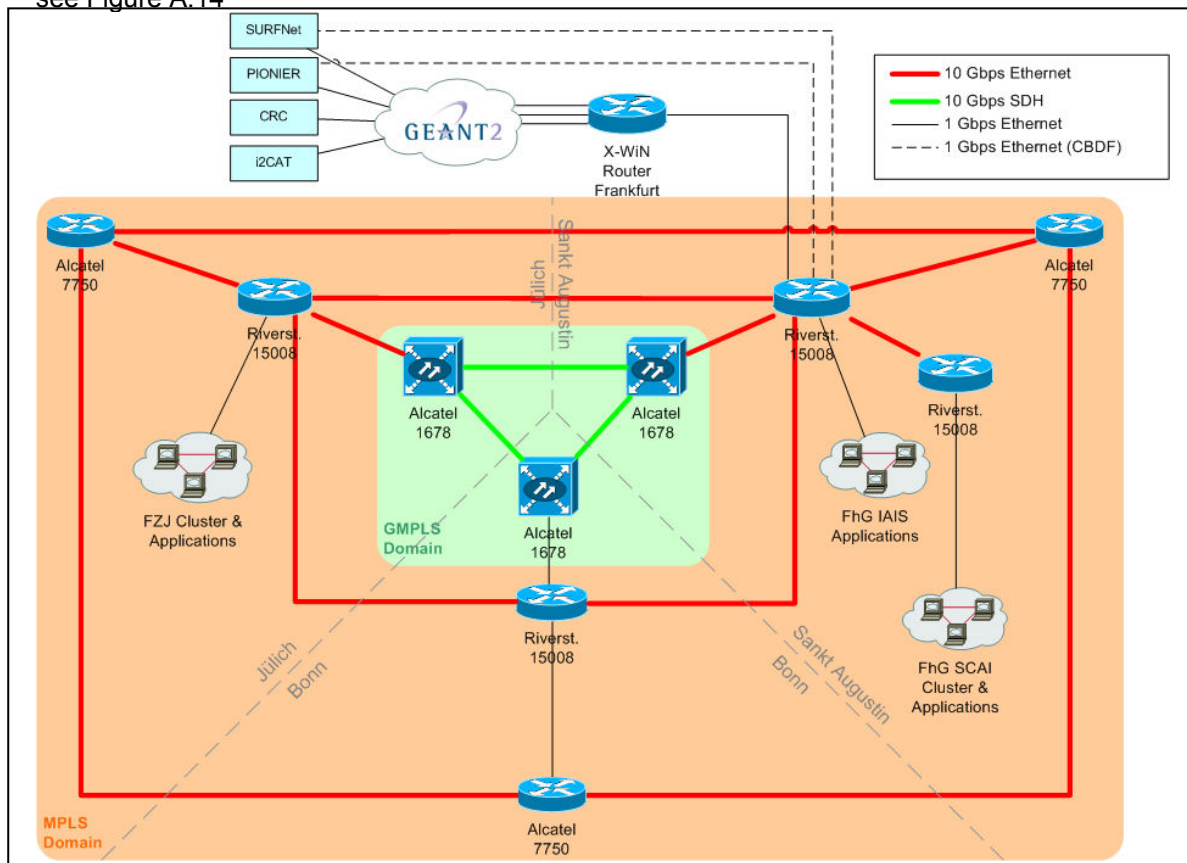


Figure A.14: Detailed view of the VIOLA testbed (simplified)

A.2.5.4 Computing and storage resources in local testbeds

Forschungszentrum Jülich (FZJ)

# Nodes	14 x 2-CPU for VIOLA and Phosphorus
CPU	Opteron, 2.2 GHz
Memory per node	4 GB per Node (2 CPU)
Home-Data	1.1 TB (global via NFS)
Temp-Data	8 GB per node local (/tmp)

Fraunhofer Institute SCAI

# Nodes	17 x 2-CPU SMP + Login/Service-nodes (2 x SMP) +File/server-nodes(2 x SMP)
---------	---

Project: Phosphorus
 Deliverable Number: D6.1
 Date of Issue: 31/01/07
 EC Contract No.: 034115
 Document Code: Phosphorus-WP6-D6.1



