Journal of Information Assurance and Security.
ISSN 1554-1010 Volume 8 (2013) pp. 057-066

Proposals for the construction of high-scalability extranet MPLS VPNs

Adão Boava[1] and Yuzo Iano[2]

[1]Federal University of Fronteira Sul,
Av Dom João Hoffmann 313, Brazil
adao@uffs.edu.br

[2]State University of Campinas,
Av. Albert Einsten, 400, Brazil
yuzo@decom.fee.unicamp.br

***Abstract*: This article presents proposals to solve the scalability problem of MPLS (Multiprotocol Label Switching) VRFs (VPN Routing and Forwarding). Because of the limitations of PEs (Provider Edges) in terms of processing and memory, the number of VRFs that can be implemented in the same equipment is limited. Due to scalability, the use of a single RD (Route Distinguisher) per VPN, rather than per site is recommended. Such recommendation is only viable in intranet construction, as there is no connection problem among sites. When it is necessary to construct an extranet, the solution is the use of one RD per site. Although this is technically feasible, it is not advisable, since it reduces the scalability of MPLS VRFs. This paper proposes some alternatives to construct extranet keeping the high scalability of MPLS VPNs. Two proposals to solve the problem of VPN scalability are presented. The first one is based on ACL (Access Control List) for sites without VPN, and import/export of routes to sites that already have VPN; the second one is founded on the implementation of firewall in the network core.**

***Keywords*:** Multiprotocol Label Switching, Intranet, Extranet, Scalability and VPN Routing and Forwarding**.**

I. Introduction

MPLS VPNs are regarded as the main elements of the architecture of convergence of Next Generation Networks and have become increasingly more accessible to users, particularly because of their high scalability and easy implementation. However, this model acts directly on VRFs (VPN Routing and Forwarding) of PEs (Provider Edges), which rapidly increase as the number of sites of VPNs connected to PEs increases. Such increase may create some problems for the MPLS VPN service provider and may also hinder scalability, thus generating some difficulties to provide new MPLS VPN services. Such problems grow in importance especially when users that do not belong to the same organization need to access the VPN, i.e. when the construction of an extranet is required.

Due to the limitations imposed by the use of one VPN identifier (RD) per site, new research lines have emerged to study a way of not hindering the high scalability of MPLS VPNs, which is their major advantage.

Aiming at better using the potential of MPLS VPNs, MPLS VPN service providers have proposed the utilization of only one route identifier per VPN, i.e. all the VPN sites would use the same route identifier (RD). This is an interesting proposal when the goal is the exclusive construction of intranet, and such a solution has been adopted by almost every MPLS VPN service provider in the world. Nevertheless, when there is a need for construction of an extranet, the creation of only one RD per VPN has shown some problems that harm the extranet construction, as we will show in the following sections. Zhangwei He and Yong Jiang have divided the PE routers into two groups: hubs and spokes. A small number of hubs maintain full reachability information for a VPN, while spokes with reduced reachability information achieve any-to-any reachability by delivering traffic via hubs [7].

This article presents two proposals. The first one is based on the creation of MPLS VPNs for big content providers (e.g. credit card companies) with access to such VPNs through import and export of routes of RT and RD attributes to users that already have their VPNs. In the case of users with no VPNs, the access to big providers will be through the creation of ACL (Access Control List). The second proposal is grounded on the implementation of a firewall centralized in the network core.

In Section II, the elements of the architecture of MPLS VPNs will be presented, with stronger emphasis on the border equipment (PE). Section III evaluates issues related to the scalability of the main element of the MPLS architecture, which is PE. In Section IV, two proposals to solve the problem of scalability are presented. Finally, Section V presents the conclusions.

II. VPN MPLS Architecture

The main element of MPLS architecture is shown in Figure 1. PE is the equipment that is in the provider environment and network border. PE routers exchange routing information with CE routers through static routing, Routing Information Protocol Version 2 (RIPv2), Open Shortest Path First (OSPF) or Border Gateway Protocol (BGP). This VPN model increases scalability because it eliminates the need for PE routers to keep VPN routes to all PEs of the MPLS VPN service provider. Each PE router keeps a VRF for each directly connected site. Multiple interfaces of PE router can be associated with only one VRF if all the access sites participate in the same VPN. After learning the local VPNs of CE routers, a PE router exchanges routing information with the other PEs through BGP. When MPLS is used to direct VPN data traffic through the backbone of the VPN service provider, ingress and egress PE routers work as ingress and egress LSRs, respectively [2, 5].

A key concept in VPN MPLS architecture is the element called Routing and Forwarding Table of PE routers (VRF). A private VRF is only accessible through interfaces that are part of the corresponding VPN. All the sites connected to the PE router should be part of a VRF. All the VPN information is reflected on the VRF and packets that travel across that site will be routed and forwarded based only on the information found in the corresponding VRF [2].

The PE router can handle thousands of sites of the MPLS VPN service providers' clients directly connected to the PE interfaces. In order to keep the connectivity among all the sites that belong to the same VPN, each PE router should have the routes to the sites belonging to the VPN in a VRF table. Therefore, the VPN routing table in PE routers rapidly increases as the number of both VPN sites and VPNs connected to PE increases. As a result of the increased number of VRFs in PE, the capacity of PE memory and processor has become a focus of study in several researches addressing issues related to both PE scalability and the increased number of VRFs.

One key benefit to MPLS VPNs is the ability to have overlapping IP address spaces between two or more VPNs. This is possible since they will never share the same routing table(i.e., the same VRF). This allows members of one VPN to use private addressing schemes in their customer networks and exchange their private routes over the service provider backbone without interfering with other customers who use exactly the same private address space in their networks. This also allows service providers allocate so-called addresses to these VPNs, since these addresses are not used for external connectivity. In cases where they are, some sort of network address translation (NAT) is required [2].

