INTERNATIONAL JOURNAL OF NETWORK MANAGEMENT
Int. J. Network Mgmt 2005; **15**: 349–362
Published online in Wiley InterScience (www.interscience.wiley.com). DOI: 10.1002/nem.581

Customer-managed end-to-end lightpath provisioning

By Jing Wu*,†, Michel Savoie, Scott Campbell, Hanxi Zhang, Gregor v. Bochmann and Bill St. Arnaud

Customer-owned and managed optical networks bring new cost-saving benefits. Two types of such networks are becoming widely used: metro dark fiber networks and long-haul leased wavelength networks. Customers may invoke a special QoS mechanism where end-to-end (E2E) lightpaths are dynamically established across multiple independently managed customer domains. The cost of bandwidth is substantially reduced since it largely becomes a capital cost rather than an ongoing service charge. Customers can optimize the overall resource consumption by utilizing resources from different suppliers. Remote peering and transit reduce the Internet connectivity cost. Bandwidth and quality of service are guaranteed because customers directly peer with each other using transport networks. An architecture for a customer-managed E2E lightpath provisioning system is presented. Integration with Grid applications is discussed and a prototype demonstration is described.

Jing Wu is a Research Scientist at the Communications Research Centre Canada. His research interests include control and management of optical networks, networking protocols and algorithms, network performance evaluation and optimization. He received his PhD degree in Systems Engineering from Xian Jiao Tong University, China.

*Michel Savoie is the Research Program Manager for the Broadband Applications and Optical Networks group at the Communications Research Centre Canada. He is the Project Manager for 'User-Controlled End-to-End Lightpath Provisioning over CA*net 4'. He holds a BSc and MSc in Electrical Engineering from the University of New Brunswick, Canada.*

Scott Campbell is a Network Researcher at the Communications Research Centre Canada. He designs and develops agent-based network management and control software for all optical networks. He has a bachelor degree in Computer Science from Dalhousie University, Canada.

Hanxi Zhang is a Research Engineer at the Communications Research Centre Canada. His research interests include distributed computing, web services, and protection/restoration for optical WDM networks. He obtained his MASc degree from the University of Ottawa, Canada.

Gregor v. Bochmann is a Professor at the University of Ottawa. Previously, he was a Professor at the University of Montreal for 25 years. He is a fellow of the IEEE and ACM, and a member of the Royal Society of Canada. He has worked in the areas of software engineering, distributed systems and optical networks.

*Bill St. Arnaud is the Senior Director of Advanced Networks for CANARIE Inc., Canada's advanced Internet development organization. He has been responsible for the coordination and implementation of Canada's next-generation optical Internet initiative called CA*net 4. He is a graduate of Carleton University, Canada.*

**Correspondence to: Jing Wu, Communications Research Centre Canada, 3701 Carling Avenue, PO Box 11490, stn H, Ottawa, Ontario, Canada K2H 8S2*
†E-mail: jing.wu@crc.ca

Introduction

There are basically two types of customer-owned and managed optical networks: metro dark fiber networks and long-haul wavelength networks. Schools, hospitals and government departments are acquiring their own dark fibers in metropolitan areas. They participate in so-called 'condominium' dark fiber networks to better manage their connectivity and bandwidth. They light up the fibers with their own equipment and interconnect their fibers to either like-minded institutions, commercial service providers or Internet Exchanges as they so choose. In the long-haul area, many providers are selling or leasing point-to-point wavelength channels. Some providers are offering 'condominium' wavelength solutions, where a number of customers share the capital costs of deploying long-haul optical networks. In return, each customer in the condominium consortium owns a set of wavelength channels. These purchased or leased wavelength channels can generally be treated as an asset rather than a telecommunication service. The institutions virtually extend their dark fiber networks many thousands of kilometers without having to purchase and maintain their own optical repeaters and associated equipment.

The advent of object-oriented and Web Service protocols such as JXTA, J2EE, CORBA, SOAP, XML and Jini indicates that optical network design may be on the cusp of a new approach in network management and in the control and management of wavelength channels. Rather than making signaling requests to a carrier to establish wavelength circuits, customers in the future may purchase wavelength channels and cross-connects from trading exchanges or from each other on a peer-to-peer basis, constructing their own optical networks and managing them directly through these object-oriented protocols.

In this paper we first discuss the benefits and technical challenges presented by customer-owned and managed optical networks. In the third section we present an architectural framework, communication protocols and software engineering technologies that are useful for building management tools for this context. In the fourth section we describe the system and software architecture of a management system developed for customer-managed lightpath provisioning. The integration with Grid applications is discussed. A prototype demonstration made at the Eighth Global Grid Forum is described in the fifth section to illustrate the deployment and service interactions. In the sixth section existing technologies for the control and management of multi-domain or inter-domain optical networks are summarized. The seventh section concludes the paper.

Benefits and Challenges of Customer-Owned Optical Networks

—Benefits—

In customer-owned optical networks, the cost of bandwidth is substantially reduced, as it now largely becomes a capital cost rather than an ongoing monthly service charge. For example, a 10-year leased fiber across Chicago costs what a carrier charges for a gigabit Ethernet service per month covering the same distance. As relatively inexpensive gigabit Ethernet and coarse wavelength division multiplexing technology catch on, exponential penetration in the market will likely occur. To expand the capacity of an existing dark fiber network, customers can add new wavelength channels. Since some all-optical networks are transparent to bit rate and data format, another option is to increase the bit rate of existing lightpaths. The latter option also applies to wavelength networks. Compared to renegotiating a new contract with service providers, customer-owned optical networks have significant cost savings when the demand for bandwidth increases dramatically over time.

Compared to renegotiating a new contract with service providers, customer-owned optical networks have significant cost savings when the demand for bandwidth increases dramatically over time.

The customer is able to optimize the overall resource consumption. The customer purchases dark fibers and/or wavelength channels from a number of independent suppliers, as well as participating in condominium wavelength networks built for some portions of their network. Therefore, the customer has more flexibility in network planning and deployment and is able to negotiate the best deal from different suppliers. The customer may fine-tune the usage of each resource from each independent supplier.

