NJIT

Advanced Networking Laboratory

Virtual Private Networking: Enhancing Security for E-Business

TR-ANL-2002-04-01

Jingdi Zeng and Nirwan Ansari

April 2002

Electrical and Computer Engineering
New Jersey Institute of Technology
University Heights
Newark, NJ 07029
U.S.A

Virtual Private Networking: Enhancing Security for E-business[1]

Abstract

With the Internet penetrating every corner of the world, e-business has found an excellent carrier which is able to fulfill the information transfer among diverse areas. Virtual private networking, a booming value-added service over the Internet, offers a low-cost, high-security, and easy-to-manage solution for e-business applications and services. The article introduces the implementation of Virtual Private Networks (VPNs), by portraying a generic deployment process following with certain details. The comparison and discussion in the article shall benefit service providers seeking for new avenues, and enterprises longing for supreme e-business services as well.

Key words: E-business, virtual private networks, deployment framework, implementation, security.

Authors:

Jingdi Zeng
jingdi.zeng@njit.edu

Prof. Nirwan Ansari
ECE department
New Jersey Institute of Technology
323 M. L. K. Blvd.
University Height
Newark, New Jersey 07102
USA
(Tel) 973-596-3670
nirwan.ansari@njit.edu

[1] This work is resulted from the effort on incorporating IP VPN services into a 10G Ethernet product for OpenCon Communication Systems.

1. INTRODUCTION

With the Internet reaching more and more homes and offices, the electronic commerce has been spreading out to every corner of the world. On one hand, e-business providers desire diverse technologies, from database management to Intranet security, to support everyday activities. On the other hand, e-business is increasingly relying on the Internet to connect remote employees, branch offices, customers, partners, and suppliers. Progressively, a variety of e-business service properties, such as cost, reliability, security, ease of management, and so forth are becoming more desirous than ever.

The effort of providing e-business services can be traced back to Public Switch Telephone Network (PSTN) and leased lines [1]. The naïve but effective leased line solution defined the basic tone of privacy, Quality of Service (QoS), reliability, etc. Its cost gave rise to a cheaper alternative targeted at data applications, frame relay. Frame relay's success, however, came with a price of variable QoS performance and complexity of integrating frame relay services into customer devices. ATM, with QoS and high speed provisioning, is designed to displace leased lines, and is supposed to provide a full range of multimedia services. Unfortunately, it is simply too complex to be used for enterprise network applications. Finally, Internet Protocol (IP)-based virtual private networking sheds the light, with all the advantages of cost, flexibility (leased line and frame relay networks require capacity specifications in advance), security, and the ability to exploit existing frame relay/ATM investments.

Virtual private networking is a rather effective solution for today's e-business applications, delivering information among multiple parties over a shared infrastructure (for example, Internet) with the same private network-like policies of security, reliability, manageability, and QoS. In addition to the firewall screening and proactive scanning for intrusion detection and protection, virtual private networking makes the network access worldwide available while protecting the information that flows across networks. As a result, Virtual Private Network (VPN) has gained intensive attention from industry vendors, standard bodies, and research communities. Not surprisingly, Cahners [2] In-Stat Group predicted that the total market for VPN gear and services would explode to $32 billion by the end of 2003.

Regarding to business applications, VPNs can be broadly classified into *remote access* VPN and *LAN-to-LAN* VPN. As shown in Figure 1, a dial-up or broadband (DSL or cable) access is a classical remote access VPN case, where a remote user or customer connects through its local Internet Service Provider (ISP) Point of Presence (POP) to the headquarters network. The other type of remote access VPN, which is relatively new, allows an individual Ethernet user connects to the headquarters LAN as well. A LAN-to-LAN VPN has two evolved scenes: an Intranet VPN brings together geographically separated branch offices of an enterprise over a public network infrastructure; an Extranet VPN enables business partners and external vendors to access specific portions of the headquarters network. There are several other classification categories; for instance, *Customer Premise Edge* (*CPE)-based* or *Network-based* VPN, *peer* or *overlay* VPN [3], etc. Further details shall be retrieved from references.

The layout of the article follows a top-to-down pattern, giving an overall picture of the VPN implementation and then following up with clarifications. The article aims to illustrate a macro-scope deployment framework from an ISP point of view, and also help customers seeking for e-business services more effectively define service criteria.

2. A PORTRAYAL OF VPN

As shown in Figure 2, VPNs are typically constructed by four fundamental building blocks: tunneling/encapsulation, authentication, encryption, and network management. What a successful e-business application desires, such as the dedicated service, secured connection, certain level of QoS, and flexibility, are fulfilled by these four functionalities.

Tunneling/encapsulation protocols provide a multiprotocol transport service and enable the use of the private IP address space within a VPN. They encapsulate VPN packets with extra headers and logically separate them from others, when VPN traffic is being conveyed over a public network (i.e., Internet or service provider networks). Packets with different headers go through different virtual path or route, just like a dedicated line. Its essential difference from a real dedicated line is that all these “dedicated” paths are sharing a

common link, or say, network pipe. Therefore, the network resource is used more efficiently, and enterprise LANs could reach the far end of the Internet in a seamless manner.

Authentication protocols verify and restrict the network access only to validated users or devices. Whoever or whatever connected to a VPN has to go through an authentication process, according to different scenarios. This is crucial because e-business services are highly and likely transporting/processing sensitive data and information, which are not supposed to be accessed by unauthorized parties.

Encryption protects data from examination or manipulation as it is transported across the Internet. With authentication, although an ineligible user or party cannot hook to a VPN network, plain packets being conveyed over the Internet are open to be attacked. Therefore, encryption protocols take place to encrypt packets with different algorithms that prevent information over a VPN from being snooped and eavesdropped.

Last but not the least, a network management infrastructure is adopted for billing, service provisioning, resource management, service level agreement (SLA) enforcement, and other related issues. The following sections will itemize major techniques of these four functionalities.

3. TUNNELING

A tunnel is a specific pathway, which crosses the Internet or service provider networks, and encapsulates traffic with new packet headers to assure the delivery to specific destinations. The destination receives encapsulated packets, strips encapsulation headers, and processes the de-encapsulated packets as if they were received on a local interface. Tunnels do not provide true confidentiality (encryption does that), but can carry encrypted traffic.