Test-bed design - Appendix

CPU	3,0 Ghz
Memory / node	2 GB ECC-RAM
Home-Data	320 GB
Temp-Data	80 GB

# Nodes	6
CPU	2 x Ultrasparc-IIIi 1,1 Ghz SUN Fire V240
Memory per node	2 GB RAM
Temp-Data	35 GB per node local (/space)

Project:	Phosphorus
Deliverable Number:	D6.1
Date of Issue:	31/01/07
EC Contract No.:	034115
Document Code:	Phosphorus-WP6-D6.1



Appendix B Requirements of Workpackage 2

WP2 will develop some adaptations of the GMPLS Control Plane to support Grid Network Services (G2MPLS).

This section summarizes the requirements of G2MPLS Network Control Plane (NCP) towards the HOSPHORUS testbed, by distinguishing requirements on hardware and software components, on connectivity for the Signalling Control Network and finally on Transport Plane. The requirements here identified derive from deliverable D2.1 "The Grid-GMPLS Control Plane architecture. Most of them could be shared with any standard GMPLS NCP instantiation, because of the numerous commonalities between G2MPLS and GMPLS. In fact, G2MPLS will be built as set of extensions to the standard GMPLS architecture in order to cope with the needs for enhanced network and Grid services by network "power" users/applications (i.e. the Grids) in addition to the needs of those standard users that just require the automatic setup and resiliency of their connections across the transport network.

B.1 G²MPLS Controller requirements

A G²MPLS controller is intended to provide the Control Plane instantiation of a G²MPLS network element.

This component is a network element participating in a G²MPLS domain. G²MPLS network elements can be at the user-network boundary of a domain (G.OUNI), in the core (G.I-NNI) or on the network-network boundary (the G.E-NNI). It is assumed that each transport network equipment (TNE) will be controlled by a corresponding G²MPLS controller in the G²MPLS topology.

The prototypes implementing the G²MPLS Controllers will be based on desktop PCs with the following hardware and software requirements.

Requirement	Description
Hardware	CPU Intel Architecture 32 bit
	RAM 512MB (min)
	HD TBD
	Network 1 Fast Ethernet card 10/100 Mbps (min)
Software	G ² MPLS is planned to be compiled and run mainly on GNU/Linux platforms. Support for BSD or Solaris platforms is under investigation.



	G ² MPLS software is planned to be delivered with a reference distribution from the open source community including all the needed applications and libraries (e.g. gcc, glibc, net-snmp, etc.).
--	---

B.2 Signalling Control Network (SCN) requirements

The G²MPLS controller will expose a number of interfaces towards peering controllers (i.e. East-West bound interfaces), management entities (i.e. Northbound interface, e.g. towards NRPS-es in some use cases), controlled TNE-s (Southbound interface). G²MPLS protocol messages will be carried through the set of DCN interfaces configured for being part of the Signalling Communication Network (SCN) [ASON-DCN].

The G²MPLS will rely on the implementation of control channel separation, i.e. the SCN to be used by G²MPLS will be a completely separate and independent network with respect to the transport one. This will allow the possibility to design different topologies, to use different infrastructure, etc.

As per IETF RFC 3945, protocol communications between two G²MPLS controllers will be possible through the SCN if an adjacency will be established between them. The existence of such adjacency implies that an IP path¹ must exist between the two communicating entities. This can be accomplished in different ways:

- IP tunnelling,
- GRE tunnelling,
- A TE-link with interface switching capability of PSC,
- A bidirectional LSP with interface switching capability of PSC.

Just the first two solutions will be considered for demonstration in the PHOSPHORUS testbed. The choice of the underlying technology for implementing this communication means is completely opaque to the G²MPLS control plane and WP2, and thus no specific requirement can be raised. For example, in the local testbeds the SCN could be implemented through a dedicated switched Ethernet infrastructure used by the IP stacks on the prototypal G²MPLS controllers, while for the interconnection of remote testbeds some mechanisms for tunnelling through the available shared connections might be needed (e.g. IPsec tunnels, some VPN solutions, etc.).