Customer-managed networks reduce Internet costs via remote peering and transit. Customers may directly peer with each other and, more importantly, set up bandwidth guaranteed connections to no-cost peering exchanges. The customers manage their peering relationships without having to contact a central management body or pay expensive Internet transit fees. There is a large incentive for many smaller ISPs and larger enterprise networks of roughly equal size to directly interconnect to each other on a settlement-free basis. Many commercial Internet service providers interconnect to each other at Internet Exchange points without charging each other money for the exchange of traffic. As long as traffic levels remain reasonably balanced, this is a very practical cost saving technique on Internet transit fees. However, when a smaller ISP or an institution connects to a larger ISP, there is generally an imbalance of traffic and the larger ISP will charge both peering and transit fees. Purchasing direct connections to other ISPs who will undertake settlement-free peering can result in significant cost savings over what would normally be paid in transit fees to a large upstream ISP.

Since the customer directly owns and manages an optical network, the bandwidth and quality of service are guaranteed.

Since the customer directly owns and manages an optical network, the bandwidth and quality of service are guaranteed. The complexity of service management at the IP layer is removed. A new opportunity for potential cost savings is introduced by eliminating expensive high-end routers in the core and replacing them with optical switches. However, there is a cost in terms of network efficiency as the functionality of IP packet multiplexing is lost. The trade-off needs to be explored regarding the bandwidth efficiency and the cost of wavelength channels versus the cost of routers.

It may be possible to convey the advantages of connectionless networks like the Internet to connection-oriented networks. Customer-empowered networks echo many of the advantages of connectionless networks. Now, circuits can be created and set up on demand by the customer without signaling a carrier or central management authority. Connection-oriented services can be deployed on a peer-to-peer basis. Connectionless packets can be transferred between computers with guaranteed throughput. Future networks may very well be hybrid facilities where attached devices will be able to transmit packets through a default connectionless facility and at the same time be able to set up end-to-end circuits on a peer-to-peer basis for high-bandwidth applications. Circuits and packets will be treated equally and managed in a seamless and transparent manner by the application or network-connected device.

Customer-managed end-to-end lightpath provisioning is a traffic engineering mechanism for inter-domain applications. The primary application for this technology is Lambda Grids to support high-end scientific research in high-energy physics, astronomy, bio-informatics, etc. Customer-managed lightpath provisioning allows these applications to invoke a special QoS mechanism where a separate optical Border Gateway Protocol (BGP) point-to-point connection is set up, operating in parallel with the normal hop-by-hop BGP route. Grid applications' data traffic is then automatically rerouted over the point-to-point optical BGP path rather than the normally congested routed path (Figure 1).

—Technical Challenges—

The first technical challenge is the management of networks with resources from different sources. Only the customer has total visibility of its own network and no provider can see all the network elements. The traditional centrally managed hierarchical networking technologies, e.g. Generalized Multi-Protocol Label Switching (GMPLS)[1] and Automatic Switched Optical/Transport Network (ASON/ASTN), assume that the provider has total visibility of all network elements and a common

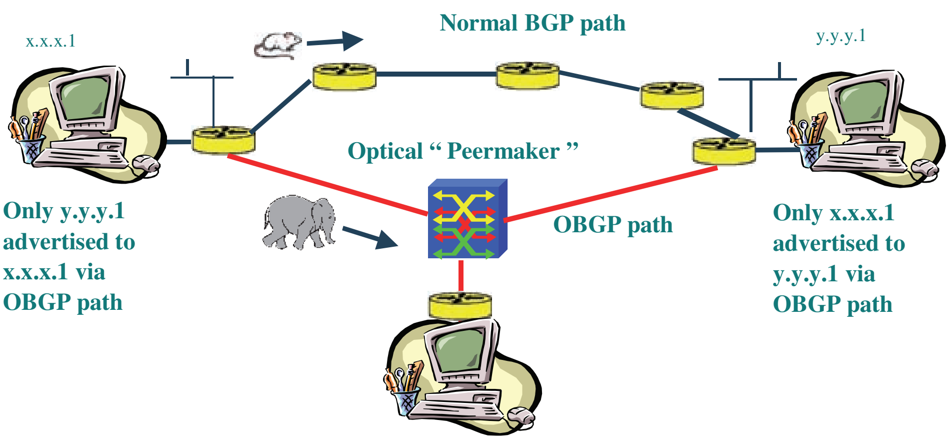

Figure 1. An application or end user controls peering of BGP optical paths for transfer of 'elephants' (i.e., vast amounts of data)

management system is used for all optical equipment. The customer-managed networks resemble some features of virtual private networks (VPNs).[2] A VPN is provisioned over public or third-party network infrastructures to provide dedicated connectivity to a closed group of customers. However, the VPN technology allows provisioning of customer networks within a single provider's domain. Clearly this type of architecture is not practical with customer-managed networks, where resources are provided by multiple suppliers. For protection and restoration, the customer, rather than any provider, is in a better position to decide the optimal solution. How to coordinate protection and restoration involving multiple providers is an open issue.

The second challenge is the collaboration among multiple independent customers without coordination through centralized management. Customer-managed networks adopt the peer-to-peer architecture, in which customers peer with each other. Each customer domain not only receives transport services from other customer domains but also contributes new transport services. During the establishment of an end-to-end connection, each segment of the connection between domains is set up on a peer-to-peer basis. Central guiding intelligence and arbitration of conflicts may be necessary, but day-to-day management and per connection control should be decentralized. An end-to-end connection from one customer to another involves at least two different customer domains, and if transit is required one or more intermediate transit domains may participate. Therefore the collaboration among multiple independent customers is critical for end-to-end connection provisioning. How to search and take control of resources in collaborative domains has to be addressed. Policy enforcement, authorization and authentication have to be applied. The organization of customer federations is a new issue.

The third challenge is the dynamic partitioning of a provider's resource to customers. VPN technology allows partitioning of a provider's resource to customers. However, VPNs are not as dynamic as some emerging applications require, e.g. Grid computing. Some customers prefer significant control and management capabilities in the provider's domain. They want a fine-grained resource allocation, which enables further optimization of the overall resource consumption. Deploying and upgrading customers' services is difficult and time-consuming in current networks owing to the closed, integrated architecture of network nodes. How to manage a provider's network element in a condominium fashion is challenging.

