Computer Networks 52 (2008) 3130–3147

Contents lists available at ScienceDirect

Computer Networks

journal homepage: www.elsevier.com/locate/comnet

Resource optimization algorithms for virtual private networks using the hose model

Monia Ghobadi *, Sudhakar Ganti, Gholamali C. Shoja

Department of Computer Science, University of Victoria, BC, Canada V8W 3P6

ARTICLE INFO

Article history:
Received 23 January 2008
Received in revised form 28 July 2008
Accepted 5 August 2008
Available online 14 August 2008

Responsible editor: J. Domingo-Pascual

Keywords:
Virtual private networks
Hose model
Quality of service
Provisioning cost
Spanning tree

ABSTRACT

Virtual private networks (VPNs) provide a secure and reliable communication between customer sites over a shared network. With increase in number and size of VPNs, service providers need efficient provisioning techniques that adapt to customer demands. The recently proposed hose model for VPN alleviates the scalability problem of the pipe model by reserving for aggregate ingress and egress bandwidths instead of between every pair of VPN endpoints. Existing studies on quality of service guarantees in the hose model either deal only with bandwidth requirements or regard the delay limit as the main objective ignoring the bandwidth cost. In this work we propose a new approach to enhance the hose model to guarantee delay limits between endpoints while optimizing the provisioning cost. We connect VPN endpoints using a tree structure and our algorithm attempts to optimize the total bandwidth reserved on edges of the VPN tree. Further, we introduce a fast and efficient algorithm in finding the shared VPN tree to reduce the total provisioning cost compared to the results proposed in previous works. Our proposed approach takes into account the user preferences in meeting the delay limits and provisioning cost to find the optimal solution of resource allocation problem. Our simulation results indicate that the VPN trees constructed by our proposed algorithm meet maximum end-to-end delay limits while reducing the bandwidth requirements as compared to previously proposed algorithms.

1. Introduction

Globalization has revolutionized the business world in the last couple of decades. Instead of simply dealing with local or regional concerns, many businesses now have to think about global markets. Many companies have facilities spread out around the world, and hence they all need a way to maintain fast, secure and reliable communications wherever their offices are. Until fairly recently, this meant the use of leased lines to maintain a wide area network (WAN)[6]. Leased lines provided a company with a way to expand its private network beyond its immediate geographic area. A WAN had obvious advantages over a public network, like the Internet, when it came to reliability, performance and security. But maintaining a WAN, particularly when using leased lines, can be quite expensive and often the cost increases with distance between the offices.

As the popularity of the Internet grew, businesses turned to it as a means of extending their own private networks. First came intranets, which are password-protected sites designed for use only by the company employees. Now, many companies are creating their own Virtual Private Network (VPN) to accommodate the needs of remote employees and distant offices.

A VPN is a group of computer systems connected as a private network that communicates over a public network. VPNs offer a cost-effective, scalable, and manageable way to create a private network over a public infrastructure such as a service provider's frame relay [7], ATM [3], or

* Corresponding author. Tel.: +1 647 899 0370.
E-mail addresses: monia@cs.uvic.ca (M. Ghobadi), sganti@cs.uvic.ca (S. Ganti), gshoja@cs.uvic.ca (G. C. Shoja).

1389-1286/$ - see front matter Crown Copyright © 2008 Published by Elsevier B.V. All rights reserved.
doi:10.1016/j.comnet.2008.08.010

IP network [20]. For this reason, VPNs are deployed by businesses to meet their networking and communication needs and have rapidly emerged as leading solutions for multi-site enterprise communication demands.

The emergence of IP technologies such as MPLS [1] and RSVP-TE [17] have made it possible to realize IP-based VPNs that can provide the end customers with QoS guarantees. Thus, an IP VPN service that replaces the traditional point-to-point connectivity between sites using legacy solutions must offer comparable performance, security and functionality.

There are two popular models for providing QoS in the context of VPNs – the *pipe* model [1] and the *hose* model [2]. The pipe model is a simple service model for an IP VPN which emulates the private line or frame relay service. As depicted in Fig. 1, in the pipe model, a VPN customer purchases a set of customer-pipes, i.e., allocations of specific bandwidth on paths between every source-destination pair of the VPN endpoints. The network provider would need to provision adequate bandwidth along the path of each pipe to ensure that the Service Level Agreement (SLA) is satisfied. The primary disadvantage of this approach is that it requires the customer to have precise knowledge of its own traffic matrix between all the VPN sites. Moreover, resources made available to a customer-pipe cannot be allocated to other traffic.

