

Resource Management for Virtual Private Networks

Satish Raghunath and K.K. Ramakrishnan¹

Abstract – Virtual Private Networks have rapidly emerged as a leading solution for multi-site enterprise communication needs. Provider-managed solutions modeled on RFC 2547 serve as a popular choice for layer-3 VPNs and the Hose Model has emerged as a common and simple service specification. It offers a “hose” of a certain contracted bandwidth to customers. With the growth in size and number of VPNs and the uncertainties in traffic patterns offered by customers, providers are faced with new challenges in efficiently provisioning and capacity planning for these networks while satisfying customer service level agreements (SLAs).

We suggest that a set of techniques can be used to help the provider build an adaptively provisioned network. These techniques involve continually processing measurement information, building inferences regarding VPN characteristics and leveraging them for adaptive resource provisioning. We have developed scalable techniques to infer VPN characteristics that are important for provisioning tasks and demonstrated the feasibility of such provisioning with existing SNMP-based measurement infrastructure from a large IP/VPN service provider. Our examination of measurement data yielded interesting new insights into VPN structure and properties. Building on our experience with analyzing VPN characteristics, in this paper we articulate an adaptive provisioning architecture that allows providers to effectively deal with the dynamic nature of customer traffic.

Index terms – IP VPNs, measurement, traffic matrix, provisioning.

I. INTRODUCTION

A. Overview of VPN technologies

A Virtual Private Network (VPN) securely connects multiple customer sites that wish to communicate amongst each other. Users demand secure, dependable, dynamic and “rich” connectivity between sets of end points from VPNs. Due to the progress in security and the overwhelming success of IP networking technologies, the number of VPNs a service provider has to support and the number of endpoints per VPN have been growing substantially. Communication patterns between endpoints are increasingly difficult to forecast. Users are often simply unable to predict loads between pairs of endpoints of the VPN in this context. It is typically difficult to specify QoS requirements on a point-to-point basis, which used to be the approach with deploying VPN services using traditional solutions such as Private Line or Frame Relay services. Private Lines isolate the performance seen by a VPN from other flows and provide guaranteed bandwidth, loss and delay characteristics. An IP VPN service that replaces the traditional point-to-point connectivity between sites that is offered by these legacy solutions must offer comparable performance and functionality. Functionality such as closed user groups, maintaining security for the data transferred over the VPN and providing network layer addressing flexibility are common. In addition, Private Lines provide guaranteed bandwidth, loss and delay characteristics. An IP-based VPN needs to offer comparable performance assurances.

B. Network Operation for IP VPNs

Service providers offering IP-based VPN services fulfill characteristics discussed above by leveraging the framework provided by RFC 2547 [1]. RFC 2547 allows a VPN to enjoy closed user groups and addressing flexibility while allowing the provider to re-use address space. While the framework outlines the mechanism for providing network-layer isolation and enables multiplexing customers on a common backbone network, the resource management components are left to the providers.

¹ Satish Raghunath is with Juniper Networks Inc, CA, rsatish@alum.rpi.edu
K.K. Ramakrishnan is with AT&T Labs – Research, Florham Park, NJ