3.1. Tunneling techniques

Popular tunneling techniques include Layer 2 Tunneling Protocol (L2TP), Generic Routing Encapsulation (GRE), IP Security (IPSec), and Multi-Protocol Label Switching (MPLS), whose positions in a packet are depicted in Fig. 3.

L2TP [4] merges the best features of two previous tunneling protocols: Point-to-Point Tunneling Protocol (PPTP) from Microsoft and Layer 2 Forwarding (L2F) from Cisco Systems. The two peers of an L2TP tunnel are L2TP Access Concentrator (LAC), a device that physically terminates a call, and L2TP Network Server (LNS), a device that terminates and possibly authenticates PPP streams. Depending on the initiation point, there are two types of L2TP tunnels: client-initiated and LAC-initiated. The former one needs the host to run L2TP client software; the latter one puts L2TP functionality in LAC, which aggregates client traffic first, encapsulates it, and finally forwards it to LNS at the other end of the tunnel. Tunneling a PPP session with L2TP consists of two steps: (1) establish the control connection for a tunnel, and (2) establish a session as triggered by an incoming or outgoing request. Control messages exchanged in step (1) are used in the establishment, maintenance and clearing of tunnels and calls. In step (2), data messages are used to encapsulate PPP frames being carried over the tunnel. Tunnel ID and Session ID indicate the identifiers of a control connection and a session within this tunnel, respectively.

GRE [5-6] provides a common mechanism for placing packets of any protocol, for example, Address Resolution Protocol (ARP), Novell IPX, AppleTalk, etc, inside any other types of protocols. A packet to be delivered is first encapsulated by a GRE header, and then encapsulated by the transportation protocol. The GRE header, therefore, is used to differentiate tunnels.

IPSec [7] defines IP-centric tunneling (and encryption) by adding headers to IP packets (and encrypted payload). Of its two modes, transport and tunnel, IPSec has a Security Parameters Index (SPI) to differentiate tunnels in its tunnel mode. There are an "outer" IP header that specifies the IPSec processing destination, and an "inner" IP header that specifies the ultimate destination for the packet. The security

protocol header (AH/ESP) appears between these two headers, and carries SPI that enables the receiving system to select a Security Association (SA) under which a received packet will be processed. SA, a set of security parameters for authentication and encryption by an IPSec tunnel, is managed by another involved standard IKE [8-9], formally known as Internet Security Association Key Management Protocol (ISAKMP/Oakley). IKE authenticates each peer in an IPSec transaction, negotiates the security policy, and handles the key exchange. More details on IKE shall be pursed from references.

Based on Cisco Tag Switching, IETF standard ***MPLS*** [10] works on a label-based paradigm, tagging packets as they enter the provider network to expedite forwarding through a connectionless IP core. Literally defined as a short, fixed length identifier that is used to identify a Forwarding Equivalence Class (FEC), a MPLS label is used to distinguish the VPN membership and separate traffic between customers. With a label inserted between the data link layer header and the network layer header as shown in Fig. 3, MPLS offers equivalent security as a trusted Frame Relay or ATM environment.

3.2. Comparison and discussion

According to several performance metrics [11], previously listed tunneling techniques are compared as shown in Table I.

There are cases where each VPN peer supports multiple customers, for instance, through a POP. Therefore, the ***multiplexing*** capability is desired by VPN tunnels. Among all tunneling techniques, L2TP (via session ID and tunnel ID), MPLS (via MPLS label), and IPSec (via SPI) all have its multiplexing mechanism. Also the GRE version defined in Request for Comments (RFC) 2701 and RFC 2702 uses its 'key' field as a multiplexing index.

To be able to have the configuration information automatically exchanged, a tunnel shall have the ***signaling*** capability. All listed protocols support certain signaling protocols: L2TP itself exchanges control messages before assigning session ID; IPSec uses IKE to negotiate parameters for the tunnel setup; GRE shares a

similar signaling mechanism as that of the mobile-IP tunneling; MPLS features Label Distribution Protocol (LDP) to distribute labels.

Besides relying on the security characteristics of the underlying IP backbone, none of these tunneling mechanisms, other than IPSec, have intrinsic ***security*** mechanisms. Although L2TP has its security concern to hide certain Attribute Value Pairs (AVPs) during the tunnel negotiation, it does not protect plain packets being conveyed over the tunnel. The 'key' field in GRE header might be used for authentication purpose; however, it has no capability to protect packet payloads. Likewise, MPLS has no real data security mechanism, although it distinguishes traffic among customers. Fortunately, all these tunnels can be further secured by the IPSec encryption, which will be discussed later.

Owing to diverse protocols utilized by different enterprise networks, a VPN tunneling technique needs to provide the ***multi-protocol*** support. L2TP is designed to transport Point-to-Point Protocol (PPP) frames, and thus can be used to carry multi-protocol traffic as inherited from PPP. GRE has certain field in the header for the identification of the protocol being tunneled, thereby is multi-protocol capable. Similarly, MPLS is applicable to any network layer protocol. IPSec, as an exception, only works for IP traffic.

To match equivalent characteristics of physical leased lines or dedicated connections, ***frame sequencing*** may be required as one of VPN service attributes. Both L2TP and GRE have a sequencing field to record the packet sequence. The 'sequence number' field of IPSec, which is designed for a receiver to perform the anti-replay check, can be possibly extended to guarantee the in-order packet delivery too. MPLS header itself has no sequence field and therefore has no frame-sequencing guarantee.

For the reliability concern, VPN peers must monitor and ***maintain*** VPN tunnels to ensure that the connectivity has not been lost, and take appropriate actions (such as back up or tear down) if there has been a failure. One straightforward approach is for the tunneling protocol itself to check in-band for the loss of connectivity, and to provide an explicit indication of failure. L2TP detects non-operational tunnels by exchanging "keep-alive" messages between peers. IPSec relies on IKE to send out 'hello' messages

periodically. MPLS has the LDP protocol to detect bad label bindings. GRE, however, has to count on routing protocols to maintain tunnels.

