Available online at www.sciencedirect.com
ScienceDirect
Optical Switching and Networking 5 (2008) 139–149
Optical Switching and Networking
www.elsevier.com/locate/osn

CHEETAH virtual label switching router for dynamic provisioning in IP optical networks

Zhaoming Li, Qiang Song*, Ibrahim Habib

City University of New York, New York, NY 10031, United States

Received 29 August 2007; received in revised form 27 December 2007; accepted 24 January 2008
Available online 13 February 2008

Abstract

Enabling new IP-based services such as triple and quad-play, as well as eScience applications at predetermined quality of service (QoS) measures, require the provisioning of guaranteed bandwidth pipes at varying granularities (e.g. from few Mbps to several Gbps and above). Dynamic provisioning of bandwidth pipes, whereby a connection is dynamically setup and released upon signalling, is a cost-effective method of enabling such services. Dynamic provisioning is a new paradigm in network control and management (NC&M) that requires the introduction of control plane (i.e. routing and signaling) capabilities within network elements such as routers, layer 2 switches and layer 1 cross-connects.

In this paper we share our experience in the design and deployment of a Generalized Multiple Protocol Label Switching (GMPLS) control plane for layer 2 switches in the experimental Circuit-Switched High-Speed End-to-End Transport Architecture (CHEETAH). We call this software engine CHEETAH Virtual Label Switching Router (CVLSR). CVLSR allows non-GMPLS devices (e.g. Ethernet switches, routers and other cross-connects) to participate in the dynamic provisioning of end-to-end bandwidth-guaranteed connections. It extends the dynamic provisioning of connections to the end-users across different administrative domains. We have successfully deployed the CVLSR in CHEETAH optical network across HOPI/DRAGON network. The interoperability of the CVLSR with commercial GMPLS SONET-based cross-connect switches has been demonstrated.

Keywords: GMPLS; Control plane; Dynamic provisioning; IP optical network

1. Introduction

Recent advances in optical networking technologies have made it possible to realize the vision of enabling new IP-based bandwidth-on-demand (BoD) services directly to end-users' devices. Examples of BoD services include video on demand (VoD), IP-TV, multimedia conferencing, as well as grid computing. The above services can be classified into two main classes: (a) those requiring high capacity pipes (e.g. multi Gbps) for file transfers of large datasets, and (b) those requiring ulow to medium bandwidth pipes (e.g. few Mbps) for visualization and steering purposes. Both classes may be further subdivided according to the holding time of the connection, i.e. short duration or long duration. Furthermore, applications may require immediate request (IR) or book-ahead (BA) service. In all cases, these emerging applications require the dynamic setup of guaranteed bandwidth granularity connections in order to ensure the delivery of the content at required quality of service measures.

* Corresponding author.
E-mail addresses: zli@ccny.cuny.edu (Z. Li), qsong@ccny.cuny.edu (Q. Song), habib@ccny.cuny.edu (I. Habib).

1573-4277/$ - see front matter © 2008 Elsevier B.V. All rights reserved.
doi:10.1016/j.osn.2008.01.005

Enabling BoD services, at cost-effective prices, requires a means of sharing the usage of expensive physical bandwidth resources among users based upon time or statistical principles. This can be easily achieved via dynamic provisioning that utilizes the control plane capabilities (i.e. routing and signalling). Obviously this requires equipping network elements with software engines that implements the control plane functionalities. Based upon users' signalling messages bandwidth pipes could then be setup/released dynamically for a few seconds or hours, thus allowing truly BoD services and avoiding the prohibitive costs of statically provisioned pipes that last for months or years.

The above clearly draws a fundamentally new picture of networking that is quite different from what is happening in the telecommunications industry today. Today, users interact with networks in a client-server model. Users provide the carrier with requirements for "static networking demands". The carrier, then, uses a centralized network management (CNM) system to provision static paths to fulfill the customer request. This provisioning process is quite complex and often requires manual intervention. In addition, the centralized approach is not scalable for "Bandwidth-on-Demand" (BoD) services where requests for service are frequent and require near real time provisioning. The CNM is suitable for large bandwidth pipes (e.g. several Gbps and above) where significant resources of the network are to be committed, thus needing prior authorization and network capacity planning from a centralized agent.

The Generalized Multiple Protocol Label Switching (GMPLS) [1] control plane is one enabling technology that could be used to set up/release end-to-end deterministic bandwidth connections across layers' 3, 2 or 1 devices. Despite the recent developments in the standards and research communities, our experience with the deployment issues of GMPLS-based control plane in either production or experimental networks has been rather limited. To overcome this shortcoming, federal agencies such as the Department of Energy (DoE) and National Science Foundation (NSF) have recently funded several new experimental high performance optical networks such as USN [2], DRAGON [3], and CHEETAH [4]. Furthermore, the internet2 community and several research and industrial partners have started a new experimental network called HOPI (hybrid optical packet infrastructure) [5]. One of the main objectives of HOPI is to foster the development of experimental control plane solutions for connection-orientated communications paths through both packet-based and optical-based networks.

The CHEETAH project is a comprehensive effort to develop the infrastructure and networking enabling technologies to support a broad class of eScience projects such as the Terascale Supernova Initiative, which is a national collaboration centred at the Oak Ridge National Laboratory (ORNL) and involves eight universities. Its data plane consists of OC-192 SONET and 10 Gbps Ethernet connections linking several enterprise networks such as the supercomputer at ORNL and the Centaur physics laboratory at North Carolina State University (NCSU) (see Fig. 1). Other connections include several 1 Gbps Ethernet links to City University of New York (CUNY) through Internet2 Hybrid Optical and Packet Infrastructure/Dynamic Resource Allocation via GMPLS Optical Networks (HOPI/DRAGON), to University of Virginia (UVa) and other enterprise networks as illustrated in the Figure. Three Point of Presences (PoPs) equipped with the Sycamore SN16000 SONET cross-connect switches (SONET-XC) form the backbone of the network. The users, in each enterprise network, gain access to the network via Gigabit Ethernet switch.

CHEETAH uses a distributed GMPLS control plane approach to enable application-initiated, on-demand dynamic provisioning of end-to-end circuits at various granularities (e.g. 10 Mbps to 10 Gbps). An important premise is the dynamic sharing of the large bandwidth pipes of backbone networks by many users, thus cutting down the expensive costs of communications at these capacities. In order to extend the connections to the end-users, each Ethernet switch is equipped with a GMPLS-enabled engine called CHEETAH Virtual Label Switching Router (CVLSR) [6]. The CVLSR enables non-GMPLS Ethernet switch to participate in the dynamic setup and release of bandwidth guaranteed connections through controlling the VLANs within the switch. Hence, it allows the users to share the network resources dynamically and makes it possible to extend the connections across the local area network to the end-users. It offers a scalable and cost-efficient solution for adding users to the network, whereas adding users directly to PoP nodes will quickly exhaust the capacity of the PoP ports, i.e. the GbE ports of the Sycamore switches in CHEETAH. The CVLSR is developed based upon the DRAGON VLSR [7], which is used for dynamic provisioning in HOPI/DRAGON network. The DRAGON VLSR, in turn, is based upon the implementation of an open source Open Shortest Path First (OSPF) engine (GNU Zebra) [8] and Resource ReSerVation Protocol (RSVP) engine (KOM RSVP) [9, 10].