

Origin of route explosion in Virtual Private Networks

Zied Ben-Houidi ‡ † zied.ben-houidi@etu.upmc.fr
Renata Teixeira † renata.teixeira@lip6.fr
Marc Capelle ‡ marc.capelle@orange-ftgroup.com
† Univ. Pierre et Marie Curie and CNRS
‡ France Telecom R&D Orange Lab

1. INTRODUCTION

Enterprises often have sites that are spread in distant locations. These sites need to interconnect with the same level of privacy as in a local-area network. Virtual Private Networks (VPNs) were introduced to serve this need. A common VPN technology uses Multiprotocol extensions for the Border Gateway Protocol (MP-BGP) and Multiprotocol Label Switching (MPLS). This technology allows a service provider to share its IP backbone among multiple VPN clients while preserving privacy. MPLS tunnels provide traffic isolation, whereas MP-BGP distributes VPN routes. Despite the wide deployment of BGP/MPLS VPNs[1], there have been only few studies to understand their behavior, mostly because of the lack of public data. Prior work focused on BGP convergence [3] and on integrity constraints to ensure connectivity [2].

This work focuses on the scalability of a VPN provider backbone in terms of the number of routes. We collect both MP-BGP routing messages and router configurations from a large European VPN service provider which has a fully dedicated network to provide the VPN service. A simple analysis shows that the BGP routing table of a core router in this network has over 680 thousand routes, which is three times more than a core Internet router. To understand the origin of such number of routes, we study the distribution of routes among VPN clients. This analysis needs a per-VPN view. Unfortunately, routing messages alone cannot distinguish VPNs. We propose a methodology to extract per-VPN information from MP-BGP messages and configurations of all routers in the network. This abstract first presents our methodology. Then, we give highlights of our preliminary results. We find a disproportion in the distribution of routes, 10% of VPN clients contribute with almost 90% of VPN routes!

2. EXTRACTING PER-VPN PROPERTIES

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.

CoNEXT'07, December 10-13, 2007, New York, NY, U.S.A.

Copyright 2007 ACM 978-1-59593-770-4/ 07/ 0012 ...\$5.00.

MP-BGP routing messages carry attributes that specify the way the route should be processed. One of these attributes is the route target (RT), it tags routes so that routers know which route belongs to which VPN. Each customer site connects to the provider backbone through a Provider Edge (PE) router. Each PE stores a Virtual Routing and Forwarding table (VRF) for each VPN connected to it. When a VRF learns a route from its directly connected VPN, it distributes it only to other VRFs in the network belonging to the same VPN. Each VRF has an import and an export list of RTs. When a VRF advertises a route, it tags it with the RTs in its export list. A VRF only imports a route if it is tagged with an RT in its import list.

Two VRFs communicate only if they import each other's routes. In this case, they are part of the same VPN client. We now formalize the definition of a VPN and explain how to extract per-VPN connectivity from router configurations. Then, we discuss how to obtain per-VPN information from MP-BGP table dumps.

2.1 Formalizing the problem

First, we define the communication relation represented by \longleftrightarrow .

Definition 1: Let VRF_1 and VRF_2 be two VRFs, $(VRF_1 \longleftrightarrow VRF_2)$ if and only if $\exists (RT_1, RT_2) |$
 $RT_1 \in imp(VRF_1) \cap exp(VRF_2)$
 $RT_2 \in exp(VRF_1) \cap imp(VRF_2)$

We define $imp(VRF)$ as the list of route targets imported by the VRF and $exp(VRF)$ as the export list. We use this basic operation to construct the graph of all VRFs. We consider VRFs as vertices and add an edge between VRF_1 and VRF_2 if $VRF_1 \longleftrightarrow VRF_2$. We define a VPN as follows.

Definition 2: Given a VRF graph, a VPN is a connected component in the graph.

This definition is flexible enough to accommodate multiple types of VPN topologies, from hub-and-spoke to full mesh connectivity.

2.2 Building VRF graph

A PE router configuration contains information about the export and import lists for each VRF that belongs to that PE. We first develop tools to parse router configurations of all PE routers in the network. The parser takes into consideration the difference in syntaxes be-




Figure 1: Distribution of routes among VPN clients

tween the two router brands used in our network. The tool first constructs what we call a VRF object file. Each entry in this file contains information about a potential vertex in our graph. Next, we apply a standard graph algorithm to extract the connected components. Each connected component is a VPN client. This list also stores information about the used RTs and graph structure.

The construction of the VRF graph using the full VRF object can give inaccurate results. Some route targets are used to administrate the network and they are imported and exported by many VRFs. They don't have to be taken into account when building the graph.

2.3 Per-VPN information from MP-BGP data

MP-BGP VPN routes are tagged with RTs. Given the set of RTs used by each VPN client, we associate routes with VPN clients.

3. RESULTS

We apply the methodology described in the previous section on a snapshot of ten days of collected MP-BGP updates and router configurations from the same period. We enumerate more than 10 thousand VPN clients (connected components on the graph). We study the distribution of routes among them. Figure 1 shows the number of routes announced by each VPN client. The vast majority of VPNs (70%) announce less than 10 routes whereas 5% of them announce more than 100 routes. If we sum up the contribution of the highest 10%, we find that it corresponds to near 590,000 routes which makes almost 90% of all the routes. Although surprising, this result has to be considered with regards to the number of sites per VPN. The number of routes per VPN has to be higher if the VPN has a high number of sites since each site needs at least one prefix. We study therefore the distribution of the number of sites per VPN client. Around 5% of VPN clients have more than 20 sites whereas 60% have less than 2. This makes that 10% of clients have 50% of the total number of VPN sites. This result implies that an important factor




Figure 2: Distribution of the number of routes per VRF

in the 10%/90% result is the disparity in VPN sizes, but it does not explain all. Figure 2 shows the distribution of routes among VRFs. We use routes announced by VRFs to estimate the number of routes announced by VPN sites, because we cannot obtain this exact number from data collected from the provider backbone. The number of routes per VRF is a good approximation because most often there is only one VPN site connected to a VRF in a PE, which implies one VRF per VPN site. Intuitively, if VPN customers allocate IP addresses to their network carefully, it should be possible to aggregate all IP addresses of a VPN site into a single IP prefix. In this scenario, there should be only one route per VRF, which was true for only 50% of the VRFs. However, around 10% of VRFs announced each more than 10 routes.

To understand the reason of this high number of routes per VRF, we study VPN address allocation among VPN clients. We find that around 30% of VPN clients have 100% of their network masks higher than /26, 55% have 80% of their network masks higher than /26. These are signs that there is a high potential for aggregation.

4. PERSPECTIVES AND FUTURE WORK

Having one route per VRF could reduce up to 75% of the total number of BGP routes in the network. We plan to study the feasibility of one route per VRF. Our future work also comprises studying how VPN structures fall on the underlying backbone topology. This would help us evaluate the BGP table sizes of each PE router.

5. REFERENCES

- [1] E. Rosen and Y. Rekhter, BGP/MPLS IP Virtual Private Networks (VPNs), in *IETF RFC 4364*, February 2006
- [2] R. Bush and T. Griffin, Integrity for Virtual Private Routed Networks, in *Proceedings of the IEEE INFOCOM*, April 2003.
- [3] D. Pei and J. V. der Merwe, BGP Convergence in Virtual Private Networks, in *Proceedings of the ACM SIGCOMM on Internet measurement*, 2006