Although the end-to-end **QoS** guarantee depends on traffic engineering issues, it is still of concern if a tunneling technique has the potential to support or cooperate with further QoS operations. Essentially, a tunneling technique shall not shield any QoS information in IP headers from intermediated routers that function on those QoS criteria. Fortunately, all tunneling techniques have one way or the other to fulfill it. L2TP works at layer two and will not affect IP headers. GRE and IPSec are both capable of having the Type of Service (ToS) information copied from original IP packets to encapsulation headers. MPLS, even better, has a preset three-bit experimental field that differentiates classes of service or per hop behavior for different traffic classes.

As a closing, there is no single best-for-all solution for tunneling/encapsulation because network engineers have to compromise to different scenarios. Generally speaking, industry vendors prefer the combination of L2TP and IPSec for the remote access VPN because L2TP is the only protocol that supports the dial-up scenario. For LAN-to-LAN VPN, IPSec alone is a good choice, with its tunnel mode and encryption added on top. Note that the IPSec suite only works for IP traffic, and thus there is another GRE/IPSec combination to compensate. In other words, GRE tunnels non-IP traffic and IPSec handles encryption. Lately, for the scalability concern, IETF is pushing the implementation of MPLS.

4. AUTHENTICATION

Authentication is the way a remote or mobile user is identified prior to being allowed the access to networks and network services. Industry vendors, with different preferences, adopt diverse authentication techniques, for example, manual key, token cards, digital certificates, challenge responses, biometrics, smart cards, RSA SecurID, Kerberos, Light Directory Access Protocol (LDAP), etc. VPN authentication involves the device and user authentication (the packet authentication is done by IPSec for a more stringent security). Among all this variety of techniques, one typical pick-up will be explained with more details for each of

them, respectively. Note that other alternatives of authentication, such as group authentication, are not included here.

Of several ***user authentication*** techniques, a challenge responses technique such as Dial-In User Service (RADIUS) is supported by a majority of service providers. Proposed by Livingston Enterprises, RADIUS [12] provides an industry-standard, client/server-based solution that lets authorized remote users to access corporate networks. A RADIUS authentication scenario has four main components: a remote user, an ISP Remote Access Server (RAS), a proxy RADIUS server, and a corporate RADIUS server.

As Figure 4 indicates, the whole authentication process [13] includes several steps:

- A user dials into its POP that often time is one of RASs of ISP, and a PPP negotiation begins.
- The RAS passes authentication information, such as username and password obtained during PPP negotiations, to the ISP's Proxy RADIUS server.
- The proxy RADIUS server parses the user name, e.g., username@companyname. It performs a translation in its database to determine the IP address of the user's enterprise RADIUS server from the "companyname". After establishing the proper remote link to the enterprise network, the Proxy RADIUS server forwards the username and password to the enterprise RADIUS server for further authentication.
- If the enterprise RADIUS server is able to authenticate the user, it issues an "accept" response to the ISP's proxy RADIUS server. The proxy RADIUS server, in turn, forwards an "accept" response to the ISP's RAS, along with the user profile obtained from the enterprise RADIUS server. If the enterprise RADIUS server is unable to authenticate this user, it issues a "reject" response to the ISP's RAS, along with a text string indicating the reason.
- With a proxy RAS between, the enterprise RADIUS server completes the PPP negotiation with the user. If the proxy RAS receives an "accept" response, it allows access to the enterprise network. If the proxy RAS receives a "reject" response, it terminates the connection and passes on the reason of termination to the user's terminal.

Depending on the implementation scenario, enterprise RADIUS servers can also be outsourced to VPN service providers. Thus remote users are authenticated by service provider VPN devices; this is so called "internal" RADIUS authentication. In addition, to convey the information between two peers, two protocol options are available: the Challenge-Handshake Authentication Protocol (CHAP) that encrypts user names and passwords, and the ubiquitous Password Authentication Protocol (PAP) that exchanges passwords in the clear.

The ***device authentication*** takes place whenever a new VPN device is added or an existing one is powered up. One type of device authentication techniques is the pre-shared key [14], including unique, group, and wildcard keys. These pre-shared keys are usually distributed through a secure out-of-band channel. When using for the device authentication for remote access VPN, unique and group pre-shared keys are tied to a specific IP address and a group name identity, respectively. Wildcard pre-shared keys, however, are not associated with any unique information to determine a peer's identity. Every device in the network uses the same wildcard key. Apparently, the former two techniques do not scale well because each device has to store all others' keys; the latter one is no longer safe even one device in the network is compromised.

Digital certificate, another technique which scales better, allows any device to authenticate any other device but do not have the security drawback of wildcard keys. It is tied to the unique, signed information on the device that is validated by a trusted third-party known as Certificate Authority (CA). In case that a hacker compromises or steals a device with a digital certificate, the network administrator is able to revoke the digital certificate and notify all other devices by broadcasting a new Certificate Revocation List (CRL) that contains a CA-signed list of revoked certificates. When a device receives a request for the tunnel establishment and uses a digital certificate for the identity proof, the device checks the peer certificate against the CRL.

By combining IKE and digital certificate techniques, a typical device authentication process [15] utilized to IPSec tunnels, can be demonstrated by the following. Essentially, all participating IPSec peers recognize one CA as an authenticating authority, each IPSec peer has its own digital certificate issued and validated

by CA, and then each peer's certificate is used to encapsulate that peer's public key. There are four steps for a device to sign on with a CA:

- A VPN Client (either a piece of software or hardware) generates a public/private key pair for the CA to sign. The Client first signs its outbound data with its private key. Then, the CA uses this Client's public key to validate that these data were originated by the VPN Client.
- The VPN Client requests the CA's public key to validate data coming from the CA.
- The VPN Client sends an enrollment request to the CA that ties the VPN Client's personal certificate to its public key, and then signs the personal certificate.
- The VPN Client accepts the signed personal certificate and validates this certificate by decrypting the signed personal certificate with its private key.

Note that the distribution of public keys is handled by the IKE protocol [9]. In addition, the success of CAs depends on the continuing deployment of Public-Key Infrastructure (PKI) [16-17] which includes directory services and key management. Both of them are beyond the scope of this article.

