A Flexible Model for Resource Management in Virtual Private Networks

N.G. Duffield, Pawan Goyal¹, Albert Greenberg

Partho Mishra, K.K. Ramakrishnan, and Jacobus E. van der Merwe
AT&T Labs-Research, 180 Park Avenue, Florham Park, NJ 07932, USA

Abstract

As IP technologies providing both tremendous capacity and the ability to establish dynamic secure associations between endpoints emerge, Virtual Private Networks (VPNs) are going through dramatic growth. The number of endpoints per VPN is growing and the communication pattern between endpoints is becoming increasingly hard to forecast. Consequently, users are demanding dependable, dynamic connectivity between endpoints, with the network expected to accommodate any traffic matrix, as long as the traffic to the endpoints does not overwhelm the rates of the respective ingress and egress links. We propose a new service interface, termed a *hose*, to provide the appropriate performance abstraction. A hose is characterized by the aggregate traffic to and from one endpoint in the VPN to the set of other endpoints in the VPN, and by an associated performance guarantee.

Hoses provide important advantages to a VPN customer: (i) flexibility to send traffic to a set of endpoints without having to specify the detailed traffic matrix, and (ii) reduction in the size of access links through multiplexing gains obtained from the natural aggregation of the flows between endpoints. As compared with the conventional point to point (or customer-pipe) model for managing QoS, hoses provide reduction in the state information a customer must maintain. On the other hand, hoses would appear to increase the complexity of the already difficult problem of resource management to support QoS. To manage network resources in the face of this increased uncertainty, we consider both conventional statistical multiplexing techniques, and a new *resizing* technique based on online measurements.

To study these performance issues, we run trace driven simulations, using traffic derived from AT&T's voice network, and from a large corporate data network. From the customer's perspective, we find that aggregation of traffic at the hose level provides significant multiplexing gains. From the provider's perspective, we find that the statistical multiplexing and resizing techniques deal effectively with uncertainties about the traffic, providing significant gains over the conventional alternative of a mesh of statically sized customer-pipes between endpoints.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
SIGCOMM '99 8/99 Cambridge, MA, USA
© 1999 ACM 1-58113-135-6/99/0008...\$5.00

1 Introduction

Virtual Private Network services have been offered in various forms over an extended period of time and have recently received considerable attention within the IP, frame-relay, MPLS, and ATM networking communities [1, 2, 3, 4, 5]. VPNs are likely to be used by customers as a replacement for networks constructed using private lines and should therefore, at the very least, provide a comparable service. Substantial progress in the technologies for IP security [6] enable us to improve on the security and privacy provided in existing VPN service offerings based on private lines or frame-relay. Other work on IP-based VPNs has mainly dealt with group membership, routing protocols and tunneling [2]. Much less attention has been paid to resource management issues related to VPNs. However, supporting a variety of mission-critical functions requires a VPN service to provide performance assurances, backed by Service Level Agreements (SLAs). Private lines isolate the performance seen by a VPN from other flows and provide guaranteed bandwidth, loss and delay characteristics. A VPN service must offer comparable performance assurances. Our focus in this paper is on the performance issues related to VPNs.

Due to the progress in security and the overwhelming success of IP networking technologies, the number of endpoints per VPN is growing, and communication patterns between endpoints are becoming increasingly difficult to forecast. We expect 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, the conventional approach. Our solution, which we call the *hose model* serves as both 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 service interface allows the customer to send traffic into the network without the need to predict point to point loads.

Though the hose model provides customers simpler, more flexible SLAs, the model appears to present the provider with a more challenging problem in resource management. Under the conventional point to point model for specifying

¹Currently with Ensim Corporation

ing QoS, there is uncertainty about temporal variation in the traffic between the two points. Under the hose model, there is also spatial uncertainty; i.e., uncertainty about traffic sinks. To cope with these uncertainties, we develop mechanisms that allow providers to use the hose model to achieve significant multiplexing gains in the network, by the use of signaling to dynamically size hose and network capacity. The flexibility offered by this approach will be crucial in building VPNs that will scale well in both the number of endpoints and capacity.

We evaluate the proposed hose VPN service model by performing a number of trace driven experiments. In particular we show that significant multiplexing gains may be achieved for both the customer and the provider when the network is capable of exploiting the hose model. Two sets of traces were used for these experiments. The first was voice traffic traces from the AT&T backbone network. The second was data traffic traces from a large corporate backbone network.

The rest of the paper is organized as follows. In the next section, the hose model for VPNs is presented. Section 3 describes implementation scenarios, and the traffic predictors we used in our experiments to estimate required capacity. After an outline of the simulation framework in Section 4, we briefly consider the variability of the traffic matrix in Section 5 based on an analysis of the data traffic traces. The section then continues by examining the benefit of the hose from the perspective of a customer. In Section 6 we look at the multiplexing benefits within the provider's network and examine the performance of alternative means of implementing a hose in the network. Of interest is the reduction in capacity as we dynamically resize the amount of resources used to adapt to changing traffic needs. We address issues of arriving at an effective bandwidth for admission control in Section 7, and conclude in Section 8.

2 The Hose Service Model

A simple service model for an IP VPN is to emulate the private line or frame relay service. This would require a customer to buy a set of **customer-pipes**, i.e., allocations of specific bandwidth on paths between source-destination pairs of endpoints of the VPN (much like virtual circuits). Figure 1 illustrates an example of the use of this kind of interface. 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 the traffic matrix between all the VPN sites. Resources made available to a customer-pipe cannot be allocated to other traffic. It is important to note the network provider may not be able to take advantage of statistical multiplexing gains across the customer-pipes.

In this paper, we propose a richer and more flexible VPN service model that we refer to as a **hose**. In the hose model, a VPN customer specifies a set of endpoints to be connected with common endpoint-to-endpoint performance guarantees. The connectivity of each endpoint to the network is specified by a hose, comprising:

- the capacity required for aggregate outgoing traffic from the endpoint into the network (to the other endpoints of the VPN)
- the capacity required for aggregate incoming traffic out

of the network to the endpoint (from the other endpoints of the VPN);

- the performance guarantee for the hose, conditioned only on the *aggregate* traffic seen at the hose interface.

Figure 2 illustrates an example of the use of hoses. Say that there are 4 VPN sites: *A*, *B*, *C*, and *D*. The customer buys 4 hoses at each of these sites and specifies the aggregate outgoing and incoming traffic for each of these hoses. The hose specification may be arrived at in a variety of ways. For example, if it is known that each of the sites *B*, *C* and *D* sends and receives at no more than 3 Mb/s to site *A*, and that each of these sites sends and receives no more than 2 Mb/s *in aggregate* to each other, then the hose capacity would be chosen as: $A_{in} = A_{out} = 9$ Mb/s and $B_{in} = B_{out} = C_{in} = C_{out} = D_{in} = D_{out} = 5$ Mb/s. Figure 2 depicts one possible realization of the connectivity from hose *A* by means of a set of provider-pipes. Different implementation options are considered in Section 3.