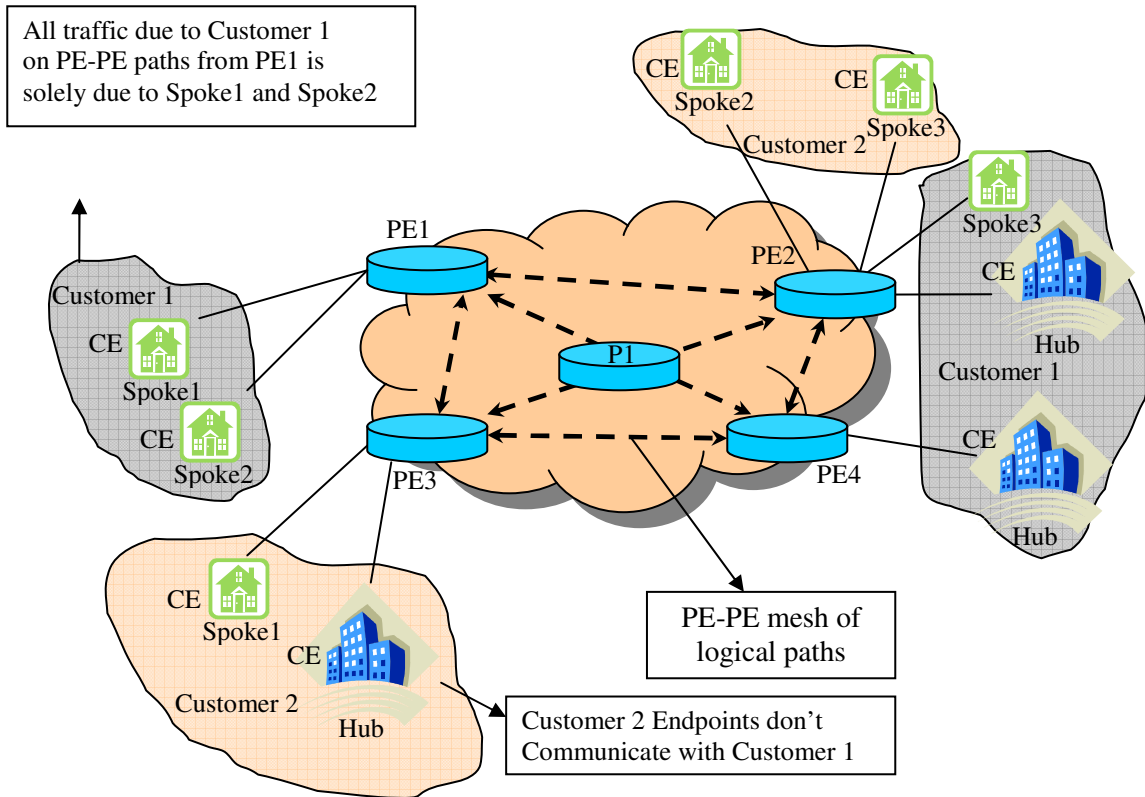


Figure 1: Schematic showing VPN properties

These resource management components can be built using per-VPN resource reservation algorithms [5][7]. In addition, there are network management tools provided by router/switch vendors that aid per-VPN provisioning. But, by and large, the task of service provisioning is left to the provider. Further, the task of managing a large number of VPNs on a common shared network while extracting statistical multiplexing gains across multiple customer VPNs still remains a challenge.

Our goal is to try and bridge this gap in resource management techniques for service providers to enable efficient multiplexing of VPN traffic while delivering QoS guarantees. We take recourse to measurement-based techniques to help complement the capabilities of the service provider network and we build on the foundation of a simple and widely accepted service specification (viz., the Hose Model). In doing so, we are able to provide an evolutionary path to an adaptive resource management architecture. Measurement-based approaches in general have been applied in the face of uncertain traffic characteristics and one good example is in the context of admission control[6].

C. Hose Model and resource management

The *hose model* [4], serves both as a VPN service interface (i.e., the way a customer thinks of a VPN) as well as a performance abstraction (i.e., the way a provider thinks of a VPN). A hose offers performance guarantees from a given endpoint to the set of all other endpoints in the VPN, and for the traffic to the given endpoint from the set of all other endpoints in the VPN. The hose is the customer's interface into the network, and is the equivalent of the customer having a "link" into the *network*. The hose model exploits the characteristics of a network and presents a much simpler service interface to the customer. It allows the customer to send traffic into the network without the need to pre-specify or predict point-to-point loads.

Though the hose model provides VPN customers simpler, more flexible SLAs, it appears to present the provider with a more challenging problem in resource management. Under the hose model, there is spatial uncertainty; i.e., uncertainty about traffic sinks. Provisioning such a VPN involves provisioning a set of point-to-multipoint connections with a traffic matrix that is unknown a-priori. The service specification

deliberately does not include details about the nature of traffic seen by the nodes in a customer’s network. However, these characteristics have an impact on the SLA offered to the customer.

The news is not all bad for the service provider. The provider can gain efficiency and scale because of the ability to multiplex VPNs on a common core network, as seen in Figure 1. In such a network, we envisage that the provider can obtain high resource utilization and simultaneously ensure SLAs using adaptive network mechanisms that exploit statistical properties of customer traffic. Such mechanisms entail understanding the temporal and spatial characteristics of VPNs to aid provisioning and capacity planning decisions. We note that Figure 1 refers to VPN endpoints that form hubs of communication as Hub nodes and their peers as Spoke nodes. Such a distinction is clear in most VPNs as we shall observe in a later section.