Due to the progress in security and the success of IP networking technologies, the number of endpoints per VPN is growing, and the communication patterns between endpoints are becoming increasingly difficult to predict. It is expected that users will be unwilling to, or simply unable to predict loads between pairs of endpoints. Similarly, it will become increasingly difficult to specify QoS requirements on a point-to-point basis, as is the conventional approach.

The hose model, introduced by Duffield et al. in [2], serves as both a VPN service interface as well as a performance abstraction. A hose offers performance guarantees at a given endpoint for the traffic to and from the set of all other endpoints in the VPN. Thus, the hose service interface allows the customer to send traffic into the network without the need to predict point-to-point loads. Fig. 2 illustrates an example of the use of the hose model. Each

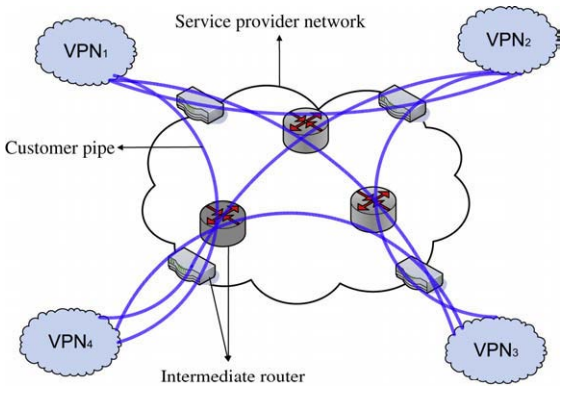

Fig. 1. VPN pipe model.

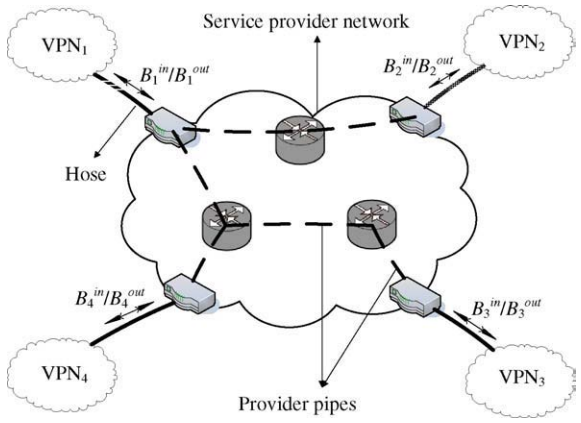

Fig. 2. VPN hose model.

VPN endpoint i is connected to the network by a hose, which is specified by its aggregate ingress and egress bandwidths (B_i^{in} and B_i^{out}, respectively). B_i^{in} is the amount of aggregate traffic from all endpoints to endpoint i and B_i^{out} is the amount of aggregate traffic from endpoint i to all other endpoints of the same VPN. Thus, in the hose model, the VPN service provider supplies the customer with certain guarantees for the traffic that each endpoint sends to and receives from other endpoints of the same VPN. The customer does not have to specify how this traffic is distributed among other endpoints. As a result, in contrast to the pipe model, the hose model does not require a customer to know its own complete traffic matrix.

Our goal is to address the resource management problem in VPNs and introduce algorithms that enable efficient resource provisioning with QoS guarantees. Our algorithms are based on the hose service model, which is a widely accepted service specification. As we will explain in Section 2, existing studies on quality of service guarantees in the hose model either deal only with bandwidth requirements or regard the delay limit as the main objective ignoring the total provisioning cost. In this work we propose a new approach to enhance the hose model to guarantee end-to-end delay limits between endpoints while optimizing the provisioning cost. Further, we introduce a fast and efficient algorithm in finding a shared VPN tree with minimum total provisioning cost compared to the results proposed previously in [23]. We connect VPN endpoints using a tree structure and our algorithm attempts to optimize the total bandwidth reserved on edges of the VPN tree. Our proposed approach takes into account the user preferences in meeting the delay limits and provisioning cost in order to find the most optimal solution with respect to user specified parameters. Our simulation results indicate that the VPN trees constructed by our proposed algorithm meet the delay limits while reducing the bandwidth requirements as compared to previously proposed algorithms [8,23].

2. VPN network model

A VPN network is modeled as a connected graph $G = (V, E)$ where V is the set of nodes and E is the set of

ID	Title	Pages
452522	Resource optimization algorithms for virtual private networks using the hose model	18