

Advance Reservation Frameworks in Hybrid IP-WDM Networks

Neal Charbonneau and Vinod M. Vokkarane, University of Massachusetts, Dartmouth
Chin Guok and Inder Monga, Lawrence Berkeley National Laboratory/ESnet

ABSTRACT

New e-Science and grid applications require the coordination of geographically distributed scientific instruments along with data and computing resources. Due to the quality of service requirements of these applications, these distributed resources can be connected by a wavelength-routed optical network, allowing each application to get dedicated bandwidth. These networks are referred to as LambdaGrids. One important service provided in these networks is advance reservation. Applications need to coordinate the use of both grid resources and the network. Advance reservation allows these applications to reserve bandwidth in advance to guarantee availability. In this article, we discuss different networks and frameworks that support advance reservation of bandwidth. We discuss the general architecture of each network and the type of advance reservation services supported.

INTRODUCTION

A grid network is a collection of geographically distributed resources, such as storage clusters, supercomputers, and scientific equipment, that are accessible to users over a network. Examples of e-Science grids include the Large Hadron Collider Computing Grid Project, the Biomedical Informatics Research Network, and the George E. Brown Network for Earthquake Engineering and Simulation. These networks typically deal with the transfer of large amounts of data in the terabytes and petabytes range. When the grid resources are connected by application-configurable optical paths (e.g., on an optical wavelength-routed network), the grid can be considered a LambdaGrid [1].

Optical wavelength-routed wavelength-division multiplexing (WDM) networks or optical circuit switched (OCS) networks are being used to interconnect computing, storage, and instruments in grid networks. An optical WDM network consists of fibers connected by switches, or optical crossconnects (OXCs). In WDM

networks, each fiber is partitioned into a number of wavelengths, each of which is capable of transmitting data at very high speeds. This allows each fiber to support data transmission rates of terabits per second. In order to transmit data over the network, a dedicated circuit is first established when a user submits a connection request. When a connection request arrives at the network, the request must be routed over the physical topology and also assigned a wavelength. The combination of a route and wavelength is known as a lightpath. Typically, the objective is to minimize the number of wavelengths used in the network in order to handle a given static traffic demand. In the dynamic traffic scenario, the objective is to minimize the cost of each lightpath request based on different cost metrics, such as path length and network load. This is known as the routing and wavelength assignment (RWA) problem. The RWA problem is NP-complete; therefore, heuristics are typically used to solve the problem. The signal may be transmitted all-optically (uses the same wavelength on every hop along the path) or converted to electronics and back to optics (O/E/O conversion) at multiple hops.

LambdaGrid networks typically support dynamic traffic. Dynamic traffic requests arrive one by one according to some stochastic process, and they are also released after some finite amount of time. The goal of the provisioning system of the network is to minimize request blocking. A request is said to be blocked if there are not enough resources available to route it. There is extensive work dealing with these problems (e.g., [2]).

We can further classify the dynamic traffic model as immediate reservation (IR) or advance reservation (AR) [3] requests. The data transmission of an IR demand starts immediately upon arrival of the request, and the holding time is typically unknown for dynamic traffic or assumed to be open ended for static traffic. AR demands, in contrast, typically specify a data transmission start time that is sometime in the future and also specify a finite holding time. Figure 1 shows the difference between AR and IR

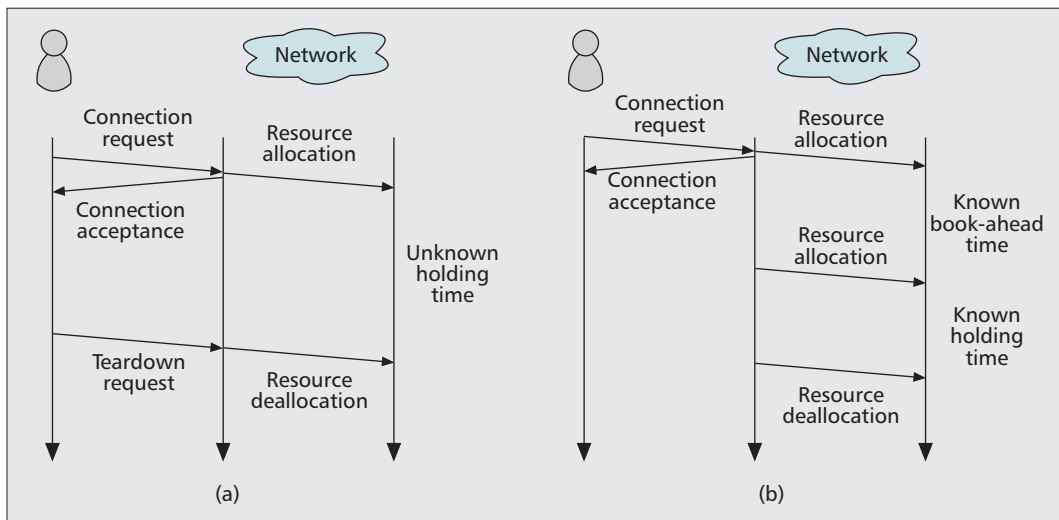


Figure 1. The request and allocation of resources for immediate and advance reservation requests. In the figure, we assume that the requested resources are available. Before reservation/allocation the network must find an appropriate lightpath. For immediate reservation (a), the allocation is at the same time as the request arrival and the duration is typically unknown. For advance reservation (b) the allocation is some (known) time after the arrival and the duration is also known.

requests. We can observe from Fig. 1a that in IR the resource allocation occurs when the request arrives at the network. In this case, the duration of the request is unknown. In AR (Fig. 1b), the actual allocation of resources does not occur until a later time. The resources are reserved when the request arrives, but they can be used by other requests before the reservation time. The difference between the arrival of the request and the beginning of transmission is the *book-ahead time*, which is specified by the request. The duration of the request is also specified in advance and known by the network. The fact that holding time and book-ahead time are known allows the network to efficiently optimize resource usage.

