Resilient Global IP/Optical Networks: DARPA CORONET

Ann Von Lehmen

Telcordia Technologies, Red Bank,NU USA,07701 avl2@research.telcordia.com

Abstract *This paper describes progress on the DARPA CORONET Program, which is focused on a 20-100 Tbps, highly dynamic, resilient IP/Optical network. The Program requires extremely rapid provisioning and restoration times, and resilience against multiple failures.*

Introduction

The past several years have seen significant activity both commercially and within R&E (Research and Education) networks in developing and demonstrating capabilities for "on-demand" optical and data services.

The recently launched DARPA CORONET program moves well beyond the current state of the art and envisions a highly dynamic, high capacity, resilient next-generation IP-over-Optical Layer (IP/OL) network[1]. The CORONET service set consists of IP/MPLS and optical connection services defined by provisioning times, restoration requirements, blocking ratios, and quality of service requirements. To illustrate: CORONET requires provisioning and restoration times as short as 100 ms for CONUS and resiliency up to 3 failures for some traffic. A significant fraction of IP and optical services have holding times of less than a few hours and 1 minute respectively, resulting in substantial network churn. There are capacity constraints imposed which limit both working and restoration capacity. These constraints essentially force solutions that are cost effective; indeed one of the Program's stated goals is transition of the technology to commercial practice. However, meeting all of these requirements simultaneously is very challenging. In this paper, we present an overview of the CORONET program requirements, and the protocols, algorithms and methodologies our team (Telcordia, AT&T, Nortel and USC) has used in addressing them. These approaches are validated through simulation on a 100-node worldwide network, the results of which will also be presented. Below we summarize our approaches in two of the key CORONET challenge areas: rapid provisioning and restoration.

Rapid Provisioning for CORONET

The most aggressive provisioning requirements within CORONET occur for a class of wavelength services that must be provisioned in <100 ms in CONUS. These services are characterized by holding times of less than a minute and service granularities that range from a single to eight wavelengths, with a blocking requirement of < 10^{-3}. These services constitute 40% of the wavelength service traffic. The challenge then essentially boils down to meeting the extremely short provisioning times and low blocking requirements in the presence of stale state information without adding substantial capacity to the network. In addition, cost considerations require that we minimize the amount of wavelength conversion required. Our approach to solving this problem is a distributed provisioning protocol called the 3-Way Handshake (3WHS)[2]. The protocol is designed to 'probe' and collect current information on available link capacities and wavelength and transponder availability for multiple possible paths between a given pair of nodes, and then to use that information to select and set up a path. The probing, resource identification and selection and optical crossconnect setup occur across three signaling passes between end nodes. This approach is very advantageous in that it avoids both the need for extremely rapid link state updates, and large amounts of excess capacity to guarantee low blocking[2]. It also enables us to identify and use all-optical connections whenever possible, thus making efficient use of transponders in the network. The paths chosen for probing are determined from link state information providing aggregate available capacity on the fiber links. The aggregate spare capacity changes more slowly in accordance with the changes in aggregate traffic intensity, so this link state information becomes stale on a much longer timescale (e.g., minutes), and therefore requires much less frequent updates.

We have performed an analysis which shows that the protocol can meet the CORONET provisioning time requirements with minimum additional capacity [2]. We have also implemented the protocol fully in an OPNET simulation platform, and will present results for provisioning times and blocking. In our presentation, we will also describe results for provisioning of IP and optical service classes requiring 2 second set-up times, and with scheduled service requirements.

Restoration for CORONET

Stringent resiliency requirements are an additional major CORONET Program challenge: all services, except IP-best effort are restorable against up to one failure; up to 20% are restorable against two failures, and up to 5% restorable against 3 failures. As noted above, there is also a an aggressive restoration time requirement: 100 ms for CONUS and 250 ms for world-wide connections. CORONET also requires that the restoration capacity in the network not exceed

978-3-8007-3173-2

75% of the working capacity. So it is imperative that the methodologies used be capacity efficient.

In our current approach, restoration for IP traffic occurs in the IP layer, and for optical services in the optical layer. It heavily leverages earlier work on restoration done by AT&T Labs[3]. We are also exploring the potential benefits of a more integrated approach to restoration.

For the optical layer, we use a variant of shared mesh restoration and a signaling protocol called ROLEX pioneered by AT&T Labs[3]. For IP services, we use a novel Two Phase Fast ReRoute (2PFRR) protocol[3]. In brief, this is how it works: upon failure of an IP link, the first phase of the 2PFRR initiates a switch-over of working traffic to pre-computed backup tunnels between the pair of routers adjacent to the failure. This approach results in extremely rapid restoration times. The second phase consists of a slower reroute of the primary tunnel over an end-to-end path optimized for the current state of the network. In both phases we employ optimized MPLS-TE. This is a critical factor which enables us to maximize sharing of bandwidth among all failure scenarios (including multiple failure scenarios), and substantially reduces the additional capacity needed for restoration.

In this talk, we will present the results of OPNET-based simulations which illustrate the performance of these restoration protocols on the CORONET network.

Conclusions

The CORONET Program presents a set of extremely demanding requirements and stringent constraints for a dynamic network with a wide variety of on-demand IP and optical services. We will report on progress towards meeting the Program challenges, as well as insights gleaned along the way.

We would like to acknowledge the support of the CORONET Program, Contract N00173-08-C-2011.

References

1 DARPA BAA06-29 “Dynamic Multi-Terabit Core Optical Networks: Architecture, Protocols, Control and Management (CORONET)
2 R.Skoog and A. Niedhardt, Proc. OFC/NFOEC NTuB5 (2009)
3 A. Chiu et al, Proc. OFC/NFOEC NThF3 (2009), and references therein.