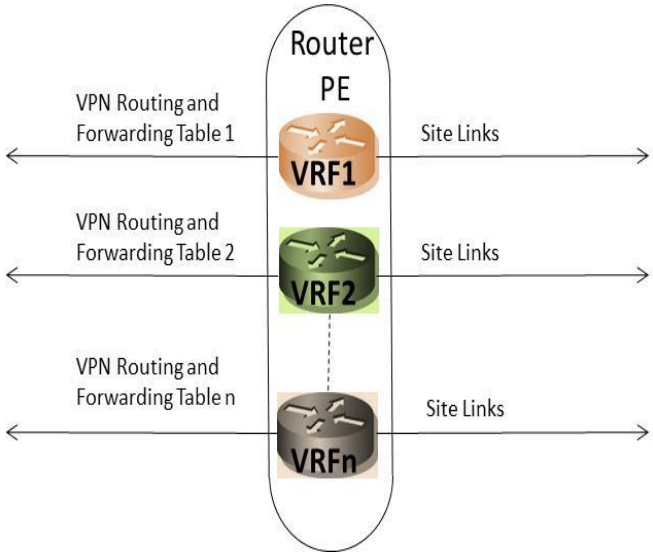

Fig. 1. PEs with several VRFs, taken from [2].

The connectionless of MPLS VPNs has many implications for scalability of the overall MPLS network, but also for security: On an ATM network, for example, a VPN customer typically will be presented with a number of virtual connections from a given router to all other routers that need to be connected. However, the customer needs to configure the router to use these virtual connections. The disadvantage here is that many virtual connections have to be configured on both the customer side and the service provider side. The advantage is that the customer has full visibility of the VPN and controls the connections. On an MPLS network, the same customer router will in most cases be presented with a single connection into the MPLS network, and it is the MPLS network itself that decides where to forward packets to. The customer loses the view of the connections through the core. The advantage of this approach is scalability: the provisioning complexity is reduced to a single connection for each customer router; but the customer does not have visibility of the core network anymore [4].

Regarding intranet VPNs, the problems above mentioned are minimized, as telecommunication operators can implement the same RD for VRFs of the same VPN, i.e. a single RD is configured per client/VPN, rather than one RD per site. As a benefit, this causes a lower consumption of memory and processors of PE routers. Figure 2 shows two MPLS VPNs

(VPN AB and VPN CD) with two sites per VPN, and RD (VPN identifier) is the only one for each VPN. Both VPN AB and VPN CD are intranet, as there is only communication between sites of the same VPN. Figure 2a illustrates VPN AB, Figure 2b illustrates VPN CD, and Figure 2c shows the implementation of MPLS VPN AB and CD. In Figure 2, it is important to observe that RD is the same one for each VRF of the same VPN; for VPN AB, RD is X, while for VPN CD, RD is Y. Hence, in order to make feasible the connectivity between sites of the same VPN, it is necessary that only the RT of all VRFs imports and exports the same RD.