Advance reservation was initially proposed for non-optical networks, focusing on circuit-switched, packet-switched, and asynchronous transfer mode (ATM) networks. Initial work focused on traffic modeling and call admission for telecommunication systems. Additional work was proposed to provide quality of service (QoS) for multimedia applications like videoconferencing. There were also extensions for Resource Reservation Protocol (RSVP) to support advance reservation. Advance reservation for optical networks was first proposed by Zheng and Mouftah in [3].

There are a number of reasons why it is beneficial to provide advance reservation services for grid applications. With user driven traffic, often bandwidth requirements and request durations are known in advance due to requests being for specific tasks. Advance reservation requests allow applications to ensure that network resources are available when certain computing resources are also available, hence enabling user access to certain grid resources at specified times in the future. In order to access these resources, the user must be able to receive guarantees of network availability. This is known as resource co-allocation.

Also, many grid applications involve delay-tolerant background or recurring tasks. For example, once a scientific instrument finishes an experiment, the data set usually must be transferred to other sites over the grid. Instead of issuing these transfers as immediate reservation requests, the user can submit them as advance reservation requests that specify a deadline or window in which the transfer must take place. By providing advance reservation for such tasks, the network can achieve higher utilization while increasing the probability that the grid applications will be able to successfully reserve the required network resources.

Collaboration is an important part of large-scale scientific computing. Advance reservation can support real-time collaboration through real-time experimentation or high-definition video conferencing. It is easier to allocate resources to advance reservation requests with booking-ahead instead of immediate reservation requests. In this article we discuss a number of real-world networks that support advance reservation over optical networks. There is a significant amount of work dealing with advance reservation of network resources in general, but this article will focus on actual implementations.

ADVANCE RESERVATION FRAMEWORKS

We now discuss the various advance reservation frameworks and architectures that have been implemented. Some projects support advance reservation across a single domain, whereas other projects focus on providing multidomain support as well. There is also work on co-allocation of network and grid resources (e.g., computing/storage) using advance reservation. Table 1 provides a comparison of the different frameworks discussed below.

Collaboration is an important part of large-scale scientific computing. Advance reservation can support real-time collaboration through real-time experimentation or high-definition video conferencing.

Framework	Region and Funding Body	Bandwidth Provision	Provisioning Layer	Network Resource Provisioning System	Grid Co-Scheduling Capabilities	Single-/Multi-Domain
OSCARS	U.S.A. (U.S. Department of Energy)	Hybrid: packet/circuit	Layer 2 & 3	Integrated (MPLS-based)	No	Single (Multidomain enabled with IDC)
DRAC	Europe (SurfNet)	Hybrid: packet/circuit	Layer 1 & 2	Integrated	No	Single
EnLIGHTened	U.S.A. (U.S. NSF)	Circuit	Layer 1	Integrated (GMPLS)	Yes	Multi-domain
G-Lambda	Japan (KDDI R&D, NTT, NICT and AIST)	Circuit	Layer 1 & 2	Integrated (GMPLS)	Yes	Multi-domain
PHOSPHORUS	Europe (E.C. FP6)	Circuit	Layer 1 & 2	ARGON, DRAC and UCLP	Yes	Multi-domain

Table 1. Comparison of advance reservation frameworks.

ON-DEMAND SECURE CIRCUITS AND ADVANCE RESERVATION SYSTEM

The On-Demand Secure Circuits and Advance Reservation System (OSCARS) is a project funded by the U.S. Department of Energy (DOE) that enables dynamic provisioning and advance scheduling of network resources within a single domain or over multiple domains between end user/customer edge devices. The basic architecture of OSCARS is shown in Fig. 2a. OSCARS is deployed in the DOE's Energy Sciences Network (ESnet) as a production service to deliver layer 2/3 virtual circuits (VCs) to users. OSCARS is implemented as a centralized service that provides a web services programmatic application programming interface (API) for peer interdomain (IDCs) and middleware clients, as well as a web-based user interface (WBUI) for users to make advance reservations (Fig. 2b). All network service requests to OSCARS require the type of service (i.e., layer 2 or 3), termination points of the VC, minimum guaranteed bandwidth, and duration of use. For advanced users, the option to specify the path via an explicit route object (ERO) is also permitted. OSCARS (version 0.6) consists of 11 distinct pluggable modules; a WBUI, a programmatic API, a lookup bridge (tied to an external lookup service), a topology bridge (tied to an external topology service), authentication, authorization, a coordinator, a notification broker, a resource manager, a path computation engine, and a path setup module. A request for network services (i.e., network resources) is first authenticated, and subsequently processed through the path computation engine (PCE) module where one or more atomic PCEs can prune the service topology based on specific constraints, such as bandwidth availability, latency, and authorization. The resulting service topology solution is subsequently stored in the resource manager as a confirmed reservation. Just prior to the reservation becoming active, the path setup module contacts the appropriate network elements and initiates the RSVP signaling to set up the VCs. OSCARS is designed to

manage resources within a single administrative domain, but supports provisioning of interdomain virtual circuits using the DICE IDC protocol¹ to communicate with other deployed IDC-compatible instances of OSCARS (e.g., Internet2's ION), or in an interoperable manner with other provisioning systems like Auto-BAHN.