Provisioning a new VPN involves satisfying a set of *point-to-multipoint* requests. So the entire set of endpoints of the VPN has to be accommodated at once. In such a scenario, we have found that understanding the VPN’s traffic matrix is useful because, in addition to describing the traffic demands, a traffic matrix also depicts the structure of interactions in the VPN. A strategy that recognizes and takes advantage of the distinct features of VPNs outperforms one that does not take these into account [3].

In the following, we first focus on building components that help us understand VPN characteristics and then outline an adaptive provisioning framework that pieces them together. These components involve extensive use of measurement and prediction to dynamically size hose and network capacity. The key elements of our approach include: (a) periodic coarse-grained measurements – we rely on existing SNMP-based measurement infrastructure, (b) scalable techniques to infer VPN characteristics in the form of traffic matrices – we obtain a VPN’s spatial characteristics (e.g., if a VPN is Hub/Spoke) and temporal characteristics (e.g., time-of-day variations) from these traffic matrices, (c) characterization of VPN change in terms of growth in number of endpoints and contracted bandwidth over time, and (d) stitching together these pieces of information to build an intelligent provisioning mechanism. We envision this ongoing cycle of measurement, inference and provisioning as being an efficient way to solve the large scale (number of endpoints and capacity) VPN resource management problem. We demonstrate the feasibility of our techniques using measurements from a large IP/VPN service provider.

II. INFERRING SPATIAL AND TEMPORAL CHARACTERISTICS OF VPNS

A. Measuring traffic matrices

To understand the temporal and spatial structure of VPNs, it is necessary to characterize the traffic exchanged between customer edge routers (CEs) (the hubs and spokes at individual customer sites in Figure 1). The traffic matrix provides estimates of traffic exchanged between every pair of CEs that belong to the VPN. Recent advances in techniques for estimation of traffic matrices have yielded fast, scalable algorithms. These techniques have been developed for a variety of scenarios including traffic matrices for pairs of IP prefixes, border routers etc [8]. While they provide a useful starting point, there are important differences in the VPN case that prevented us from directly employing these pre-existing traffic matrix estimation techniques: (a) the scale of the network taken as a whole results in a computationally expensive and infeasible formulation (e.g., with the data we examined, we were dealing with 2.8×10^6 non-zero elements in a $(18 \times 10^3, 950 \times 10^3)$ sparse matrix); (b) per-VPN traffic information is not available for core network links, resulting in insufficient measurement information; (c) a shared core network infrastructure with only aggregate traffic counts for core network links introduced dependencies among the many VPNs that share those links.

The size of the VPN in terms of the number of CEs and the number of provider edge routers (PEs) that service the VPN greatly influence the scale of our problem. If there are N PEs that the CEs of a VPN communicate with, there can be $O(N^2)$ PE-PE paths that have to be factored in the estimation formulation. Figure 2 shows the distribution of number of endpoints per VPN. While there are a lot of small VPNs, there is a significant fraction with sizes in the tens to hundreds of endpoints. In the absence of per-VPN traffic information on a per-link basis (as is the case here - the traffic counts available for PE-PE logical links are based on aggregated traffic across VPNs), the estimation had to account for all pairs of CEs as potentially communicating peers.

Each of these issues is magnified further when we observe that there is continual growth in the number and size of VPN customers. Obtaining fine-grained reliable measurement information becomes much harder. We need a scalable technique to compute VPN traffic matrices – one that works well even with coarse-grained measurement information. We built such a technique and examined what characteristics of VPNs can be reliably estimated with existing information. In so doing, we were able to arrive at deployable techniques for improving the existing provisioning infrastructure.

B. Scalable VPN traffic matrix estimation

Our key observations are based on certain distinct properties of VPNs which allow us to break the network-wide problem down into many per-VPN traffic matrix formulations. The nature of VPNs is such that although they may share a common core network, no two endpoints belonging to different VPNs communicate with each other. This lends separability to our problem and suggests a strategy to reduce its size. Instead of solving the problem for all VPNs as part of a single network, we compute the traffic matrices for each VPN independently.