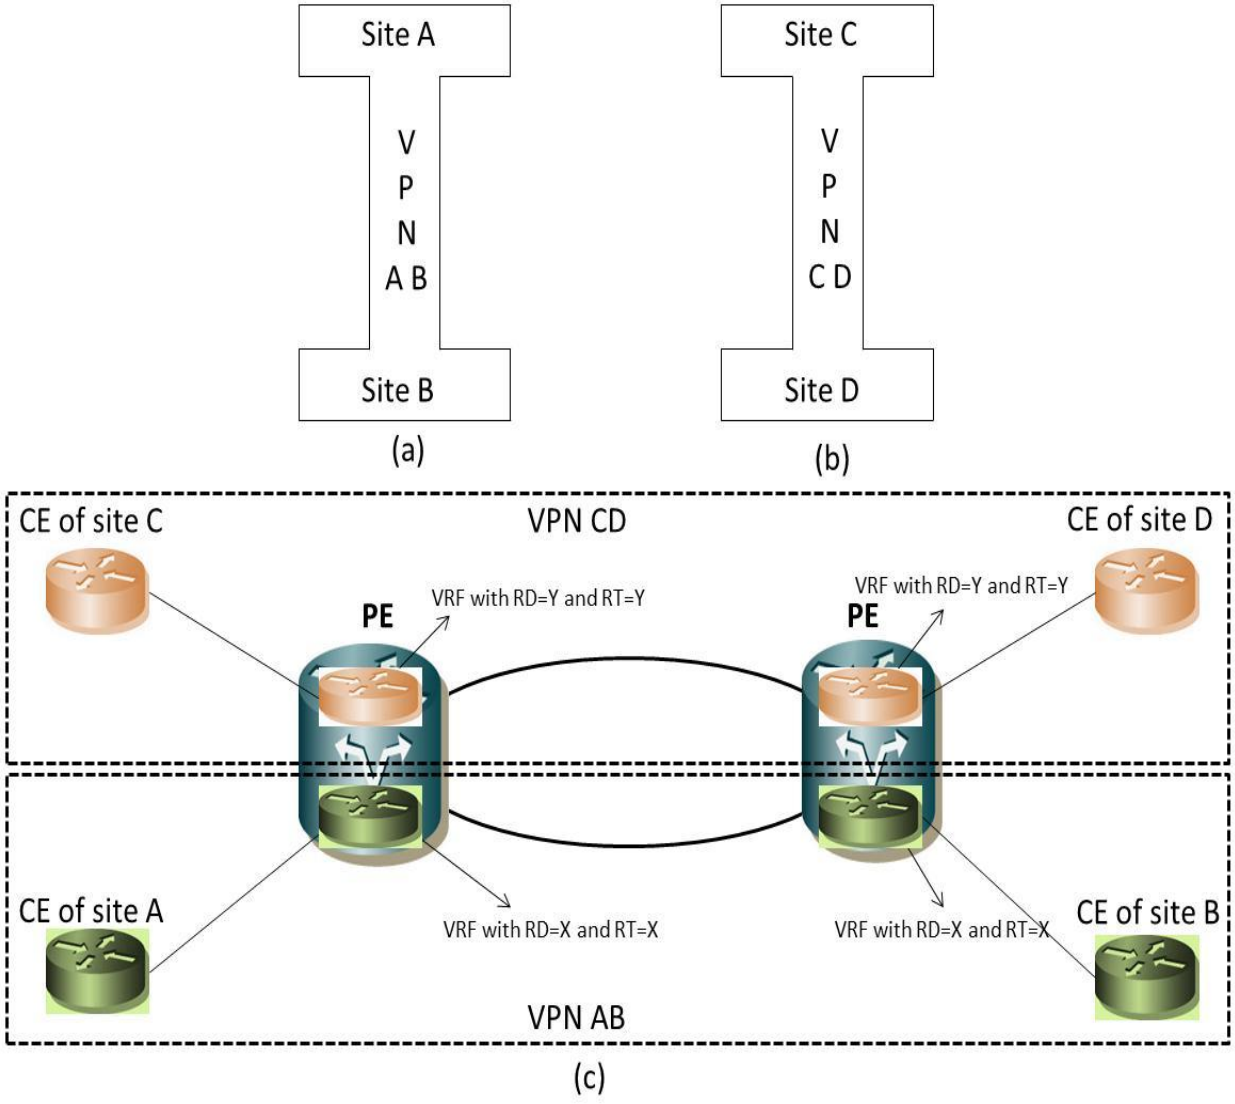

Fig. 2. Implementation of intranet MPLS VPN

The decision to create RD per VPN rather than per site aims to improve scalability, but the problem becomes critical when it is necessary to construct an extranet, e.g. when a site must be part of more than one VPN at the same time, as it is shown in Figure 3. In this case, two sites (B and C) belong to two VPNs simultaneously. These sites form what has been known as extranet. We will consider that the VPN AB identifier (RD) is X, and the VPN CD identifier (RD) is Y. In VPN BC, only site B is connected to site C. If in site C RT is equal to X, not only will site B be connected to site C, but site A will also present connectivity. This is due to the fact that one RD is configured per VPN, rather than per site. Therefore, a theoretically viable solution would be configuring RD per site, not per VPN, as Figure 5 shows.

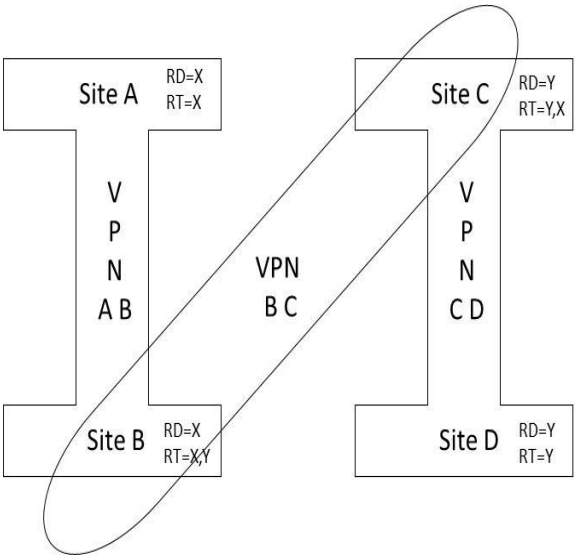

Fig. 3. Implementing intranet VPN with RD per VPN.

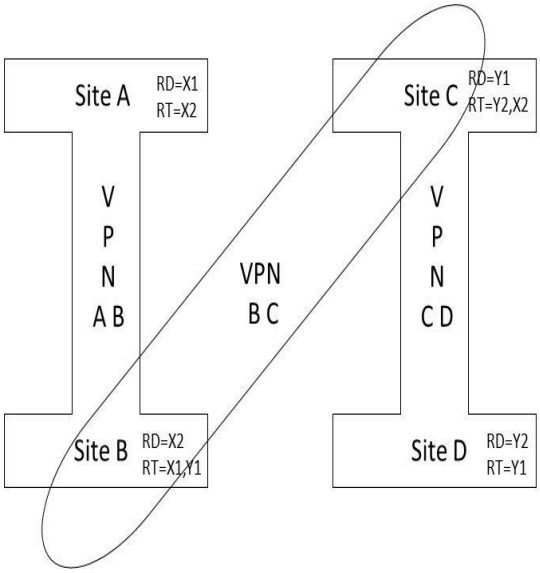

Fig. 5. Implementation of extranet VPN with RD per site.

The creation of RD per site is a theoretically feasible alternative desired by every user, as it is safe and allows for the use of the same IP address in all sites, as long as RDs are different. However, this is not the solution chosen by telecommunication operators because the creation of RD per site would cause high consumption of PE memory and processing, as shown in Figure 4.

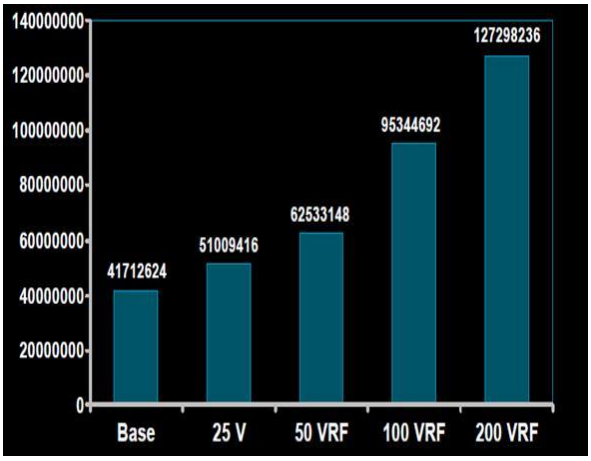

Fig. 4. Number of VRFs x PE Memory Use, taken from [1].

Figure 5 presents three VPNs (AB, CD and BC). Site B of VPN AB and site C of VPN CD belong simultaneously to VPN BC. The solution implemented should allow only site A of VPN AB to connect only to site B, site B to connect to site A and site C, site C to connect to B and D, and site D to connect only to C. As each site has a different RD, the implementation of extranet is totally possible. However, due to the recommendation not to use RD per site, as this hinders the scalability of MPLS networks, in the next section we propose some alternatives concerning this issue. Before presenting the proposal of scalability, we will present some issues related to routing protocols that are configured between CE and PE, as the selection of a particular routing protocol may influence PE scalability.

A major issue in MPLS VPN architecture is to assess how the routes of VPN sites are directed and routed in the MPLS network through VPN Routing and Forwarding (VRF). It is important to know the process of creation and propagation of routing tables in PE routers. Depending on the routing protocol used between CE and PE, there will be either a big or a small impact on processing and memory of PE routers, thus considerably affecting MPLS VPN scalability [3].

PEs learn which CE routes are linked to VRFs through the most common routing protocols, which are the following: Static Routes, RIPv2, OSPF and BGP. Such routing protocols must insert the routes learned in VRF through an interface between CE, which is the user's equipment, and PE, which is the provider's network border equipment. After the connection of CE of each site of VPN with a VRF in PE is defined, it is necessary to choose which routing protocol will be used to advertise the routes to the VRF of PE [63]. The main factors that determine the utilization of a certain type of protocol are: security, user's CE processing limitation, and required control [3, 4, 6].