How to manage a provider's network element in a condominium fashion is challenging.

Architectural Principles for Customer-Managed Networks

—Network Architecture and Management Domains—

We assume a general mesh type of network architecture where the nodes consist of optical cross-connects and the edges consist of optical fibers over which multiplexing is provided through wavelength division multiplexing. The optical lightpaths are mostly used as links between packet switches and IP routers that are connected at the endpoints of these lightpaths.

In analogy to the architecture of the Internet, we assume that the optical network is partitioned into a number of administrative domains. In customer-managed optical networks, each domain represents an individual customer or a group of customers. Within each domain, we assume that complete management information would be available, which could be consulted when a new lightpath is to be established. However, between different domains, only partial information would be exchanged.

Customer domains may choose to interconnect to each other using condominium dark fibers or leased wavelength channels, or peer with each other at condominium peering points. Figure 2 shows these two scenarios. The peering points are usually non-blocking optical cross-connects. Rather than having one central organization manage these peering points, their management is partitioned amongst different customers.

Each cross-connect on the peering points can be managed independently by the individual customer. Interconnections between independent customers then must be done on a bilateral peering basis. The example in Figure 3 illustrates the principles of peering points. The peering point is partitioned into four separate management domains representing four customers. Associated with each management domain are the essential functions: grooming, switching and control plane services. Instead of having a single management interface for all these functions as in a traditional optical network, agents or objects are associated with each respective function for every customer on the peering point.

—Advertising Network Resources—

To establish an end-to-end lightpath across multiple customer domains, it is sometimes necessary

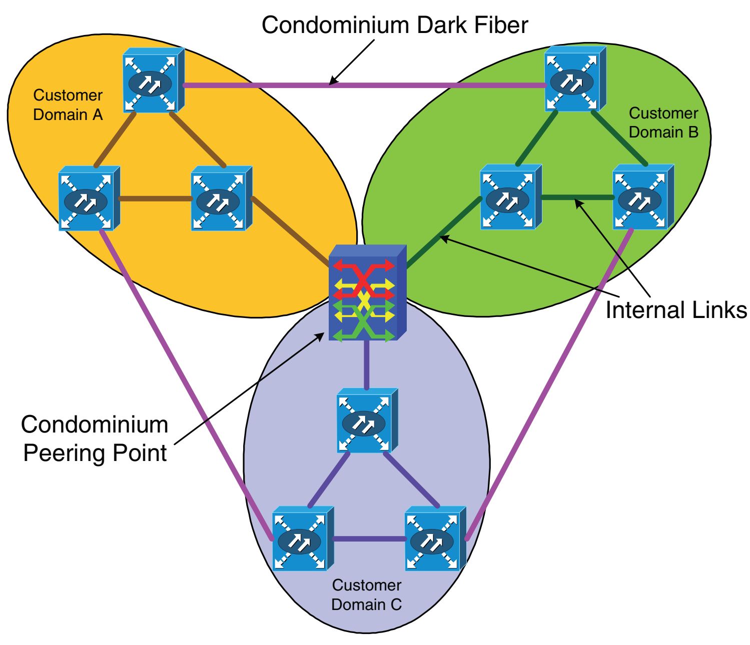

Figure 2. Customer domains may interconnect to each other using condominium dark fibers or leased wavelength channels, or peer with each other at condominium peering points

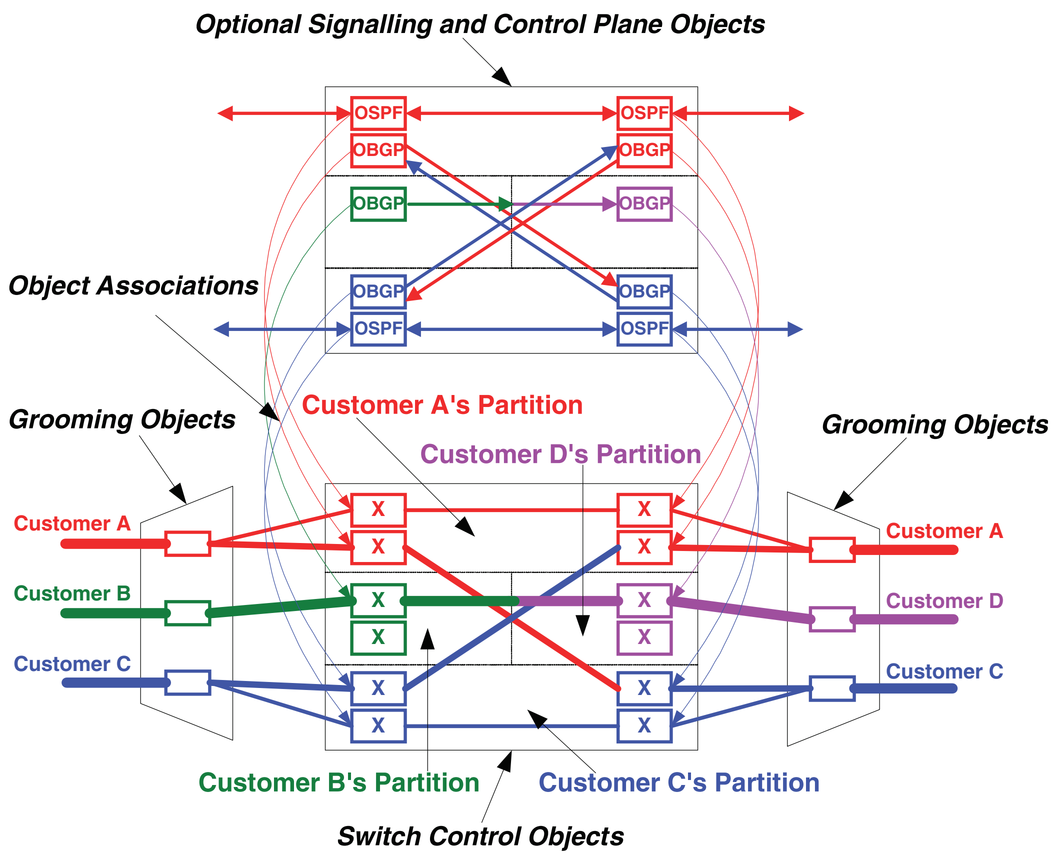

Figure 3. A peering point with customer control of ports and cross-connects

to concatenate lightspans that belong to different customers. A lightspan is a single physical link or a logical link composed of several concatenated physical links. The following two cases may be considered:

- *Peering*. Two lightspans from node A to B and from B to C are to be interconnected at the peering point B to create a new lightspan from A to C; the lightspans from A to B and from B to C may belong to different customers.
- *Leasing*. To establish a lightspan from A to C, a given customer P1 may own a free lightspan from A to B, and needs another free lightspan from B to C to be interconnected with the former; another customer P2 may own such a lightspan and be willing to sublease it to P1.