The path from a CE to another CE consists of two segments: a) an access segment (between the PE and the CE) where there is traffic from this VPN alone, and b) a core network segment (link between two PEs) which carries traffic multiplexed across multiple VPNs. In the network we examine, we assume that there exists a full mesh of paths (e.g., MPLS label switched paths (LSPs) between every pair of PEs). Typically, we have aggregate SNMP information for each of these segments. Thus we need to infer what fraction of the PE-PE aggregate traffic is attributable to the particular VPN for which we are estimating the traffic matrix. But there is not enough information to deduce this quantity. Instead, we introduce a bound on the contribution of a particular VPN to the measured PE-PE link traffic. For example in Table 1, the portion of traffic measured on the link PE1-PE2 due to Customer 1 must be less than the sum of traffic from the CEs, ‘Spoke1’ and ‘Spoke2’, for that customer. In order to improve this bound, we observe that the fraction of traffic on the PE-PE link attributed to a VPN cannot exceed the total observed traffic on the link. This constraint is also important when we observe that there are a lot of cases where certain PE-PE links are very lightly loaded. Consequently, it is likely that the sum of the traffic coming out of the CEs of a VPN attached to a PE might indeed be greater than the total traffic on some of the PE-PE links (in as many as 30% of the cases in many large VPNs [2]).

Using these observations we built a per-VPN estimation problem. The per-VPN formulation is efficient and scalable – it allows us to evaluate traffic matrices for hundreds of VPNs over many months of measurement data. In the following paragraphs we discuss some of the key insights we obtained by analyzing five months of SNMP measurements.




Figure 2: CDF of VPN Size

C. Evaluating the accuracy of traffic matrices

Evaluation of the reliability of the traffic matrix estimation technique assumes special significance in the case of VPNs, partly due to the scale of the problem as well as the approximations we used to break it down to manageable pieces.

We conducted two sets of evaluations – first, with synthetic inputs, and second, with real SNMP measurement inputs. In the case of synthetic inputs, we generate a traffic matrix and then deduce the link traffic counts that are fed to the estimation technique. It is straightforward in this case to measure the accuracy of the estimation technique. With SNMP data we do not know what the traffic matrix was. So we take recourse to properties of the estimation result that we know have to hold true. Specifically, for each PE-PE link we compute the estimated traffic for every VPN that uses it and then sum up these estimates. These estimates are then compared with the measured SNMP aggregate traffic for verification.

Using synthetic verification techniques [2], we found that our estimation techniques do well in identifying hubs of communication and the traffic matrices themselves are most accurate in the case of Hub/Spoke VPNs (which we describe below) featuring a single hub or multiple hubs. We examine the spatial structure of VPNs and find that this indeed forms a significant set, confirming general intuition. We refer the reader to [2] for a detailed discussion of the accuracy of our traffic matrix estimates.

D. Spatial Structure

VPN type	% VPNs where node with max peers had:	
	Max CIR	Max PIR
Pure Hub/Spoke	72.2	66.1
Dual Hub/Spoke	77.4	77.4
Hybrid	49.5	42.3

Table 1: Relation between CIR and number of peers

Traffic matrix measurements provide important insights in to the way CEs communicate with each other. We are able to classify CEs based on how they interact with other CEs in the VPN. For example, in a Hub/Spoke VPN, a hub node communicates with most or all other CEs (the spokes) of that VPN. It is

necessary to estimate traffic matrices in order to arrive at such properties since the Committed Information Rate (CIR) and Peak Information Rate (PIR) (which are typical static parameters specified for a VPN in an agreement with the service provider) or access link traffic measurements are not always reliable indicators. Moreover, the routing tables are not setup to constrain traffic between hub and spoke nodes. Any node in the VPN can communicate with any other node within the same VPN. Thus, routing tables do not reflect any consistent patterns that could help us infer the structure of the VPN.

In fact as Table 1 indicates, in many instances, the node with the highest CIR need not be the hub node (e.g., 28% of the nodes in a pure hub/spoke environment had the max. number of peers but did not have the max. CIR). On the other hand if we rely on the traffic measured on the access link between the CE and the PE, it is not clear whether we should use the traffic from CE to PE or in the opposite direction to judge if the node is a hub. If we observe traffic entering and leaving a hub, they need not be symmetric (Figure 3) – this means we cannot use one or the other as a definite indicator for a CE to be a hub node.