III. PE Scalability

The physical properties of border routers (PEs) of MPLS network, such as central processing unit (CPU) processing capacity, memory size and time of convergence of routing protocols, together define the potentiality of the PE router to work with several VRFs simultaneously. The amount of physical and logical interfaces that a PE can handle is the first important factor in the evaluation of scalability of a MPLS VPN, as this restrains the number of VRFs that a PE can handle. There may be a situation in which the number of VRFs is smaller than the number of logical interfaces, leading PE suppliers to specify two basic characteristics, namely, the amounts of VRFs and the amounts of routes per VRF. The

amount of routes per VRF depends on the routing protocol that is configured between CE and PE. The amount of routes handled by VRFs of a PE that is connected to CE through RIP, for instance, is smaller than the amount of routes per VRF if the routing protocol between CE and PE is a static route [3].

IV. Proposal for scalability

In this section, we present two proposals aiming at mitigating the scalability problem in VRFs of MPLS VPNs configured in PE routers. For better understanding the need for scalability in MPLS VPN service, consider the need for a certain telecommunication service provider to offer a solution to three credit and debit card companies (A, B and C) that intend to provide financial transactions to shops. Figure 6 shows three shops. Two of them already have MPLS VPN (VPN 1 and VPN 2) for corporate use, and the other one does not have VPN yet. All the shops need to have access to VPNs A, B and C, in which the database of credit and debit cards is located. It is important to consider that all the sites connected to each VPN already have natural connectivity among themselves, since they are either in VPN 1 or VPN 2.

It is also possible to see that users of shops that have only one site and no VPN (e.g. shops with only one commercial site) need to connect to credit and debit card companies. A scenario in which connectivity to several content providers is required has become a common situation. This case should be carefully assessed, as it has a strong impact on scalability. As we have already mentioned, the possibility of creating VPNs per site is currently out of the scope of telecommunication operators. Therefore, an alternative should be developed. The main premises of the proposals to be here presented are the following.

A. Users/Shops with VPNs

Users that already have their corporate a VPN routing/forwarding instance (VRF) implemented (for example, VPN 1 and VPN 2) and need to connect to card companies. Today, all intranet VPNs basically need to connect to huge databases, such as credit card companies. In this situation, in which companies already have VPN, i.e. they have an route distinguisher identifying their VPN, the most immediate solution is to configure exportation and importation routes in content VPNs for that route distinguisher.

The purpose of the route distinguisher is to allow the entire IPv4 space to be used in different contexts (for VPNs, in our example). On a given router, a single route distinguisher can define a VPN, in which the entire IPv4 address space may be used independently [4,5].

All security mechanisms explained in this article work only when configured correctly and when the network is correctly implemented.

Technically, extranets are constructed by using route-targets to determine which routes get included into which VRF. In this example, the VPN routes from VPN 1 and VPN 2 are imported into the VRF of the extranet. This is achieved through the route-target import statement within the VRF configuration. This way, the extranet receives the routes of the two VPNs. In the other direction, the extranet VPN routes are imported into VPNs 1 and 2 with the same command [4].

B. Users/Shops without VPNs

Users with only one site and no VPN routing/forwarding instance (VRF) that want to connect to VPNs of card companies.

Users with such characteristics (without VPN) are often found in the current communication scenario, as small shops must be connected to the content provider databases, which in turn are connected to the MPLS network. The problem is that such small shops do not have any connectivity to the MPLS backbone.

These small shops (users) usually have only one site, therefore, they have neither VPN nor VRF configured in the MPLS backbone of the telecommunication operator enabling their connection to content providers. This situation ends up hindering the implementation of a solution for access to the content providers' VPN.

In the past, small shops used to rely upon the public switched telephone network to access large databases through dial-up access (several shops still do). However, the average response time required by this kind of access was about 15 seconds to perform debt and credit card transactions. Such response time will be significantly reduced by using MPLS VPNs with dedicated access, falling from 15 seconds to approximately 2 to 4 seconds per transaction.

C. MPLS provider configures RD per VPN, rather than per site

MPLS VPN routing/forwarding instance (VRF) service operators provide MPLS VPN service, and the route identifier (RD) is configured per VPN, not per access/site. That means that all the accesses of the same VPN have one RD for all VRFs.

This environment in which one RD is configured per VPN is the solution that offers better scalability to the telecommunication service provider, although it is not often preferred by users.

Technically, from the perspective of MPLS VPN users, it would be more interesting to have one RD per site, as it could

repeat the private IP address schemes in each site, which means that different IP address configurations are not necessary in VPN access; also, this could favor a great economy of IP addresses, since each site would be with a different RD.

Nonetheless, the solution of one RD per VPN has shown to be the most effective in terms of optimization of performance of PE routers.

If the Service Provider assigns a customer interface to the wrong VPN or commits some other configuration errors which mean that it assigns RT attributes to some illegal VPN sites,unauthorized parties might join a VPN.In the case of attack,invaders may get right RT attributes by some ways to make his site join a VPN and gain access to important resources[6].

