Lightpaths on Demand: A Web-Services-Based Management System

Raouf Boutaba, Wojciech Golab, and Youssef Iraqi, University of Waterloo
Bill St. Arnaud, CANARIE Inc.

ABSTRACT

User-controlled optical networks play a key role in supporting electronic transfer of the enormous volumes of data generated in emerging e-science experiments. The ability of users to manage their own resources enables provisioning of bandwidth-guaranteed tunnels on demand without the costs associated with conventional managed services offered by network providers. However, building high-performance user-controlled networks has only become feasible in the last few years, as trends in the telecommunications industry have made it possible for users to purchase installed optical fiber and light it using their own premises equipment. Consequently, suitable network management technologies have not yet evolved. In particular, there is presently no means for users to easily provision bandwidth-guaranteed tunnels across multiple independent management domains. In this article we present a user-controlled lightpath management system that addresses this problem. We begin by reviewing the high-level functionality of the system. Then we examine the software architecture. Finally, we discuss design challenges faced while building the system and propose future extensions.

INTRODUCTION

The Internet, while interconnecting a large part of the globe and supporting key services like the World Wide Web and email, is being outpaced by the data transport needs of the scientific community. Emerging experiments in the physical sciences call for a high-speed network capable of gigabit or even terabit per second aggregate data rates in order to enable remote collection and sharing of data. These include the ATLAS particle collider under construction at CERN, and the LOFAR radio telescope that consists of tens of thousands of radio antennas distributed throughout the Netherlands.

Attempts have been made in recent years to increase the throughput of Internet data transfers through enhancements to existing software technologies. These include new transport layer protocols optimized for high-performance networks [1], and improved file transfer software that supports parallel data streams and failure recovery [2]. Such techniques can reduce the inefficiencies in the transport layer congestion control mechanism and improve its stability, as well as increase throughput severalfold by circumventing the fairness mechanism of TCP. However, there remains the performance bottleneck associated with using a shared best effort delivery system.

In order to achieve the quality of service called for by emerging e-science applications, it is necessary to provision end-to-end bandwidth-guaranteed tunnels through the network. For example, in an optical network these can take the form of synchronous optical network (SONET) circuits, or wavelengths in a wavelength-division multiplexing (WDM) system. The latter are known as lightpaths, although we shall use the same term in reference to all types of circuits. Circuit emulation and packet switching technologies such as aysnchronous transfer mode (ATM), multiprotocol label switching (MPLS), and differentiated services (DiffServ) can also be used, but at high data rates the simple circuit switching mechanisms of optical networks prove to be the most economical. Of particular interest are emerging WDM systems that use all-optical switching based on microelectromechanical system (MEMS) technologies instead of costly high-speed electronics. For further information on the general topic of physical layer virtual private networks, the reader is referred to Recommendation Y.1312 recently released by the International Telecommunication Union (ITU) [3].

Provisioning of bandwidth-guaranteed tunnels has traditionally required the service of the network provider due to the fact that state must be set up in routing or switching elements in the core network. Until the late 1990s, this was the

This research is partially supported by CANARIE Inc. and Cisco Canada Inc.

only option for network users, who tolerated the costs associated with managed services. However, after the downturn of the telecommunications industry around 2000, which was due in part to the over-provisioning of optical infrastructure with respect to the near-term demand, a new alternative has emerged. Network users such as schools and research institutions now have the option of purchasing strands of dark fiber from a fiber bundle planted by a third party, and lighting it up with their own premises equipment [4]. Despite having to share in the ongoing costs of maintenance, users benefit from significant long-term savings. Moreover, having control of the electronic equipment, they are in a position to manage their interconnections to other users in the region, or to a carrier that provides transport over a wide area network. This type of peering can be achieved using switching hardware such as a SONET digital crossconnect system (DCS), jointly owned by users in a condominium fashion to reduce infrastructure costs. This raises the important possibility of provisioning lightpaths across multiple independent management domains, through collaborative trading of lightpaths and shared control over user-owned switching equipment.

The final barrier faced by users in provisioning wide-area bandwidth-guaranteed connectivity on demand and without the overhead of services from network providers is the lack of suitable management software. Conventional optical network management systems operate as isolated domains, with interdomain forwarding delegated to the IP layer. The Border Gateway Protocol (BGP) is the de facto Internet standard for exchanging interdomain routing information, focusing on simple connectivity and peering policies. It is not suitable for provisioning or advertisement of bandwidth-guaranteed tunnels in lower layers of the network. A new technology is needed that allows individual users to construct end-to-end lightpaths across multiple independent management domains on demand, enabling the same kind of peer-to-peer interaction that fostered the success of the IP-based Internet.