DYNAMIC RESOURCE ALLOCATION CONTROLLER

The Dynamic Resource Allocation Controller (DRAC) [4], developed by Nortel (now part of Ciena), enables automated and dynamic provisioning of network resources to meet applications' network requirements. Applications leverage DRAC to request and control their share of network resources, brokering and configuring on an end-to-end basis all the necessary pieces of the network, regardless of the type of network — circuit or packet, wireless or wireline. A key application that DRAC enabled in SURFnet, a research and education network in the Netherlands, was “cut-through” capabilities by steering very large flows of packets or low-latency applications dynamically over layer 1 circuits instead of layer 3. For example, instead of routing high-bandwidth point-to-point traffic transfers over intermediate layer 3 hops (i.e., IP routers), which can incur additional jitter and latency due to buffering, or setting up long-term dedicated and costly high-bandwidth optical connections to address transient bandwidth demands, DRAC simply sets up and takes down “ephemeral” optical circuits as needed by the application, minute by minute, hour by hour.

The DRAC core framework (Fig. 3a) includes authentication, authorization, and accounting (AAA) services, a policy engine, a topology discovery engine, resource reservation services, workflow utilities, interdomain routing facilities, and smart bandwidth management capabilities. Its interface provides applications with the means to directly drive their share of network resources within a policy-defined enve-

¹ <http://www.controlplane.net/>

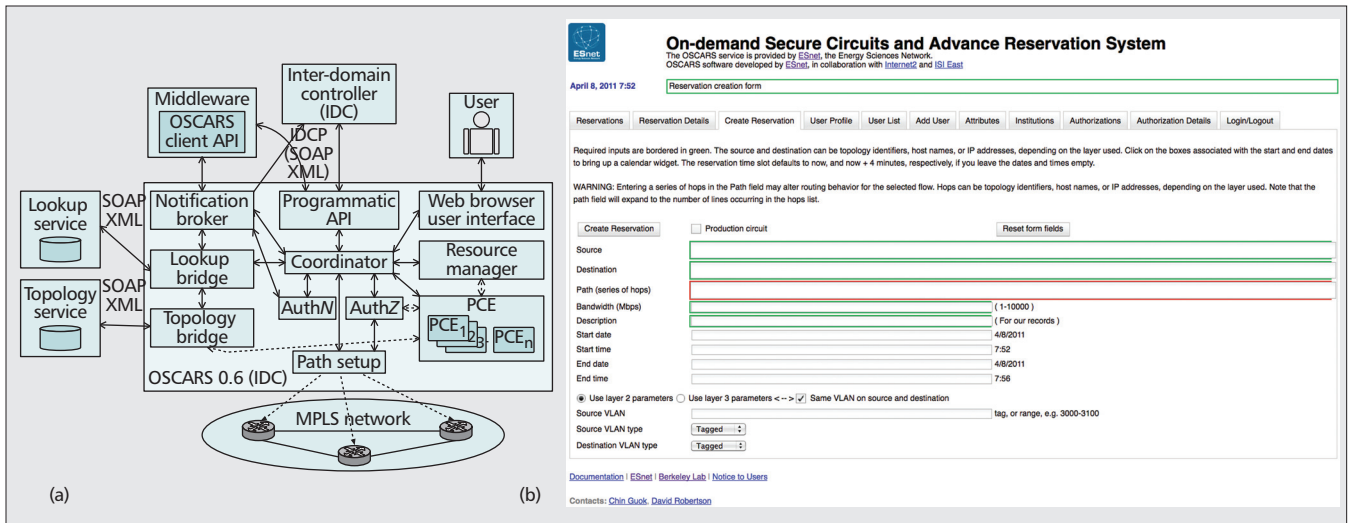


Figure 2. a) Overview of OSCARS architecture; b) screenshot of OSCARS reservation system.