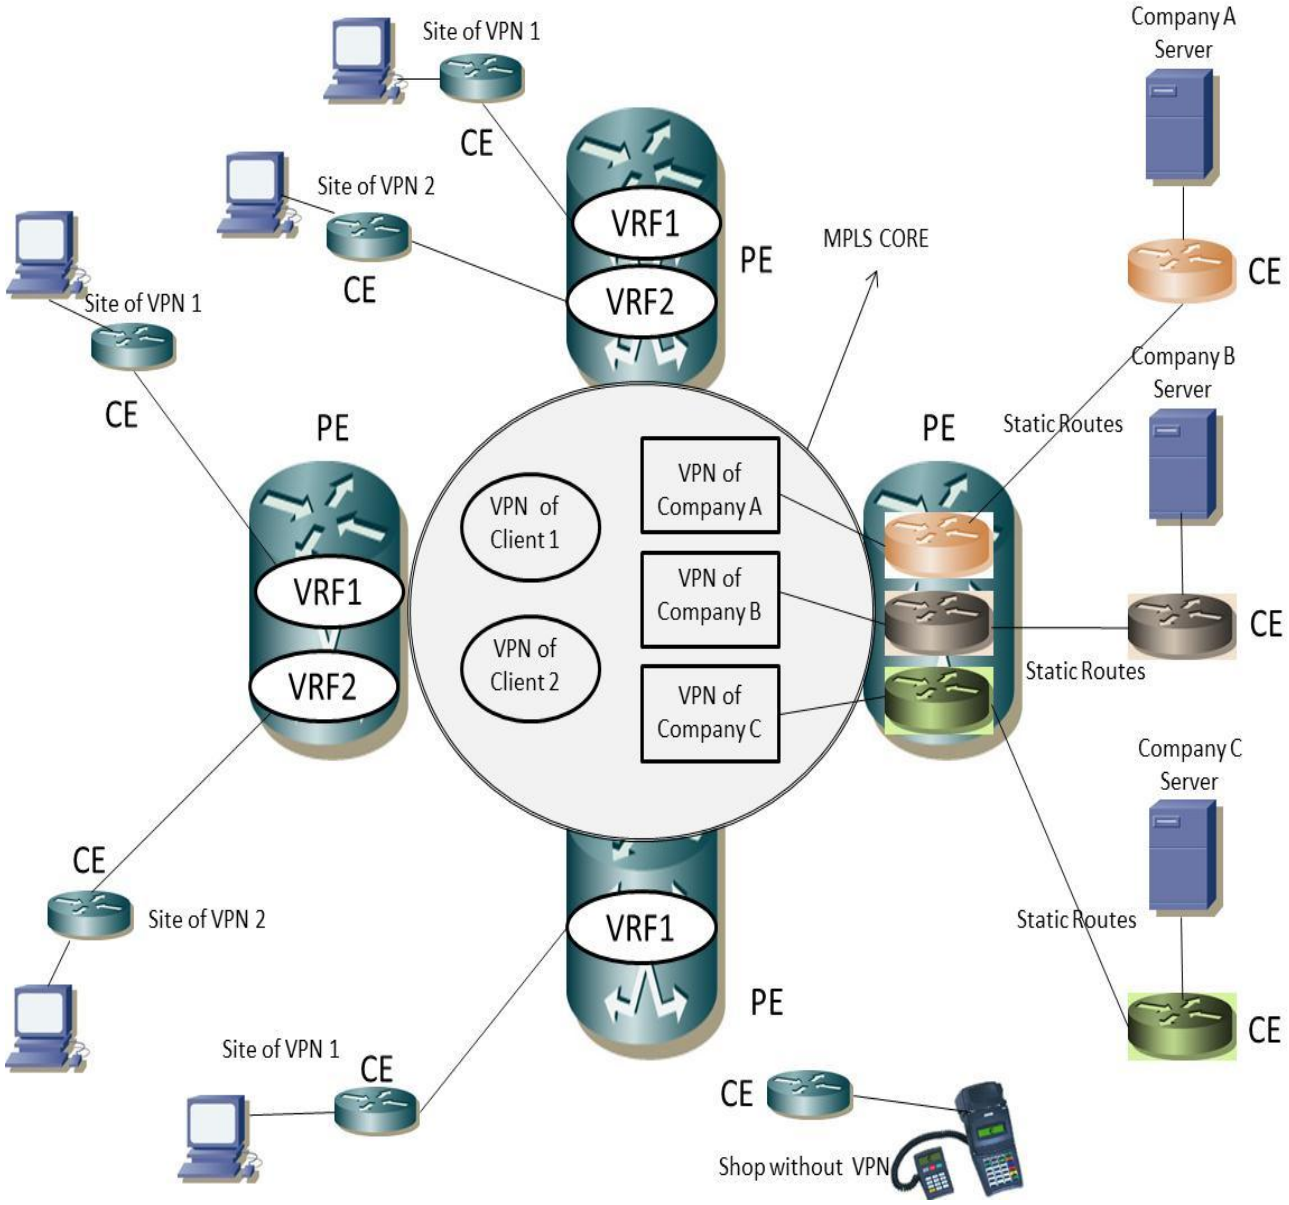

Fig. 6. Topology for VPN integration.

1) PROPOSAL 1

The first proposal suggests the creation of a Permanent Virtual Circuit (PVC) from the shop that does not have VPN to the card companies' VPN. Such solution requires a PVC for each card company. The need of a PVC for each content

server is one of the problems that lead to Proposal 2. Besides the PVC configuration to establish connectivity, it is necessary to configure an Access Control List (ACL) to allow the access to a particular shop to have connectivity only with the card company server, thus making impossible the connectivity to other shops. It is recommended the static routing protocol between (CE routers) customers' edge routers and VPN routing/forwarding instance (VRF) through a static route between the IP address of CE and the IP of the card company server.

All the shops must be connected to VRF of card companies interested in being connected, as depicted in Figure 7. As there will be several shops accessing the same VRF of a particular card company and it is not possible to have CE with the same IP addresses in the same VRF, the solution proposed to tackle this problem is the configuration of NAT (Network Address Translation) in the customers' edge routers (CE routers) of each shop.