In this article we present a solution to the problem of managing next-generation user-owned networks [5]. Through virtualization of user-owned hardware devices distributed across multiple management domains, our system empowers users to construct end-to-end bandwidth-guaranteed tunnels by composing multiple inter-domain lightpaths.

Our prototype implementation, demonstrated at the 2003 CANARIE Advanced Networks Workshop [6], is based on Web services technology as the service discovery and remote procedure call mechanism. This technology takes advantage of the universality of XML to bridge the gap between applications and services implemented using different platforms and programming environments. It also allows the use of existing tools for browsing XML-based service registries, which facilitates the development of applications that dynamically establish lightpaths. In addition, our implementation is compliant with emerging Grid services standards, which build on Web services to support functionality needed in exposing distributed computational and data systems as shared commodities. This feature facilitates integration of our system with Grid infrastructures, which are gaining popularity in disseminating and processing data from e-science experiments.

This article is organized as follows. We discuss the main functionality of our user-controlled lightpath management system, which allows users to dynamically provision lightpaths for large data transfers. We offer details of the architecture of the system. We then describe the design challenges faced. Finally, we conclude the article with a discussion of possible extensions to the system and research directions.

LIGHTPATH MANAGEMENT FUNCTIONALITY

OVERVIEW

The functionality of our lightpath management system is centered on the notion of user ownership and control of network resources. The two types of resources in question are the bandwidth occupied in a series of physical links by each lightpath, and the switching equipment that allows lightpaths to be crossconnected to each other or to local area networks (LANs). Bandwidth resources are represented in the system as lightpath objects, subsequently referred to as LPOs, each of which is owned by a user of the system such as a research institution, government department, or large corporation. The bandwidth and crossconnect devices associated with each LPO are controlled through a set of high-level operations, which are described in detail in the next section. We use the term *crossconnect device* in reference to circuit switching or emulation hardware such as SONET DCSs, all-optical WDM switches, and ATM switches.

The functionality of the system aims to empower users to construct high-performance virtual networks by obtaining the necessary resources from other users, and combining those resources to form the desired logical topology. Consequently, the operations supported on LPOs are concerned with transfer of ownership, partitioning of bandwidth, and forming crossconnections among lightpaths. Operations are exposed to users based on their access level. Ordinary users have the ability to lease lightpaths from others, to partition and crossconnect them as needed, and to advertise unused bandwidth for lease. Domain administrators are technically competent users who enlist bandwidth resources within the system by specifying low-level hardware parameters, and who create accounts for ordinary users.

The functionality of the system aims to empower users to construct high-performance virtual networks by obtaining the necessary resources from other users, and combining those resources to form the desired logical topology.

SUPPORTED OPERATIONS

In this section we describe each of the major operations supported by the system. Note that for each operation discussed below, there is an analogous inverse operation.

Create Root LPO — Root LPO creation is a fundamental process during which bandwidth resources are enlisted within the system. Root LPOs represent lightpaths between pairs of physically adjacent crossconnect devices. They

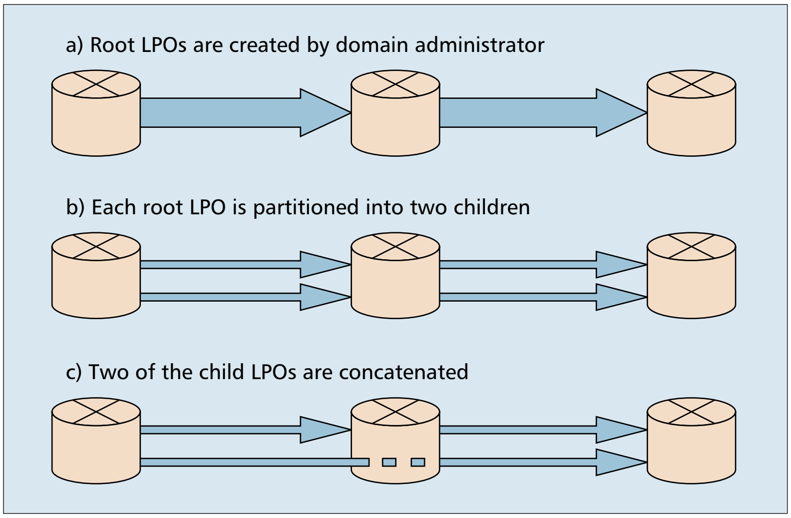

Figure 1. *A sequence of operations demonstrating the partitioning and composition of lightpaths from the hardware perspective.*

