CUSTOMER-MANAGED END-TO-END LIGHTPATH PROVISIONING

Jing Wu, J. Michel Savoie, Scott Campbell, Hanxi Zhang

Communications Research Centre Canada, 3701 Carling Avenue, Ottawa, Ontario, Canada K2H 8S2, {jing.wu, michel.savoie, scott.campbell, hanxi.zhang}@crc.ca

Gregor v. Bochmann

School of Information Technology & Engineering, University of Ottawa, bochmann@site.uottawa.ca

Bill St. Arnaud

CANARIE Inc., bill.st.arnaud@canarie.ca

Abstract: 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 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. The bandwidth and quality of service are guaranteed because the customers directly peer with each other using transport networks. The technical challenges include distributed resource management, collaboration among independent customers without co-ordination through centralized management and dynamic partitioning of a provider's resource to customers. An architecture for a customer-managed end-to-end lightpath provisioning system is presented. The distributed network resource management in a multi-domain environment is realized as service directories. Jini and JavaSpaces are used to develop the management system for customers to establish end-to-end lightpaths.

1. 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 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 telecom 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.

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 setup between a client host and a destination server, operating in parallel with the normal hop-by-hop BGP route. Grid applications' data traffic is then automatically re-routed over the point-to-point optical BGP path rather than the normally congested routed path.

In this paper we first discuss the benefits and technical challenges presented by customer owned and managed optical networks. We then propose an architectural framework, communication protocols and software engineering technologies that are useful for building management tools for this context. In Section 4, we describe the system and software architecture of a management system developed for customer-managed lightpath provisioning.

2. Benefits and Challenges of Customer Owned Optical Networks

2.1 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 the existing lightpaths. The latter option also applies to the wavelength networks. Compared to re-negotiating a new contract with service providers, customer owned optical networks have significant cost-saving 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 participates 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 relationship without having to contact a central management body or pay expensive Internet transit fees.

Since the customer directly owns and manages the 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.

2.2 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 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 multiple suppliers provide resources. For the protection and restoration, the customer, rather than any provider, is in a better position to decide the optimal solution. How to co-ordinate the protection and restoration involving multiple providers is an open issue.

The second challenge is the collaboration among multiple independent customers without co-ordination 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. So 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 customer's services is difficult and time-consuming in today's networks due to the closed, integrated architecture of network nodes. How to manage a provider's network element in a condominium fashion is challenging.

2.3 Related work

2.3.1 Status of control techniques for multi-domain optical networks

The end-to-end lightpath provisioning in the customer-managed environment involves multiple independent customers without co-ordination through centralized control or management. This feature has some similarities to multi-domain control or management. In this section, we will summarize the status and existing techniques for the multi-domain or inter-domain control and management for optical networks.

The motivations for dividing a network into domains are identified as [3]:

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 [4]: 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. At early stages, the interoperability and performance of control protocols may not be the primary concern for customer-controlled or customer-managed optical networks.

2.3.2 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 [5], 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 have full customization of 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.

A Virtual Active Network (VAN) was proposed to transform a multi-domain environment into a single domain view for a customer [6]. A VAN is a generic service offered by the VAN provider to the customer. The customer can install, run and supervise active services into a VAN, without interaction with the provider. Active networks using the VAN concept give customers new possibilities and benefits. The customers benefit from rapid new service installation without standardization. The customers may perform customized service management without interaction with their service providers and control fine-grained resources from multiple suppliers.

The management of VPNs and active networks 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 carrier. These management solutions meet the requirements for operation of network switching element in a condominium fashion in the customer-managed networks, which were defined by the IETF in [7]. However, the collaboration among autonomously operated customers is not addressed. When an end-to-end lightpath is to be established in a customer-managed optical network, multiple independent customer domains participate collaboratively.

3. Architectural Considerations for Customer-Managed Networks

3.1 Network Architecture and Management Domains

We assume a general mesh-type of network architecture where the nodes consists 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 end-points 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 1 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 bi-lateral peering basis. The example in Figure 2 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.