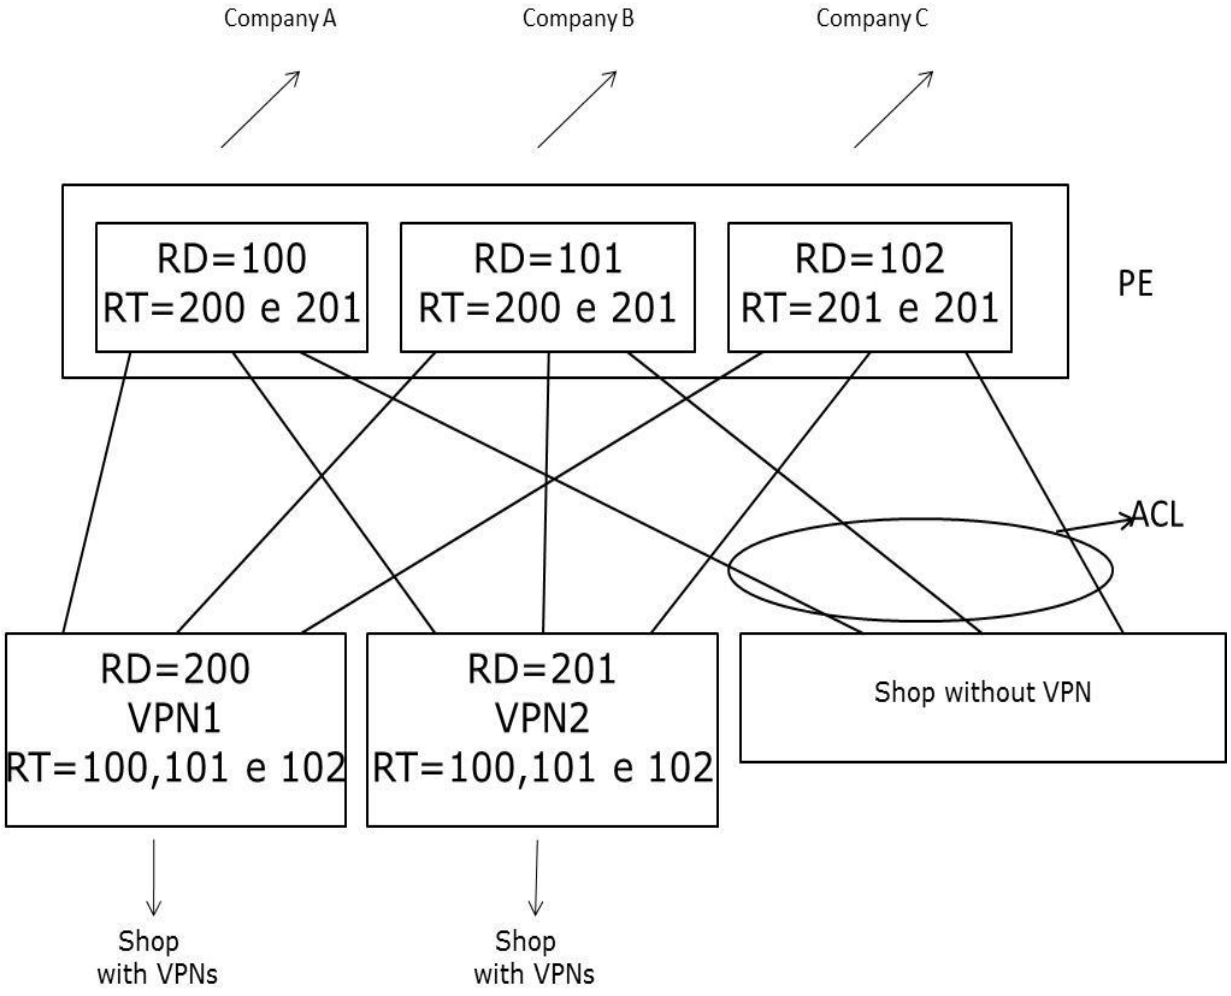

Fig. 7. Improving scalability through ACL.

In Figure 7, in the case of shops that have VRF (VPN routing/forwarding instance) for their corporate intranets, only the import and export of routes of the existing VPNs will be configured in VPNs of the card companies, as mentioned in previous sections. For example, consider the VPN of Company A. In order to allow the access of shops that have their VPNs with route identifier (RDs) 200 and 201, respectively, the route target (RT) parameter equal to 200 and 201should be configured in the VRF of Company A, thus enabling the import of routes from the shops to the VPN of the company. The same configuration procedure applies to VPNs of Company B and Company C.

If the extranet VRF would accidentally export also routes from VPNs, connectivity between VPNs would result through the extranet. Therefore, it is important to operationally control the route-targets: misconfigurations may result in a break of separation of VPNs. This applies to all route-target configurations . In practice, many service providers use automated provisioning tools, which make such misconfigurations unlikely. So the more realistic threat is coming from deliberate misconfigurations of an operator. On

the other hand, operational tools control running configurations and compare them against a correct configuration, such that even malicious changes would normally be detected [4].

Figure 8 shows the points in which Access Control List (ACL), Network Address Translation (NAT) and the implementation shown in Figure 7 must be configured.

It is important to observe that each shop without VPN will need the VPNs of Company A, B and C of several permanent virtual circuit (PVC) in order to establish connection (Figure 8).

Hence, one permanent virtual circuit of each company will have to be configured for each shop, i.e. if there were a demand by 2,000 shops in a particular area for connection to 30 companies, 60,000 permanent virtual circuit configurations would be required. This would practically cause the technical unfeasibility of such a solution (Proposal 2 presents a solution to this problem).