Comparison of traffic entering and leaving hubs

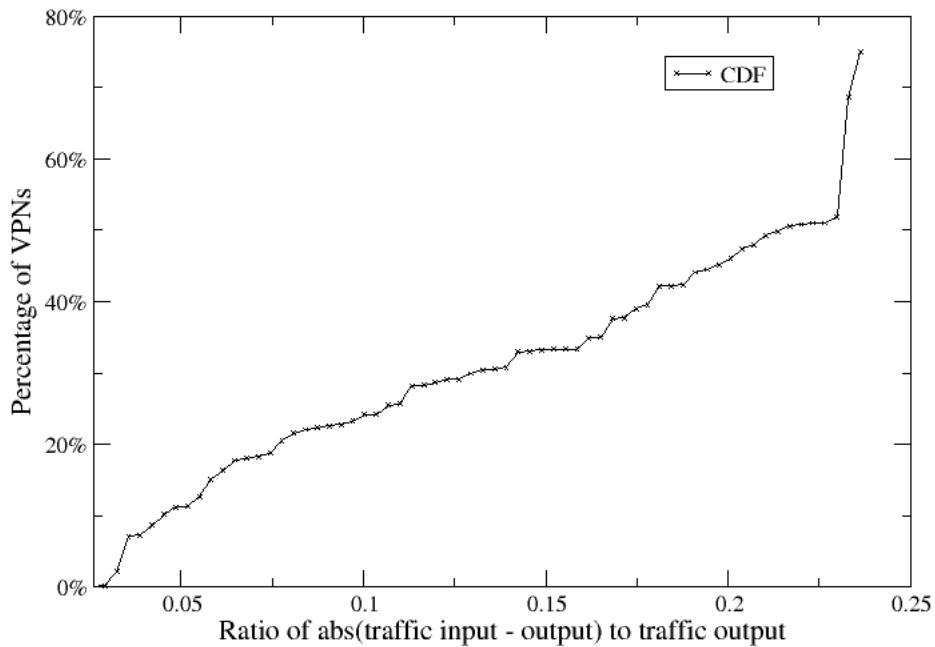


Figure 3: Traffic at a hub is often asymmetric – roughly 60% the VPNs that were analyzed featured a difference of more than 15% (0.15 on X-axis) between traffic entering and leaving the hub.

An important result that emerges from a structural analysis of VPNs is the significance of the Hub/Spoke structure. We find that irrespective of the size of the VPNs being analyzed, there is a sizable percentage of VPNs that conforms to this structure (Figure 4). We categorize VPNs into Hub/Spoke, Multi-hub, Meshed and Hybrid VPNs.

To characterize a VPN’s spatial structure, we use a two-step procedure that involves traffic matrix measurements for the VPN. First, we examine each CE in the VPN and it’s communicating peers. After appropriate pruning to mitigate measurement errors, we categorize a CE as a hub if it communicates with more than 50% of the CEs in the VPN. The CE is a spoke node, if it communicates with just one other node. If we are unable to put a CE in either of these classes, we put the CE node in the ‘hybrid’ class.

The second step looks at the VPN as a whole in light of the estimated properties for the CEs. If there is one hub node and the rest of the nodes are spokes, we have a pure Hub/Spoke VPN. When there is more than one hub, the VPN is classified as a multi-hub VPN. VPNs that feature one or more hybrid CEs are

classified into hybrid VPNs. Finally, if every node has every other node in the VPN as its peer, it is judged to be a meshed VPN.

Many of the multi-hub VPNs feature just two hubs that communicate with most of the other CEs. Very often, these two hubs are co-located and serve the purposes of redundancy and load-balancing. This means that taken together with pure Hub/Spoke VPNs, these structures cover well over half of all VPNs (Figure 4), irrespective of the size of the VPN.




Figure 4: Structural classification of VPNs

E. Temporal Properties

VPN traffic matrices analyzed over a period of many months confirm stability in temporal characteristics [2]. We find that there are patterns not only in the aggregate SNMP measurements, but also with CE-CE traffic. This has important implications for adaptive provisioning strategies. It means that service providers can deploy measurement-based provisioning strategies that rely on traffic patterns that are learnt on timescales of many weeks or months.

III. “CHANGE CHARACTERISTICS” OF IP VPNS

A. Changing nature of IP VPNS

As part of our vision for resource management, we anticipate that a service provider would adapt its resource allocation to VPN customer demands and traffic characteristics by provisioning, learning through measurement, and subsequently re-provisioning as needed. Thus, it is useful to examine the growth trends for VPNs, both in terms of the number of new VPN customers added and the changes in individual VPNs. We examine changes in the number of customer endpoints (size of VPN), the CIR and the PIR. By understanding these factors in addition to spatial and temporal trends in customer traffic, we can then envision our adaptive provisioning architecture.