There are several advantages of the hose model from a customer's perspective:

Ease of Specification: Only one inward and outward rate per hose endpoint needs to be specified, as compared with that for each customer-pipe between pairs of endpoints.

Flexibility: Data to and from a given hose endpoint can be distributed arbitrarily over other endpoints provided the aggregate conforms to the hose size.

Multiplexing Gain: Due to statistical multiplexing gain, hose rates can be less than the aggregate rate required for a set of customer-pipes.

Characterization: Hose requirements are easier to characterize because the statistical variability in the individual source-destination traffic is smoothed by aggregation into hoses.

The nature of the service level agreement between a customer and a service provider is driven by the traffic characteristics and QoS requirements of the (customer) applications that make use of the VPN. For example, an IP voice VPN service might require tight bounds on the per-packet loss rates, delay and possibly jitter. On the other hand, a data-only VPN service might have relatively less stringent delay requirements. To ensure that the appropriate requirements can be met, it is essential for the customer to provide a description of the traffic characteristics.

The complexity of the traffic specification provided by a customer is expected to vary depending on the needs and sophistication of the customer. A reasonable service offering would indeed be for customers to start off with a fairly simple specification and to then refine this specification based on operational experience and service provider feedback.

Coming up with an initial estimate of the VPN specification might be a service offered by the provider to customers, as part of a VPN characterization phase. The customer (or the provider) can then monitor the performance of the VPN in terms of loss and delay to determine whether it satisfies the needs of the customer's applications and can then negotiate a different capacity and the corresponding SLA.

During the characterization phase, the SLA for the VPN might be undefined or might be defined as some best-effort service. Alternatively, provider may act conservatively and over-provision in terms of the resources it allocates for the VPN during this phase, and provide QoS assurances.

Figure 1: A VPN BASED ON THE CUSTOMER-PIPE MODEL. A mesh of customer-pipes is needed, each extending from one customer endpoint to another. A customer endpoint must maintain a logical interface for each of its customer-pipes.

Service level agreements following the characterization phase might be based on the current traffic load with provisions made for expected gradual growth as well as expected drastic traffic changes that the customer might foresee (or protect against).

Both the customer and the provider may play a role in testing whether the SLAs are met. The provider may police (and possibly shape) the incoming traffic to a hose from the customer's access link to ensure that it stays within the specified profile. Similarly, traffic leaving the VPN at a hose egress (i.e., traffic potentially generated from multiple sources that has traversed the network) may have to be monitored and measured at the hose egress point, to ensure that such traffic stays within the specified profile and that the provider has met the SLA. The customer might also be required to specify a policy for actions to be taken should egress traffic be more than the specified egress hose capacity.

2.1 Capacity Management

From a provider's perspective, it is potentially more challenging to support the hose model, due to the need to meet the SLAs with a very weak specification of the traffic matrix. To manage resources so as to deal with this increased uncertainty, we consider two basic mechanisms:

Statistical Multiplexing: As a single QoS assurance applies to a hose, the provider can consider multiplexing all the traffic of a given hose together. Similarly, the set of hoses making up the VPN have a common QoS assurance, and the provider can consider multiplexing all the traffic of a given VPN together. These techniques can be applied on both access links and network internal links.

Resizing: In order to provide tight QoS assurances, the provider may use (aggregate) network resource reservation mechanisms that allocate capacity on a set of links for a given hose or VPN. A provider can take the approach of allocating this capacity statically, taking into account worst case demands. Alternatively, a provider can make an initial allocation, and then resize that allocation based on online measurements. Again, such techniques can be applied on both access and network internal links. Resizing is allowed only within the envelope defined by the SLA. Resizing would occur at a substantially finer time scale than the time scale over which SLA's might be renegotiated.

Figure 2: A VPN BASED ON THE HOSE MODEL. A customer endpoint maintains just one logical interface, a hose, to the provider access router. In the Figure, we show the implementation of one hose (based at A) using provider-pipes.

These two resource management mechanisms can be used separately or in combination.

Some more remarks are in order on resizing. Provisioning decisions normally have an impact over fairly long timescales. Within the context of a VPN framework, measurements of actual usage can be used on much shorter timescales to enable efficient capacity management. Underlying this is an assumption that within the network boundaries will exist between resources that might be used by different classes of traffic to ensure that performance guarantees are met. For example, traffic from different VPNs might be isolated from each other, and from other classes of traffic. In the context of this paper, resources available for VPN traffic cannot be used by other traffic requiring performance guarantees. We assume that this perspective holds whether the boundaries reflect reservation of resources, such as in the case of Intserv, or whether it represents some allocation in a bandwidth broker in a Diffserv environment.

If we can use the measurements of actual usage to resize the boundary for a given VPN's traffic, more bandwidth will be made available to other traffic and we can make better use of available capacity. In reality, measurements of current usage would be used to make a prediction about near term future usage, and this prediction will be used to resize the share of resources allocated.

In the hose model, this approach can be realized by allowing customers to resize the respective hose capacities of a VPN. Presumably there will be some cost incentive for customers to resize their hose capacities. While we envisage this mechanism to be mainly used to track actual usage, by exposing this interface to the customer, it would also enable the customer to resize its hose capacities based on local policy decisions.

How frequently hoses may be resized will depend on implementation and overheads for resizing and measurement. More important, however, is whether frequent resizing is beneficial and whether it is possible to make predictions with sufficient accuracy. Finally, short timescale resizing is not a replacement for provisioning and admission control and the appropriate relationship between these resource management approaches is important.

Figure 3: An example network to illustrate various implementation possibilities. 1-3 represent customer routers, and A-G represent provider routers.

3 Implementation Scenarios

The flexibility offered by the hose interface presents a challenge for its realization. In this section we examine the various alternatives that a provider has for implementing a hose.

Provisioned VPNs Consider the example network shown in Figure 3. Let 1, 2, and 3 be the customer routers that are part of a VPN, each originating a hose of size 1 Mb/s. To simplify the discussion, let us assume that capacity is reserved for hoses based on a worst-case traffic split, i.e., the traffic from each hose can be directed entirely to just one other endpoint. To begin with, a provider may not want to employ any hose specific state in the network, relying simply on default routing and not making use of any resource sharing (in this case the hose model is only a service interface for the customer.) Then, the hose originating from, say, router 1, can be implemented by reserving two “provider-pipes” from ingress router A to egress routers E and G¹. Since a worst-case traffic split is being assumed, the capacity of both the provider-pipes are 1 Mb/s. To determine the total resources reserved in the network, let us assume that the default shortest paths from A to E and G are A-C-B-E and A-C-D-G, respectively. Thus, the total capacity required (i.e. summed across all links) when a hose is implemented using provider-pipes is 6 Mb/s.

Observe that on link A-C, 2 Mb/s is reserved even though from the hose specification we know that at most 1 Mb/s worth of traffic may enter the network from source 1 and thus be present on that link at any time. To reduce the reservation to 1 Mb/s on link A-C, the provider can utilize the *source tree* rooted at A and constructed using the default shortest paths to destinations E and G. On each edge that belongs to this tree, 1 Mb/s is reserved. Since 5 links belong to the source tree rooted at A, the total capacity required is reduced to 5 Mb/s. However, to achieve this reduction, the provider is required to know about and make use of hose specific state in the network in the form of a source tree. Since default shortest path routing is being assumed, this hose state is limited to control of *resource sharing*.