Although the authentication technique surely continues to evolve with the demand of new techniques and threats, it is well known that firewall products alone are way too weak to assure the e-business safety. Two major issues of the authentication technique, scalability and affectivity, are under intensive development. For instance, as compared to pre-shared keys, digital certificate enjoys the advantage of scalability and flexibility. For the effective concern, a combination of multiple authentication techniques, for example, token card and password, has been suggested to ensure an effective and strong authentication. As a price to pay, the additional administrative burden becomes significant when the size of VPN or the requirements for a strong device authentication increases.

5. ENCRYPTION

Even without making use of cryptographic security measures, Layer 3 VPNs, such as VPNs utilizing GRE, IPSec (tunnel mode), or MPLS, are intended to provide a level of security equivalent to what layer two backbones (e.g., ATM and Frame Relay) can obtain [18]. That is, in the absence of misconfiguration or

deliberate interconnection of different VPNs, it is not possible for systems in one VPN to access those in another VPN. However, transmitting encapsulated packets in a clear and plain text, VPNs solely utilizing tunneling techniques do not ensure the privacy. Likewise, authentication only verifies the identity of a user or device. Data integrity is still an issue if an e-business transaction wants to deliver confidential data over the Internet. Therefore, IPSec, a framework of open standards for ensuring secure private communications over IP networks, comes into play along with other cryptographic protocols for the network management such as *Secure Shell Protocol* (SSH) and *Secure Sockets Layer* (SSL). For the data integrity and confidentiality, the IPSec suite provides an ***authentication header*** **(AH)** and an ***encapsulating security payload*** **(ESP)** protocol, which can be used separately or collaboratively.

The ***AH*** protocol [19] provides the connectionless data integrity and data origin authentication for IP packets. Connectionless data integrity means the original IP packet was not modified in transit from the source to the destination, and data origin authentication verifies the source of the data. Inserted into the IP packet between the IP header and the rest of the packet, AH contains a cryptographic checksum of the packet contents, including part of the IP header itself. It uses cryptographic algorithms, such as *Hashed-based Message Authentication Code* (HMAC) coupled with the *Message Digest 5* (MD5) hash function and HMAC coupled with the SHA-1 hash function, to calculate the checksum. A hash algorithm is a one-way mathematical function that takes a variable-length message and produces a unique fixed-length value. By taking a received message, calculating the same cryptographic checksum, and comparing it with the value received, the receiver can verify that the message has not been altered in transit. Since AH does not keep the packet content confidential, it is not widely used alone for IPSec implementations across the Internet.

ESP [20] provides the confidentiality for IP traffic, as well as authentication and anti-replay capabilities. Its confidentiality is achieved through encryption that is a process of taking a message, referred to as clear text, and passing it through a mathematical algorithm to produce what is known as cipher-text. Decryption is a reversed process. Encryption algorithms, such as *Data Encryption Standard* (DES) and *Triple DES* (3DES), typically rely on a value, called a key, in order to encrypt and decrypt the data. ESP encrypts the higher-

level protocol information (the TCP header, for instance) and the actual data itself. Different from AH, the authentication functionality of ESP does not protect the IP header.

The major concern for encryption techniques is the processing speed, facing many high-speed applications in the Internet. In addition to software-based deployments, ASIC is utilized for faster solutions. For instance, a single remote user could use the software-based encryption to lower the cost, while a VPN gateway or remote access server probably needs a hardware boost.

6. NETWORK MANAGEMENT

After picking up appropriate tunneling, authentication and encryption techniques, a network infrastructure is desired to manage VPN services. Many standard network infrastructures, for example, Policy-Based Network (PBN), Telecommunications Management Network (TMN), etc., are tailored for the VPN management. Omitting details on provisioning, billing, Service Level Agreement (SLA), fault management, resilience, etc., this section focuses on the management infrastructure of different solutions.

As indicated in Fig. 5, ***PBN*** was originally [21] designed for security management purposes, in particular for the access control. Driven by new technologies including VPN and voice over IP [22], a PBN system is tailored to monitor and manage a network based on rules that define how and when to handle network applications. Most vendors' policy management products consist of three major components [23]: a directory, a policy server known as the Policy Decision Point (PDP), and a VPN device referred to as Policy Enforcement Point (PEP).

Directories have multiple functions: store global settings, coordinate and synchronize multiple policy servers, and provide information about users, file servers, and other resources where the policy server wants to apply policy. All interfaces on VPN devices are assigned a series of roles, which are defined in the policy server. For instance, there may be a role called "if the traffic goes from gateway A to gateway B, it uses GRE tunneling and DES encryption". If a role is changed, the modified information is automatically pushed to involved VPN devices. This is why VPN devices are called the policy enforcement points. The policy

server assigns roles to interfaces by using Command Line Interface (CLI) commands, Simple Network Management Protocol (SNMP), or Common Open Policy Service (COPS). As policy enforcement points, VPN devices ensure the given policy is carried out on clients via specific hardware and software functions, such as packet filtering, bandwidth reservation, traffic prioritization, and port configuration.

One of the protocols supporting the communication between policy servers and VPN device interfaces, SNMP [24-26] defines a structure to monitor and push configuration or policy information flowing among network entities. SNMP management messages are represented by instances of all object types defined either in Internet-standard Management Information Base (MIB) or in another Internet-standard Structure of Management Information (SMI). While the standard body is making efforts to standardize it in finer details, industry vendors have utilized SNMP by defining proprietary MIBs such as IPSec MIB, L2TP MIB, and VPN MIB.

Comparing to traditional network management infrastructures, PBNs implement the policy by centralizing the storage of defined roles instead of by centralizing control functions into a single software application [27]. Different from configuring individual devices, it focuses on setting policy for the network in aggregate and controlling device behaviors through these policies. However, an accurate and clear policy definition can be a problem for large-scaled PBNs with different VPN devices. Likewise, breaking the service functionality into device-specific functions that are outlined in related MIB can be very time-consuming and error-prone. Ongoing work, both in standards forums and research communities, focuses on the element management problem, i.e., on the specification of policies for managing multiple devices, and end-to-end QoS [28] across the Internet.