To facilitate the search for free lightspans that belong to other customers, free resources should be publicly advertised. The concept of service directories has been proposed in distributed computing and realized as the Common Object Request Broker Architecture (CORBA), Jini Lookup Services and JavaSpaces,[3,4] or Web Services directories.[5] Each resource or service that is to be made available to other customers must be registered in the service directory. Potential customers of these resources or services may query the directory to find a resource that meets their requirements. The object-oriented paradigm is used to support meaningful queries. Each registered object instance belongs to a class that defines the properties of the object instances. In addition, the class defines a certain number of attributes, and each instance is characterized by the values of its attributes. A customer's search for a service will therefore indicate the class of service desired and possibly some attribute values. For instance, customer P1, in the example above, will search the directory for an object of class *lightspan* with the attributes *source* = *B* and *destination* = *C*.

It is also conceivable that a given customer, e.g. P1, has leased a lightspan from customer P2 to build a longer lightspan, e.g. from A to C. Now P1 may subdivide the bandwidth of the lightspan and create multiple low-bandwidth lightspans from A to C. P1 may use a few of these lightspans and may readvertise the others as available in one of the service directories. The advertised leases could be associated with prices. In this way, the leasing mechanism may be used as a basis for establishing a broker market for optical networking.

—Distributed Resource Management—

Even if we have access to all the information about available lightspans and the possibility of leases, the actual resource reservations necessary for the establishment of a new lightpath will generally involve several customers. Therefore, several databases will be maintained for the state information about the different resources. To avoid inconsistencies due to concurrent access to these resources by several customers, it is important to foresee appropriate mechanisms for mutual exclusion of access. Since persistent storage is required in the presence of occasional crashes of the computers that maintain the databases, the transaction concept developed for centralized and distributed databases appears to be useful here. A transaction is a sequence of actions, such as reading the status of resources and requesting resource reservations, that are all executed as specified (i.e., the transaction commits) or not executed at all (i.e., the transaction aborts). This applies even in the case where one of the computers managing certain resources crashes during the operations (in this case the transaction aborts).

It is interesting to note that JavaSpaces[8] represents a service that provides persistent storage of object instances, retrieval of object instances and transaction management involving actions taking place in different JavaSpaces possibly residing on different computers. This technology is therefore an interesting platform for implementing distributed resource management tools.

—Using Jini and JavaSpaces in the Implementation of a Management Tool—

Jini and JavaSpaces provide a number of advantages for developing distributed applications. Jini runs on top of Java and uses the Remote Method Invocation (RMI) to access remote services. Jini also provides a set of application programming interfaces that hide the underlying complexity of distributed computing from the customer.[9] The Jini Lookup Service (JLS), a distributed service registry, allows customers to find services without having to know where they are located. The whole network provides a lookup service. Multiple lookup services may be introduced for redundancy and fault tolerance. In the start-up process, customers obtain the lookup service either by preconfiguration or through a discovery protocol that uses multicast throughout the network. Through the federation of JLSs, a customer can find any service in any domain. Services that are registered in the JLS are persistent and will be maintained as long as the service is alive. Jini also provides mechanisms for distributed events, distributed leasing and transactions.

JavaSpaces provides a distributed data store for Java objects. Objects stored in JavaSpaces are loosely coupled; anyone can take an object from a space without knowing (or caring) the details about the person who put it there. Operations on JavaSpaces are transactionally secure if the transaction mechanism is used. All the service calls in a transaction are committed, or none of them is committed. Transactions are supported for a single operation on a single space as well as multiple operations over many spaces. Like the JLS, JavaSpaces are also persistent; an object will remain in a space until it is explicitly removed. It also includes the search facilities of Jini and its mechanisms for distributed leases.

As mentioned below, our prototype system provides a management interface that conforms to the Web Services standard, which uses the XML (Extended Markup Language) coding scheme for client–server communication. Since the Web Services registries provide a service quite similar to the Jini Lookup Service, we were contemplating the possibility of using the Web Services standard also within our distributed prototype system; however, we decided to use the Jini technology internally. The JLS is more powerful and mature than the XML-based service registries. Since Jini passes Java objects via RMI, there is no need to have XML schema definitions for all remote service calls. Although using Jini/JavaSpaces limits us to the Java programming language (while XML is language independent), the internal Jini service calls are transparent to the customer and to other applications.

Design of a Management Tool Based on Jini and JavaSpaces

—System Architecture—

The architecture of a lightpath management system,[6] which we call a User-Controlled Lightpath Provisioning (UCLP) system, is shown in Figure 4. Our UCLP system is primarily designed to support Grid applications defined in the Open Grid Service Infrastructure (OGSI). The UCLP system enables Grid applications to dynamically set up end-to-end lightpaths across multiple independently managed customer domains to transfer vast amount of data. The figure shows the generic architecture and does not show the replication of the system components in the different parts of the network. Typically, one instance of each component shown would exist in each federation (i.e., a collaborative group of customer domains). However, they may also be shared. The JLS, JavaSpaces, Switch Communication Service (which interfaces to a single switch or a cloud of tightly coupled switches) and the Grid Service Access Point (SAP) may run on different computers. The Jini SAP and the LPO Service are downloaded to the process using them, in this case the Grid SAP.