3.2 Advertising Network Resources

To establish an end-to-end lightpath across multiple customer domains, it is sometimes necessary 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 [8, 9], or Web Services Directories [10]. 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 searching for a service will therefore indicate the class of service desired and possibly some attribute values. For instance, the 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 re-advertise the others as available in one of the service directories. The advertised leases could be associated with a price. In this way, the leasing mechanism may be used as a basis for establishing a broker market for optical networking.

3.3 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 status 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] represent a service which 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.

3.4 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 Remote Method Invocation (RMI) to access remote services (Figure 3). 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 pre-configuration or through a discovery protocol that uses multicast throughout the network. Through the federation of JLS's, 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. All the service calls in a transaction are committed, or none at all. 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 conform to the Web Services standard, which uses the XML (Extended Markup Language) coding scheme for client-server communication. Since the Web Services registries provides a service quite similar to the Jini Lookup Service, we were contemplating the possibility of using Web Services standards 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.

4. Design of a Management Tool Based on Jini and JavaSpaces

4.1 System Architecture

The architecture of a lightpath management system, we call it a User-Controlled Lightpath Provisioning (UCLP) system, is shown in Figure 4. The 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. Figure 4 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 a 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 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. Firstly, internally, the JLS is used to find the different instances of switches and JavaSpaces in the network. Secondly, Grid SAP advertises its service instantiation through a well-known process described by the OGSI implementation such as a Web Services Inspection Language (WSIL) pointer or Universal Description, Discovery and Integration (UDDI) database. The client communicates the customer requests to the Grid SAP using the Simple Object Access Protocol (SOAP) protocol 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.

4.2 Lightpath Management Services

Figure 5 shows the system architecture in more detail and indicates the main service methods provided 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, 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. LPOs are created through the Admin functions. Supported customer operations on LPOs include: concatenating two LPOs, partitioning one LPO into many LPOs sharing common start and end points 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, the 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.

5. Conclusions

With the availability of dark fibers in some metropolitan areas and leased or condominium long-haul point-to-point wavelength networks, customers have an opportunity to construct and manage their own networks. Thus, the end-to-end 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.

The 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 prototype of a management system for customers to provision end-to-end lightpaths across multiple independent customer domains. This tool can be used as a traffic engineering mechanism for inter-domain applications, where separate BGP point-to-point lightpaths are set up operating in parallel with the normal hop-by-hop BGP routes. Interfaces to Grid applications are designed so that this tool can be easily integrated into Grid applications.

The architecture for customer-managed optical networks is the foundation for the future broker market of optical networking, where dynamic configuration of the network resources and 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.

6. References

1. Z. Zhang, J. Fu, D. Guo and L. Zhang. "Lightpath Routing for Intelligent Optical Networks", IEEE Network, Vol.15, No.4, July/August 2001, pp.28-35.
2. R. Isaacs and I. Leslie. "Support for Resource-Assured and Dynamic Virtual Private Networks", IEEE Journal on Selected Areas in Communications, Vol.19, No.3, Mar. 2001, pp.460-472.
3. G. M. Bernstein, V. Sharma and L. Ong. "Interdomain Optical Routing", Journal of Optical Networking, Vol.1, No.2, Feb. 2002, pp.80-92.
4. J. Strand, et al. "NNI Requirements and Framework Resource Document", Optical Internetworking Forum (OIF) technical document OIF2001.535, Jan. 2002.
5. R. Boutaba and A. Polyrakis. "Projecting Advanced Enterprise Network and Service Management to Active Networks." IEEE Network, Vol.16, No.1, Jan./Feb. 2002, pp.28-33.
6. M. Brunner. "A Service Management Toolkit for Active Networks", Proceedings of IEEE/IFIP Network Operations and Management Symposium (NOMS), Honolulu, Hawaii, USA, April 2000, pp.265-278.
7. T. Anderson and J. Buerkle. "Requirements for the Dynamic Partitioning of Switching Elements", IETF RFC 3532, May 2003.
8. E. Freeman, S. Hupfer and K. Arnold. JavaSpaces Principles, Patterns, and Practice, Addison-Wesley Pub Co, 1999.
9. S. Li, et al. Professional Jini, Wrox Press Inc, 2000.
10. F. Curbera, et al. "Unraveling the Web Services Web: an Introduction to SOAP, WSDL, and UDDI", IEEE Internet Computing, Vol.6, No.2, March/April 2002, pp.86-93.