If PBN aims to handle network elements, ***TMN*** [29] illustrated in Fig. 6 intends to support a wide variety of management areas that cover the planning, installation, operations, administration, maintenance and provisioning of telecommunication networks and services. A general Logical Layer Architecture (AAL) of TMN is defined to organize management functions into a grouping called “logical layers” as well as to

describe the relationship between layers. A TMN reference model includes four layers: business layer, service layer, network layer, and element layer.

Supporting an abstraction of functions provided by the network management layer, the roles of the *element management layer* include the control and coordination of a subset of network elements on an individual basis and an collective basis as well, and maintenance of statistical, log and other data about elements within its control field. The *network management layer* is responsible for the management of a network supported by the element management layer. Complete visibility of the whole network and, as an objective, a technology independent view will be provided to the service management layer. *Service management layer* is responsible for all negotiations and resulting contractual agreements between a customer and services offered to this customer. Its principal roles consist of interactions between services, interactions with the service provider, the maintenance of statistical data (e.g. QOS), and the customer contact and interfacing with other administrations/Recognized Operating Agencies (ROAs). Finally, the *business management layer* has responsibility for the total enterprise. The main function of business management layer is to optimize the investment and usage of new resources, while that of service and network management layers is to maximize the utilization of existing resources.

In a nutshell, by defining the functionality of each layer and interfaces between components in the same layer and in the successive layers, the logical layer architecture of a TMN system is clearly defined. Besides the relationship between management systems in the same layer, an ongoing effort is to define service layer management specifications and interface points between different TMNs [30].

To be comprehensive, a VPN management system needs to address several issues; for instance, heterogeneous hardware platform support, local configuration consistency across the network, multi-vendor support, management functionality outsourcing, end-to-end QoS guarantee, etc. Given IP DiffServ architecture, a ***hybrid*** VPN management architecture [31] that brings together the advantages of other infrastructures is depicted in Fig. 7.

As indicated in Fig. 7, lines expanding through the network elements management and network management layers represent SNMP message flows. SNMP allows monitoring of network elements and push of configuration information into all kinds of networking devices, which solves heterogeneous hardware platform problem and in some degree enforces configuration consistency from forwarding point of view. SNMP MIBs are used to represent device management information. A device driver translates user requests and pseudo policy rules into device-specific rules to configure network devices, such as VPN and DiffServ aware routers and switches.

Moving up to the service management layer, consistency among network configurations and running services is enforced through a centralized software agent located in Service Management System (SMS), which is essentially a policy server. Before defining policies for services such as VPN, IPSec, and QoS, SMS needs to check their availability, which is supported by a collection of databases managed by Network Management System (NMS). For instance, *SLA database* contains user's identification, maximum amount of traffic for a tunnel, etc, and defines the boundary of the VPN. *Resource database* holds resource information for all the devices in its field. *Connection database* keeps a list of currently active VPN connections. *Interface database* records edge routers used as VPN tunnel endpoints, and a tunnel map attached to the outbound interface of the router. Finally, *pricing and billing database* contains the details of terminated connections and their prices.

The Business Management layer handles the negotiation between customers and service providers, including the establishment of SLA. Also, the billing information coming from the service management layer is collected here and sent to customers. In a multi-ISP scenario, the business management system incorporates the functionality of setting up an SLA between service providers that enables one provider to access the management of its peer providers; this is inherited from the standard TMN infrastructure.

Setting up a VPN end-to-end trail [31] involves user request, admission control, and service activation. First, a customer requests a SLA by negotiating with Business Management System (BMS), and then BMS connects SMS. By consulting related policies, SMS finds out the combination of different network

functionalities to provide the requested end-to-end services (for instance, a LAN-to-LAN secured tunnel requires GRE and IPSec functionalities). Given these parameters, SMS contacts the designated NMS components and consult databases for resources availability, such as bandwidth and the encryption processing power. If the other end of the peer is located in another ISP, SMS should contact the related service component of that ISP. If SMS admits the customer request, it guides NMS to reserve resources and modify databases. Meanwhile, SMS triggers certain components to delegate configurations and policies all the way down to the network element layer. Finally, SMS notifies BMS to inform the customer that the requested VPN service is set up.

7. OTHER IMPLEMENTATION ISSUES

Owing to a fundamental property of the VPN technology, that is, it tends to shield or change packet information such as headers and checksums, there are several associated implementation issues. First, no matter a VPN device is stand-alone or coupled with other network devices such as routers, firewalls, and layer three switches, the VPN functionality has to coexist with other networking functionalities. The implementation order of these functionalities, accordingly, is of intensive concern because tunneled/encapsulated or encrypted VPN packets may not go through firewalls or may be falsely dropped by routers. Second, the Network Address Translation (NAT) and Network Port Address Translation (NPAT) applications may bring up the same problem without careful design. Third, for the QoS issue, either the VPN implementation provides a means of copying the QoS information into tunneling/encapsulation headers or the traffic is classified before tunneling and encryption; otherwise, the tunnel header appended to an IP packet would make the original QoS markings invisible to intermediate routers/switches. Last but not the least, some VPNs probably need Internet connectivity as well. The meaning of the Internet connectivity is two-fold: being able to reach Internet destinations, and being reachable from any Internet source. Accordingly, the VPN deployment shall have the Internet connection functionality included, or implement corresponding security mechanisms if the Internet connection is stand-alone.

Other issues, related to network management, include network resilience (provide hot swap or standby VPN devices), load balancing (distribute loads among VPN devices), web caching (relieve Internet access burden), service provisioning (manage multiple VPNs), etc.