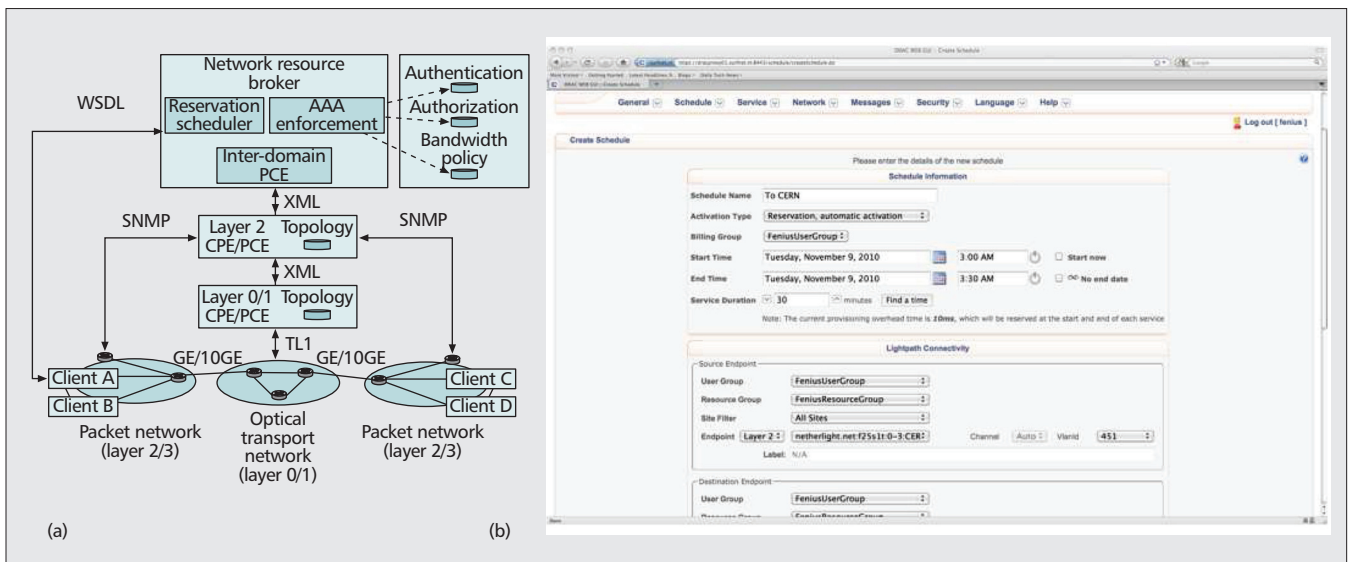


Figure 3. a) Overview of DRAC architecture; b) screenshot of DRAC reservation system.

lope of flexibility while receiving abstracted availability and performance information from the network. Network resources include bandwidth, acceleration appliances, and sensors, to name a few. The DRAC interface to the network uses existing management protocols and standards like SNMP/TL1 and toolsets for interfaces (equipment CLI's), which greatly simplifies deployment in multivendor multitechnology environments. DRAC is now available as an open source implementation called Open DRAC.²

ENLIGHTENED

EnLIGHTened (initially funded by the U.S. National Science Foundation [NSF]) is a project that focuses on advance reservation of both grid and network resources (lightpaths) [5]. The goal of the project is to allow grid applications to request in advance or on-demand compute, storage, instrumentation, and network resources. These resources have to be co-allocated and may be spread across multiple domains. The architec-

ture includes a resource broker (ERB), resource monitoring, and network scheduling (in a network domain manager). The architecture utilizes the Highly-Available Resource Co-allocator (HARC) system, which allows clients to reserve multiple distributed resources in one step. HARC consists of acceptors that manage co-allocation and resource managers, which are the interfaces used to make reservations. Custom acceptors and resource managers were implemented specifically for the project. The resource broker provides the primary interface to the client, and the ERB uses HARC internally. The architecture of the ERB is shown in Fig. 4a. The ERB accepts requests from the user, who is authenticated through the AAA module. From the user request, the broker then gets availability information from each of the local managers (e.g., the storage managers and network managers). With this information, it uses its own scheduler to find a resource set able to satisfy the request. Given a set of resources, it then uses HARC acceptors to co-allocate the required