A provider can achieve further reduction in the capacity by employing hose specific state not only in the sharing of resources but also by making use of non-default or *explicit routing*. For example, let the forwarding tables be changed

¹We use the term provider-pipe to clearly indicate that these pipes are not exposed in the customer interface.

such that traffic from A to E and G is forwarded along paths A-C-F-E and A-C-F-G, respectively. Though these are also shortest paths, these were not considered the default shortest paths in our example. In this case, by employing resource sharing as before, as well as explicit routing, the total reserved capacity reduces to 4 Mb/s.

The above implementation alternatives only exploit hose specific state. A provider can achieve further reduction by additionally exploiting *VPN specific state*. To observe this, assume that the routing occurs as in the previous scenario for hose 1 and shortest path routing is employed for hoses 2 and 3, i.e., for hose 1 the tree is A-C-F-E and A-C-F-G, for hose 2 it is E-F-C-A and E-F-G and for hose 3 it is G-F-E and G-F-C-A. If we limit ourselves to hose specific state only then the total reserved capacity in the network for the VPN is 12 Mb/s (4 Mb/s for each hose). Now, consider the link FG. On link FG, capacity of 2 Mb/s (1 Mb/s each for hoses 1 and 2) is being reserved for traffic that is destined for node 3. However, from the specification of hose 3, we know that at most 1 Mb/s may be received by 3. Hence, by recognizing that hoses 1 and 2 belong to the same VPN, a provider can reduce the capacity reserved on link FG to 1 Mb/s. A similar reduction of 1 Mb/s can be achieved on links FE, FC, and CA. Thus, by exploiting VPN specific state, the aggregate reserved capacity can be reduced from 12 Mb/s to 8 Mb/s.

The above implementation alternatives are applicable even when a network does not accommodate the capability to signal the traffic requirements, e.g., in the case where VPNs are provisioned. This case is of practical importance as initial VPN offerings might be realized in this fashion. We investigated this in our simulation experiments. In this type of realization, however, a provider has no alternative but to allocate for the peak rate and to assume the worst case traffic split from the hose, because the hose model explicitly allows changes in the traffic matrix. Without the capability to signal the traffic requirements, this can lead to significant underutilization of resources in a network. To address this inefficiency, we propose an extension to the above set of implementations, which could lead to a better utilization of network resources.

Dynamically Resized VPNs The essential idea is to use on-line measurements to determine the capacity requirements of hoses and to then dynamically adapt the amount of resources reserved based on such measurements. We assume the existence of appropriate mechanisms to achieve these two functions. The implementation alternatives are as follows.

Resized Provider Pipes: In this case, a hose is implemented by a mesh of provider-pipes between the ingress and egress routers of a VPN. The resizing of these pipes is done from the ingress edge router: it measures the traffic for each provider-pipe, and based on these measurements predicts the size of the pipes required. It then signals the reservation for each of the provider-pipes. Since this implementation does not use any hose specific state in routing or resource reservation (i.e., no sharing), it can be realized in any IP network that can support dynamically resized pipes, e.g., IntServ, DiffServ or MPLS networks.

Resized Trees: A hose is realized by a source based tree. The aggregate hose traffic is measured at each link and resources are reserved for the aggregate. Although, as we traverse down the tree, we measure and reserve resources for progressively smaller aggregates.

Instead of employing a source based tree, we can also employ a sink based tree. For example, in Figure 3, we can employ a sink tree rooted at A with E and G being the leaves. In this case, the reservation is for traffic flowing out on the hose into router 1. Reservation of resources occurs in a manner analogous to the source based tree.

Since this implementation does not have any hose specific state in routing, it can be realized in an IP network that has only default routing. However, compared to the first case, the signaling protocols that enable reservation of resources have to make use of hose specific information. We will use the terms hose and tree interchangeably in this implementation alternative.

Resized Trees with Explicit Routing: In this case, a hose may be realized using a Steiner tree. Measurement and signaling for resources occur as in the previous case. But, the main advantage of this approach over the previous one is that it increases the number of links on which resources are shared. This implementation, however, can only be realized in networks that support explicit routing (for example, MPLS) and also requires the signaling protocol to be hose aware.

Resource Aggregation across a VPN: In this case, we first route each of the hoses using the default shortest path routing or Steiner trees (i.e., either of the previous two approaches). Then on the links of the graph resulting from the union of the hose trees, we measure the aggregate traffic for the VPN and reserve resources appropriately. This requires the signaling and measurement mechanisms to associate different hoses to the VPN they belong to.

The realizations described above all require: (1) a method for measuring the traffic and based on such measurements, predicting the required capacity, and (2) signaling protocols for dynamically reserving resources based on predicted capacity requirements. In the next section, we present techniques for predicting the required capacity, while the design of appropriate signaling mechanisms is the subject of future work.

3.1 Prediction of Traffic Rates

In this section we describe schemes for predicting future traffic rates of a traffic flow from measurements. As explained in the previous section, such predictions are used to dynamically change or renegotiate the resources associated with a VPN. Specifically we are interested in flows comprising traffic aggregated at either the pipe, the hose, or the VPN level. We assume that the measurements comprise samples gathered at (regularly spaced) instants during a window of duration T_{meas} . The samples are themselves some function of the traffic rates in the interval between sampling instants; in this paper we use average rate over the inter-sample interval. The measurements are used to predict an effective bandwidth for the traffic flow over some window of duration T_{ren} following the measurement window. Such predictions have the locality property that they depend only on measurements over the window of duration T_{meas} into the past. Specific examples that we employ are:

Local Maximum Predictor: The renegotiated rate is the maximum of the rate sampled during the measurement window.

Local Gaussian Predictor: The renegotiated rate is equal to $m + \alpha\sqrt{v}$, where m and v are respectively the mean and

variance of the rates sampled during the measurement window, and α is a multiplier that controls the extent to which the negotiated rate accommodates variability in the samples. The interpretation of α is that in a Gaussian approximation to the rate distribution, we expect the bandwidth $m + \alpha\sqrt{v}$ to be exceeded with probability $1 - G(\alpha)$, where G is cumulative distribution of the standard normal distribution.

The class of predictors is robust with respect to systematic variability (i.e. non-stationarity) in the demand provided that T_{meas} is smaller than the timescale at which demand systematically varies. For example, it is well known that telephone traffic exhibits diurnal variation; the call arrival process can be accurately modeled as a time dependent Poisson process. On the other hand, the call arrival rate is relatively static over intervals of a few minutes; robustness to systematic variations requires choosing T_{meas} to be no larger than this timescale.