8. FUTURE TRENDS

As a flourishing value-added service, VPN has been experiencing intensive changes. In March 2002, Cisco announced [32] its unified VPN suite to deliver new VPN technology for IP and MPLS backbones. This comprehensive delivery includes new protocols such as Any Transport over MPLS (AToM) and Layer 2 Tunneling Protocol version 3 (L2TPv3), new deployment capabilities for integrating IPSec with MPLS, and new VPN provisioning tools for both service provider and enterprises. In January 2002, Aleron [33] became the first major Internet backbone provider that fully implements MPLS technology across its entire core network; this enables IP packets to travel in a "switched" fashion directly over optical networks, providing customers with decreased network latency and router hops. MPLS, when associated with resource reservation and traffic engineering technologies, permits the construction of highly configurable IP-based VPNs as well as QoS-defined applications. As a general trend, MPLS VPN, along with DiffServ and IPSec, is considered as a promising solution for LAN-to-LAN VPNs, where MPLS breaks up the scalability bottleneck, DiffServ guarantees QoS through the core network, and IPSec secures e-business information up to the packet level.

9. CONCLUSIONS

Bearing ultimate goals of low cost, high security, easy management, real scalability, and end-to-end QoS, virtual private networking will benefit the booming e-business in years to come. This article pictures the implementation of VPN services (mainly layer three), highlighting its capability of supporting low-cost and high-security e-business services. Other practical implementation issues are briefly addressed for the purpose of completeness. For ISPs seeking for new avenues, the article provides a light but comprehensive guide of the VPN deployment framework; for enterprises looking for e-business services, the article explains considerations behind different implementation decisions, thereby helping enterprise customers define their e-business service criteria more effectively and successfully.

REFERENCES

[1] Network-based IP VPNs: the role of MPLS, Ennovate networks white paper, 1999, http://www.tradespeak.com/docdetails.asp?docid=1232.

[2] R. Younglove, Virtual private networks - how they work, *Computing & Control Engineering Journal*, Vol. 11, No. 6, pp. 260–262, 2000.

[3] W. Yurcik and D. Doss, A planning framework for implementing virtual private networks, *IT Professional*, Vol. 3, No. 3, pp. 41 –44, 2001.

[4] W. Townsley, A. Valencia, A. Rubens, G. Pall, G. Zorn, and B. Palter, Layer two tunneling protocol "L2TP", Internet Engineering Task Force (IETF) RFC 2661, August 1999, *http://www.ietf.org*.

[5] Hanks, Li, Farinacci, and Traina, Generic routing encapsulation (GRE), IETF RFC 1701, October 1994.

[6] Hanks, Li, Farinacci, and Traina, GRE over IPv4 networks, IETF RFC 1702, October 1994.

[7] S. Kent and R. Atkinson, Security architecture for the Internet protocol, IETF RFC 2401, November 1998.

[8] D. Maughan, M. Schertler, M. Schneider, and J. Turner, Internet security association and key management protocol (ISAKMP), IETF RFC 2408, November 1998.

[9] D. Harkins and D. Carrel, The internet key exchange (IKE), IETF RFC 2409, November 1998.

[10] E. Rosen, A. Viswanathan, and R. Callon, Multiprotocol label switching architecture, IETF RFC3031, January 2001.

[11] B. Gleeson, A. Lin, J. Heinanen, G. Armitage, and A. Malis, A framework for IP based virtual private networks, IETF RFC 2764, February 2000

[12] C. Rigney, S. Willens, A. Rubens, and W. Simpson, Remote Authentication Dial In User Service (RADIUS), IETF RFC 2865, June 2000.

[13] ISPs simplify remote access for enterprises, Bay Networks white paper, *http://www.firstvpn.com*.

[14] SAFE VPN: IPSec virtual private networks in depth, Cisco white paper, 2001, *http://www.cisco.com*.

[15] Case study for layer 3 authentication and encryption, Cisco white paper, August 2000.

[16] C. Adams, S. Farrell, Internet X.509 Public Key Infrastructure Certificate Management Protocols, IETF RFC 2510, March 1999.

[17] S. Chokhani, W. Ford, Internet X.509 Public Key Infrastructure Certificate Policy and Certification Practices Framework, IETF RFC 2527, March 1999

[18] E. Rosen and Y. Rekhter, BGP/MPLS VPNs, IETF RFC 2547, March 1999.

[19] S. Kent and R. Atkinson, IP authentication header, IETF RFC 2402, November 1998.

[20] S. Kent and R. Atkinson, IP Encapsulating Security Payload (ESP), IETF RFC 2406, November 1998.

[21] D. C. Blight and T. Hamada, Policy-based networking architecture for QoS interworking in IP management-scalable architecture for large-scale enterprise-public interoperation, *Proceedings of the Sixth IFIP/IEEE International Symposium on Integrated Network Management*, pp. 813--826, 1999.

[22] End-to-end quality of service for WAN and campus networks, Cisco tech. note, August 1999.

[23] S. J. Shepard, Policy-based networks: hype and hope, *IT Professional*, pp. 12--16, 2000.

[24] K. McCloghrie, An administrative infrastructure for SNMPv2, IIETF RFC 1909, February 1996.

[25] B. Wijnen, D. Harrington, and R. Presuhn, An architecture for describing SNMP management frameworks, IETF RFC2571, April 1999.

[26] J.D. Case, M. Fedor, M.L. Schoffstall, and C. Davin, Simple Network Management Protocol (SNMP), IETF RFC 1157, May 1990.

[27]C. K. Wang, Policy-based network management, *Proceedings of International Conference on Communication Technology (WCC-ICCT)*, pp. 101--105, 2000.

[28] Rajan, R.; Chiu, A.; Civanlar, S., A policy based approach for QoS-on-demand over the Internet, *Proceedings of the Eighth International Workshop on Quality of Service (IWQOS)*, pp. 167--169, 2000.

[29] ITU Rec. M3010, Principles for a telecommunications management network, Geneva, 1996.

[30] R. Larsson and L. Ferrari, Use case driven integration of TMN interface specification and application design, *Proceedings of IEEE Network Operations and Management Symposium*, pp. 434--443, 1996.

[31]T. Braun, M. Guenter, and I. Khalil, Management of quality of service enabled VPNs, *IEEE Communications Magazine*, Vol.39, No. 5, pp. 90--98, 2001.

[32] *http://newsroom.cisco.com/dlls/prod_030402.html*.

[33] *http://www.businesswire.com/cgi-bin/f_headline.cgi?bw.011402/220142533*.

Captions for illustrations

Fig. 1. VPN application scenarios.

Fig. 2. A VPN block architecture.

Fig. 3. Tunneling/Encapsulation protocols.