are created by domain administrators, who specify hardware parameters such as the card and port on which the lightpath terminates at each endpoint. In the case of SONET technology, the starting channel of the corresponding synchronous transport sgnal (STS) is also specified. Once created, root LPOs serve as a basis for constructing other lightpaths through the partitioning and concatenation operations described below. This scheme is similar to the virtual networking concept discussed in [7, 8].

Concatenate LPOs — The purpose of LPO concatenation is to crossconnect a series of constituent lightpaths having uniform bandwidth in order to form a single longer lightpath of the same bandwidth (Fig. 1). The process is used mainly in the establishment of end-to-end lightpaths, which is described later. The resultant compound lightpath is represented by a new LPO. The constituent LPOs store the LPO ID of the corresponding compound LPO, while the compound LPO stores a list of the IDs of its constituents.

Partition LPO — LPO partitioning refers to the process of dividing the bandwidth of a lightpath, referred to as the parent, in order to form multiple child lightpaths of smaller bandwidth between the same pair of endpoints (Fig. 1). A new LPO is created to represent each child lightpath.

The partitioning process is done in several steps. After selecting the LPO to be partitioned, the user chooses the bandwidth of the desired child lightpath. A list of valid choices is provided, representing any granularity constraints imposed by the hardware. For example, when using SONET technology, the child size must be chosen from a hierarchy of bandwidths differing by integer multiples. The user's choice is then validated and the system returns a list of additional children (if any) that must be created to satisfy hardware constraints (i.e., if the remaining bandwidth cannot be represented as a single unit). Finally, the user confirms, and the system updates its state by reconfiguring the affected crossconnect devices and updating the database.

Compound lightpaths present a challenge during the partitioning operation since a set of child bandwidths must be negotiated that is supported by all the constituents, which may themselves be compound. To this end, a recursive algorithm is used that iteratively builds a list of child bandwidths acceptable to each crossconnect device involved. Once a suitable combination of children is found, all crossconnections within the compound parent lightpath are partitioned accordingly.

Access/Reconfigure LPO — Accessing an LPO refers to the configuration that must be done before data can be transmitted over the corresponding lightpath. Specifically, each endpoint of the lightpath is crossconnected to an access point (e.g., an Ethernet port). The particular pair of access points used is chosen from a set of valid access points, which is associated with each LPO and inherited during partition and concatenate operations. This set can be restricted before an LPO is advertised for lease. Thus, an LPO represents not only bandwidth available to a user but also a policy concerning its usage.

Advertise/Lease LPO — The advertise and lease operations empower users to share resources. Unused lightpaths can be advertised, which makes them available for other users to lease. Thus, a large pool of unused resources is maintained, consisting of contributions from multiple domains. This is a peer-to-peer resource sharing model, where each user is both a customer and a provider. In contrast, in a conventional client-server scenario a customer must negotiate separately with a series of providers, and has no means to integrate resources from multiple sources.

Advertisement and lease are subject to time limits defined in terms of expiry dates. An LPO can be leased until the end of the advertisement period, and can be advertised until the end of the period for which it was leased. The lease expiry date is inherited during partitioning, and in concatenation it is set to the minimum of the expiry dates of the constituents. The owner of an LPO can manually cancel or extend an advertisement or lease, provided that this does not violate the constraints mentioned above.

End-to-End LPO Establishment — The establishment of an end-to-end LPO is a high-level operation used to produce a single lightpath between a particular pair of crossconnect devices. This process occurs in two phases. The first phase is the end-to-end path computation phase. It involves the formation of and search over a logical topology consisting of a subset of lightpaths that are accessible to the requesting user and eligible. A lightpath is accessible to a user if it is owned by that user or advertised for lease. Eligibility is defined according to two criteria. First, the bandwidth of an LPO must meet or exceed the user's requirement (we assume that a lightpath with excess bandwidth can be partitioned as needed). Second, the user must be able to obtain or retain ownership of an LPO for the desired period of time. The path computation phase fails if there are insufficient accessible and eligible resources to form the desired end-to-end path. In our prototype implementation, shortest path routing is used where distance is defined as hop count (i.e., in terms of the number of crossconnect devices traversed by a lightpath). Excess bandwidth, which is the bandwidth of an LPO less the user's requirement, can be considered to break ties among multiple shortest paths in order to prevent unnecessary fragmentation of high-capacity lightpaths. However, the system supports multiple routing engines and path selection strategies, which allows users to choose among various resource allocation policies. This feature makes the software extensible, and able to serve as a basis for traffic engineering in addition to mere circuit provisioning.