If the interval between renegotiations encompasses periods of systematic variation the local estimators above can become inaccurate. Remedies for this include:

- (i) Use a worst case predictor over the largest time scale of variation, e.g., the maximum rate over a day for telephone traffic. This is robust but wasteful of resources.
- (ii) Use historical data to predict trends, e.g., when average telephone call arrival rates are a known non-constant process $Q(t)$, then instead of using the predicted bandwidth $S(t)$ directly, we may use instead $S(t)Q(t + T_{ren})/Q(t)$ as our predictor. Here the ratio $Q(t + T_{ren})/Q(t)$ is used to model the systematic change of the arrival rate upwards or downwards.

Apart from these systematic errors, there are two statistical effects which, if uncorrected, may cause the prediction to underestimate bandwidth requirements. The first is sampling error, the second short-time scale burstiness. Estimation of mean and variance is subject to inherent sampling error since the estimates are themselves random variables. This additional variability can lead to violation of target quality metrics if estimated parameters are assumed to be the true ones; see e.g. [7]. We can show that sampling error for the local Gaussian the sampling error from n samples can be avoided by increasing the multiplier α to

$$\alpha' := \left((1+n)(e^{\alpha^2/n} - 1) \right)^{1/2} > \alpha$$

For example, for $n = 60$ and $\alpha = 3$, $\alpha' = 3.14$.

Burstiness at multiple timescales has been observed in Internet data traffic [8, 9]. The variability of the window-averaged rate of such traffic over a given window may increase for smaller window sizes. A predictor based on a given sampling window can underestimate the bandwidth required to satisfy QoS guarantees specified at shorter timescales. In the case of the local Gaussian predictor, a priori knowledge of the scaling relations between rate variance at different time scales [8, 10] can be used to correct the boost factor α in order to accommodate short timescale variability.

4 Simulation Experiments

We use trace-driven simulation to examine the effectiveness of using hoses. For all the simulated experiments reported in this section, we used an approximation of the AT&T Worldnet IP backbone topology comprising 12 core routers spanning the continental U.S., shown in Figure 4. We use two sets of traces, one for voice traffic and one for data traffic.

Figure 4: Physical Topology of IP Backbone

Voice Traffic We use the call detail records for telephone calls offered to the AT&T switched network in the domestic United States for the August-December 1998 time period. Each call detail record enumerates the originating, dialed and terminating number, along with the origination time and duration of a call. The origination time and duration are captured at the granularity of seconds. One may view this as the potential load to be carried over an IP telephony VPN.

To determine the routing of each of these voice calls, we use the following rules. A given ten-digit telephone number is associated with an area code (based on the first three digits of the phone number). Each of the 168 area codes is assumed to funnel all its traffic through an access link to one of the 12 backbone nodes of Figure 4, that is geographically closest to the centroid of the geographical region corresponding to an area code. Traffic between the backbone nodes follows the shortest path route in the network topology, following normal IP intra-domain routing.

To simulate the network traffic, we convert the call detail records into call counts on a minute by minute interval between every pair of area codes. These call counts are then used to compute the aggregate call statistics on every access and backbone link. Most of the results are derived from simulations corresponding to call counts over a 24 hour interval on Monday, November 9 1998. However, we also use call statistics gathered over the longer period (from August-December, 1998) to motivate how admission control decisions could be made for this class of VPN traffic.

Data Traffic For our data experiments we used NetFlow traffic records [11] gathered from a set of Cisco routers in a corporate network during a 12 hour period. NetFlow provides records at the level of traffic flows. A flow is a logical grouping of packets that share common properties and which are localized in time. The properties used for grouping may include source and destination IP addresses or subnets and port numbers. Localization is achieved by terminating flows through criteria such as timeouts (a packet is deemed to be the last in the flow if no additional packets are received with the same property within a given subsequent period) or through protocol specific packet detail (e.g., a TCP FIN packet). The range of IP addresses present in the NetFlow trace data was divided into 12 groups, each of which was assigned to one of the nodes in the physical topology of Figure 4. Each flow record contained the start and end time of the flow, together with the number of packets and bytes in the flow. Each such flow was mapped into a flow of constant rate transferring the same total number of bytes between the

start and end times. Start and end times were given at the granularity of 1 second.

Experiments For all the simulations using the data traffic, we assume that there are 12 VPN access points, one each for the 12 nodes in the example topology. For the voice simulations, we allow additional levels of traffic aggregation. At one extreme, each area code is considered as a separate VPN access port. In this case, the traffic generated by a hose corresponds to all of the calls originating from a particular area code (independent of the destination) resulting in a total of 168 hoses. Analogously, with a customer-pipe service interface, all the calls made from one area code to another would be a distinct pipe, resulting in a total of 168×168 customer-pipes. At the other extreme, all of the area codes funneling traffic into a particular backbone node are considered to constitute a single VPN access port. Thus, there are a total of 12 hoses and 144 customer-pipes. We also present results for two intermediate levels of aggregation with 24 and 48 hoses respectively.

Using these two sets of data traces we conducted a number of experiments to evaluate the effectiveness of the hose model.

- We investigate the stability of VPN traffic matrices.
- We evaluate the usefulness of the hose model from the point of view of a customer, compared to the customer using a set of customer-pipes, and the benefit of resizing hoses.
- We compare two different mechanisms a service provider may use to realize a hose: a mesh of provider-pipes in the network vs. a source based tree. The provider can manage capacity in a couple of ways:
 - aggregating the required capacity across a hose or across an entire VPN.
 - renegotiate the required capacity, either for the mesh of provider-pipes or for the source based tree.

We quantify the bandwidth savings from either one or using a combination of both techniques.

- Finally, we examine the relationship between short term capacity management by resizing and the longer term admission control algorithms within the context of the hose model.

5 Hoses From a Customer's Viewpoint

One of the primary benefits of the hose model for the customer is the ability of the hose to accommodate changes in the traffic matrix. The use of hoses can also lead to more efficient use of access-link capacity relative to customer-pipes. We explore these issues experimentally in this section.

5.1 Traffic Matrix Variability

We examine the variability of the traffic matrix for data traffic in Figure 5, for a given source hose to all the hose destinations. Out of the 12 hoses in the configuration, we chose one that was representative of the high variability in the traffic matrix, observing the traffic matrix changing across five 2-hour intervals. We observe from the figure that there is considerable variability in the rate of traffic from

Figure 5: VARIABILITY IN THE DATA TRAFFIC MATRIX: Capacity for each destination of a selected hose

the source to the different destinations. The proportion of traffic going from the source to the different destinations changes across the 2-hour intervals. The variability in the traffic matrix indicates the benefit to the customer from the hose interface. However, from the perspective of the network provider, this variability poses a challenge. Indeed, observe in Figure 5 that the most significant variability is for the source-destination flows that dominate the hose.