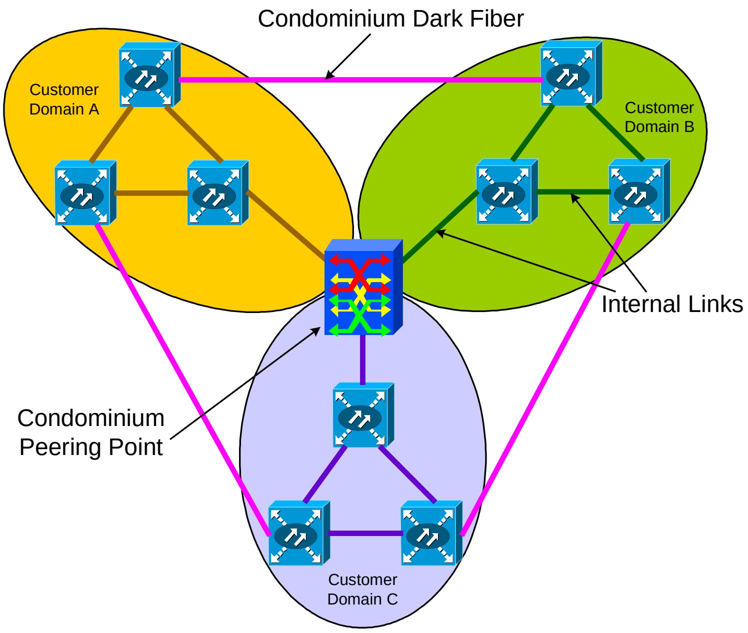

Figure 1. Customer domains may choose to interconnect to each other using condominium dark fibers or leased wavelength channels, or peer with each other at condominium peering points.