We observed the daily logs and the SNMP data for the IP/VPN service which had several thousand VPN customers, tens of thousands of distinct CEs, connected to a few hundred PE routers over a 30 month period. We found significant changes in VPN size over the observed period. While there are changes with CIR and PIR as well, we found those changes to be less frequent.

Examining the growth in VPN size (the number of CEs in the VPN) over a period of 1 month (a typical manual provisioning time-scale), Figure 5 shows that there is quite a significant percentage (almost 60%) of VPNs whose size remains nearly constant – close to the 0% change on the X-axis. However, we do see that the change is more frequent, going from one year to the next. Furthermore, there are a small percentage of VPNs that see a large amount of change in the number of endpoints over the same interval of one month – these clearly pose challenges to the service provider in terms of provisioning capacity. Our measured traffic matrices also show an increase in traffic toward hubs over many weeks in the case of VPNs that add new CEs.

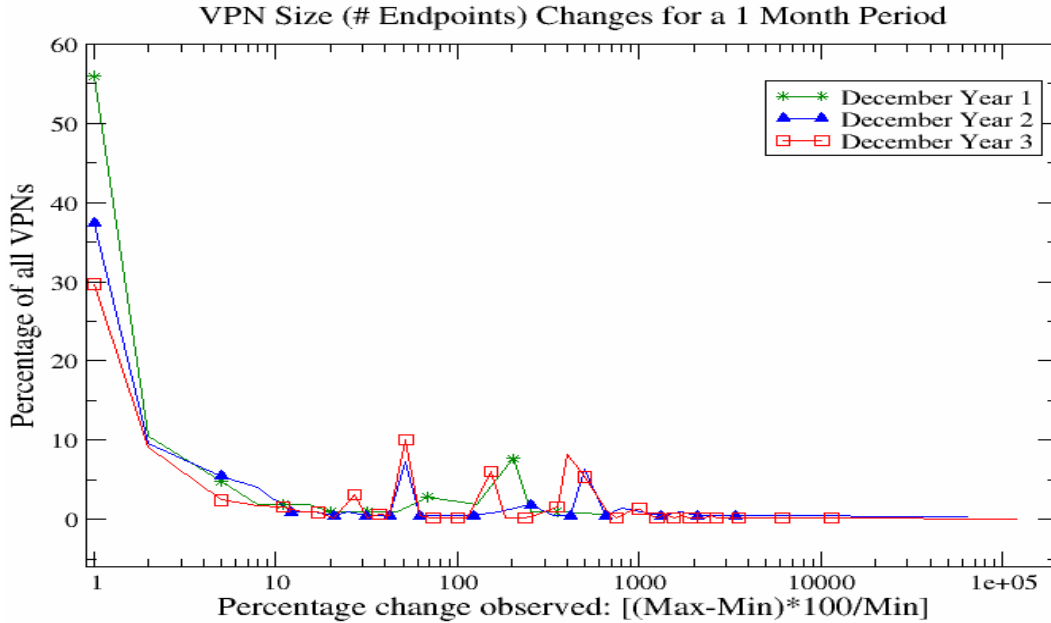


Figure 5: Growth in VPN size

Over a longer period (e.g., 6 months, which is typical of the timescales for upgrading the service provider’s facilities), almost a majority of VPNs (over 50%) see change in size, with some seeing very significant change. Thus, it is important for the service provider to continually monitor and adapt to changes in the customer characteristics – it may not be sufficient to admit an individual VPN customer and provide the resources at the beginning with no further adaptation in provisioned resources. After all, the value of the VPN service for the customer is when good service is provided to all the endpoints in the VPN, including the new endpoints that are added as the customer’s business evolves. Thus, we believe that it is important for the service provider to adopt a strategy of “provision, learn, re-provision” as an integral part of their VPN resource management approach.

B. An adaptive provisioning architecture

We have made three important observations with supporting measurement evidence:

1. We have built a scalable estimation technique for VPN traffic matrices. This means we can compute periodic traffic matrices even for a large service with several hundreds or thousands of VPNs.
2. Traffic matrices deduced by our estimation technique are reliable for a significant percentage of VPNs and hence can be a basis for provisioning decisions.
3. VPN characteristics are continually subject to change, but the time-scale and pace of change in temporal trends are slow enough to imply that providers can employ tools to learn traffic trends.

The important observation here is that in the time-scales important to providers in terms of provisioning, we can deploy traffic matrix estimation techniques to reliably learn VPN characteristics. This leads us to an architecture that adapts to change.