5.2 Performance Benefit of Hoses for the Customer

In this section we evaluate and compare the benefits to the customer of provisioning with hoses (as opposed to customer-pipes) on the access link, and the benefits of dynamic re-provisioning for either the hose or the customer-pipe model. To define our experiments more precisely, let $E_{\ell}(i)$ denote the set of destination endpoints for traffic sent from source endpoint i and which is routed through link ℓ . Let $r_{ij}(t)$ denote the time series of traffic from endpoint i to endpoint j . Let $S(r)$ denote the capacity required for traffic with a given time series r . With static provisioning, we calculate $S(r)$ as the maximal rate attained over the time series. With dynamic resizing, we calculate $S(r)$ as in Section 3.1, using the Gaussian predictor. In this case, the required capacity is the time series of the predictor. On an access link carrying traffic sourced on hose i , we define the customer-pipe and hose requirements as follows:

- **Customer-Pipe Requirement** $= \sum_{j \in E_{\ell}(i)} S(r_{ij})$
- **Hose Requirement** $= S\left(\sum_{j \in E_{\ell}(i)} r_{ij}\right)$

We calculate the access hose-gain as the ratio of the customer-pipe requirement to the hose requirement. We refer to this quantity as the **statically provisioned access hose-gain** when the capacity for both the customer-pipe and the hose are determined based on their maximum rates. It is called the **dynamically resized access hose-gain** when the customer-pipes and hoses are dynamically resized, calculated as the ratio of their time averages.

5.3 Provisioning the Access Link

The capacity required by a customer on each access link depends on the service model being offered to the VPN cus-

Hose Source	Static Requirement (kB/sec)		Static prov. hose gain
	Customer-Pipe	Hose	
1	2229	1164	1.92
2	2873	1379	2.08
3	13379	12538	1.07
4	4925	2031	2.43
5	619	255	2.42
6	86	79	1.08
7	112	59	1.88
8	3104	2538	1.22
9	1483	416	3.57
10	752	251	2.99
11	778	303	2.56
12	1606	771	2.08

Table 1: STATICALLY PROVISIONED ACCESS HOSE GAIN FOR DATA TRAFFIC: static requirements for customer-pipes and hoses.

tomers. If the customer is presented the abstraction of a frame-relay style independent customer-pipes, adequate capacity would need to get provisioned independently on the access link for each such customer-pipe. A conservative approach would provision adequate capacity per pipe to equal the maximum traffic demand for that pipe, over all time. The aggregate capacity needed on the access link would then be the sum of the maximum capacities needed for each of the customer-pipes on that link.

On the other hand, when the customer's service interface into the network is a hose, then the capacity needed on the access link is the maximum traffic demand for that hose, i.e., the maximum of the aggregate traffic transmitted from the source to all destinations.

Using the voice and data traffic sets described above, we performed an evaluation of the access link capacity required to be statically provisioned for each of these two service models and use this to compute the statically provisioned access hose-gain.

The results of this comparison for the data traffic is shown in Table 1. The statically provisioned access hose-gain varies between 1.07 and 3.57 for this experiment. Examining the actual trace data showed that in the cases where the hose-gain was close to 1, there was a very large burst of traffic between the source and a particular destination which completely overshadowed the traffic contribution to other destinations. Such a burst would naturally dominate both the sum of customer-pipe capacities and the hose capacity, leading to a small gain.

Figure 6 shows a similar comparison for the telephony traffic. The cumulative distribution function (CDF) of the access hose-gain for different levels of source traffic aggregation is shown, i.e., for a given value of the gain on the horizontal axis, the fraction of hoses whose gain was less than this value. (For clarity, we display here and in subsequent figures a smoothed version of the CDF, namely a quantile-quantile plot of the experimental distribution interpolated against 20 points distributed uniformly on the unit interval.) The distribution is across all the access links for each level of aggregation.

When each area code is treated as a separate VPN access point, this would result in a total of 168x168 customer-pipes. On the other hand, with a hose interface, we would have 168 hoses. In this case the statically provisioned access hose-

Figure 6: STATICALLY PROVISIONED ACCESS HOSE-GAIN FOR VOICE TRAFFIC: CDF of hose-gain for different aggregation levels

	number of hoses			
	12	24	48	168
Max. resizing gain	2.15	2.17	2.22	3.08
Mean resizing gain	1.94	1.95	1.98	2.06
Min. resizing gain	1.80	1.79	1.77	1.82

Table 2: HOSE RESIZING GAIN FOR VOICE TRAFFIC: maximum, mean and minimum across hoses at different aggregation levels. Renegotiation at 1 minute intervals .

gain varies between 1.22 and 16.39, with a mean of 1.78. In absolute terms, the aggregate capacity required for *all* of the access links is 45 Gb/s with the hose interface, versus requiring 73 Gb/s with the customer-pipe interface.

5.4 Resizing the Access Link

When there are significant fluctuations over time in the offered traffic, it is useful to provide customers with the capability to renegotiate hose capacities. This renegotiation may be based on demand predictions derived from measurements that track the fluctuations in the offered traffic. Erroneous predictions would result in either wastage of network capacity or the inability to accept all of the offered traffic (resulting in blocked calls for the telephony service and packet losses for the data service). The next set of experiments evaluate the utility of renegotiating hose capacity and quantify how well the predictors described earlier deal with the deterministic and statistical variations in the traffic.

To quantify the utility of renegotiation, we compute the **hose resizing gain**: this is the ratio of the statically provisioned hose requirement (i.e. the maximum offered traffic over the length of the experiment) to the time-average of the renegotiated hose requirement.

5.4.1 Benefit of Resizing the Access Link for Voice Traffic

The next set of experiments explores the utility of hose resizing for the telephony traffic. In these experiments, we use the variance based predictors to compute the capacity required for hoses at each of the 4 different levels of aggregation (12, 24, 48, 168). We use a moving window of 10

Resize Freq.	12 hoses	24 hoses	48 hoses	168 hoses
1 minute	2.34, 0.40	2.37, 0.39	2.53, 0.39	3.55, 0.47
5 minute	2.72, 0.88	2.73, 0.86	2.88, 0.86	3.96, 0.98
10 minute	3.11, 1.42	3.11, 1.38	3.25, 1.37	4.36, 1.49
30 minutes	4.72, 3.54	4.64, 3.40	4.72, 3.36	5.76, 3.48

Table 3: BLOCKING ABOVE HOSE REQUIREMENT FOR VOICE TRAFFIC: Impact of resizing interval on performance. Each table entry represents the percentage overallocation and percentage of calls blocked.

minutes for the measurement and traffic predictions and a renegotiation interval of 1 minute.

Table 2 shows the minimum, maximum and average values for the hose resizing gain, computed across all hoses, over a 24 hour interval, at each of the 4 different levels of aggregation. The mean resizing gain of nearly 2 indicates that there is a significant benefit to be derived by resizing for telephony traffic, even when there is considerable aggregation.

Figure 7: TIME SERIES OF VOICE TRAFFIC: actual traffic and hose prediction for a single hose.

The dynamic behavior of the variance based traffic predictor is illustrated in Figure 7. This is a plot of the time series over the interval [10 am - 1 pm] (busy hour) for the aggregated traffic and the corresponding capacity predictor for a hose. We assume an aggregation level of 168, i.e. each area code constitutes a VPN hose endpoint; this particular hose trace corresponded to telephone traffic originating from an area code in the New York area.