This architecture uses the concept of service directories at two levels. First, internally, the JLS is

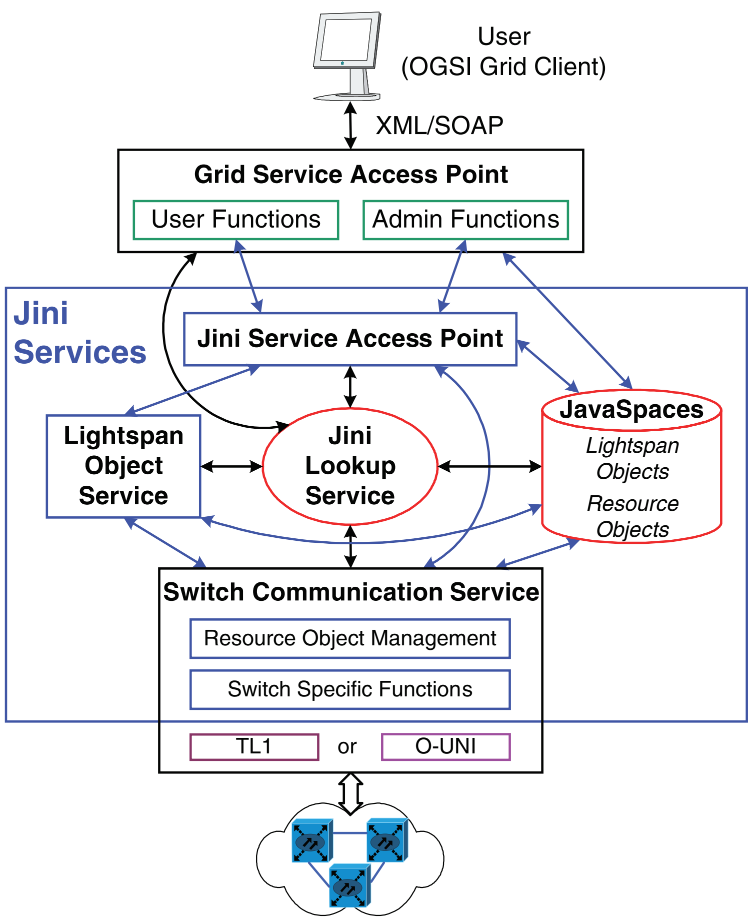

Figure 4. Overview of UCLP architecture

used to find the different instances of switches and JavaSpaces in the network. Second, a Grid SAP advertises its service instantiation through a Web Services Inspection Language (WSIL) pointer or an entry in the Universal Description, Discovery and Integration (UDDI) database. The client communicates the customer requests to the Grid SAP using the Simple Object Access Protocol (SOAP) adopted for Web Services. These requests are converted into Java procedure calls within the Grid SAP which then performs these calls on its local Jini SAP. Then, the Jini SAP executes these commands with the help of the other components within the system.

In a typical system configuration, each federation has its own set of services supported by the UCLP system (for short, we call them UCLP services), including its own JavaSpaces and JLS. Even though many services are accessible across federations, they are maintained independently of those in other federations. It should be noted that although resources are shared among different domains, it is still important to maintain the administrative boundaries between each domain to avoid confusion about the ownership of assets and administrative privileges.

—Lightpath Management Services—

Figure 5 shows the system architecture in more detail and indicates the main service methods pro-