² <https://www.opendrac.org/>

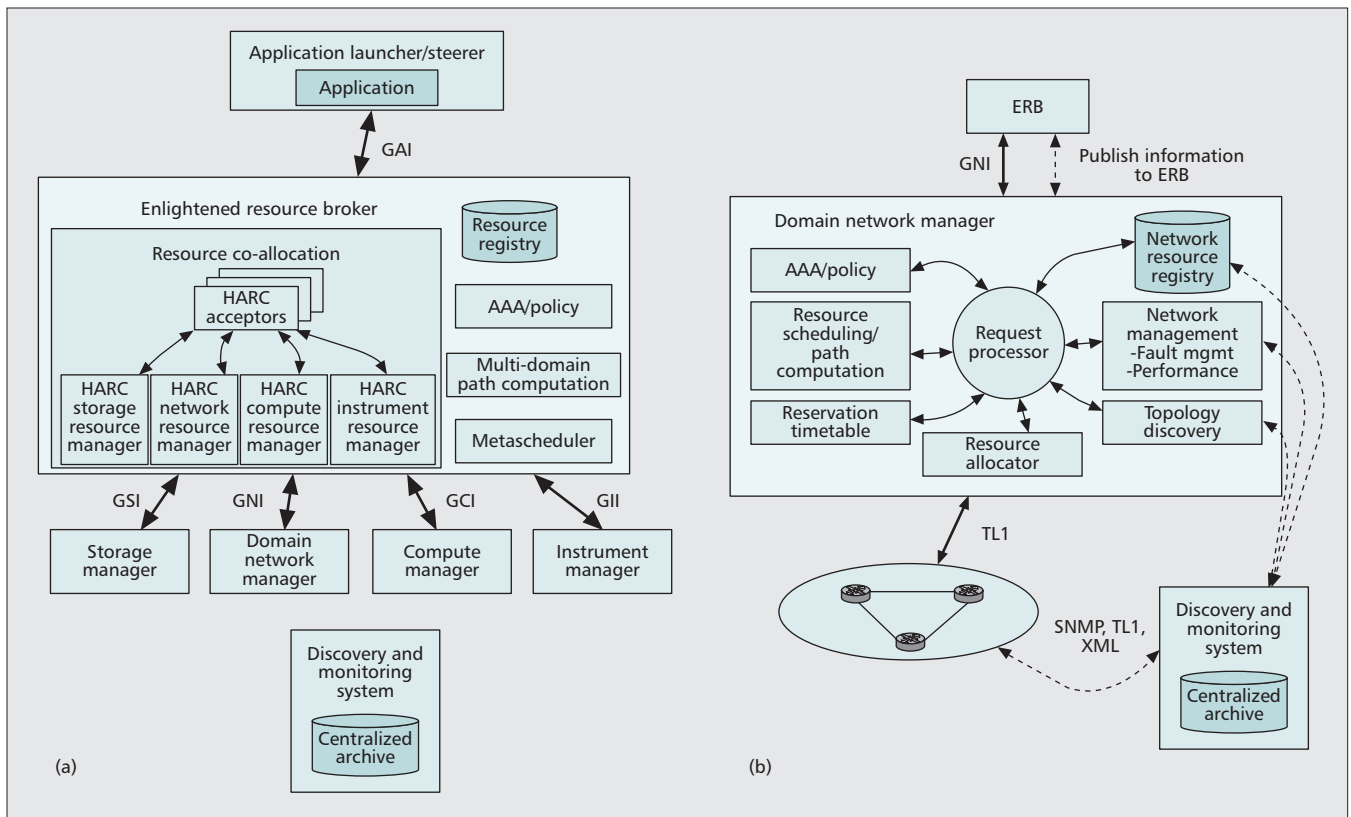


Figure 4. Overview of EnLIGHTened architecture: a) ERB architecture; b) the DNM.

resources. The acceptors in turn use the HARC resource managers to schedule the final resources with the local managers. The local managers are then responsible for activating and deactivating resources at the proper times.

The domain network manager (DNM) is responsible for controlling network resources, resource reservation and path computation, restoration, and dynamic teardown/setup of lightpaths for a single administrative domain. Each domain will have a DNM. An example of a DNM is shown in Fig. 4b.

The discovery and monitoring system (DMS) is required to support the dynamic addition or removal of resources. It is also responsible for monitoring the performance and reliability of resources in real time. There are three types of resource monitoring. At-reservation-time monitoring is used to verify lightpath establishment. During-reservation monitoring is used to assess whether QoS/service level agreement (SLA) requirements are being met. Ongoing grid resource status monitoring is used to collect relevant performance metrics from grid resources in real time. All the monitoring types use the perfSONAR framework. Details of this module can be found in [6].

Lastly, EnLIGHTened also emphasizes the importance of standardized interfaces. Interaction with all local managers is done through well defined interfaces, such as grid security infrastructure (GSI) and grid network infrastructure (GNI). The project promotes collaboration with other groups to define a GNI as well as an application interface, the grid application interface (GAI).

G-LAMBDA

Another project with similar goals was the Japanese project G-Lambda [7]. The main goal of the project was to define a standard web interface between a grid resource scheduler (similar to the ERB) and network resource management systems (similar to the DNM). Network operators could then implement this interface to allow the grid resource coordinators to make advance reservations of network resources. The grid resource coordinators are implemented as middleware, and coordinate both network and computation/storage resources at the request of grid applications.

An overview of the G-lambda architecture is shown in Fig. 5a. A detailed view of the architecture can be seen in Fig. 5b showing the modules discussed. As introduced above, the main focus on the project was to define the grid network service-web service interface (GNSWSI) shown in the figure. The GNS-WSI allows a number of commands to be sent to the network resource manager (NRM). These include commands to reserve paths, modify previously reserved paths, release paths, and query path or status information. An example architecture of an NRM is shown in Fig. 5c.

The grid resource scheduler (GRS) provides a web service interface to grid clients using the web service resource framework (WSRF). The GRS is implemented using the Globus Toolkit 4. The modules include a client API, the web service module, and a scheduling module that uses the CRMs and NRMs to co-allocate and reserve resources. The scheduling module consists of a co-allocator and planner. The planner takes the