B.3 Transport Plane requirements

The G²MPLS NCP developed by the WP2 will be used to control and configure connections of the underlying TNEs being part of the testbed.

¹ Since GMPLS and consequently G²MPLS are designed natively for IP networks, the SCN is expected to support IP at Layer 3.



Test-bed design - Appendix

The interfacing of a G²MPLS controller to its respective TNE will be implemented through the southbound interface, which will provide the technology specific discovery, updates and configurations for transport resources. For this purposes, it is needed that the controlled TNEs expose this information via some interfaces (e.g. SNMP, CLI, TL1, etc.) in order to allow G²MPLS operations.

B.4 Configuration requirements

Along with the requirements on infrastructures, G²MPLS demonstration relies also on a number of configurations that must be under the control of the Network Operator, i.e. the administrator of the local testbed. In the following table a preliminary list of these Traffic Engineering parameters is provided for reference. The final list of the configuration parameters of the G²MPLS controllers will be delivered along with the G²MPLS prototypes at the planned dates.

Parameter	Description	Origin	Values	TE LSA Status	TE LSA
Router Address	specifies a stable IP address of the advertising router that is always reachable if there is any connectivity to it; this is typically implemented as a "loopback address". This is known as the "router ID"	RFC3630	IP address	Optional	Exactly once
Link type	defines the type of the link	RFC3631	1 - Point-to-point, 2 - Multi-access	Mandatory	Exactly once
Local interface IP address	specifies the IP address(es) of the interface corresponding to this link. If there are multiple local addresses on the link, they are all listed in this sub-TLV	RFC3633	IP address	Optional	Exactly once
Remote interface IP address	specifies the IP address(es) of the neighbour's interface corresponding to this link. If the Link Type of the link is Multi-access, the Remote Interface IP Address is set to 0.0.0.0; alternatively, an implementation MAY choose not to send this sub-TLV.	RFC3634	IP address	Optional	Exactly once
Traffic engineering metric	specifies the link metric for traffic engineering purposes. This metric may be different than the standard OSPF link metric.	RFC3635	32-bit value	Optional	Exactly once



Test-bed design - Appendix

Parameter	Description	Origin	Values	TE LSA Status	TE LSA
Maximum reservable bandwidth	specifies the maximum bandwidth that may be reserved on this link, in this direction (units are bytes per second). This may be greater than the maximum bandwidth (oversubscription).	RFC3637	32-bit IEEE Floating Point format	Optional	Exactly once
Admin. group	bit mask (Resource Class/Colour) in which each set bit corresponds to one administrative group assigned to the interface. A link may belong to multiple groups.	RFC3639	32-bit value	Optional	Exactly once
Link Local/Remote Identifiers	specifies the link identifiers (local and remote) in case of unnumbered links. If the Remote Identifier is unknown, it is set to 0.	RFC4203	2 * 32-bit value	Optional	Exactly once
Link Protection Type	represents the protection capability that exists for a link	RFC4203	0x01 Extra Traffic, 0x02 Unprotected, 0x04 Shared, 0x08 Dedicated 1:1, 0x10 Dedicated 1+1, 0x20 Enhanced, 0x40 Reserved, 0x80 Reserved	Optional	Exactly once
Max LSP Bandwidth [prio]	specifies the smaller of the unreserved bandwidth at priority p	RFC4203	8 * 32-bit IEEE Floating Point format	Optional	More than once
Min LSP Bandwidth	specifies the minimum bandwidth an LSP could reserve -- PSC & TDM ONLY	RFC4203	32-bit IEEE Floating Point format	Optional	More than once
Shared Risk Link Group	unordered list of 32 bit numbers that are the SRLGs that the link belongs to.	RFC4203	N * 32-bit values	Optional	Exactly once
BER estimate	specifies the exponent from the BER representation	RFC4209	32-bit value	Optional	Exactly once
Span length	represents the total length of the WDM span in meters	RFC4209	32-bit value	Optional	Exactly once
OSNR	specifies the OSNR measurement/estimation for the optical link	N/A	32-bit value	Optional	Exactly once