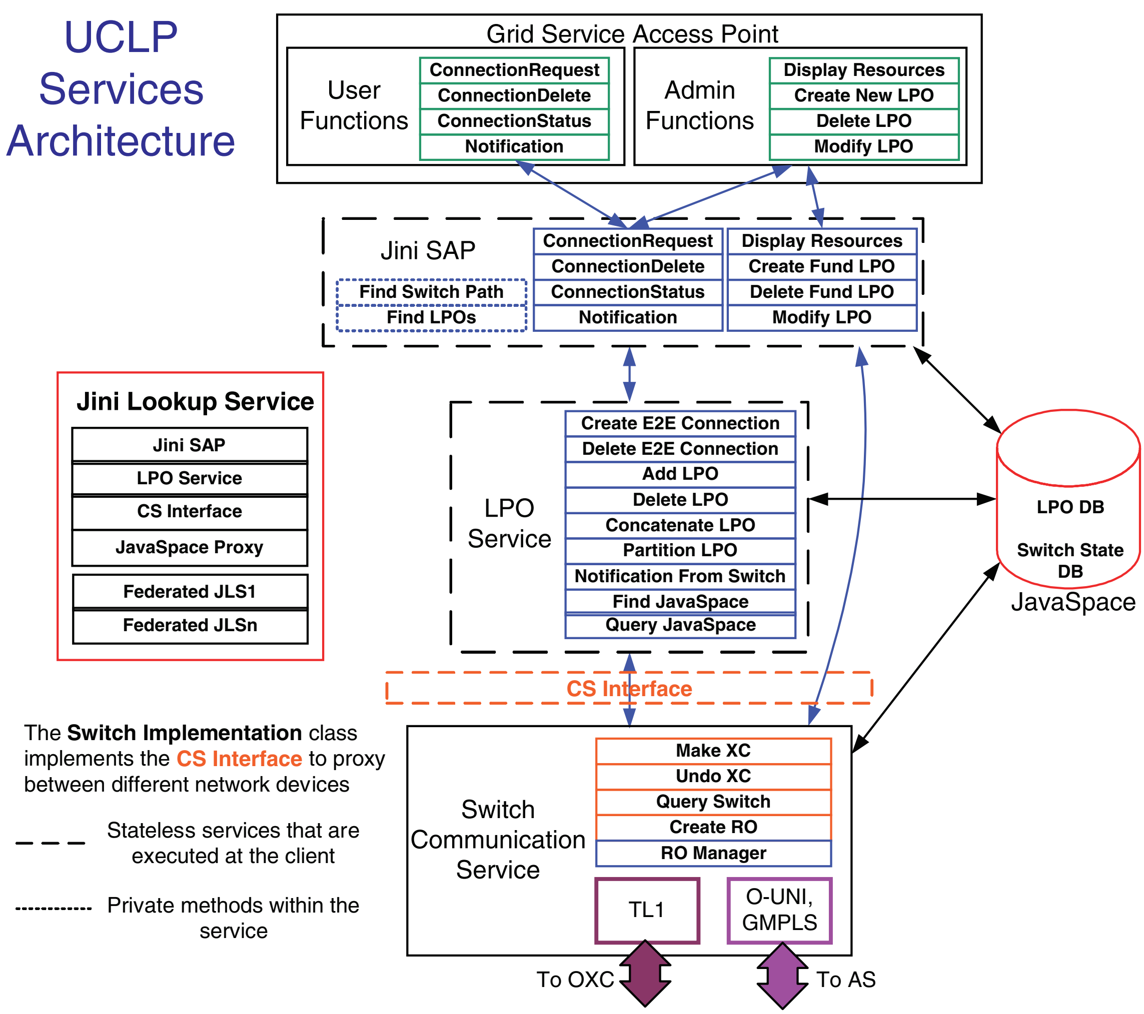

Figure 5. Detailed UCLP services architecture

vided by the component interface. The management services provided by our system are classified into two groups: those only available to administrative customers and those available to all customers. The latter include in particular 'ConnectionRequest', by which a customer can request the establishment of an end-to-end connection from a given entry port of a given switch to a given exit port on another given switch, possibly belonging to a different federation. One of the functions reserved to administrative customers is the addition of new physical links to the available optical network.

In our design, Light Path Objects (LPOs) are objects stored in JavaSpaces. An LPO is an abstraction of a lightspan. It is associated with a set of attributes and methods that enable possible peering to other LPOs at a switch to create an end-to-end connection or a longer lightspan. Supported customer operations on LPOs include: concatenating two LPOs, partitioning one LPO into many LPOs sharing common start and endpoints but with smaller bandwidth allocations, and reserving/using/releasing LPOs. The administrative operations include adding new LPOs and deleting LPOs corresponding to changes in the physical layer and the allocation of new resources.

For the execution of the ConnectionRequest, Jini SAP uses the internal methods FindSwitchPath and FindLPOs. The latter searches through the pertinent JavaSpaces to look for LPOs with attributes suitable for the end-to-end connection to be built. It also uses the functions provided by the methods of the LPO Service.

—Integration with Grid Applications—

To make UCLP services easy to utilize by conventional network applications, the interface to UCLP services is IP oriented. For example, the 'create E2E connection' operation takes two public IP addresses of source and destination hosts as inputs and returns information about the newly created connection, including the IP addresses of the UCLP endpoint network interface cards of the two hosts. Therefore, the integration becomes a simple IP replacement. Taking GridFTP as an example, a UCLP-enabled GridFTP client is actually a simple wrapper of the original GridFTP client. The UCLP service is called first to create an E2E connection. Then the original GridFTP application is performed for clients with the UCLP endpoint IP addresses to transfer data through the established UCLP connection.

However, this simple solution won't work if the transfer is initiated by a third party (Figure 6). In this case, the UCLP endpoint IP addresses of the source and destination host are unknown to the client that initiates the transfer. Although the FTP/GridFTP protocol is general enough to achieve this kind of transfer, the existing GridFTP client doesn't support third-party initiated file transfers yet. It assumes that FTP control messages and data flows always go through the same network interface. Therefore a more sophisticated GridFTP client is needed to realize a third-party initiated file transfer over a UCLP connection.

We use the Reliable File Transfer (RFT) service as an alternative solution. The RFT service is an

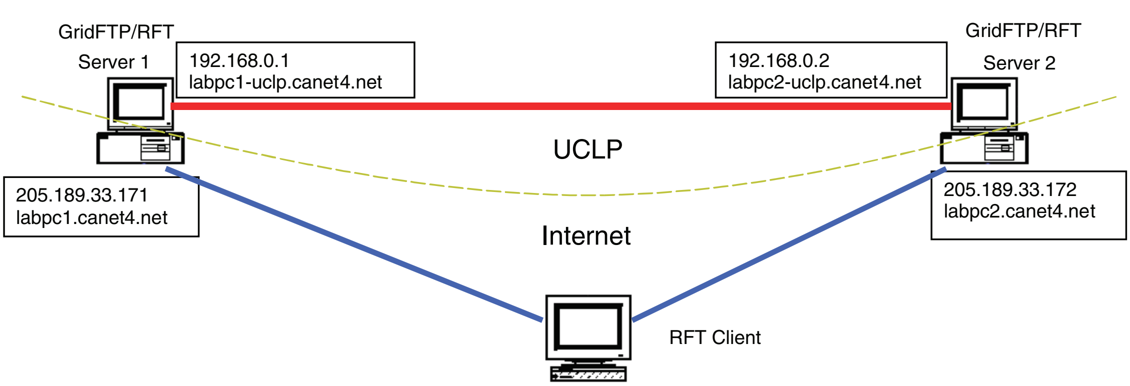

Figure 6. A third-party initiated file transfer through UCLP

OGSI-compliant Grid service that provides interfaces for controlling and monitoring third-party file transfers between GridFTP servers. Using the RFT service, clients can connect to the RFT service using one IP address and ask it to perform transfers between other IP addresses. Therefore, third-party transfers can be realized easily by a simple wrapper of the RFT client. The only thing is to make sure UCLP end points are reachable by the RFT service. This can be easily done by starting the RFT services on the hosts with a UCLP end point.

Prototype Demonstration at the Eighth Global Grid Forum

The prototype of our UCLP system was demonstrated at the Eighth Global Grid Forum (GGF8) held in Seattle on June 24–27, 2003. An RFT application was used to initiate the setup of an E2E lightpath (in the demonstration, a SONET STS-3 connection was established to simulate a lightpath). An RFT client used the established E2E lightpath to perform a server-to-server file transfer. The Grid FTP client ran at the site of the GGF8, while the UCLP system ran remotely in our labs in Ottawa, Canada. The UCLP deployment architecture for the demonstration is shown in Figure 7.

The purpose of this demonstration was mainly to show the system architecture and the interactions among service modules. Therefore, only one customer domain was used in the demonstration. An E2E lightpath was successfully set up dynamically on the demand of the Grid client. The demonstration proved that the design of our UCLP system is technically feasible. Service architecture and interactions are shown in Figure 8. The major steps of interactions are listed as follows:

1. Grid Service receives the request and invokes the Jini SAP.
2. User makes 'connectionRequest' call via the Grid client (user enters IP address of source and destination machines that are to be connected).
3. Jini SAP queries the JLS to get the proxy to the JavaSpaces the resource objects and LPOs reside in.
4. Jini SAP uses the source and destination IP addresses entered by the user to search the JavaSpaces for the resource objects for that connection.