The figure illustrates that the predictor tracks the actual capacity requirements quite closely. A measure of effectiveness of the predictor is the overallocation of capacity, which is 1.78% for this trace. Another measure, which is blocking of offered calls, is nearly 0.3% over a 24 hour interval for this trace.

The results in Table 2 were for a resizing interval of 1 minute. Table 3 illustrates the effect of reducing the resizing frequency on the performance of the predictor as measured by the values of the tuple {%overallocation, %blocking}. The results show that the blocking probabilities become unacceptably high when the resizing interval increases beyond about 5 minutes. However, the amount of overallocation

Figure 8: HOSE PREDICTORS FOR DATA TRAFFIC: variance-based predictors for data over 1 second, 1 minute and 10 minute windows. Bandwidth is shown in nominal units.

Figure 9: DYNAMICALLY-RESIZED ACCESS HOSE GAIN FOR VOICE TRAFFIC. CDF (across access links) of the gain in capacity in going from dynamically resized customer-pipes to dynamically resized hoses.

Figure 10: DYNAMICALLY-RESIZED ACCESS HOSE GAIN FOR DATA TRAFFIC. CDF (across access links) of the gain in capacity in going from dynamically resized customer-pipes to dynamically resized hoses.

Resizing Interval	1 min.	5 min.	10 min.	30 min
Max. resizing gain	62.9	49.8	44.4	37.8
Mean resizing gain	15.5	11.9	10.7	9.23
Min. resizing gain	1.95	1.84	1.80	1.75

Table 4: HOSE RESIZING GAIN FOR DATA TRAFFIC: maximum, mean and minimum across hoses for different resizing intervals T_{ren} .

remains reasonable for such resizing intervals.

5.4.2 Benefit of Resizing the Access Link for Data Traffic

In Table 4 we display the hose resizing gains for data traffic on access links. As we resize the hose more frequently, the higher the gain. The mean gain improves from around 9 to 15 as the resizing interval changes from 30 minutes down to 1 minute.

The behavior of the variance based predictor used in these experiments is illustrated in Figure 8. We plot a sample path of the aggregate traffic over a hose (at one second granularity) together with variance-based predictors with 1 and 10 minutes windows. Observe how the 1 minute predictor follows the trace more closely than the 10 minute predictor. The variance-based predictor is responsive to both upward and downward trends in the trace. This is an advantage for prediction during upward trends; on the other hand it can lead to periods of low utilization after downward trends; see e.g., the central portion of the 10 minute predictor in Figure 8.

The data traffic exhibits far greater short-term variability than the voice traffic. We see later (see the second column of Table 7) the proportion of bytes that exceed the dynamically resized hose rate is quite small, less than 1 or 2%.

5.5 Comparison of Benefits of Resized Hoses and Customer-Pipes

In the next set of experiments, we compare the access link capacity required with resized hoses versus resized customer pipes. As before, we express the ratio of these capacities as a “dynamically resized” access hose-gain. Figure 9 provides results for the voice traffic comparing the access link capacity required using resized hoses instead of resized customer-pipes. The average access-hose gain varies between 10 and 20% depending on the degree of aggregation, with the greater benefits at lower levels of aggregation (observe the case of 168 hose endpoints). The comparison is based on commensurate renegotiation intervals for both alternatives. Thus, on the access link, there is a benefit from the natural aggregation that the hose provides.

Next, consider the data traffic. The CDF for the dynamically resized access-hose gain is shown in Figure 10 for renegotiation windows $T_{ren} = T_{meas}$ of 1, 5, 10 and 30 minutes. The gain is less pronounced for smaller time windows, ranging on average from about 1.4 for the 30 minute window down to about 1.2 for the 1 minute window. The predictor follows the trace more closely for smaller windows; hence the pipe predictor becomes closer to the hose predictor.

Figure 13: PROVIDER-PIPE RESIZING GAIN FOR VOICE TRAFFIC: CDF (over all provider-pipes) of static to dynamically sized provider-pipe requirements, for different levels of aggregation.

Figure 14: DYNAMICALLY RESIZED TREE GAIN FOR VOICE TRAFFIC: CDF (over all network links) of the ratio of dynamically resized provider-pipe requirement to dynamically resized tree requirement, according to aggregation level.

Resizing Freq.	1 min.	5 min.	10 min.	30 min.
max	1291	592	422	272
mean	48.9	35.2	30.6	25.0
min	2.11	1.91	1.80	1.69

Table 5: PROVIDER-PIPE RESIZING GAIN FOR DATA TRAFFIC: Maximum, mean, and minimum of static to dynamically resized provider-pipe. The mean is calculated by weighting each provider-pipe gain by its static provider-pipe capacity requirement.

	min	mean	max
source-destination-pairs per link	2	5.81	13
source-destination-pairs per tree	1	1.85	8
trees per link	1	3.14	8

Table 6: TOPOLOGICAL FACTORS FOR AGGREGATION: provider-pipes vs. trees vs. VPN

gain provided by the tree implementation is between 1% and 17%. While this ratio is smaller than the corresponding numbers for the data traffic, this still represents a fairly substantial bandwidth savings - 540 Mb/s - across all of the links in the backbone network topology, with an aggregation level of 48 hose endpoints.

When we increase the level of aggregation and only have 12 distinct hose endpoints, the additional multiplexing gain for a resized tree beyond that due to resizing provider-pipes becomes progressively smaller. For this reason, for the voice traffic we do not explore the further gains due to aggregation over an entire VPN.

Data Traffic

Dynamically Resized Provider-Pipes: In Table 5 we summarize the statistics of provider-pipe gains when resizing; i.e. the ratio of maximum data rate on the provider-pipe to the time averaged resized rate, for renegotiation windows $T_{ren} = T_{meas}$ of 1, 5, 10 and 30 minutes. Some extreme gains are evident due to localized peaks in the data rate. Mainly these extremes occurred for a provider-pipe with low maximum rate and so did not dominate the weighted mean in Table 5. Broadly, the gains decrease as T_{ren} increases; this can be understood from the dynamics of rates shown in Figure 8; long term variability can lead to increased variance and hence increase predicted requirement at longer timescales.

Dynamically Resized Trees: The CDF over internal links of the tree gain (for resized trees vs. resized provider-pipes) is shown in Figure 15. In nearly 40% of links the gain is 1. To understand the relatively small gain attained by the tree implementation, we also investigated the topological characteristics of the trees. We found that the average number of distinct endpoints for a tree on a link was 1.85. Aggregation has little impact especially as we near the destination.

Dynamically Resized VPNs: Aggregating all flows present from a VPN on a given link potentially gives more multiplexing gain. This is evident from Table 6 where we compare the statistics of the number of source-destination-pairs that are carried on a link, the number of source-destination-pairs per tree on the link, and the number of trees on the link. In contrast to an average of less than 2 source-destination-pairs per tree, there are on the average nearly 6 source-destination-pairs per link for the VPN. We believe that aggregating per VPN is feasible in that hoses within a VPN are likely to be "aggregatable" in that they will be statistically similar and to have the same QoS requirements. This motivates sharing reservations by VPN's. The CDF (over network links) of the VPN gain with dynamic resizing is shown in Figure 16. We observe greater gain than for resized trees. There are also increased gains for longer resizing intervals T_{ren} . We

Figure 15: DYNAMICALLY RESIZED TREE GAIN FOR DATA TRAFFIC: CDF (over all network links) of the ratio of dynamically resized provider-pipe requirement to dynamically resized tree requirement, according to resizing window.