The path establishment phase completes the end-to-end LPO establishment process. Here, a sequence of operations is performed in order to construct an end-to-end LPO given the list of constituents selected earlier. First, the necessary LPOs are leased. Next, LPOs with excess bandwidth are partitioned, and in the case of leased LPOs the excess children are automatically re-advertised. Finally, the modified set of constituent LPOs is concatenated, unless of course the set contains only one lightpath. The end-to-end LPO cancellation operation reverses the changes performed during end-to-end LPO establishment to the extent permitted by the current state of resources.

SAMPLE SCENARIO

Here we consider a complete scenario that illustrates how the operations supported by our system can be composed in order to create an end-to-end lightpath for a large data transfer. Suppose that research institution A wishes to share 1 Tbyte of data from a particle physics experiment with research institution B. Suppose also that A and B reside in different administrative domains, and each is connected via Gigabit Ethernet LAN to lightpath crossconnect devices registered in our lightpath management system.

The process begins with user A performing end-to-end lightpath establishment between the crossconnect devices to which A and B are connected. The specifications of the request are as follows: routing will be done using the shortest path metric; advertised LPOs of other users (B's in particular) will be considered in addition to A's own LPOs; the bandwidth is 1 Gb, corresponding to the limit imposed by the LAN technology used to access the lightpath; and the expiry date is five hours in the future, which gives A and B enough time to set up the connection and transfer 1 Tbyte of data. Assuming that the necessary resources are available in the system, the system returns a candidate sequence of lightpaths that can be concatenated to form the desired end-to-end lightpath. A accepts, at which time the suggested constituents are leased (if necessary), partitioned, and concatenated. Finally, A accesses the end-to-end lightpath and begins transmitting data. After the data transfer is complete, A cancels the end-to-end LPO, or allows the system to automatically tear it down when it expires.

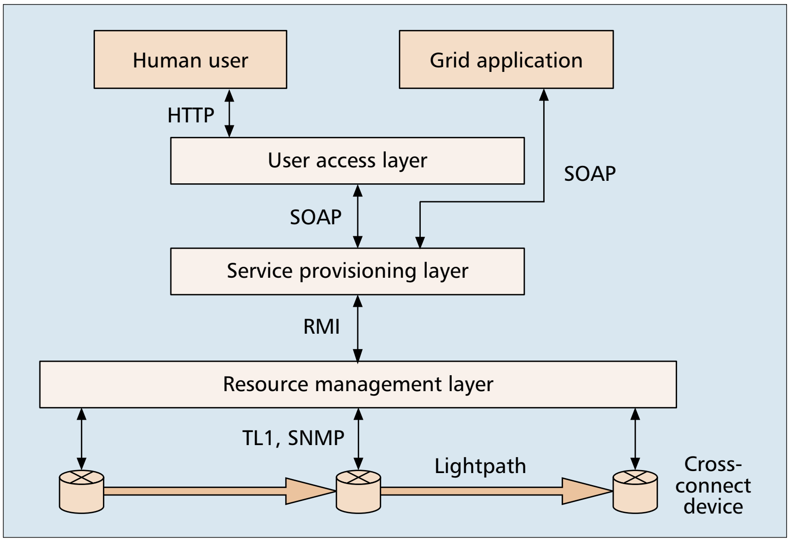

■ **Figure 2.** *High-level architecture of the lightpath management system.*

ARCHITECTURE FOR A WEB-SERVICES-BASED LIGHTPATH MANAGEMENT SYSTEM

This section discusses the architecture of our lightpath management system. We begin with a discussion of the overall architecture, and then explore each layer individually. Finally, we explain the nature and distribution of the data maintained by the system.