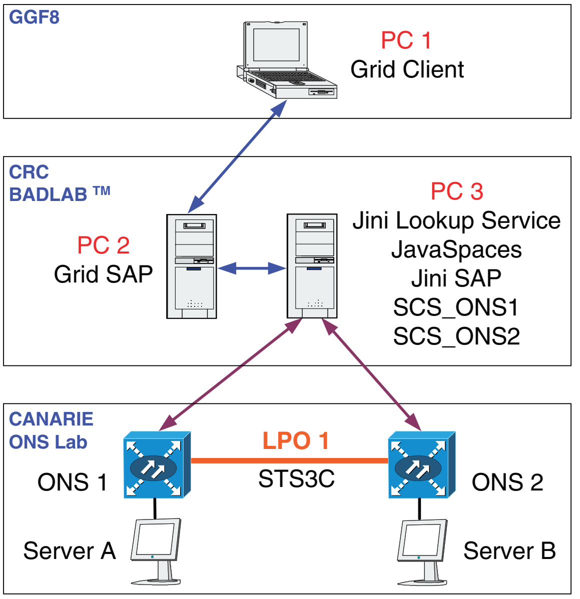

Figure 7. UCLP deployment architecture for the demonstration at the Eighth Global Grid Forum

5. Once the specific switch information is retrieved, it can be used to find the LPO.
6. Jini SAP then uses the source switch ID from the LPO to download the proxy for the SCS controlling Switch 1 (ONS 1).
7. Jini SAP does the same for Switch 2 (ONS 2).
8. Jini SAP sends a 'Cross-connect' command to SCS 1.
9. SCS 1 initiates a TL1 session with the switch to make cross-connection at Switch 1.
10. Jini SAP sends a 'Cross-connect' command to SCS 2.
11. SCS 2 initiates a TL1 session with the switch to make cross-connection at Switch 2.
12. The two servers are now connected via LPO 1.

Related Work

—Status of Control Techniques for Multi-Domain Optical Networks—

The E2E lightpath provisioning in the customer-managed environment involves multiple independent customers without coordination through centralized control or management. This feature

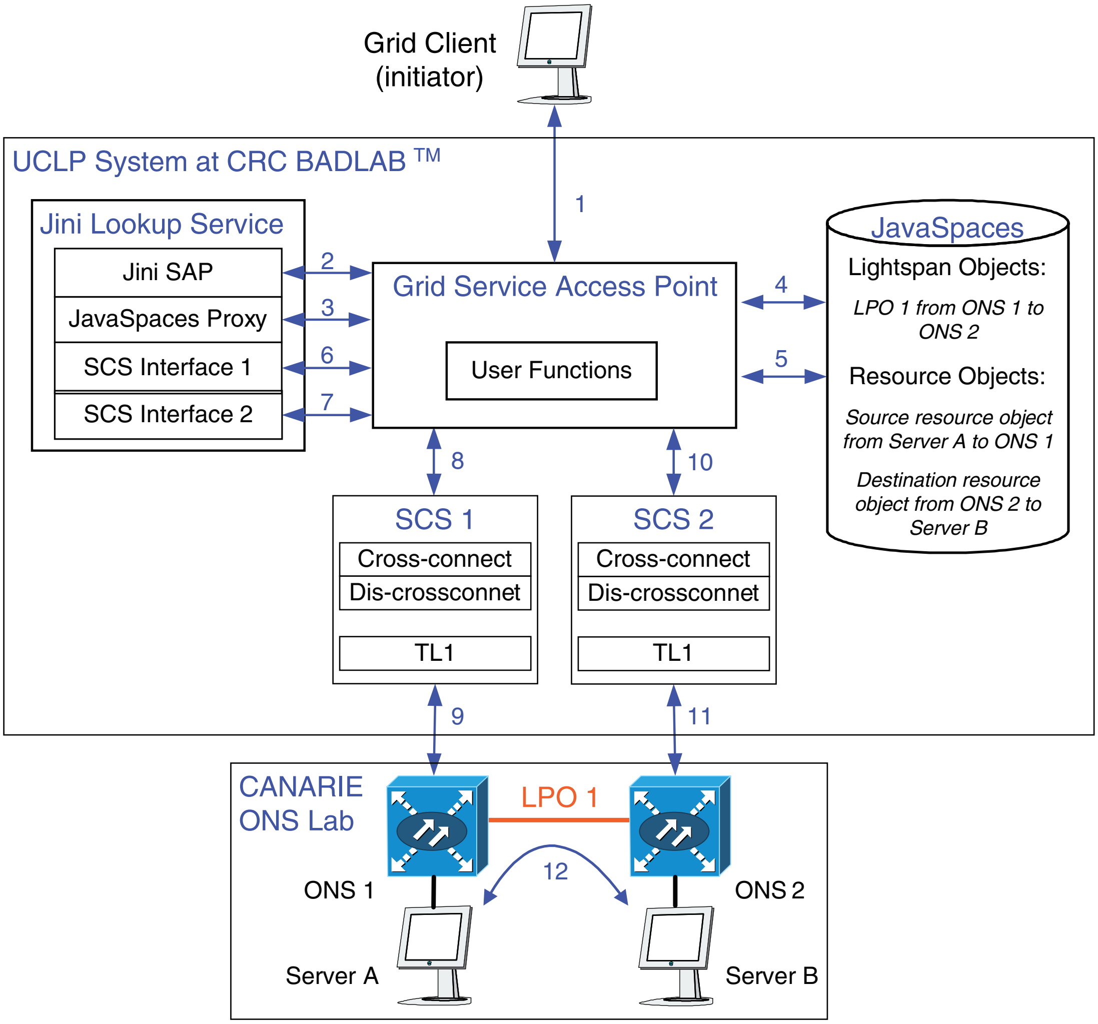

Figure 8. Service architecture and interactions for the demonstration at the Eighth Global Grid Forum

has some similarities to multi-domain control or management. In this section, we will summarize the status and existing techniques of multi-domain or inter-domain control and management for optical networks.