Project: Phosphorus
 Deliverable Number: D6.1
 Date of Issue: 31/01/07
 EC Contract No.: 034115
 Document Code: Phosphorus-WP6-D6.1



Appendix C Requirements of Workpackage 3

The requirements of the applications regarding connectivity between participating sites and testbeds, the hardware, middleware, tools and applications have been acquired and described in Deliverable D3.1 “Use-cases, requirements and design changes of the applications and middleware”. This appendix summarizes these requirements with a focus on the testbed:

- What are the connectivity requirements between the local testbeds?
- Which hardware resources have to be provided in the local testbeds to support the applications?
- Which middleware and which other software tools have to be provided on the above mentioned hardware resources?

C.1 Network requirements

In the initial set-up, the applications will be deployed and tested at the partner sites that participate in the respective application as described in the technical annex. In order to achieve a broader involvement of the local testbeds, a “call for participation” has been issued. The table of participating local testbeds includes both the partners listed in the Technical Annex and the partners that have already responded to the call for participation. Therefore it is likely that the number of involved local testbeds will grow during the course of the project.

Application	Participating local testbeds	Bandwidth / latency requirements	Other network requirements
WISDOM	VIOLA, PSNC, UESSEX, (CNRS/IN2P3 in France, others later)	No constraints	None
KoDaVis	VIOLA, UoA, UESSEX, PSNC	700 Mbit/s / 50 msec from all testbeds to	IP-Multicast, Jumbo frames preferred



Test-bed design - Appendix

		VIOLA	
TOPS	VIOLA, SARA, later Barcelona, Prague	1 ... 10 Gbit/s / 100 msec from all testbeds to SARA	Jumbo frames required
DDSS	PSNC, VIOLA, UESSEX	500 Mbit/s / -	Jumbo frames preferred
INCA	To be determined	No constraints	Jumbo frames preferred

c.2 Local testbed hardware and software requirements

The hardware and software requirements in the following table have also been collected from the application questionnaire (see D3.1) and the call for participation. The application partners have made sure that these requirements can be fulfilled by the participating sites.

The applications that use UNICORE as underlying Grid middleware (WISDOM, KoDaVis) will make use of the advance reservation and co-allocations capabilities of the Phosphorus testbed. Therefore, a local resource management system (RMS) with advance reservation capabilities for the compute resources has to be available at the participating sites. For DDSS, which relies on Globus Toolkit, and TOPS, which does not rely on a particular Grid middleware up to now, there are no such requirements.

Application	Participating local testbeds	Hardware-Platform	Middleware, application and firewall requirements
WISDOM	VIOLA (data server, results database, workflow execution)	Linux-Cluster, min. 500 MB main memory, 30 GB shared disk space, 1 GbE Interfaces	UNICORE, FlexX/Autodock, access to FLEXIm- and Oracle servers needed, RMS with advance reservation capability, (e.g. EASY, Torque+MAUI, PBSpro) firewall: 2 open TCP-ports (for UNICORE and MSS)

Project:	Phosphorus
Deliverable Number:	D6.1
Date of Issue:	31/01/07
EC Contract No.:	034115
Document Code:	Phosphorus-WP6-D6.1