A HIGH-LEVEL OVERVIEW

Our user-controlled lightpath management system uses a three-layer architecture that separates the human user interface, the business logic and data management functionality, and the hardware resource management functionality. Accordingly, the system is divided into the user access layer (UAL), service provisioning layer (SPL), and resource management layer (RML). The high-level architecture is shown in Fig. 2.

The UAL exposes a Web interface to human users. Its role is to translate incoming requests into XML-based Simple Object Access Protocol (SOAP) messages that are understood by the SPL and present the results back to the user in HTML format. The SPL comprises a set of Grid services that carry out the operations defined in the system, a service registry, as well as a database that tracks the state of all lightpaths managed by the system. Each lightpath here is represented as an LPO. The SPL also serves as an access point for Grid applications. The SPL delegates hardware configuration tasks to the RML using the Java Remote Method Invocation (RMI) protocol. The RML handles low-level communication with network hardware devices while exposing a high-level abstract interface to the SPL.

The UAL does not maintain persistent data, and can be either deployed at a centralized location, or customized and deployed separately by each user. The SPL is logically centralized in our prototype, and uses a single database. This layer can be physically distributed over a cluster of

The Service Provisioning Layer tracks the states of all lightpaths managed by the system. Lightpaths are created, modified, and destroyed in response to requests from human users via the User Access Layer, or directly from Grid applications.

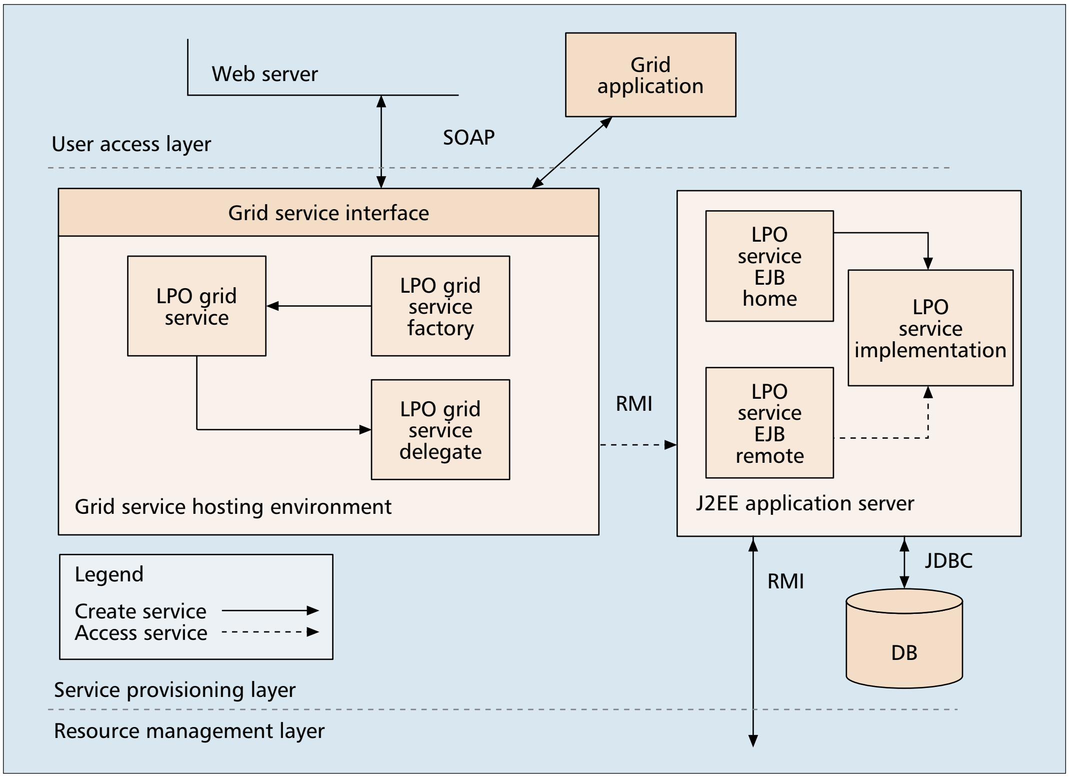

Figure 3. *Detailed architecture of the service provisioning layer.*

machines for increased fault tolerance. The RML is logically and physically distributed. It comprises a set of resource agents, one per crossconnect device managed by the system.