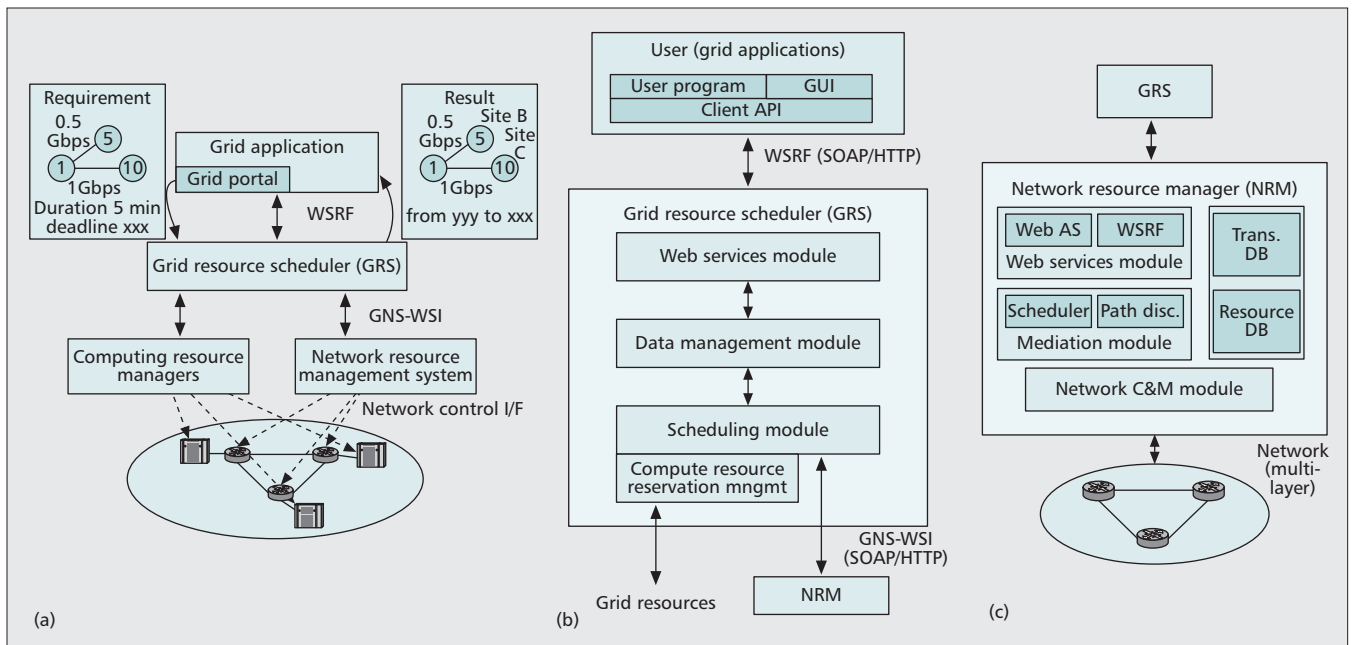


Figure 5. Overview of G-lambda architecture: a) detailed view of the G-lambda architecture; b) grid resource scheduler; c) network resource manager.

user requests and selects a candidate set of resources, which are then reserved simultaneously by the co-allocator. A two-phase commit protocol is used to reserve resources in parallel. A detailed description of the modules and protocols can be found in [8].

The CRMs are essentially wrappers around existing software like Plus and GridEngine. The project implemented two NRMs (one by KDDI and another by NTT). The NRMs are responsible for path virtualization between endpoints, local scheduling, and activation/de-activation of lightpaths. Paths are virtualized to hide implementation details. NRMs have a web service module (interface), a mediation module that virtualizes the generalized multiprotocol label switching (GMPLS) optical network and performs the scheduling based on requests from the GRS, a network control module for the GMPLS switches to manage state information. Details about NRMs can be found in [9].

EnLIGHTened and G-Lambda established a collaboration to prove network co-allocation across the two network domains. Software wrappers around components in the architectures were used so that the systems could interoperate. The technical details can be found in [10].

PHOSPHOROUS

The EU's PHOSPHOROUS project also incorporates advance reservation of grid and networking resources [11]. The goal of PHOSPHOROUS is to provide on-demand and in-advance end-to-end provisioning of network and grid resources across multiple domains and multiple vendors. The PHOSPHOROUS project comprises two phases. In the first phase each independent domain is controlled by an existing network resource provisioning system (NRPS), while in the second phase interoperability is added with other existing networks and grid resources through stan-

dardized interfaces. The NRPSs are similar to the NRM/DRM of the previous two projects.

PHOSPHOROUS defines an architecture composed of three planes. The first is the *service plane*, which consists of middleware extensions to allow applications to make advance reservations, co-allocate resources, provide AAA across multiple domains, and so on. The next layer is the *network resource provisioning plane*. This is an adaptation layer between the service plane and existing NRPSs. Lastly, there is the *control plane*, which is used to control physical resources.

The general architecture is shown in Fig. 6. The project started with a centralized architecture with only one service plane (global broker). The service plane (network service plane, NSP) is responsible for finding end-to-end paths, managing AAA, and so on. The individual NRPSs, shown underneath the NSP, manage their own domains and intradomain routing. Each NRSP has its own control plane, such as GMPLS. All interfaces are deployed as web services.

At present, PHOSPHOROUS architecture works with three existing NRPSs: ARGON [12] from Germany, DRAC from Nortel, and User Controlled Lightpaths (UCLP) from CANARIE in Canada. The NRSPs publish border endpoints to the NSP, which can then do interdomain routing while NRSPs perform intradomain routing based on the selected endpoints. Each NRPS provides the ability to make a reservation request, cancel a reservation, get the status of a request, and bind and activate requests. This functionality is provided by an NPRS adapter which abstracts each specific NRPS. A number of interfaces (NBI, EWI, SBI), shown in the figure, for communication between all the components are defined.

The general flow when a grid request is received is as follows. The grid application or grid middleware sends a request to the reserva-

tion web service of the NSP. Once it is validated by the AAA module and available resources are found, the NSP finds an interdomain path. With the interdomain path selected, each independent domain is notified, and lookup of an intradomain path is initiated.

Harmony [13] is the network service provisioning system in PHOSPHORUS. Harmony assumes a group of independent NRPSs and provides an abstract service plane to do end-to-end co-allocation. It defines three types of architectures: centralized, hierarchical, and distributed. In each case, a Harmony interdomain broker sits above each NRPS. In the case of a centralized system there is only one, whereas in the case of a distributed system there is one for each NRPS. A Harmony service interface (HSI) is used to exchange abstract topology and other information between the brokers and the NRPSs.

As discussed previously, Harmony is responsible for interdomain paths while each NRPS is responsible for intradomain paths. The NRPSs provide border endpoints and abstract links to

Harmony. A reservation specifies one or more services, where each service requires one or more connections. The requests can be fixed or malleable advance reservation requests. A simple iterative search over the interdomain paths is used for malleable requests, trying different start times and different bandwidths. For each path, Harmony checks whether or not each individual domain can create an intradomain path.

