

A Grid Oriented Lightpath Provisioning System

Jing Wu, Hanxi Zhang, Scott Campbell
and Michel Savoie

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

Gregor v. Bochmann
School of Information Technology
& Engineering
University of Ottawa
Ottawa, Ontario, Canada
bochmann@site.uottawa.ca

Bill St. Arnaud
Canarie Inc.
Ottawa, Ontario, Canada
bill.st.arnaud@canarie.ca

Abstract— Multi-domain resource sharing is the fundamental feature of computing and data Grid applications. In this paper, we address the issue of multi-domain optical network resource sharing and present a network management system designed for end-to-end lightpath provisioning across multiple independently managed domains. Our network management system, called user-controlled lightpath provisioning (UCLP) system, is based on the Grid services concept and built on the Jini and JavaSpaces technologies. The UCLP system provides a traffic engineering approach to control network infrastructure for computing and data Grid applications.

Keywords— *Network Management Systems; User-Controlled Networks; Lightpath Provisioning Systems; Grid Applications*

I. INTRODUCTION

The communication infrastructure is important for distributed Grid computing. For bandwidth demanding Grid computing applications, bandwidth guaranteed connections are preferable compared to the best-effort Internet. Therefore, circuit-switched optical networks are a natural choice for Grid computing. With the emerging trend of dark fibre in metro areas, directly owned and managed by the user, and leased wavelength services in inter-city and long-haul areas, Grid computing may take advantage of such less expensive network infrastructures. In this paper, a Grid computing user is referred to as a user for short. Also a lightpath should be thought of as any bandwidth guaranteed connection because our management approach is independent of the underlying transmission technology.

The fundamental feature of Grid computing is multi-domain resource sharing, which distinguishes Grid computing from conventional distributed computing [1]. Such resource sharing potentially involves a large number of computing, data storage and communication resources that are distributed across multiple independently managed domains. A domain is a management entity that controls the inventory and configuration of resources. Within a domain, detailed management information is available, while across domain boundaries, only abstracted management information is disseminated to partner domains. Such multi-domain characteristics are challenging for network management and has not been solved in existing management technologies for a single domain [2].

In this paper, we address the issue of multi-domain optical network resource sharing and present a network management system designed for end-to-end lightpath provisioning across multiple independently managed domains. Our network management system, called user-controlled lightpath provisioning (UCLP) system, is based on the Grid services concept and built on the Jini and JavaSpaces technologies. In Section 2, we outline the emerging user-owned-and-managed optical networks, which are becoming a cost-effective network infrastructure for Grid computing. In Section 3, some technical challenges are summarized for network management we faced when we designed the UCLP system. The architectural design of the UCLP system is presented in Section 4. Section 5 contains the conclusions.

II. USER OWNERSHIP AND MANAGEMENT OF TRANSPORT NETWORKS

There are basically two types of user-owned and managed optical networks: metro dark fibre networks and long-haul wavelength networks. Schools, hospitals and government departments are acquiring their own dark fibres in metropolitan areas. They participate in so-called “condominium” dark fibre networks to better manage their connectivity and bandwidth. They light up the fibres with their own equipment and interconnect their fibres 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 users share the capital costs of deploying long-haul optical networks. In return, each user 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 fibre networks many thousands of kilometers without having to purchase and maintain their own optical repeaters and associated equipment.

In user-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 fibre 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 fibre network, users can add new wavelength channels. Since 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, user-owned optical networks have significant cost-savings when the demand for bandwidth increases dramatically over time.

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

User-managed networks reduce Internet costs via remote peering and transit. Users may directly peer with each other and more importantly set up bandwidth guaranteed connections to no-cost peering exchanges. The users manage their peering relationships without having to contact a central management body or pay expensive Internet transit fees.

Since the user 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.