Test-bed design - Appendix

	PSNC, UESSEX, (CNRS/IN2P3 in France, others later) (workflow execution)	Linux/UNIX system, 1 GbE Interface	UNICORE, FlexX/Autodock, access to FLEXIm- and Oracle servers needed, C/C++ compiler, MPI, RMS with advance reservation capability, (e.g. EASY, Torque+MAUI, PBSpro) firewall: 2 open TCP-ports (for UNICORE and MSS)
KoDaVis	VIOLA (data server)	Linux-Cluster system, 1 GbE Interfaces	UNICORE, C/C++ compiler, RMS with advance reservation capability, (e.g. EASY, Torque+MAUI, PBSpro) firewall: 3 open TCP-ports (for UNICORE, MSS and data streaming)
	UoA, UESSEX, PSNC (client)	Linux/UNIX system, 1 GbE Interface	UNICORE, C/C++ compiler, AVS/Express would be useful, RMS with advance reservation capability, (e.g. EASY, Torque+MAUI, PBSpro) useful firewall: 3 open TCP-ports (for UNICORE, MSS and data streaming)
TOPS	SARA (2/3D data source)	Linux-Cluster, 20 GB disk space, nvidia GPU in each node (for 3D), 1 or 10 GbE Interfaces	Various Linux software packages: swig, python, libz-dev, libtiff4-dev, sdl-dev, libxxf86dga-dev, firewall: some open ports, ssh access
	VIOLA, later Barcelona, Prague (visualization end point)	Tiled display panel, 1 or 10 GbE Interface	
DDSS	PSNC, VIOLA (GridFTP- and TSM- server + client)	Linux-system, dual- CPU, 512 MB main memory, 400 GB disk space, 1 GbE Interface(s) SAN or FC-disk or tape system (for TSM)	Globus Toolkit 4.0, TSM-server + client 5.2 or higher, firewall: 50-100 high ports open (GridFTP)

Project: Phosphorus
 Deliverable Number: D6.1
 Date of Issue: 31/01/07
 EC Contract No.: 034115
 Document Code: Phosphorus-WP6-D6.1



Test-bed design - Appendix

	UESSEX (GridFTP- and TSM- client)	Linux-system, single or dual-CPU, 512 MB main memory, 100 GB disk space, 1 GbE Interface(s)	Globus Toolkit 4.0, TSM client 5.2 or higher, firewall: 50-100 high ports open (GridFTP)
INCA	To be determined (data nodes)	Linux-PCs, 4 TB disk space for data nodes, 1 GbE Interface(s)	C/C++ compiler, Firewall: some ports to be opened
	To be determined (client nodes)	Linux/UNIX-PCs, 1 GbE Interface(s)	C/C++ compiler, Firewall: some ports to be opened



Appendix D **VPN-solutions for realizing a network for Control Plane traffic in Phosphorus**

Carsten Rosche, Jochen Schon

Fraunhofer IAIS

D.1 Requirements for the Control Plane Data Communication Network (DCN)

Within the Phosphorus project, the network control plane will be the major topic of research. Typically, this control plane will consist of several network elements and workstations running services. Therefore, these devices need to be able to exchange protocol messages, especially RSVP and OSPF, but possibly also others.

Each partners' local testbed might already have a control plane. Most probably, this control plane will be realized through an Ethernet based LAN, which is called Data Communication Network (DCN). For the control plane research in Phosphorus, these local DCNs have to be interconnected. This means a connection between different Ethernet segments has to be established. This could either be realized by deploying routers which will forward messages between the different sites or by bridging the segments with Hubs or switches.

As the different DCNs are geographically located at various places, there is no point where all networks could physically be stitched together. Therefore, an approach is needed, which allows to connect the DCNs via the public Internet and also protects the DCNs from unauthorized access through the Internet by encrypting the traffic. A solution for this are Virtual Private Networks (VPNs).

Basically, VPNs can be set up through Routers. But as there is no funding for network hardware in Phosphorus, and cheap routers are normally not capable of setting up encrypted VPNs, this does not seem very feasible. Another solution is a software based approach. In this case, an ordinary workstation can run a software which will spawn the VPN. As there will be workstations needed to implement the control plane feature developments of Phosphorus, this could even be done without a dedicated machine, but of course, it might still be a dedicated one.

Project:	Phosphorus
Deliverable Number:	D6.1
Date of Issue:	31/01/07
EC Contract No.:	034115
Document Code:	Phosphorus-WP6-D6.1



Test-bed design - Appendix

Besides security, robustness is also an issue. Software VPNs are normally realized either as client/server or peer-to-peer architecture. While both architectures will offer the needed functionalities, the client/server approach might lead to a single point of failure at the server.