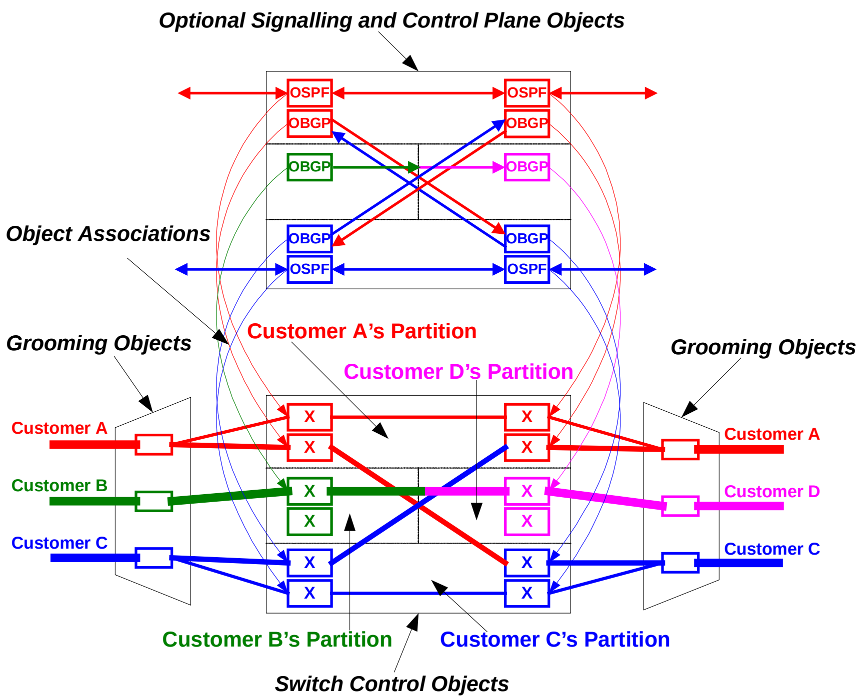

Figure 2. A peering point with customer control of ports and cross-connects

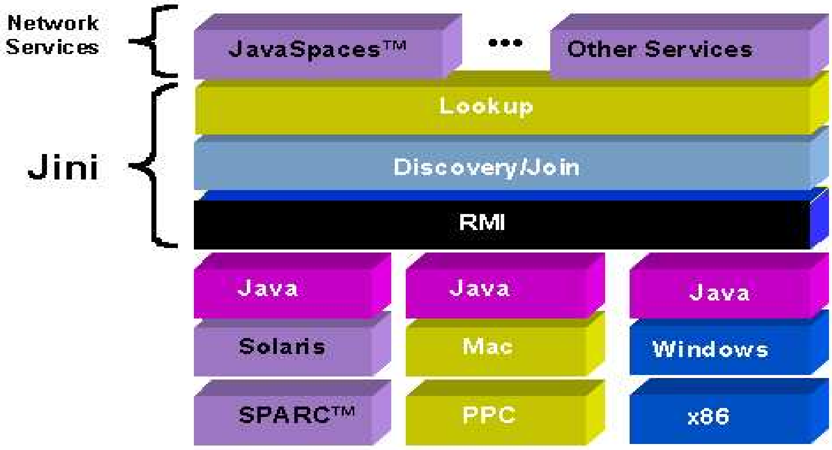

Figure 3. Jini Architecture

UCLP Architecture

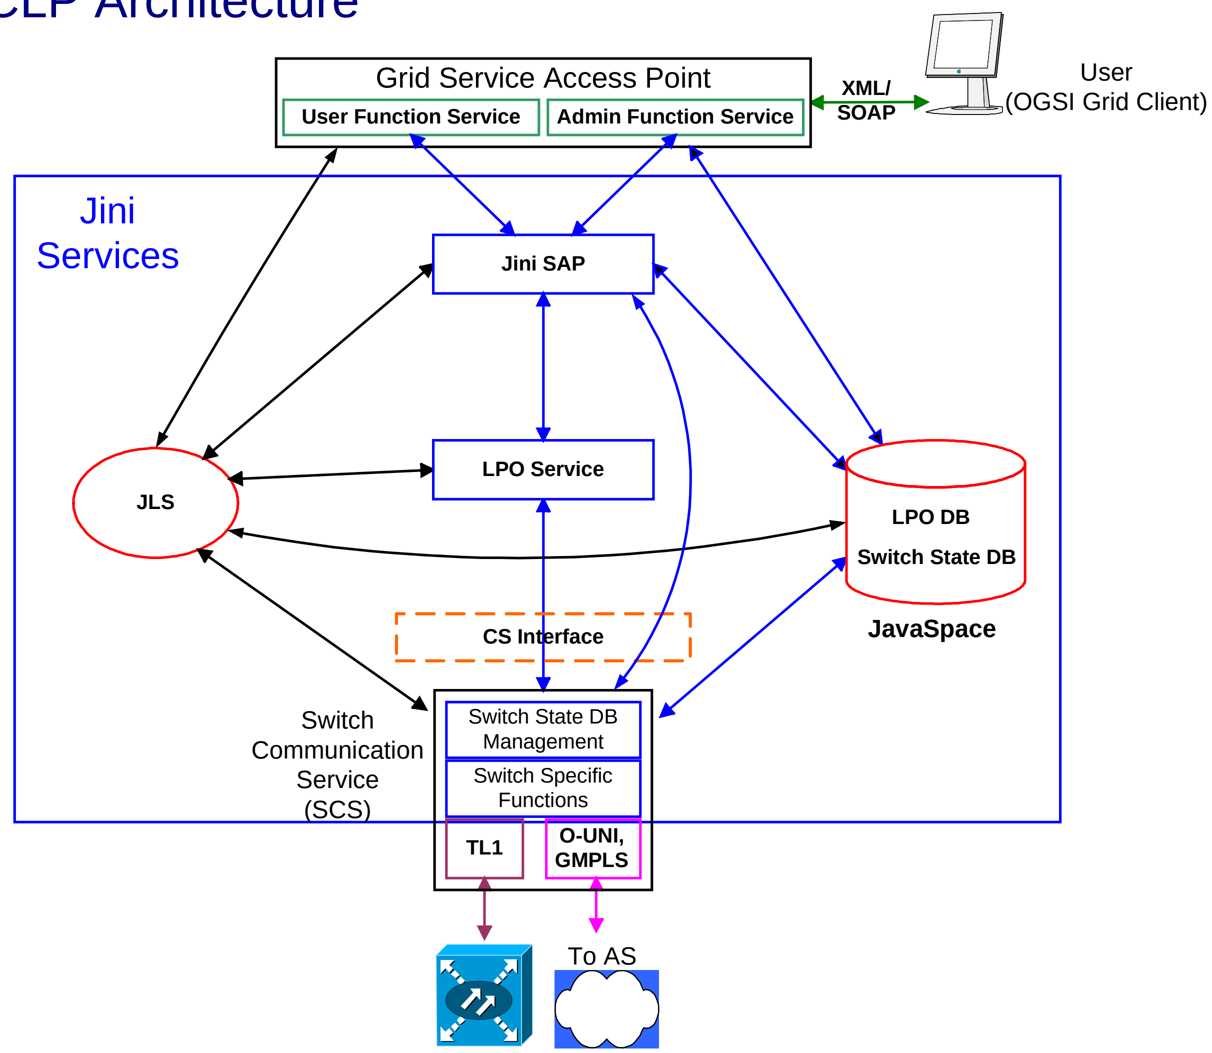

Figure 4. Overview of UCLP Architecture

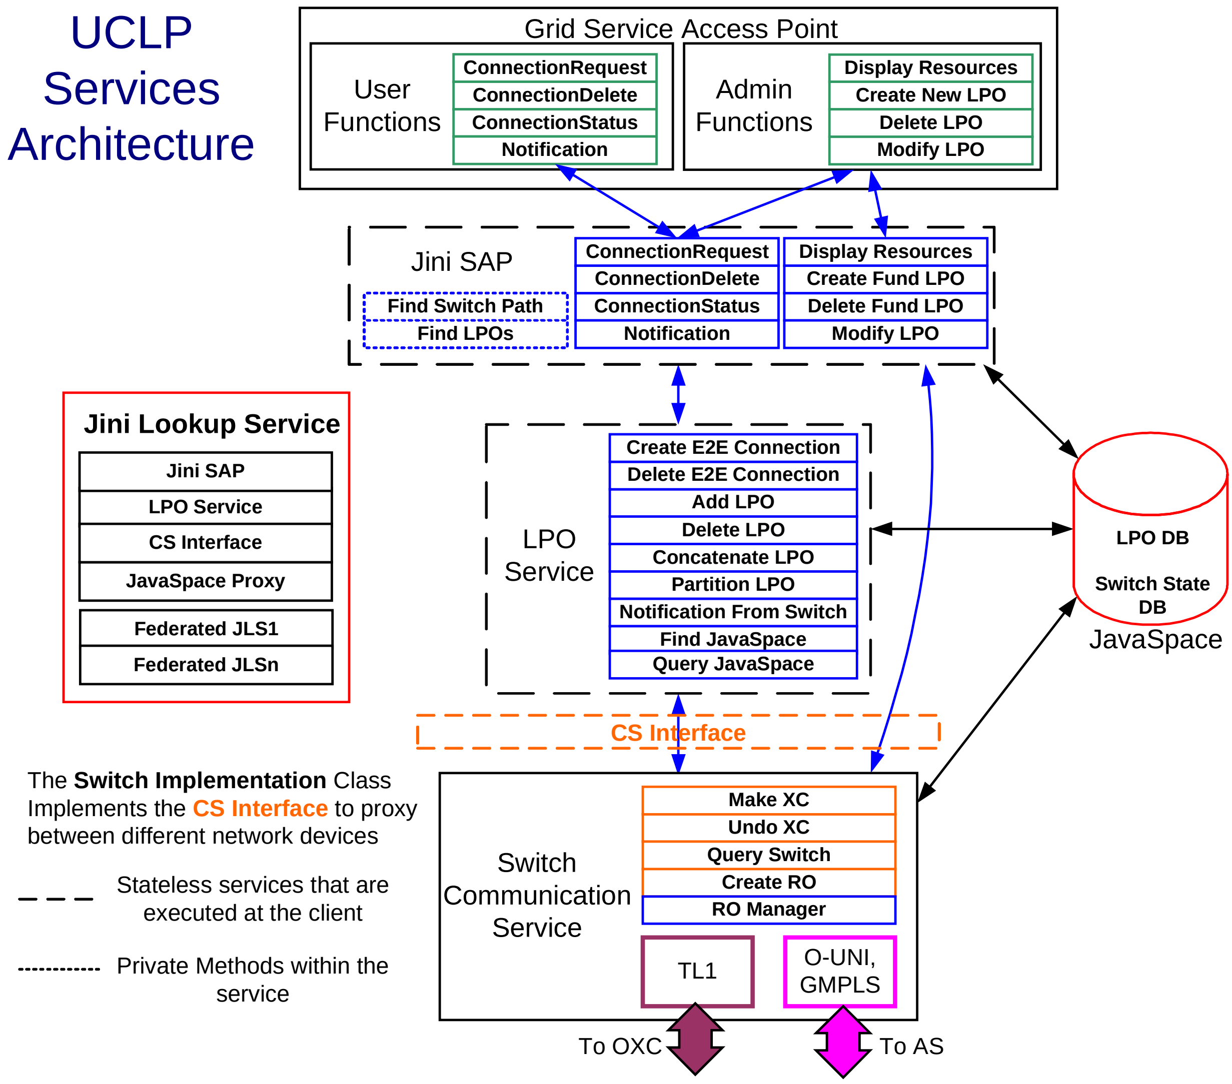

Figure 5. Detailed UCLP services architecture

Jing Wu (M'98) has been a Research Scientist in the Communications Research Centre Canada since 2001. His research interests mainly include control and management of optical networks, protocols and algorithms in networking, network performance evaluation and optimization, etc. In the past, he worked at Beijing University of Posts and Telecommunications, Beijing, China, as an Assistant Professor, Queen's University, Kingston, Ontario, Canada, as a Postdoctoral Fellow, and Nortel Networks Corporate, Ottawa, Ontario, Canada, as a System Design Engineer. He received his Ph.D. degree in systems engineering from Xian Jiao Tong University, Xian, China, in 1997.