III. TECHNICAL CHALLENGES

The first technical challenge is the management of networks with resources from different sources. Only the user 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) [3] 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 user-managed networks resemble some features of Virtual Private Networks (VPNs) [4]. A VPN is provisioned over public or third party network infrastructures to provide dedicated connectivity to a closed group of users. However, the VPN technology allows provisioning of user networks within a single provider's domain. Clearly this type of architecture is not practical with user-managed networks, where multiple suppliers provide resources. For the protection and restoration, the user, 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 users without co-ordination through centralized management. User-managed networks adopt the peer-to-peer architecture, in which users peer with each other. Each user domain not only receives transport services from other user 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 user to another involves at least two different user domains, and if transit is required, one or more intermediate transit domains may participate. So the collaboration among multiple independent users 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 user federations is a new issue.

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

IV. ARCHITECTURE OF THE NETWORK MANAGEMENT SYSTEM FOR GRID-ORIENTED LIGHTPATH PROVISIONING

A. Overview of the Architecture

The management system has three functional layers: resource management layer, lightpath provisioning layer and Grid application interface layer (Figure 1). With the resource management layer, individual users can directly access their portion of resources while being isolated from other users' resources on the same network element. Therefore, the resource management layer is a proxy providing a condominium type of resource management. The lightpath provisioning layer offers resource searching, peering, leasing and provisioning of end-to-end lightpath across multiple independently managed domains. With the lightpath provisioning layer, users may collaboratively share resources on a peer-to-peer basis. The Grid application interface layer allows users to manually set up end-to-end lightpath through a Graphic User Interface (GUI). A service access point is also designed for Grid applications to initiate the establishment of end-to-end lightpaths to transfer large amount of data.

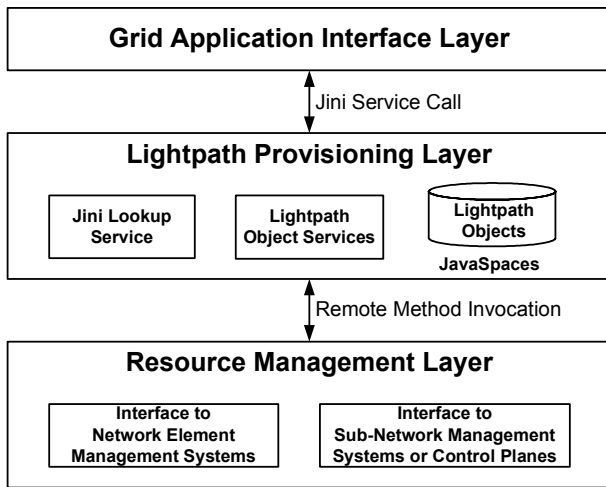


Figure 1. Architecture of the Grid-Oriented Lightpath Provisioning System

B. Lightpath Provisioning Layer

Each domain has a module for lightpath provisioning. The lightpath provisioning layer is designed in the Jini and JavaSpaces environment. The collaboration among domains is implemented as Jini services. A service directory is used for a domain to publicly advertise available resources to its users. In our implementation, we use Jini Lookup Services as service directories and JavaSpaces as distributed databases for available resources.

Jini and JavaSpaces provide a number of advantages for developing distributed applications [5]. Jini runs on top of Java and uses Remote Method Invocation (RMI) calls to access remote services. Jini also provides a set of application programming interfaces that hide the underlying complexity of distributed computing from the user. The Jini Lookup Service (JLS), a distributed service registry, allows users to find services without having to know where they are located if the JLS and users are within the same multicast domain. Each domain provides a JLS. Multiple JLS's may be introduced for redundancy and fault tolerance. By federating JLSs together, a user can find any service in any other domain. Jini 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 is chosen as a distributed data store for Java objects [6]. Communications through a JavaSpaces is loosely coupled. Anyone can take an object from a space without knowing (or caring) the details about 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 JavaSpaces as well as multiple operations over many JavaSpaces. Like the JLS, JavaSpaces is also persistent; an object will remain in a JavaSpaces until explicitly removed.