Fig. 4. A general authentication process.

Fig. 5. A Policy-Based Network (PBN) model.

Fig. 6. A general TMN model.

Fig. 7. A hybrid VPN management model.

Table I. The comparison of tunneling/encapsulation techniques

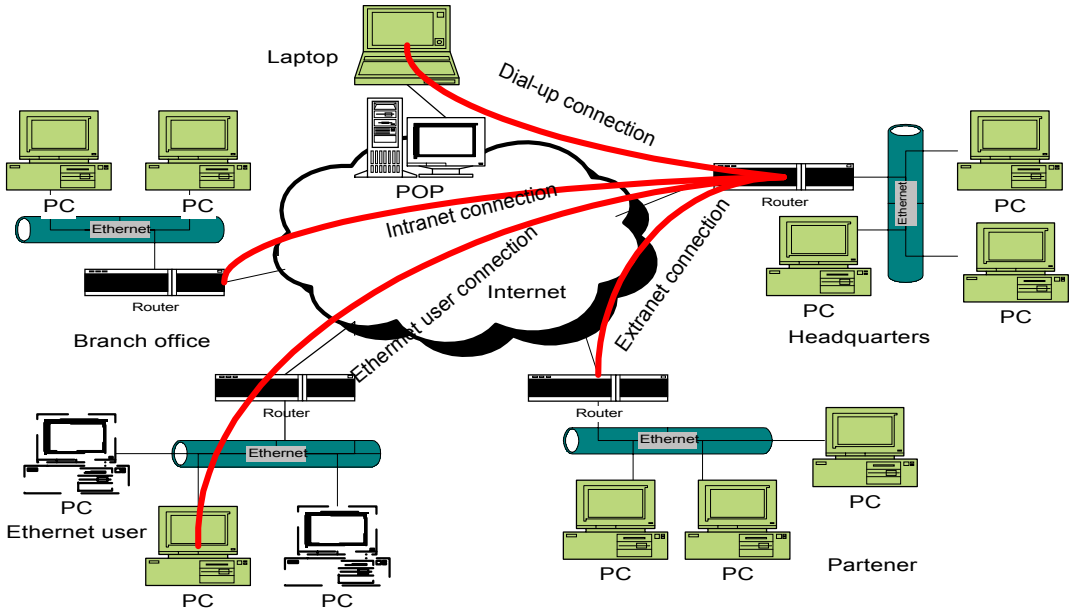

Fig. 1. VPN application scenarios.

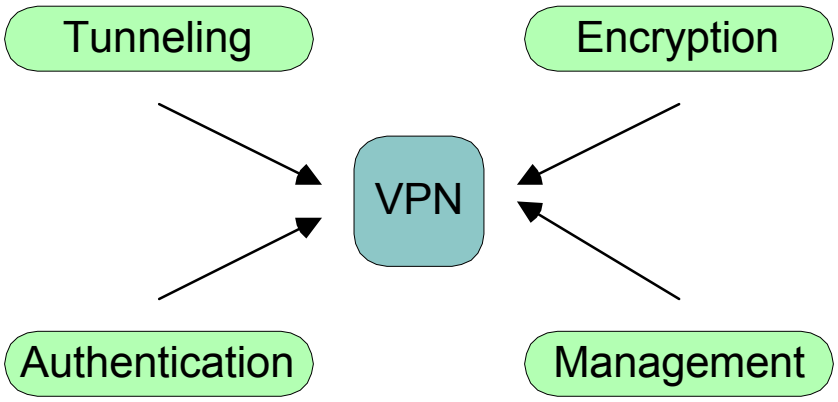

Fig. 2. A VPN block architecture.

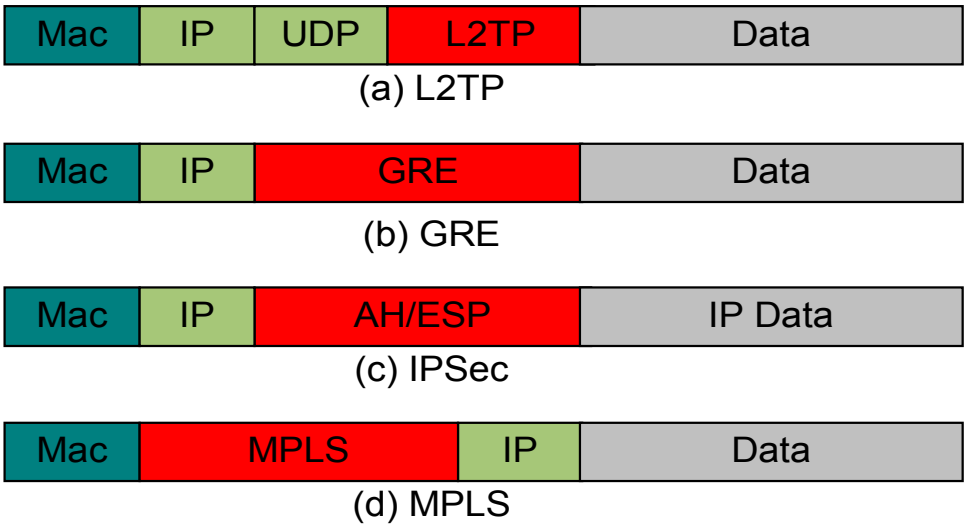

Fig. 3. Tunneling/Encapsulation protocols.