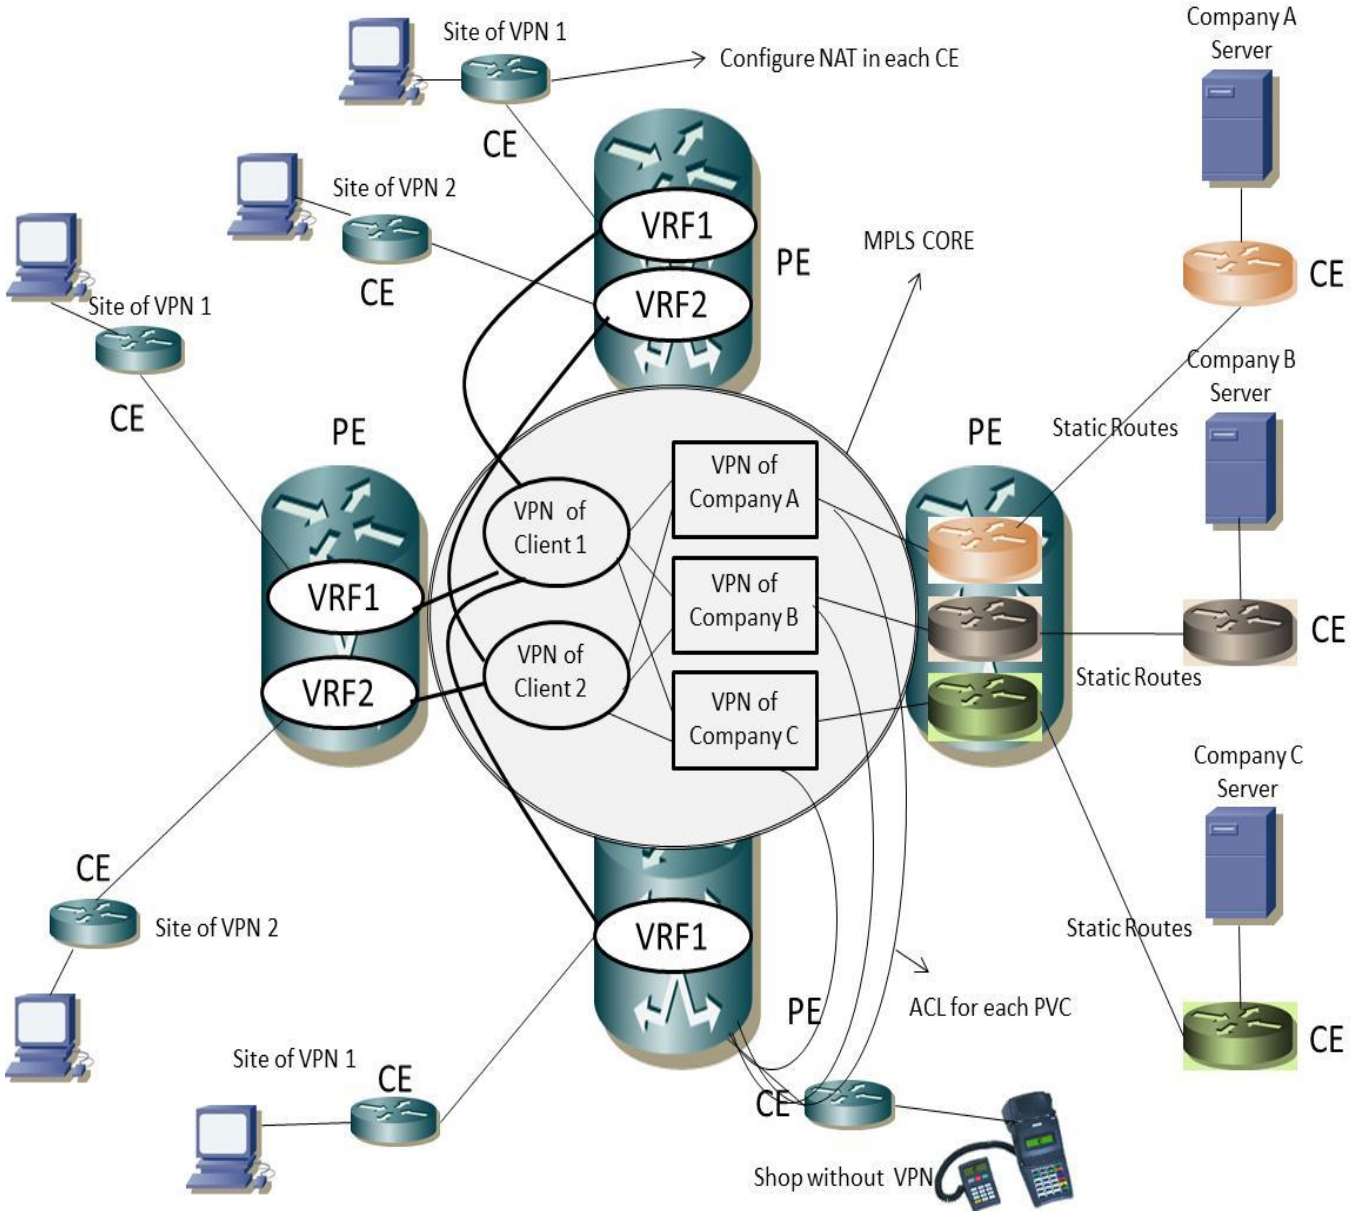

Fig. 8. Implementing scalability through ACL.

The solution illustrated in Figure 8 has some limitations:

The use of NAT functionality minimizes the possibility of IP address conflict, but cannot avoid such conflict in 100 per cent of the cases, especially when the number of accesses to VRFs is high.

The need for habilitating a PVC and an ACL between VPNs A, B and C of the card companies and the shop interested in connecting weakens the scalability offered by this solution.

2) PROPOSAL 2

Due to the limitations of Proposal 1, we suggest the use of a new model, which is compounded of an external firewall as described below. A security policy should be implemented in CE, limiting the access to the content servers of VPNs A, B and C.

Now, with the new proposal, each corporate VPN or a shop with access without VPN needs only one connection to the firewall, which is connected to all the content VPNs. Such procedure increases scalability, as it considerably reduces the amount of configuration needed. In the preceding proposal, *n* connections of permanent virtual circuit between the shop customers' edge routers and *n* content servers/card companies were necessary. In Proposal 2, only one connection is necessary. The new proposal suggests the utilization of the functionality of virtual context or virtual systems.