THE USER ACCESS LAYER

The UAL consists of a Web server, which exposes a Web interface to human users and communicates with the SPL using SOAP messages. Our implementation is based on the Java Web Services Developer Pack (Java WSDP) from Sun Microsystems [9]. This software is available free of charge and provides a set of tools for building and deploying Web services and Web applications. One of these tools is Tomcat, a commercial-quality Web server from the Apache Software Foundation. Tomcat is also the official reference implementation for Java Servlet and JavaServer Pages technologies, which makes it a suitable solution for serving dynamic Web content. Another tool provided by the Java WSDP is the Java application programming interface (API) for XML Messaging (JAXM), which enables creation and parsing of SOAP messages.

THE SERVICE PROVISIONING LAYER

The middle layer of our system is concerned with supporting the logic and data involved in the management operations offered by the system. It is implemented using the JBoss application server, MySQL relational database, and Globus Toolkit 3 Grid hosting environment.

Each of the high-level operations described earlier is supported as a service in the SPL. Each service consists of an Enterprise JavaBean (EJB), which contains the business logic and is hosted in JBoss, as well as a separate delegate object that exposes the actual Grid interface and is supported by the Grid hosting environment (Fig. 3). The delegate object forwards requests to the appropriate EJB, and separates the service interface from the service implementation. The EJB itself comprises a home interface used to create instances of the bean, a remote interface used to communicate with these instances, and the implementation of the bean. The Grid interface of each service is described using the standard Web Service Description Language (WSDL), and is listed in the Grid service registry. Globus Toolkit 3 provides a tool that automatically generates an OGSA-conformant interface given an EJB.

The SPL tracks the states of all lightpaths managed by the system. Lightpaths are created, modified, and destroyed in response to requests from human users via the UAL, or directly from Grid applications. SOAP messages are used in both interactions, which makes it possible to treat the UAL in the same manner as a Grid application.

Communication with the RML, on the other hand, is done using the more lightweight RMI protocol. Messages between the SPL and RML are exchanged whenever hardware configuration changes are needed (e.g., when forming crossconnections among lightpaths).

THE RESOURCE MANAGEMENT LAYER

The RML consists of a set of resource agents that control the underlying network resources. Each resource agent is associated with exactly one lightpath crossconnect device. In the CA*net4 network [10], each device is a Cisco ONS 15454 SONET Multiservice Provisioning Platform. The high-level role of the resource agent is to enable shared user control over cross-

```
<CrossConnectDevice num.slots = “17”>
    <card full.name = “OC192 LR/STM64 LH 1550” type = “SONET” num.ports = “1” num.channels = “192”>
        <common.name name=”OC192-LR”/>
        <common.name name=”OC192LR”/>
        <supported.circuit ID=”STS1”/>
        <supported.circuit ID=”STS3C”/>
        <supported.circuit ID=”STS6C”/>
        …
    </card>
    <circuit ID = “STS1” num.channels = “1”  bandwidth = “51840”    start.channel.granularity = “1”/>
    <circuit ID = “STS3C” num.channels = “3” bandwidth = “155520” start.channel.granularity = “3”/>
    <circuit ID = “STS6C” num.channels = “6” bandwidth = “311040” start.channel.granularity = “3”/>
    …
</CrossConnectDevice>
```

Figure 4. *An example of XML-encoded SONET hardware information.*

connect devices while abstracting the details of the underlying network hardware. This is accomplished using a pluggable switch interface component, whose implementation provides support for a particular type of crossconnect device. Its role is to perform hardware configuration and performance monitoring using the communication protocols spoken by the device. For example, in a concatenation operation on a SONET switch, the SPL instructs the agent to crossconnect two lightpath endpoints specified using abstract data structures, and in turn the switch interface issues the appropriate commands to the cross-connect device using Transaction Language 1 (TL1) messages, which contain hardware-specific details. The switch interface also plays a key role in lightpath partitioning, whereby it exposes resource allocation constraints such as bandwidth granularity. This topic is discussed further later.

DESIGN CHALLENGES AND SOLUTIONS

The design of a user-controlled network management system is a challenging problem, due not only to the distributed nature of the software, but also to the need to hide the complexity of network configuration from the user. This section discusses specific problems that arise and describes solution techniques.

HARDWARE PARAMETERS AND RELATED CONSTRAINTS

An important challenge in the design of a user-controlled lightpath management system is making the critical logic of the SPL independent of the particular network hardware technology in use. This allows the SPL to accommodate heterogeneous network hardware, and evolve gracefully as new technologies are deployed. Consequently, the RML must expose an interface that is generic with respect to the hardware.