Figure 16: DYNAMICALLY RESIZED VPN GAIN FOR DATA TRAFFIC: CDF (across all network links) of ratio of resized VPN requirement to resized provider-pipe requirement, according to resizing interval.

Resizing Frequency	Access Links		Internal Links		
	Pipe	Hose	Pipe	Tree	VPN
1 minute	0.041	0.020	0.041	0.035	0.023
5 minute	0.029	0.015	0.030	0.026	0.018
10 minute	0.021	0.012	0.021	0.020	0.016
30 minutes	0.016	0.011	0.015	0.014	0.011

Table 7: EXCURSIONS ABOVE RATE REQUIREMENTS FOR DATA TRAFFIC: Impact of resizing interval on the proportion of offered traffic that exceeds the allocated capacity, according to level of aggregation.

interpret this as indicating that potentially higher variability at longer timescales (see e.g. Figure 8) is reduced due to smoothing by aggregation across source-destination pairs.

The ability of aggregation to maintain quality is borne out by our data traffic example; in Table 7 we show statistics of the proportion of offered traffic that exceeds the allocated capacity, when using dynamically resized provider-pipes, trees and VPNs. In general, increasing the degree of aggregation reduces the level of underallocation.

7 Effective Bandwidths for Admission Control

For admission control for hoses of a fixed capacity it suffices to perform a test at each network link over which a hose's traffic is routed. The test verifies that the total capacity allocated for active hoses on that link plus the capacity required for the new hose does not exceed the available bandwidth.

However, when hoses can be resized, the computation of the total capacity allocated to active hoses at a link, cannot be done based on the currently requested capacity for each of these hoses. This is because the decision to admit a hose impact resource usage for the lifetime of the hoses, not just in the short term. Instead, it is desirable to use a more

stringent measure of the bandwidth requirement of each active hose - we will refer to this as an effective bandwidth. Here we propose that the effective bandwidth for an admitted hose be the maximum over its lifetime of a shorter term bandwidth requirement. Candidates for algorithms to characterize of the latter include algorithms proposed for bandwidth estimation for statistical QoS; see e.g. [12] for a review. A natural choice available in the present framework is the predicted bandwidth used for renegotiation. For the local maximum predictor, the effective bandwidth reduces to the maximum traffic rate seen at the sampling timescale.

In the experiments reported in this section used at the effective bandwidth the maximum over the hose lifetime of the local Gaussian predictor. We compare the properties of the effective bandwidths for the different levels of aggregation: pipe, hose and VPN. The CDF of the ratio of effective bandwidths for customer-pipes vs. hoses on access links in Figure 17, for provider-pipes versus hoses on internal links in Figure 18 and for provider-pipes vs. VPN in Figure 19. Observe greater gain for link than for hose aggregation; however, the gain is relatively insensitive to the size of the prediction window. Table 8 shows that the proportion of bytes above the effective bandwidth is exceedingly small, and decreases with the degree of aggregation.

The Gaussian predictor admits some tuning for sensitivity through the parameter α that determines the target number of standard deviations above the mean at which the predictor should lie. Tuning this parameter downwards enables the provider to overbook resources.

8 Conclusions and Future Work

VPNs are undergoing dramatic change owing to at least three interrelated factors:

- rapid progress in IP network technologies (in overall capacity and the development of diverse network access technologies);

Figure 11: STATICALLY PROVISIONED TREE GAIN FOR DATA TRAFFIC FOR 1 TREE: Ratio of provider-pipe requirements to tree requirements, for each link used by a hose connected to node 1.

6 Performance Benefit of Hose Realizations for a Provider

For the topologies and voice and data traffic described in Section 4, we evaluate the benefits for a provider of the three implementations of the hose model described in Section 3: a set of provider-pipes, trees, and aggregated VPNs. For all of these alternatives, we assume default (shortest path) routing.

Using the same notation as in Section 5 we define:

- **Provider-pipe Requirement** = $\sum_i \sum_{j \in E_\ell(i)} S(r_{ij})$
- **Tree Requirement** = $\sum_i S\left(\sum_{j \in E_\ell(i)} r_{ij}\right)$
- **VPN Requirement** = $S\left(\sum_i \sum_{j \in E_\ell(i)} r_{ij}\right)$

We evaluate the **tree gain** for a link ℓ as the ratio of the provider-pipe requirement to the Tree requirement. For voice traffic we perform this evaluation at different levels of hose aggregation; see Section 4. For data traffic we evaluate the requirements for two levels of aggregation. First, we calculate the tree requirement for trees based on hoses from 12 VPN endpoints, one accessing the network at each of the 12 nodes of the network topology. Second, we calculate the VPN requirement: that of the aggregate of all provider-pipes routed over link ℓ ; the **VPN gain** is the ratio of the provider-pipe requirement to the VPN requirement. We calculate these gains under both static provisioning and dynamic resizing.

6.1 Benefits of Statically Provisioned Trees

Consider the requirement for a tree-based implementation of a single hose. Moving from the root of a tree corresponding to a given hose towards a leaf, progressively fewer flows are aggregated together and hence we expect the benefit of sharing reservations in the tree to decrease. This is illustrated

Figure 12: STATICALLY PROVISIONED TREE GAIN FOR DATA TRAFFIC FOR ALL TREES. Ratio of provider-pipe requirements to tree requirements for each link. Each link is labeled with two ratios, as the links are bidirectional. The ratio corresponding to the link from node A to node B is placed nearest to node A.

for the data traffic in Figure 11, which pertains to the hose with the ingress being node 1. At the access link, the gain (the ratio of the requirement of the hose to the corresponding sum of the requirements of customer-pipes) is 1.92. As we trace a particular path towards a leaf, for example the path between nodes 1 and 12 via links 1 to 2, 2 to 5 and 5 to 12, the tree-gain is respectively 1.73, 1.94 and 1. The full tree and provider-pipe requirements are shown in Figure 12. It turns out that the tree gain varies between 1 and 2.57 with an average of 1.38. A tree gain of 1 occurs on links where each tree present on the link leads toward a single destination.

6.2 Benefits of Dynamical Resizing

Voice Traffic

Dynamically resized provider-pipes: Figure 13 shows the CDF (over all provider-pipes) of the resizing gain for voice traffic. By resizing gain, we mean the ratio of the peak capacity of the static provider-pipe to the peak capacity of the dynamically resized provide pipe. When there is a medium level of aggregation (where the 168 area codes are aggregated to only 48 distinct hose endpoints), there is a reasonable benefit from resizing. The average gain for an internal link capacity is about 1.27, with a minimum of 1.08 and a maximum of 1.71, using a resizing interval of 1 minute. However, as we increase the level of aggregation and only have 12 distinct hose endpoints, the gain due to resizing becomes much smaller. The mean gain is only about 7%. This should be expected though, because there is progressively less variability to be exploited by resizing as aggregation increases. With 12 endpoints the gain results from time of day variation in the traffic, seen on cross-continental links that carry traffic for different time-zones with different time-of-day peaks.