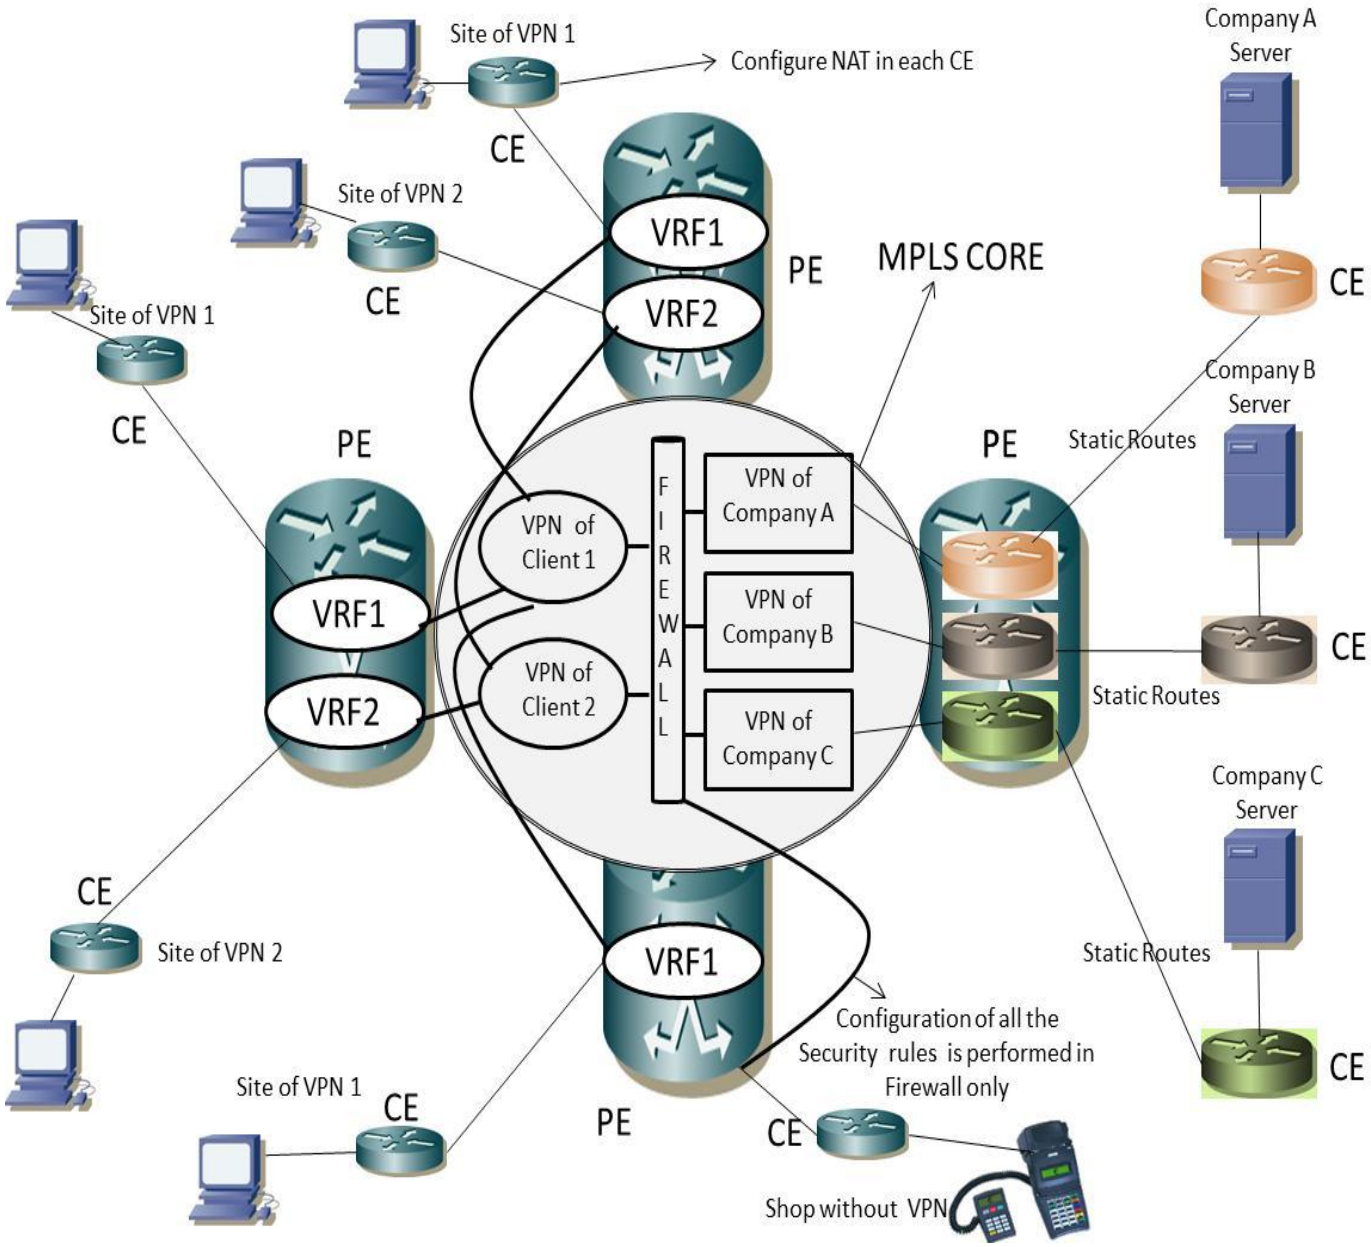

Fig. 9. Increasing scalability with the use of Firewall.

The investment to implement Proposal 2 is different from that required by Proposal 1. While the latter does not require any additional investment in the network, the former requires an additional investment in a firewall. However, in terms of scalability and security, the benefits provided by Proposal 2 are greater than those provided by Proposal 1, and such

solution is recommended to telecommunication providers that want to offer extranet VPN service.

V. Conclusion

The issues related to MPLS VPN scalability have been addressed in this article, particularly those concerning the PE router. Due to scalability, the telecommunication companies have configured only one RD identifier per VPN, rather than per site to construct intranet, as there is no problem in connecting one site to another. However, when there is the need for extranet construction, the alternative of creating an RD per site significantly decreases MPLS VPN scalability. Due to this problem, two alternatives have been proposed in order to enable the extranet construction keeping MPLS VPN scalability high. Proposal 1 is an alternative to improve scalability without a high investment in the network, but it has presented some problems, such as high likelihood of conflict of IP addresses, amounts of PVCs and NAT configuration in each CE. Considering the deficiencies found in Proposal 1, Proposal 2 has been presented. Despite requiring new investments in equipments in the network core, Proposal 2 is the solution that best meets the need for construction of high-scalability extranet MPLS VPNs. In general, it is possible to conclude that operators should implement firewall architectures for extranet construction, if they are to keep the high scalability of MPLS networks.

References

[1] Cisco Systems, Brett. Chapman “Deploying Large Scale VPN with MPLS”, Presentation, 2001.

[2] P. Tonsu, G. Wieser, “MPLS-Based VPNs” Prentice Hall series 2001.

[3] Monique Morrow, Azhar Sayeed. “MPLS and Next-Generation Networks: Foundations for NGN and Enterprise Virtualization”,2006.

[4] Michael H. Behringer, Monique J. Morrow. "MPLS VPN Security", Cisco Press 2005.

[5] RFC 4364 – BGP/MPLS IP Virtual Private Networks, http://www.ietf.org/rfc/rfc4364.txt

[6] Yi Ji;Yaping Deng, “A scheme to enhance the security of BGP/MPLS VPN”. IEEE, 2006

[7] Zhangwei He, Yong Jiang, “IMPROVE RELIABILITY OF SCALABLE VPN ROUTING VIA RELAYING” IEEE, Proceedings of IC-NIDC2010, pp 1061-1066, 2010

Author Biographies

Adao Boava received his B.S. degree in Electrical Engineering in 1991 from Federal University of Santa Catarina (UFSC), Florianopolis-SC, Brazil. He received his M.S. degree from the State University of Campinas (Unicamp), and M.B.A. from Foundation Getulio Vargas (FGV), São Paulo, Brazil. In 2011, he received his PhD in Telecommunication Engineering and Telematics from Unicamp, São Paulo, Brazil. Currently, he is a professor at Federal University of Fronteira Sul (UFFS), Santa Catarina, Brazil. He has worked in Brasil Telecom and OI for 16 years with MPLS product development. In addition, he has worked as a consultant in different data communication projects for Monsanto, Santander, Itau, Visa, Redecard and others.

Yuzo Iano received his PhD. in electrical engineering in 1986. Currently he is an Associate Professor in Electrical Engineering at Unicamp (State University of Campinas, Brazil). He also works at the Visual Communication Laboratory in the same University and is responsible for digital signal processing (sound and image) projects. His research interests include video and audio coding, digital video and audio compression and digital signal transmission. He is a member of IEEE.