For example, in the partitioning operation the RML provides information concerning the valid child bandwidths and exposes functionality that allows the SPL to determine what additional children must be created in order to satisfy the constraints of the hardware. This way, agents associated with devices based on time-division multiplexing (TDM) technology, such as SONET, can ensure that the bandwidth hierarchy is respected, while agents based on statistical multiplexing technology, such as ATM, can allow more flexible partitioning of bandwidth.

In SONET technology, hardware constraints are not limited to a simple bandwidth hierarchy. One must also take into account the set of circuit sizes supported by a network interface card, as well as the set of valid STS channel numbers for each circuit size [11]. This information is naturally encoded in an XML file, allowing external configuration of each resource agent. We have devised an XML schema for this purpose, as illustrated in Fig. 4.

As an example of how this information is used, consider a request to partition an STS192C lightpath provisioned between two OC192LR cards. The agent can infer that STS6C is a valid child circuit size, and can only start at channel numbers of the form $1 + 3k$, where k is an integer. Note that because actual starting channel constraints do not follow simple formulas, the XML file only specifies necessary conditions. The resource agent uses this information and also polls the hardware device itself in order to obtain a complete set of constraints. In the interest of performance, channels are polled only when needed, and any constraints learned from the device are retained in memory using a special cache.

SECURITY

The security mechanisms employed in our system are mostly based on the features available in Globus Toolkit 3. These features include an implementation of the Grid Security Infrastructure, which provides authentication using X.509 public key certificates, delegation of credentials through special proxy certificates, and message-level security based on the WS-Security specification. The Globus Toolkit also provides an authorization scheme that assigns privileges to certified identities.

In addition to the security mechanisms provided by the Grid hosting environment, our system uses an internal authorization scheme that enforces different levels of access according to the roles described earlier. Such a design was chosen to ensure consistent access control across all three layers of the system, only one of which is based on Grid technology.

> *Presently, we are investigating the problem of distributing the logically centralized Service Provisioning Layer in order to allow each management domain to operate its own customized instance of the system.*

Our authorization scheme is integrated with the security mechanisms of Globus Toolkit as follows. A Grid application first authenticates with the SPL using a public key certificate, which is verified by the Grid hosting environment. Next, the Grid application contacts a special authentication Grid service in order for the internal authorization mechanism of the SPL to assign privileges to the session. The SPL maps the identity of the user running the Grid application to a user ID in the database, and returns a randomized session ID. The session ID is then presented by the Grid application with each request, allowing the SPL to make authorization decisions. Interaction with human users is handled similarly, by treating the entire UAL as a single user at the level of the Grid hosting environment. After accepting a user name and password from the human user, the UAL authenticates to the SPL using its own certificate, and forwards the user's credentials to the authentication Grid service. The SPL then authenticates the human user and returns a session ID to be used by the UAL in future requests on behalf of the same user.

CONCLUSION AND FUTURE WORK

The lightpath management system presented in this article provides the essential services needed in order for users to create bandwidth-guaranteed tunnels across multiple management domains.

The functionality of the system can be extended readily by deploying new services in the SPL. New network hardware technologies can also be accommodated thanks to the abstract interface and the modular design of the resource agent component.

Extensions to the system are possible along several directions, ranging from customizability to fault tolerance. Presently, we are investigating the problem of distributing the logically centralized SPL in order to allow each management domain to operate its own customized instance of the system. The XML-based service description, discovery, and invocation mechanisms of the Web services technology used in our system simplify the problem of interoperability between instances. Another problem that remains to be addressed is how to disseminate the availability of interdomain lightpaths, and how to perform route computation in the end-to-end lightpath establishment operation. A hierarchical solution analogous to BGP is a good candidate here.

Another possible extension is to add support for provisioning survivable lightpaths. This functionality could be incorporated into the end-to-end lightpath establishment operation. For example, users could be given the choice of using path protection, if it is supported by the network hardware, a slower but more resource-efficient restoration mechanism, or no protection at all, depending on the requirements of the traffic.

ACKNOWLEDGMENTS

The following persons are hereby acknowledged in alphabetical order for their contribution to the development of the lightpath management system: Shivani Aggarwal, Bassam Aoun, Mohamed Dadi, Acila Derbal, Adel Ghlamallah, Boris Jabes, Dennis Tianshu Li, Darcy Quesnel, Basem Shihada, Wayne Szeto, Thomas Tam, and Iban Touchet.