Finally, the software should be available for different platforms, at least Windows and Linux, but possibly also others. It should also not be too difficult to be set up and maintained.

D.2 Tested solutions

With regard to the requirements mentioned above, some software based VPN solutions have been investigated. The advantages and disadvantages will be described briefly in this chapter.

D.2.1 Hamachi²

Hamachi is a very popular VPN solution. It was originally designed to be able to play games over the Internet, which were only meant to be played over a LAN. Of course, as it is offering a VPN, it could also be used for other purposes. It is available for several platforms and offers very easy installation and configuration.

Unfortunately it turned out to not support Ethernet bridging which renders it useless for the purposes of Phosphorus.

D.2.2 OpenVPN³

OpenVPN is an open source VPN solution licensed under the GPL. It is also very popular, available for different platforms (Linux, Solaris, Windows, Mac OS X and some BSDs) and is not too hard to configure. It allows for Ethernet bridging with encrypted traffic and IPv4 based routing.

The major drawback for this solution is that it is based on a client/server architecture and therefore very vulnerable to server failures. Even though, traffic can be forwarded directly between clients without passing the server, the server is still needed for the VPN to work.

D.2.3 tinc⁴

Tinc is another GPL based free software VPN. It runs on several systems (currently Linux, FreeBSD, OpenBSD, NetBSD, MacOS/X, Solaris, Windows), is quite easy to configure and works with encrypted traffic. It offers

² <http://www.hamachi.cc/>

³ <http://openvpn.net/>

⁴ <http://www.tinc-vpn.org/>

Project:	Phosphorus
Deliverable Number:	D6.1
Date of Issue:	31/01/07
EC Contract No.:	034115
Document Code:	Phosphorus-WP6-D6.1



Test-bed design - Appendix

a routed mode, a switched mode with MAC-learning and a Hub-mode where all packets are forwarded to all destinations. Changing between the different modes can easily be done.

Tinc works as a peer-to-peer model. Therefore, if a node fails, the VPN will still continue to work. The mesh between the nodes can be freely designed.

D.3 Conclusion

Hardware based VPNs require compatible routers at each partner and are therefore expensive. A software based approach could minimize the costs and possibly offer even more flexibility and functionalities.

Three freely available software VPNs have been tested. The Hamachi VPN unfortunately does not offer bridging of Ethernet segments which is needed for Phosphorus. Therefore it cannot be used.

OpenVPN and tinc both offer very similar functionalities. The installation and configuration is not as easy as with Hamachi, but still not very difficult. Both offer encrypted communication and can be run in a routed or switched mode. Furthermore, tinc offers a virtual Hub mode, where all traffic can be seen at every client. This could be useful for protocol analysis in some cases.

The main difference between tinc and OpenVPN is their architectural model. Tinc is implemented as a peer-to-peer-model with freely configurable mesh in between, while OpenVPN is a client/server model which makes it vulnerable to server failures.

The additional features of tinc (namely the Hub mode and enhanced redundancy through the peer-to-peer-model) made tinc the most valuable and feasible solution for Phosphorus in this test.

Project:	Phosphorus
Deliverable Number:	D6.1
Date of Issue:	31/01/07
EC Contract No.:	034115
Document Code:	Phosphorus-WP6-D6.1



Disclaimer

The Phosphorus project is funded by the European Commission under the FP6 contract no. 034115. This document contains material which is the copyright of Phosphorus contractors and the EC, and may not be reproduced or copied without permission. The information herein does not express the opinion of the EC. The EC is not responsible for any use that might be made of data appearing herein. The information in this document is provided as is and no guarantee or warranty is given that the information is fit for any particular purpose. The user thereof uses the information at its sole risk and liability.

Project:	Phosphorus
Deliverable Number:	D6.1
Date of Issue:	31/01/07
EC Contract No.:	034115
Document Code:	Phosphorus-WP6-D6.1