Grid enabled GMPLS (G²MPLS) [14] is the new control plane defined in PHOSPHORUS which extends GMPLS to work with both grid and network resources. The goal is to incorporate grid network services (GNS) into GMPLS. GNS is a service that allows provisioning of network and grid resources in a single step. G²MPLS allows selection, co-allocation, and maintenance of grid and network resources. Some of the extensions include:

- Discovery and advertisement of grid capabilities (e.g., number of CPUs) and network resources (e.g., bandwidth)
- Service setup, which involves coordination with grid middleware, configuration of network connections, advance reservation, etc.
- Service monitoring

G²MPLS provides two models of operation. One is an overlay model designed for legacy systems that only supports standard GMPLS. In this case, NRPSs still exist and can use G²MPLS in the same way they would use GMPLS. The other model is the integrated model, where G²MPLS provides an interface directly to the NSP. It can then be used to co-allocate grid and network resources for a given domain. The basic unit of work is a grid job.

CONCLUSION

In this article we discuss a number architectures for supporting advance reservation of network resources. Supporting advance reservation is becoming increasingly important due to the requirements of grid and e-science applications. These applications must coordinate geographically distributed resources and be able to reserve network bandwidth in advance.

The various frameworks discussed in the article follow very similar architectural concepts of abstraction and orchestration to reserve, allocate, and provision dynamic circuits across heterogeneous network resources. These concepts and high-level design can be leveraged to implement similar techniques to meet requirements of commercial cloud computing providers.

ACKNOWLEDGMENT

We would like to thank John MacAuley for providing access to a DRAC test system for screen captures, and the anonymous reviewers for their insightful comments and suggestions to improve the quality of the article.

REFERENCES

- [1] T. DeFanti *et al.*, "Optical Switching Middleware for the OptiPuter," *IEICE Trans. Commun.*, 2003, pp. 2263–72.
- [2] H. Zang, J. P. Jue, and B. Mukherjee, "A Review of Routing and Wavelength Assignment Approaches for Wavelength-Routed Optical WDM Networks," *SPIE Optical Networks*, vol. 1, 2000, pp. 47–59.

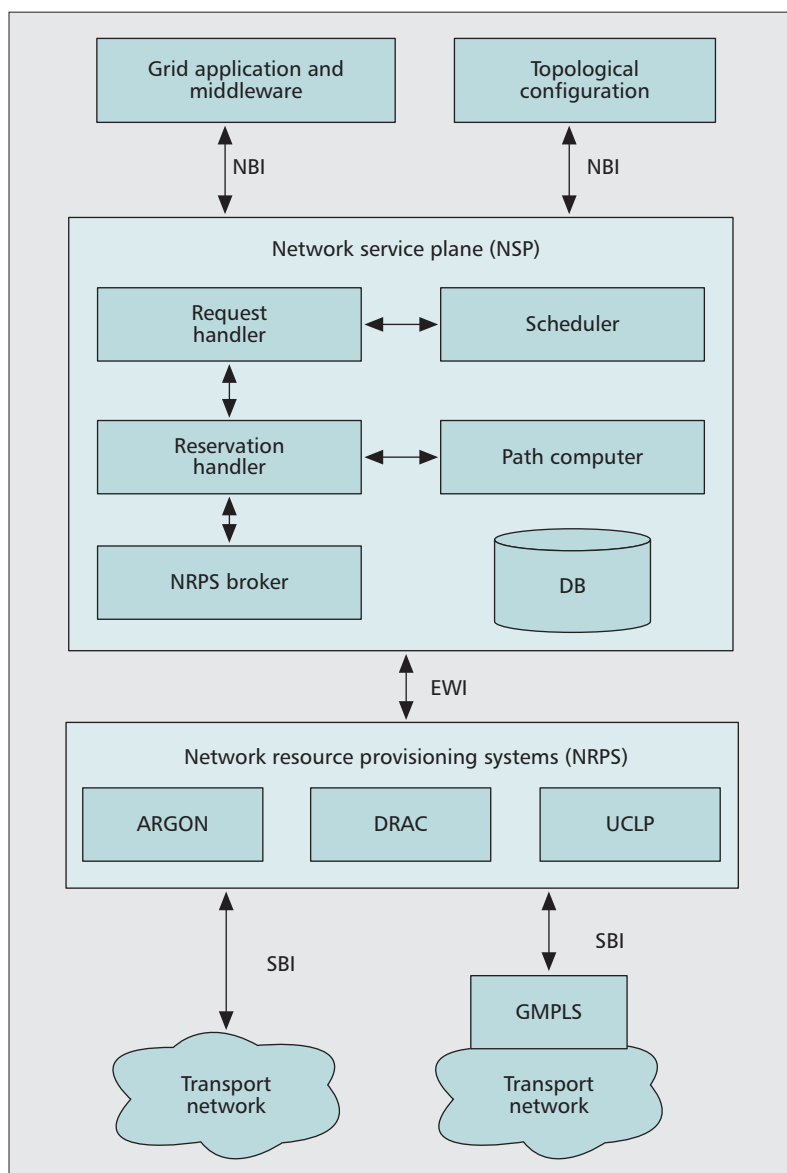


Figure 6. Overview of PHOSPHORUS architecture.

