Quantifying Trade-offs in Resource Allocation for VPNs *

Satish Raghunath
Dept. of ECSE, RPI
raghus@rpi.edu

Shivkumar Kalyanaraman
Dept. of ECSE, RPI
shivkuma@ecse.rpi.edu

K.K. Ramakrishnan
AT&T Labs - Research
kkrama@research.att.com

ABSTRACT

Virtual Private Networks (VPNs) feature notable characteristics in structure and traffic patterns that allow for efficient resource allocation. A strategy that exploits the underlying characteristics of a VPN can result in significant capacity savings to the service provider.

There are a number of admission control and bandwidth provisioning strategies to choose from. We examine trade-offs in design choices in the context of distinctive characteristics of VPNs. We examine the value of signaling-based mechanisms, traffic matrix information and structural characteristics of VPNs in the way they impact resource utilization and service quality. We arrive at important conclusions which could have an impact on the way VPNs are architected. We show that the structure of VPNs profoundly influences achievable resource utilization gains with various admission control and provisioning schemes.

Categories and Subject Descriptors: C.2.1 [Computer-Communication Networks]: Network Architecture and Design; C.4 [Performance of Systems]: Design Studies

General Terms: Performance, Design

Keywords: Point-to-multipoint, Virtual Private Networks, Hose Model, Point-to-Set

1. INTRODUCTION

Traditional models for VPNs have involved building point-to-point links provisioned at peak bandwidth demand. Recent proposals like the Hose Model [1] and the Point-to-Set [2, 3] architecture provide a superior alternative that allow the provider to exploit multiplexing gains by leveraging a common core network infrastructure shared among all customers. These models propose admission control and adaptive provisioning mechanisms to preserve Service Level Agreements (SLAs). Our work is complementary to such solutions in that we provide guidelines to a designer in architecting the components of these models.

*Supported in part by NSF contracts ANI9806660, ANI9819112, Intel Corp. and Nortel Networks Limited.

SIGMETRICS/Performance'04, June 12–16, 2004, New York, NY, USA.
ACM 1-58113-873-3/04/0006.

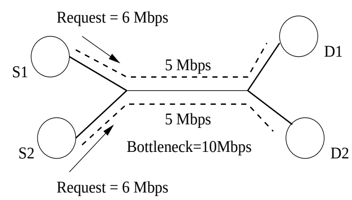

Figure 1: By statically apportioning link capacity among paths, we can substitute the function of signaling. E.g., each network edge sees $5Mbps$ as available and rejects the second $6Mbps$ request

2. DESIGN CHOICES

In order to assure SLAs are not violated, a provider needs a mechanism to ensure that an admitted request can be appropriately serviced. This involves an admission control algorithm and a bandwidth provisioning module. These components can be built on a per-link basis with the help of signaling protocols or with edge-based algorithms that do not assume core network support. An additional dimension is the customer traffic information these components exploit in terms of the traffic matrix and structure of VPNs. We briefly describe these choices below.

Statistical admission schemes typically evaluate the probability of violating a given QoS metric; e.g., if L denotes the random variable for loss, the admission condition could be $Pr\{L > 0\} \leq \epsilon$ where $\epsilon \in (0, 1)$ is a pre-specified parameter. A deterministic strategy is simpler in that it usually just involves a peak bandwidth requirement that is reserved inside the network. A statistical condition exploits variations in traffic and delivers far higher resource utilization but is complex in implementation.

In the presence of signaling support, admission control and bandwidth reservation decisions can incorporate information from each hop of a path along which a flow is admitted (e.g., link capacity, utilization etc.). The downside is the complexity of the system; changes in routing and link failures can cause considerable overhead. In order to build a mechanism that is edge-based we need to substitute the functionality achieved by signaling, viz., ensuring that link capacities are not overbooked and are optimally utilized.

A simple means of achieving this is to statically apportion link capacities among various paths in the network

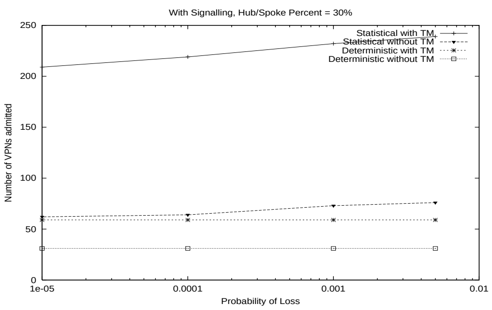

Figure 2: Number of Admitted VPNs with 30% of the generated VPNs being of the Hub/Spoke type