With adaptive resource management, a service provider starts with a conservative allocation, based on the static information a customer provides (e.g., CIR) and coarse-grained available capacity information. Throughout the life of the VPN, the provider measures and learns the traffic characteristics of the customer. As more measurement samples become available, the provider’s view of customer VPN’s characteristics as well as the available capacity on the network gets more refined. The result is that the provider is now in a position to anticipate changes in customer bandwidth needs, enabling timely remedial actions. For example,

the provider could augment the capacity of a PE-PE path if the customer aggregates sharing the path are expected to grow their demands.

To quantify the benefits of such an adaptive architecture, we built a simulation framework that modeled VPN characteristics as seen in the measurement data-set analyzed in earlier sections. Figure 6 captures the benefits of exploiting traffic matrix information for provisioning, a key component of the “provision-learn-re-provision” strategy.

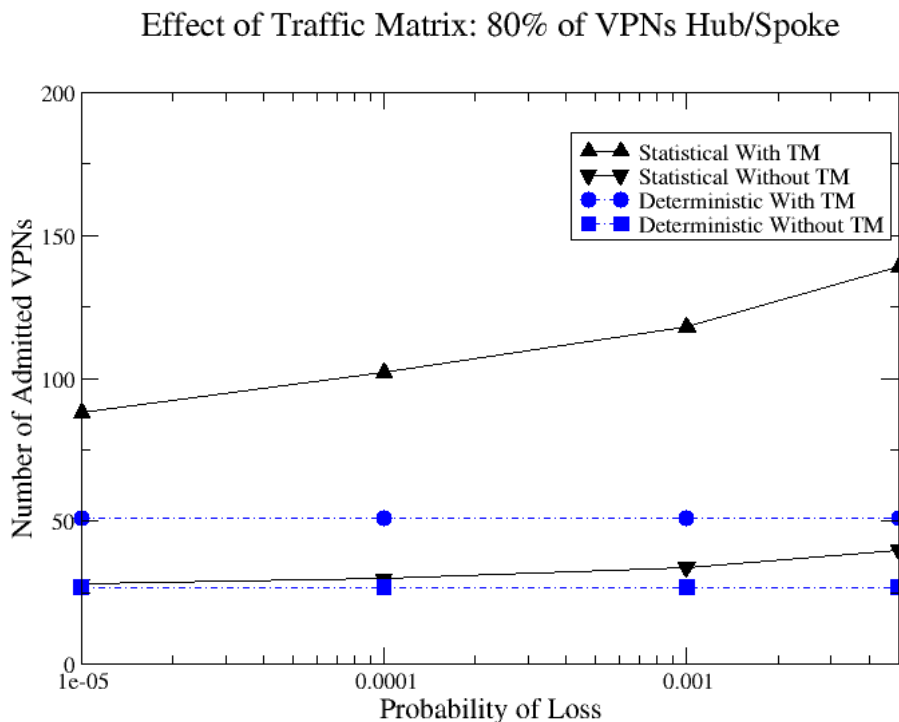


Figure 6: Impact of using traffic matrix information

Figure 6 provides a measure of the multiplexing gains that a provider can attain with both deterministic and statistical bandwidth allocation. With statistical allocation, the provider seeks to achieve a QoS objective with a pre-set probability. In Figure 6, we use the probability of packet loss as the QoS objective – higher the tolerance for loss, higher the multiplexing gain with a statistical allocation technique. In contrast, a deterministic scheme statically partitions bandwidth, and is a constant (horizontal line in Figure).

The key observation here is the dramatic improvement delivered by incorporating VPN traffic matrices. In the presence of traffic matrix information, even if the provider chooses a simple deterministic allocation scheme (that does not oversubscribe capacity), it might do better than a statistical allocation mechanism (which would oversubscribe the capacity with some probability) without traffic matrices. The gains obtained by employing traffic matrices suggest that incorporating them in a continual learning process to fine-tune provisioning can be beneficial. We anticipate that the initial provisioning in our adaptive strategy can be somewhat coarse, ensuring that there is sufficient capacity just based on the customer contracted CIR. Subsequent measurement and learning may then improve provisioning to take advantage of statistical multiplexing gains.