- [3] J. Zheng and H. T. Mouftah, "Routing and Wavelength Assignment for Advance Reservation in Wavelength-Routed WDM Optical Networks," *Proc. IEEE ICC*, 2002, vol. 5, pp. 2722–26.
- [4] D.B. Hoang et al., "DWDM-RAM: An Architecture for Data Intensive Services Enabled by Next Generation Dynamic Optical Networks," *Proc. IEEE GLOBECOM*, Nov. 2004, pp. 400–09.
- [5] L. Battestilli et al., "enLIGHTened Computing: An Architecture for Co-Allocating Network, Compute, and Other Grid Resources for High-End Applications," *Int'l. Symp. High Capacity Optical Networks and Enabling Technologies*, Nov. 2007, pp. 1–8.
- [6] S. Tanwir et al., "Monitoring and Discovery for enLIGHTened Computing," *High-Capacity Optical Networks and Enabling Technologies*, 2007.
- [7] A. Takefusa et al., "G-lambda: Coordination of a Grid Scheduler and Lambda Path Service over GMPLS," *Future Generation Comp. Sys.*, vol. 22, no. 8, 2006, pp. 868–75.
- [8] A. Takefusa et al., "GridARS: An Advance Reservation-based Grid Co-Allocation Framework for Distributed Computing and Network Resources," *Int'l. Wksp. Job Scheduling Strategies for Parallel Processing*, 2007, pp. 152–68.
- [9] M. Hayashi, H. Tanaka, and M. Suzuki, "Advance Reservation-based Network Resource Manger for Optical Networks," *Proc. OFC*, 2008, pp. 1–3.
- [10] S. R. Thorpe et al., "G-lambda and enLIGHTened: Wrapped in Middleware Co-Allocating Compute and Network Resources Across Japan and the US," *GridNets '07*, 2007.
- [11] S. Figuerola et al., "PHOSPHORUS: Single-Step On-Demand Services Across Multi-Domain Networks for e-Science," *SPIE*, 2007.
- [12] M. Pilz et al., "ARGON: Reservation in Grid-Enabled Networks," *DFN Forum Commun. Tech.*, 2008.
- [13] A. Willner et al., "Harmony — Advance Reservations in Heterogeneous Multi-Domain Environments," *Proc. Int'l. IFIP-TC 6 Networking Conf.*, 2009.
- [14] G. Zervas et al., "Phosphorus Grid-Enabled GMPLS Control Plane (GMPLS): Architectures, Services, and Interfaces," *IEEE Commun. Mag.*, vol. 46, no. 6, June 2008, pp. 128–37.

BIOGRAPHIES

NEAL CHARBONNEAU [S'08] (ncharbonneau@ieee.org) received B.S. and M.S. degrees in computer science from the University of Massachusetts, Dartmouth in 2008 and 2010, respectively. His interests include computer networks and software design and development.

CHIN P. GUOK (chin@es.net) joined ESnet in 1997 as a network Engineer, focusing primarily on network statistics.

He was a core engineer in the testing and production deployment of MPLS and QoS (Scavenger Service) within ESnet. He is the technical lead of the ESnet OSCARS project, which enables end users to provision guaranteed bandwidth virtual circuits within ESnet. He also serves as a co-chair to the Open Grid Forum On-Demand Infrastructure Service Provisioning Working Group.

INDER MONGA [M'00] (imonga@ES.NET) is developing new ways to advance networking services for collaborative and distributed science by leading research and services within ESnet. He contributes to ongoing ESnet research projects, such as Advanced Resource Computation for Hybrid Service and Topology Networks (ARCHSTONE) and ANI-Testbed as well as to the OSCARS and Advanced Network Initiative (ANI) 100G projects. He also serves as co-chair of the Network Services Interface working group in the Open Grid Forum. His research interests include network virtualization, network energy efficiency, grid/cloud computing, and sensor networking. He currently holds 10 patents, and has over 15 years of industry and research experience in telecommunications and data networking at Wellfleet Communications, Bay Networks, and Nortel. He earned his undergraduate degree in electrical/ electronics engineering from the Indian Institute of Technology, Kanpur, before coming to the United States to pursue his graduate studies at Boston University's Electrical and Electronics Communication Systems Department.

VINOD M. VOKKARANE [S'02, M'04, SM'09] (vvokkarane@ieee.org) received a B.E. degree with Honors in computer science and engineering from the University of Mysore, India, in 1999, and M.S. and Ph.D. degrees in computer science from the University of Texas at Dallas in 2001 and 2004, respectively. He is currently an associate professor of computer and information science at the University of Massachusetts, Dartmouth. His primary areas of research include design and analysis of architectures and protocols for optical and wireless networks. He is the co-author of a book, *Optical Burst Switched Networks* (Springer, 2005). He is a recipient of the Texas Telecommunication Engineering Consortium Fellowship, 2002–2003, and the University of Texas at Dallas Computer Science Dissertation of the Year Award, 2003–2004. He was a recipient of the Best Paper Award at the IEEE GLOBECOM 2005 Symposium on Optical Systems and Networks. He currently serves on the Editorial Board of *IEEE Communications Letters*. He has been on the Technical Program Committees of several IEEE conferences, including INFOCOM, ICC, GLOBECOM, ICCCN, HSN, and ANTS, and served as TPC Co-Chair for the Optical Networks and Systems Symposia at ICCCN 2007 and 2010, GLOBECOM 2011, and ICC 2012, and the INFOCOM High-Speed Networks Workshop 2011.

Supporting advance reservation is becoming increasingly important due to the requirements of grid and e-science applications. These applications must coordinate geographically distributed resources and be able to reserve network bandwidth in advance.