The motivations for dividing a network into domains are identified as:[7] (1) to define administrative boundaries between network operators; (2) to allow for scalability of routing and/or signaling; (3) to enable isolation of partitions for security or reliability; (4) to accommodate technology differences between systems, for example, by means of partitioning a single carrier's network into separate single-vendor sub-networks.

Standards organizations are active to define the requirements, framework and standards of the inter-domain control plane for optical networks. Three classes of functionality have been identified:[8] discovery, signaling and routing. The discovery process obtains information about the connectivity to and capabilities of neighbor network elements and sub-networks. The signaling or connection control is to set up, tear down, modify and maintain connections across networks. The routing is to exchange reachability, topology and resource status information.

So far, the control architecture and protocols for customer-controlled optical networks have not been well defined. This situation will remain for a period of time before the control architecture and protocols for inter-domain optical networking are well developed. The market demands play a key role in the process of the standardization of the control architecture and protocols for customer-controlled optical networks. In the early stages, the interoperability and performance of control protocols may not be the primary concern for customer-controlled or customer-managed optical networks.

—Existing Management Techniques for Customer-Managed Optical Networks—

The ability of active networks to program the network enables new protocols and innovative cost-effective technologies to be easily deployed at intermediate nodes. An important application of active networks is the distribution of network management functionality,[9] e.g. network configuration management. The configuration management performs several tasks such as discovering new devices, maintaining topology information, partitioning resources, controlling (e.g. installing, updating, reconfiguring) the software of the network elements remotely, setting up VPNs, etc. A programmable controller in a resource agent is the key building block that allows customers to customize network control and provides the resource agent with the intelligence for autonomous operations. Each customer network management system has its own programmable controllers in the resource agents. The programmable controller can download software packages with different functions. An example is the signaling and routing control software required for each logical subset of a switch.

Although a virtual active network was proposed to transform a multi-domain environment into a single domain view for customers,[10] the management of VPNs and active networks mainly concentrates on the partitioning of carriers' network resources and the autonomous operation of resources partitioned for individual customers without or with minimal interaction with the carriers. These management solutions meet the requirements for the operation of network switching elements in a condominium fashion, which were defined by the IETF in Anderson and Buerkle.[11] However, the collaboration among autonomously operated networks is not addressed. In a customer-managed optical network, to establish an E2E lightpath, multiple independent customer domains participate collaboratively. This motivates us to use the Jini technology to design a management system for customer-managed E2E lightpath provisioning. The federation concept and associated functions in the Jini technology facilitate our system design.

Conclusions

With the availability of dark fibers in some metropolitan areas and leased or condominium long-haul point-to-point wavelength networks, customers have the opportunity to construct and manage their own networks. Thus, the E2E lightpath provisioning involves multiple independently managed customer domains. The distributed resource management, collaboration among independent domains and partitioning of provider's resources to customers are major issues for this new network architecture.

A management system for such applications may be built based on the concept of service directories. Jini and JavaSpaces provide a number of advantages for developing such applications. We presented a design of a management system for customers to provision E2E lightpaths across multiple independent customer domains. This tool can be used as a traffic engineering mechanism for inter-domain applications, where separate point-to-point lightpaths are set up operating in parallel with the normal hop-by-hop routes. Interfaces to Grid applications are designed so that this tool can be easily integrated into Grid applications. The prototype of our management system was successfully demonstrated at the Eighth Global Grid Forum to dynamically establish an E2E lightpath on the demand of a Grid client to do a Grid FTP server-to-server file transfer.

The architecture for customer-managed optical networks is the foundation for the future broker trading market of optical networking, where the dynamic configuration of network resources and the management of partnership and leasing are essential. It also finds significant opportunities in non-profit research and educational networks, where reduced operational cost is observed.

Acknowledgements

We thank Jun Chen, Wei Zhang, Ling Zou and Mathieu Lemay for their contributions to the implementation of the prototype system for customer-managed end-to-end lightpath provisioning. We thank Sergi Figuerola, Eduard Grasa, Joaquim Recio and Albert López from the Technical University of Catalonia, Barcelona, Spain, for their participation in the discussion. This research is partially supported by CANARIE, Inc. under a directed research program.

References

1. Zhang Z, Fu J, Guo D, Zhang L. Lightpath routing for intelligent optical networks. *IEEE Network* 2001; **15**(4): 28–35.
2. Isaacs R, Leslie I. Support for resource-assured and dynamic virtual private networks. *IEEE Journal on Selected Areas in Communications* 2001; **19**(3): 460–472.
3. Freeman E, Hupfer S, Arnold K. *JavaSpaces Principles, Patterns, and Practice*. Addison-Wesley: Reading, MA, 1999.
4. Li S *et al*. *Professional Jini*. Wrox Press Inc., 2000.
5. Curbera F *et al*. Unraveling the Web Services Web: an introduction to SOAP, WSDL, and UDDI. *IEEE Internet Computing* 2002; **6**(2): 86–93.
6. Wu J, Campbell S, Savoie M, Zhang H, Bochmann G, St Arnaud B. User-managed end-to-end lightpath provisioning over CA*net 4. In *Proceedings of National Fiber Optic Engineers Conference (NFOEC 2003)*, Orlando, FL, 7–11 September 2003.
7. Bernstein GM, Sharma V, Ong L. Interdomain optical routing. *Journal of Optical Networking* 2002; **1**(2): 80–92.
8. Strand J *et al*. NNI requirements and framework resource document. *Optical Internetworking Forum (OIF) technical document OIF2001.535*, 2002.
9. Boutaba R, Polyrakis A. Projecting advanced enterprise network and service management to active networks. *IEEE Network* 2002, **16**(1): 28–33.
10. Brunner M. A service management toolkit for active networks. In *Proceedings of IEEE/IFIP Network Operations and Management Symposium (NOMS)*, Honolulu, HI, April 2000; 265–278.
11. Anderson T, Buerkle J. Requirements for the dynamic partitioning of switching elements. *IETF RFC 3532*, May 2003. ■