Dynamically Resized Trees: Figure 14 shows the CDF of the tree gain over all the internal network links. The additional

Figure 17: HOSE GAIN FOR EFFECTIVE BANDWIDTH ON ACCESS LINKS FOR DATA TRAFFIC: CDF (over access links) of ratio of maximum pipe requirement to maximum hose requirement, according to renegotiation interval.

Figure 18: HOSE GAIN FOR EFFECTIVE BANDWIDTH ON INTERNAL LINKS FOR DATA TRAFFIC: CDF (over internal links) of ratio of maximum pipe requirement to maximum hose requirement, according to renegotiation interval.

Resize Freq.	Access Links		Internal Links		
	Pipe	Hose	Pipe	Hose	VPN
1 minute	1.4e-5	5.9e-6	1.6e-5	1.3e-5	1.0e-5
5 minute	1.1e-4	3.9e-5	1.0e-4	9.1e-5	4.1e-5
10 minute	2.1e-4	6.3e-5	1.9e-4	1.6e-4	9.9e-5
30 minutes	5.8e-3	1.7e-4	5.4e-4	4.1e-4	3.5e-4

Table 8: EXCURSIONS ABOVE EFFECTIVE BANDWIDTH FOR DATA TRAFFIC: Impact of resizing interval on the proportion of byte underallocation of requirement, according to level of aggregation

- progress in IP security (in flexible, dynamic methods for establishing secure associations);
- rapid change in the diversity and dynamics of communication and collaboration patterns at work and at home.

As a result, communication patterns are evolving from fairly static, predictable flows between endpoint *pairs* to dynamic, difficult to forecast traffic between *sets* of endpoints.

Accordingly, we have proposed a more powerful, easier to specify service model for the customer, termed a *hose*. A hose is characterized by the aggregate traffic to and from one endpoint in the VPN to the set of other endpoints in the VPN, and by an associated performance guarantee. A hose allows a customer to simply buy a logical *access* link and use it to send traffic to any one of the remote hose endpoints, with reliable QoS, and with the rates of the customer access links the only limitation. In addition, hoses naturally allow the customer to take advantage of aggregation of the flows to and from access links, reducing required access link capacities. Though it would appear that hoses present greater resource management challenges for the provider, these difficulties can be addressed by statistical multiplexing or resizing techniques, applied separately or in combination.

Using trace-driven simulations, we examined the effectiveness of the hose interface and the benefit of resizing. An IP telephony service may be one of the major customers of the type of VPN service that we propose here. Therefore, we looked at traces of telephone calls over the AT&T national long distance network, as well as traces of data traffic on a large corporate private network. Our simulations showed significant capacity savings by using techniques to improve statistical multiplexing and resizing. Under conditions where there is significant fluctuation in the customer demands:

- On access links, we find there is considerable potential for statistical multiplexing. It achieves a factor of 2 to 3 in capacity savings over statically provisioned customer-pipes. On network internal links, we found that statistical multiplexing, by hose, typically provides small benefit because the number of distinct destinations reachable on the link for a given hose is small. We suspect that with a richer network topology, especially with less aggregation in each customer-pipe, the provider may achieve greater benefits even on network internal links. On the other hand, statistical multiplexing by VPN, provides significant benefit.
- Similarly, resizing provides about a factor of 2 in savings in access link capacity when resizing the hose just once a minute. A simple predictor based on a local Gaussian approximation was used to calculate the expected load. The gains are much higher with data traffic because of the higher variability. On network internal links, resizing is similarly effective.
- Combining statistical multiplexing and resizing, provides additional benefit, over applying either separately.

We believe the VPN service model presented here will naturally and economically fit emerging business practices in an increasingly IP networked environment. In this paper, we have just begun to address VPN performance issues

Figure 19: VPN GAIN FOR EFFECTIVE BANDWIDTH ON INTERNAL LINKS FOR DATA TRAFFIC: CDF (over internal links) of ratio of maximum pipe requirement to maximum VPN requirement, according to renegotiation interval.

within the IP context. It is very important to understand the detailed technical specifications of the QoS attributes for hoses, i.e., the assurances in the SLA related to delay, loss, and jitter. In general, we expect these specifications to be looser for hoses than for customer-pipes. Thus, we expect both paradigms will play important roles. It is our hope that the step toward problem formulation and analysis taken in this paper will encourage further research on VPN performance.

References

- [1] S. Fotedar, M. Gerla, P. Crocetti, and L. Fratta, "ATM Virtual Private Networks," *Communications of the ACM*, vol. 38, pp. 101–109, Feb 1995.
- [2] B. Gleeson, A. Lin, J. Heinanen, and G. Armitage, "A Framework for IP Based Virtual Private Networks." draft-gleeson-vpn-framework-00.txt, September 1998. Work in progress.
- [3] K. Muthukrishnan and A. Malis, "Core IP VPN Architecture." draft-muthukrishnan-corevpn-arch-00.txt, October 1998. Work in progress.
- [4] D. Jamieson, B. Jamoussi, G. Wright, and P. Beaubien, "MPLS VPN Architecture." draft-jamieson-mpls-vpn-00.txt, August 1998. Work in progress.
- [5] S. Rooney, J. E. van der Merwe, S. Crosby, and I. Leslie, "The Tempest, a Framework for Safe, Resource Assured, Programmable Networks," *IEEE Communications Magazine*, vol. 36, pp. 42–53, October 1998.
- [6] S. Kent and R. Atkinson, "Security Architecture for the Internet Protocol." RFC 2401, November 1998.
- [7] M. Grossglauser and D. N. Tse, "A Framework for Robust Measurement-Based Admission Control," in *Proceedings ACM SIGCOMM 97*, (Cannes, France), 1997.
- [8] W. E. Leland, M. S. Taqqu, W. Willinger, and D. V. Wilson, "On the Self-Similar Nature of Ethernet Traffic (Extended Version)," *IEEE/ACM Transactions on Networking*, vol. 2, no. 1, pp. 1–15, 1994.
- [9] V. Paxson and S. Floyd, "Wide-Area Traffic: The Failure of Poisson Modeling," *IEEE/ACM Transactions on Networking*, vol. 3, no. 1, pp. 226–244, 1995.
- [10] A. Feldmann, A. C. Gilbert, W. Willinger, and T. Kurtz, "Looking Behind And Beyond Self-similarity: On Scaling Phenomena in Measured WAN Traffic," in *Proceedings of the 35th Allerton Conference on Communication, Control and Computing*, 1997.
- [11] Cisco NetFlow. For more information see <http://www.cisco.com/warp/public/732/netflow/>.
- [12] E. Knightly and N. Shroff, "Admission Control for Statistical QoS: Theory and Practice," *IEEE Network*, vol. 13, no. 2, pp. 20–29, 1999.