J. Michel Savoie is the research program manager for the Broadband and Optical Networks group of the Broadband Network Technologies Research Branch at the Communications Research Centre, an Agency of Industry Canada. He maintains expertise in broadband systems and related technologies such as advanced IP, ATM and WDM based Optical Networks in order to provide advice on important national initiatives and to demonstrate the application of CRC technologies in a real operational environment. He is the project manager for the "User-Controlled End-to-End Lightpath Provisioning over CA*net 4" joint project with the University of Ottawa as part of CANARIE's new directed research program. Mr. Savoie holds a B.Sc. and M.Sc. in Electrical Engineering from the University of New Brunswick.

Scott Campbell is a network researcher in the Broadband and Optical Networking group at the Communications Research Centre. He has been working there since 2001 when he graduated from Dalhousie University in Halifax, Nova Scotia, Canada with a bachelor of Computer Science Degree. At CRC he is involved in the design and development of agent based network management and control software for all optical networks.

Hanxi Zhang obtained his M.A.Sc degree from the University of Ottawa in 2002. He is now with the Communications Research Centre Canada, as a research engineer on broadband and optical networks. His research interests include distributed computing, web services, and the protection/restoration for optical WDM networks.

Gregor v. Bochmann is a professor at the School of Information Technology and Engineering at the University of Ottawa since January 1998. Previously, he was 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 programming languages, compiler design, communication protocols, and software engineering and has published many papers in these areas. He has also been actively involved in the standardization of formal description techniques for communications protocols and services. His present work deals with methodologies for the design of distributed systems, quality of service negotiation for distributed multimedia applications, and optical networks.

Bill St. Arnaud is Senior Director Advanced Networks for CANARIE Inc., Canada's Advanced Internet Development Organization. At CANARIE Bill St. Arnaud 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 School of Engineering.