All functional components in the lightpath provisioning layer are implemented as Jini services. They all register themselves in the JLS. When a component calls others for the

first time, the component looks up the JLS to get the proxies for the others. For subsequent calls, the component will interact with the others directly by using the proxies without having to look up the JLS.

The lightpath provisioning utilities provide a GUI for administrators to create, delete or modify root lightpath objects, which represent the configured channels between network elements in the transport plane. Administrators may use a graphic tool to display available network resources for a domain.

C. Discovery and Federation of the Lightpath Provisioning Layers

The discovery and federation of the lightpath provisioning layers that belong to independently managed domains are critical to the UCLP system. It is the bootstrap phase for the UCLP system. JLS's find each other either by unicast or by multicast discovery. Unicast-discovery does not work well in a dynamic environment where JLS's come and go, or when there are a large number of JLS's. On the other hand, multicast-based discovery only works within the IP multicast range, and is usually not feasible in a wide area network environment. A special design is required to discover and federate the lightpath provisioning layers that belong to different domains.

A domain federation manager is introduced, in order for JLS's of different domains to be able to find and register with each other dynamically, without being bound by the IP multicast range. The domain federation manager consists of one or a cluster of JLS's hosted by a third party. We extend the standard JLS that comes with Jini, such that when a domain JLS starts up and registers with the domain federation manager via Jini unicast discovery, it will automatically get registered with all other domain JLS's, and at the same time all other domain JLS's will register themselves with the new domain JLS.

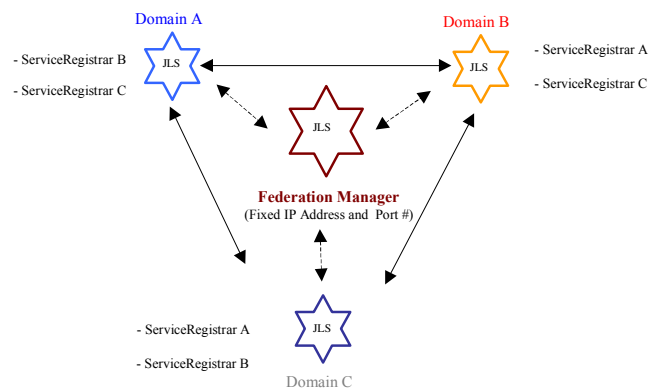


Figure 2. Discovery and federation of the lightpath provisioning layers that belong to independently managed domains

Figure 2 shows three domains A, B and C which discover each other and federate together via the domain federation manager. The IP address and port number of the JLS for the

domain federation manager is known to all users participating in the same Grid application. The result of the discovery and federation process is that the JLS for a domain contains the service registrars for all other JLS's.

D. *Lightpath Provisioning Operations*

A user may access network resources by a number of approaches. In the first case, if all the required network resources are within one domain that serves this user or within several domains that this user subscribes to, the lightpath is local. The resource searching and lightpath provisioning do not pose new technical challenges compared to the traditional single domain management.

In the second case, part of the required network resources are beyond the domains that the user requesting the lightpath subscribes to. Therefore, the resource searching and lightpath provisioning rely on the collaborations and partnerships that domains have established. The domain that serves the user requesting the lightpath will delegate the action to search the required resources in the collaborative domains that this domain has partnership relations with, then control the switching actions in other domains.

The third case is more complicated, where even the collaborative group of domains do not acquire sufficient resources for the lightpath request. A peering of two segments of lightpaths is necessary to provision the entire end-to-end lightpath. The originating and terminating ends of the lightpath need to establish a segment of lightpath to a common peering point. Then a cross-connection at the peering point makes the end-to-end lightpath.

Our UCLP system is to provide a solution for these operational scenarios. The lightpath provisioning layer offers resource searching, peering, leasing and provisioning of end-to-end lightpath across domains. With the lightpath provisioning layer, users may collaboratively share resources on a peer-to-peer basis. A user may access the JavaSpaces in its domain and search for available lightpath objects. Its search may also be delegated by its domain to the JavaSpaces of other collaborative domains.