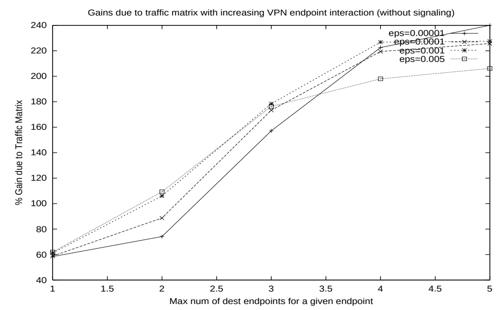

Figure 3: Higher the number of endpoints with which a node communicates more important the TM

Algorithm 1 Static Path Capacity of path p

Denote capacity of link l as C_l and that of path p as C_p
Input: $\mathcal{L}_p \leftarrow \{$ Set of links in p $\}$
Input: $\mathcal{P}_l \leftarrow \{$ Set of paths traversing link l $\}$
for each link $l \in \mathcal{L}_p$ **do**
  $|P_l| \leftarrow$ Number of paths traversing link l
  $S_p(l) \leftarrow \frac{C_l}{|P_l|}$
end for
$C_p \leftarrow \min_{l \in \mathcal{L}_p} S_p(l)$

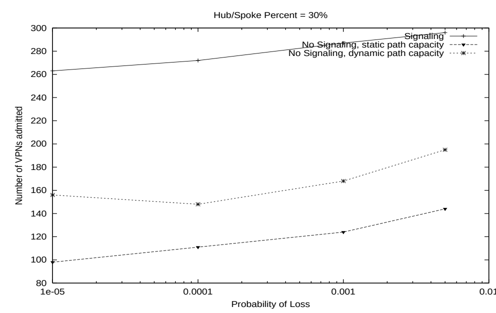

Figure 4: Dynamic capacity allocation considerably improves performance compared to static scheme

to create virtual edge-to-edge links as illustrated in Fig. 1 and outlined in algorithm 1. Although this method allows distributed admission control at the edges of the network, it results in wastage of bandwidth. We implement a measurement-based dynamic path capacity scheme (algorithm 2) that attempts to obtain performance that is closer to that provided by signaling, yet simpler in deployment. This scheme involves using the static path capacity algorithm to start with but adapts allocations depending on the utilization along the path (computed using popular traffic matrix computation algorithms [4]).

Algorithm 2 Dynamic Path Capacity of path p

For a path p, compute capacity C_s form Algorithm 1
Path Capacity $C_d \leftarrow \beta C_s$
Adapt C_d based on measurement information

An important attribute of provisioning decisions is the amount of customer traffic information that is incorporated. In the context of VPNs, the amount of traffic exchanged between endpoints and the manner of this communication are both important. Since the SLA is a point-to-multipoint QoS assurance, it is important to know how much traffic is exchanged between any given pair of customer endpoints, so that over-provisioning is avoided. Further, there are commonly observed VPN structures that can simplify traffic engineering. E.g., a hub/spoke VPN has a central node called the "hub" and rest of the endpoints communicate only with the hub. Recognizing such a structure can lead to bandwidth savings and help handle network events like failures in a more elegant manner. The trade-off in using such information is the additional complexity in the network.

3. COMPARATIVE ANALYSIS

We conducted NS-2 based simulations to help understand how to mix and match these design choices for optimal performance. We briefly report some of the important findings here: (a) Traffic matrix information has a significant effect on achievable resource utilization. Fig. 2 demonstrates that even a deterministic admission scheme can match a statistical scheme when enhanced with traffic matrix information; (b) The structure of VPNs controls the importance of traffic matrix information. As complexity increases, gains from traffic matrices increase dramatically (Fig. 3). This means that a provider can make network mechanisms simpler depending on the structure of VPNs serviced. If a large fraction of customers are of the hub/spoke nature, the provider might choose to de-emphasize such traffic matrix information; (c) The Dynamic algorithm is easy to deploy and is closer in performance to signaling compared to the static scheme (Fig. 4). Observe that an edge-based mechanism requires only the path capacity as its input. This means that as long as the network edges have the appropriate capacity information, admission control and SLAs are shielded from routing and topology changes.

4. REFERENCES

[1] N.G. Duffield, P. Goyal, A. Greenberg, P. Mishra, K.K. Ramakrishnan, J.E. van der Merive, "Resource management with hoses: point-to-cloud services for virtual private networks", *ACM SIGCOMM 99*

[2] S. Raghunath, K. Chandrayana, S. Kalyanaraman, "Edge-based QoS Provisioning for Point-to-Set Assured Services," *IEEE ICC 2002*, Vol. 2, pp. 1128-1134, 2002.

[3] S. Raghunath, S. Kalyanaraman, "Statistical Point-to-Set Edge-Based Quality of Service Provisioning," *QoFIS 2003* , Springer Verlag LNCS 2811, pp. 132-141, 2003.

[4] Y. Zhang, M. Roughan, C. Lund, D. Donoho, "An Information-theoretic Approach to Traffic Matrix Estimation,"*ACM SIGCOMM 2003*