REFERENCES

[1] C. Jin *et al.*, "FAST TCP: From Theory to Experiments," Mar. 2003; http://netlab.caltech.edu/pub/papers/fast-030401.pdf

[2] The Globus Alliance, "The GridFTP Protocol and Software"; http://www.globus.org/datagrid/gridftp.html

[3] ITU-T Rec. y.1312, "Layer 1 Virtual Private Network Generic Requirements and Architectures," Sept. 2003; http://www.itu.int/itudoc/itu-t/aap/sg13aap/recaap/y131/

[4] B. St. Arnaud, "Frequently Asked Questions about Customer Owned Dark Fiber, Condominium Fiber, Community and Municipal Fiber Networks," Mar. 2003; http://www.canarie.ca/canet4/library/customer/frequentlyaskedquestionsa boutdarkfiber.pdf

[5] U. Waterloo UCLP Team, "User Controlled Lightpaths Project," Mar. 2004; http://bcr2.uwaterloo.ca/~canarie/

[6] "CANARIE Advanced Networks Workshop 2003"; http://www.risq.qc.ca/ risq2003-canw2003/en/index.html

[7] A. Leon-Garcia and L. G. Mason, "Virtual Network Resource Management for Next-generation Networks," *IEEE Commun. Mag.*, vol. 41, no. 7, July 2003, pp. 102–09.

[8] R. Boutaba, W. Ng, and A. Leon-Garcia, "Web-Based Customer Management of Virtual Private Networks," *Int'l. J. Networks and Sys. Mgmt.*, special issue on Web-Based Management, vol. 9, no. 1, 2001, pp. 67–87.

[9] Sun Microsystems, "Java Web Services Developer Pack"; http://java.sun.com/Webservices/downloads/Webservicespack.html

[10] CANARIE Inc., "CA*net4 Connected Institutions Map"; http://www.canarie.ca/ canet4/connected/canet4_map.html

[11] Bellcore, "GR-253-CORE: Synchronous Optical Network (SONET) Transport Systems Common Generic Criteria," *Industry Reqs. and Stds.*, no. 03, Sept. 2000.

BIOGRAPHIES

RAOUF BOUTABA (rboutaba@bbcr.uwaterloo.ca) received his M.Sc. and Ph.D., both in computer science, from the Université Pierre et Marie Curie in France in 1990 and 1994, respectively. He is currently an associate professor in the School of Computer Science of the University of Waterloo. He teaches networks and distributed systems, and conducts research in network management, service engineering, and wired and wireless multimedia networking. In these areas he has published more than 100 papers in refereed journals and conferences. He is chairman of the IFIP working group on network and distributed systems management and vice chair of the IEEE Communications Society Technical Committee on Information Infrastructure. He HAS received several awards including the Premier's Research Excellence Award.

WOJCIECH GOLAB (wgolab@bbcr.uwaterloo.ca) received an Honors B.Sc. degree from the University of Toronto in 2002. Upon graduation, he was awarded the Governor General's silver academic medal. Currently he is completing a Master's degree in mathematics at the University of Waterloo. He is interested in optical networks, customer-controlled networks, and network resource allocation problems.

YOUSSEF IRAQI (iraqi@bbcr.uwaterloo.ca) received a B.Sc. in computer engineering, with high honors, from Mohamed V University, Morocco, in 1995. He received his M.S. and Ph.D. degrees in computer science from the University of Montreal in 2000 and 2003. He is currently a research assistant professor at the School of Computer Science at the University of Waterloo. From 1996 to 1998 he was a research assistant at the Computer Science Research Institute of Montreal, Canada. His research interests include network and distributed systems management, resource management in multimedia wired and wireless networks, and wireless ad hoc networking.

BILL ST. ARNAUD (bill.st.arnaud@canarie.ca) is senior director, sdvancednNetworks for CANARIE Inc., Canada's advanced Internet development organization. At CANARIE he has been responsible for the coordination and implementation of the world's first customer controlled optical network CA*net 4. He was recently featured in a *Time* magazine article as the engineer who is "wiring together advanced Canadian science". He is a frequent guest speaker at numerous conferences on the Internet and optical networking, and is a regular contributor to several networking magazines. He is a graduate of Carleton University School of Engineering.