C. Future directions

Analyzing measurement data has helped us understand important VPN characteristics and yielded practical strategies for fine-tuning provisioning. These insights open up a range of research directions:

- **Traffic Engineering:** Traffic matrices provide us an estimate of the size of the customer aggregate in the core network. This allows us to perform traffic engineering on a per-customer aggregate basis. Without this information, traffic engineering would have to handle PE-PE aggregates as a whole.
- **Bandwidth Allocation:** Exploiting spatial characteristics can lead to simplified provisioning and efficient resource allocation, especially in the case of endpoints which communicate with just one or two other peers.
- **Customer Differentiation:** Since traffic matrices provide an estimate of the size of the customer aggregate in the core network, the provider can choose to give preferential treatment to a selected set of customers more efficiently. The temporal characteristics of traffic matrices indicate that the aggregate characteristics vary slowly and can be learnt through measurement.
- **Managing network failures:** The additional knowledge of customer traffic can lead to elegant management of network failure and maintenance events. E.g., the aggregates leading to a hub node can be mapped on to a new path which has more available capacity.
- **Anticipating VPN changes:** By understanding typical trends in VPN growth, a provider can better prepare for future provisioning needs.

IV. CONCLUSIONS

Resource management is an important issue faced by service providers offering IP VPN services. The hose model provides a simple service specification for the customer and allows the provider to multiplex multiple VPNs on a shared network. However efficiently provisioning such a network, while satisfying SLAs, is not simple. We envisage an environment where providers use measurement information to build an adaptive resource management framework that can endure changes in traffic characteristics.

We demonstrated the utility of VPN traffic matrices in provisioning and outlined scalable techniques to compute such matrices. We are able to arrive at these traffic matrices with coarse-grained SNMP measurement of traffic aggregated across VPNs in the core of the network. This is very useful in the light of difficulties in obtaining per-VPN measurements with large-scale service deployments. Analyzing traffic matrices allowed us to draw important conclusions about spatial and temporal characteristics of VPNs. We found that Hub/Spoke VPNs formed a significant percentage of the prevalent structures. The temporal trends in customer traffic show stability over periods of a few weeks to a month, suggesting the feasibility of measurement-based learning of traffic characteristics.

A provider that can accommodate VPN growth and still provide good service (SLA) obviously stands to gain. Our fundamental observation is that an approximate “provision – learn – reprovision” approach can be very beneficial to the provider. We have presented important tools to enable this approach and have laid out an evolutionary path toward such a system.

V. ACKNOWLEDGEMENTS

The authors gratefully acknowledge guidance and support from Prof. Shivkumar Kalyanaraman (ECSE, RPI) and Dr. Chris Chase (AT&T Labs).

REFERENCES

- [1] E. Rosen and Y. Rekhter, “BGP/MPLS VPNs,” RFC 2547, Mar. 1999.
- [2] S. Raghunath, K.K. Ramakrishnan, S. Kalyanaraman, and C. Chase, “Measurement based characterization and provisioning of IP VPNs,” in *Proc. of IMC 2004*, pp. 342–355.
- [3] S. Raghunath, S. Kalyanaraman, and K. K. Ramakrishnan, “Trade-offs in resource management for virtual private networks,” in *Proc. IEEE INFOCOM 2005*.
- [4] N. Duffield, P. Goyal, A. Greenberg, P. Mishra, K. K. Ramakrishnan, and J. van der Merive, “Resource management with hoses: point-to-cloud services for virtual private networks,” *IEEE/ACM Trans. Networking*, vol. 10, no. 5, pp. 679–692, Oct. 2002.
- [5] A. Kumar, R. Rastogi, A. Silberschatz, and B. Yener, “Algorithms for provisioning virtual private networks in the hose model,” in *Proc. of ACM SIGCOMM 2001*, pp. 135–146.
- [6] S. Jamin, P.B. Danzig, S. J. Shenker, and L. Zhang, “A Measurement-based Admission Control Algorithm for Integrated Services Packet Networks,” *IEEE/ACM Trans. Networking*, vol. 5, no. 1, pp. 56-70, Feb. 1997.

- [7] A. Gupta, J. M. Kleinberg, A. Kumar, R. Rastogi, and Bulent Yener, "Provisioning a Virtual Private Network: A network design problem for multicommodity flow", *ACM Symposium on Theory of Computing*, pp. 389-398, 2001
- [8] Y. Zhang, M. Roughan, C. Lund, D. Donoho, "An information-theoretic approach to traffic matrix estimation," *ACM Sigcomm 2003*, pp.301-312