Our UCLP system provides the following operations on lightpath objects: add, concatenate and partition LPOs, as well as the reverse operations [7]. A user may take available lightpath objects from its domain. Moreover, a user may create new derived lightpath objects and put them into the JavaSpaces in its domain and make them available for other users to use. For example, a user may concatenate several lightpaths into one longer lightpath. Also, a user may partition a lightpath with large bandwidth into several lightpaths with smaller bandwidth.

E. *Grid Application Interface Layer*

The basic application interface is a GUI for a user to request an end-to-end lightpath from a given entry port of a given switch to a given exit port on another given switch, possibly belonging to a different domain. A user may also specify a peering point as an intermediate switch to peer with another incoming lightpath request.

In addition, our system is designed to support Grid applications defined in the Open Grid Service Infrastructure (OGSI). Our system enables Grid applications to dynamically set up end-to-end lightpaths across multiple domains to transfer vast amount of data. The primary application for this technology is Lambda Grids to support high-end scientific research in high-energy physics, astronomy, bio-informatics, etc. Lightpath-aware Grid applications may invoke a special QoS mechanism where a separate optical Border Gateway Protocol (BGP) point-to-point connection is set up 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 a point-to-point bandwidth guaranteed path rather than a best-effort path.

The Grid application interface layer uses a Grid service access point to advertise its service instantiation through a well-known process described by the OGSI implementation such as a Web Services Inspection Language (WSIL) pointer or a Universal Description, Discovery and Integration (UDDI) database [8]. End users communicate their requests to the Grid service access point module using the Simple Object Access Protocol (SOAP) adopted for Web Services. These requests are converted into Java procedure calls within the Grid service access point which then performs these calls on its local lightpath provisioning module as a Jini service. Detailed functional description can be found in [7].

V. CONCLUSIONS

The widely deployed user-owned and managed optical networks have shown significant cost-saving, which is a promising infrastructure for emerging computing and data Grid applications. However, to fully achieve seamless integration with computing and data Grid applications, new network management architectures are required. In this paper, we presented a new architecture based on the Grid service concept. It uses three functional layers for resource management, lightpath provisioning and Grid application interface. The key software components in our network management system are implemented using the Jini and JavaSpaces technologies. Our prototyping has demonstrated the design is feasible [9].

ACKNOWLEDGMENT

We wish to 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.

REFERENCES

- [1] Z. Nemeth and V. Sunderam. "Characterizing Grids: Attributes, Definitions, and Formalisms", *Journal of Grid Computing*, Vol.1, No.1, 2003, pp.9-23.
- [2] R. Boutaba, W. Golab, Y. Iraqi, T. Li and B. St. Arnaud. "Grid-Controlled Lightpaths for High Performance Grid Applications", To appear in *Journal of Grid Computing*, 2004.
- [3] 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.

- [4] 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.
- [5] S. Li, et al. Professional Jini, Wrox Press Inc, 2000.
- [6] E. Freeman, S. Hupfer and K. Arnold. JavaSpaces Principles, Patterns, and Practice, Addison-Wesley Pub Co, 1999.
- [7] J. Wu, S. Campbell, M. Savoie, H. Zhang, G. Bochmann and B. St. Arnaud, "User-Managed End-To-End Lightpath Provisioning Over CA*net 4", Proc. of the National Fiber Optic Engineers Conference (NFOEC 2003), Vol.1, pp.275-282, Orlando, Florida, USA, Sept.7-11, 2003.
- [8] 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.
- [9] H. Zhang, J. Wu, S. Campbell, M. Savoie, G. Bochmann and B. St. Arnaud, "A Distributed Testbed for a User-Controlled Lightpath Provisioning System Using Jini/JavaSpaces Technologies", To be presented at the 30th European Conference on Optical Communications (ECOC 2004), Workshop on Optical Networking for Grid Services, Stockholm, Sweden, Sept. 5-9, 2004.