Table I. The comparison of tunneling/encapsulation techniques

	Multiplexing	Signaling	Data Security	Multiprotocol	Frame Sequencing	Tunnel Maintenance	QoS capability
L2TP	yes	Yes	no	yes	Yes	yes	yes
GRE	yes	yes	no	yes	Yes	no	yes
IPSec	yes	yes	yes	no	Yes	yes	yes
MPLS	yes	yes	no	yes	No	yes	yes

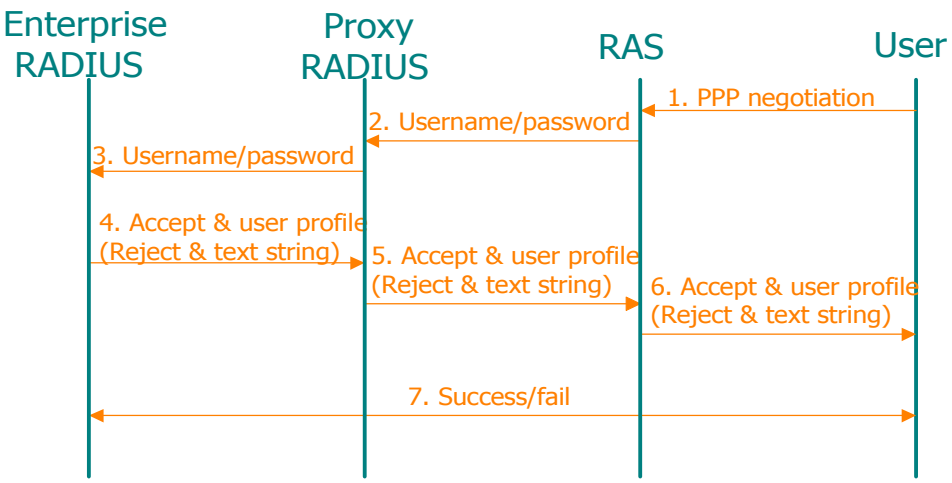

Fig. 4. A general authentication process.

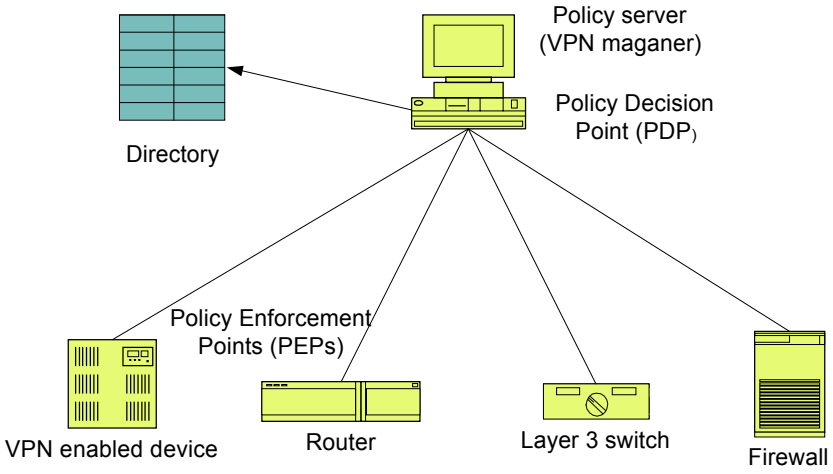

Fig. 5. A Policy-Based Network (PBN) model.

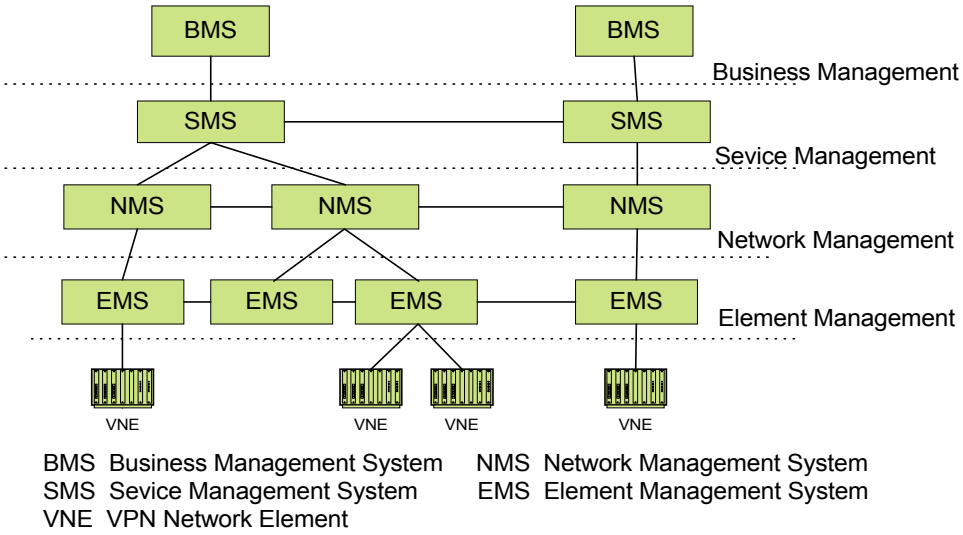

Fig. 6. A general TMN model.

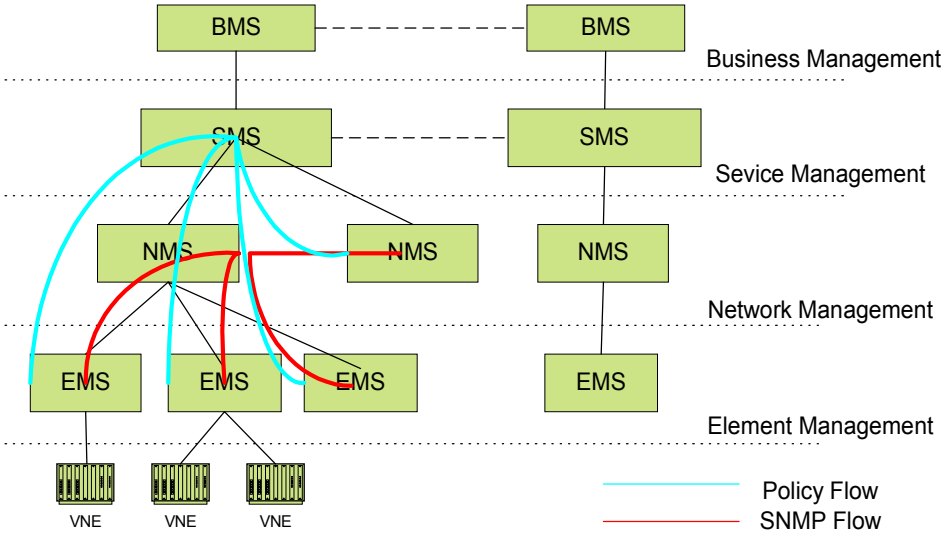

Fig. 7. A hybrid VPN management model.