Development of a Commercially Viable Best-effort Layer-2 Virtual Private Network System

Akira Saito, Masakazu Miyamoto[†]*, and Kouichi Suto*

Abstract

This article describes the development of edge routers and settings servers for adapting the existing FLET'S network to provide a Layer-2 virtual private network (L2-VPN) service to small and medium-sized businesses and SOHO (small office, home office) users. This service can be easily provided by installing compact edge routers in users' premises and settings servers within the service provider network. In contrast to earlier Wide-Area Ethernet services, which were designed mainly for use by large corporations, this system can provide smaller users with an L2-VPN service at a low cost.

1. Virtual private networks

Virtual private network (VPN) technology is designed to provide a virtual environment that appears to the user as a dedicated line. Until now, the main design consisted of overlay model networks in which user nodes were connected by asynchronous transfer mode (ATM) or frame relay (FR) technologies. Networks of this type have a number of advantages in that they support redundancy, a high degree of reliability, and detailed quality of service (QoS) control. However, they also have a number of problems. Routers are needed between each pair of nodes, resulting in higher device costs and the need for advanced installation skills, and it becomes extremely difficult to manage and maintain paths when the number of nodes increases and the network becomes larger. This is because nodes must be linked by one-to-one connections. Moreover, the processing load is placed mainly on the core of the communications carrier network. To address such problems, new VPN technologies were developed using IP or virtual LAN (VLAN) communications. These are classified as Layer-3 (L3) or Layer-2 (L2) VPNs according to the packet transmission methods used within the network.

2. L3-VPN

In an L3-VPN, the VPN is controlled on the basis of IP packet header information. A common example of an L3-VPN service is an IP-VPN service using virtual router technology, or a combination of border gateway protocol (BGP) and multiprotocol label switching (MPLS) technology. A network of this type can reduce the need to specify settings for individual node routers and reduce the processing load placed on such routers because the path control from each VPN node is terminated by the edge routers (ERs) located within the network. The processing load placed on internal network routers is also reduced because no path control information is released to the network from individual user nodes, and dynamic path control protocols can be used to provide flexible redundancy for transmission paths. Such a network also presents problems, however, in that ERs must assume the majority of the processing load and only the IP protocol may be used.

3. L2-VPN

In an L2-VPN, the VPN is controlled on the basis of Ethernet frame packet header information. A common example of an L2-VPN service is a wide-area Ethernet service using VLAN-VPN technology. In a network of this type, the entire communications carrier network is composed of Layer-2 switches, with each switch and its associated VPN being identified

[†] NTT Access Network Service Systems Laboratories
Chiba-shi, 261-0023 Japan
E-mail: miyamoto@ansl.ntt.co.jp

by a VLAN tag. It can use protocols other than the IP protocol (e.g., the Microsoft Windows NetBIOS or Macintosh AppleTalk protocol) because the network as a whole appears as a single switching hub to each user node. Furthermore, there is no need for any advanced skills in specifying operating settings because there is no need to install a router on the user side. However, it provides very limited QoS control capabilities and one problem is that the transfer tables for each switch become extremely large due to non-hierarchical media access control (MAC) addresses being used to control transmission paths. Care must be taken in specifying switch path control settings because there is no limit on the number of transfer hops that can occur in the MAC layer, and it is difficult to build large-scale networks even when VLAN tags are stored in stacks because there are restrictions on the number of VPNs that can be identified.

4. Best-effort L2-VPN system

This system uses an Ethernet over IP-VPN (Ether/IP-VPN) protocol, which makes it possible to provide Ethernet connectivity over a best-effort IP-VPN network. An Ether/IP-VPN network offers all of the advantages of both IP-VPN and VLAN-VPN. More specifically, while the communications carrier network appears as a single switching hub from the viewpoint of a user node, IP path control is actually performed within the network. This design has several advantages that result from eliminating the need to install routers at individual user nodes: it is easier to specify operating settings, protocols other than the IP protocol can be used, and it is possible to use existing IP networks because the network itself need not be constructed as an Ethernet network. There are other advantages too. The MAC address, which must be assigned to each user device to perform L2-VPN control, can be automatically learned by the L2-VPN edge routers installed at each user node, and the learned addresses are handled separately by each ER. This makes ERs easy to operate and maintain and reduces the load placed on internal routers. This in turn makes this type of system superior for constructing large-scale networks.

5. Packet transfer method

The way in which packets are transferred in an Ether/IP-VPN system and the format of ER transfer packets used in such a system are shown in **Fig. 1**. In this system, an ordinary IP network is used for transferring data between ERs, and Ethernet traffic is transmitted along the IP network via tunneling connections[*1]. When an Ethernet frame is delivered to an ingress ER from its associated device, the ER gives the Ethernet packet a VPN header to identify the associated VPN, together with an IP header addressed to the egress ER. It then transfers the packet to the network. Network-internal IP routers then transfer the packet to the egress ER, handling IP path control independently of any user nodes. The egress ER then removes the attached VPN and IP headers

[*1] Tunneling connection: A technique whereby path control headers are added and deleted when packets are transferred across network borders so as to make it possible to perform end-to-end packet transfers over two networks with differing transfer protocols. For instance, if two IPv6 sites are connected via an IPv4 network, then IPv6 packets are encapsulated in IPv4 packets for transfer within the IPv4 network.

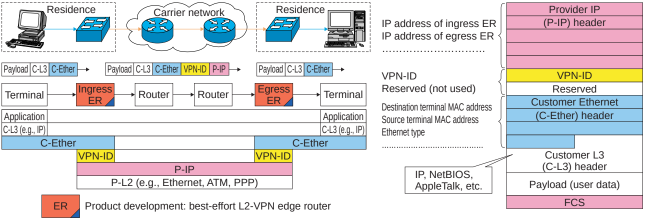

(a) Packet transfer

(b) Packet format

Fig. 1. Ether/IP-VPN packet transfer and packet format.

and transfers the Ethernet frame to its own associated device.

6. Automatic learning of ER MAC addresses

Each ER can automatically learn information about which devices exist at any given ER node. This means that traffic is never directed anywhere except to the appropriate node within the same VPN, making it possible to reduce the usage of network bandwidth and reduce the processing load placed on individual ERs. There is also the advantage that the network as a whole may be treated as a single switching hub from the viewpoint of the user.

The automatic MAC address learning method is shown in **Fig. 2**. **Figure 2(a)** shows the processes that automatically build a transfer table for incoming packets from outgoing packets transmitted from a home node, and **Fig. 2(b)** shows the processes that build a transfer table for outgoing packets from incoming packets from the network.

First, when an outgoing packet from home node terminal *A* within a local area network (LAN) arrives at remote node terminal *B* within a wide area network (WAN), an entry for the MAC address of the sender of the packet (terminal *A*) is added to the incoming packet transfer table. By storing table entries in this way, the ER can recognize that terminal *A* exists at its own home node when a packet from a remote node within the WAN addressed to terminal *A* arrives at the ingress ER. Once the VPN-ID has been checked and the IP header that allowed the packet to be transferred along the network has been removed, the packet can be transferred within the LAN.

Conversely, at the egress ER, upon the arrival of an incoming packet from remote WAN terminal *A* addressed to LAN terminal *B*, the VPN-ID from within the packet header is checked and an entry is added to the outgoing packet transfer table listing the MAC address of the sender (terminal *A*) and the IP address of the sending (ingress) ER. By storing table entries in this way, the egress ER can recognize that terminal *A* exists at the ingress ER node when a packet from a home node of the egress ER addressed to terminal *A* arrives. This information is used to attach an IP header for transmitting the packet along the network, and

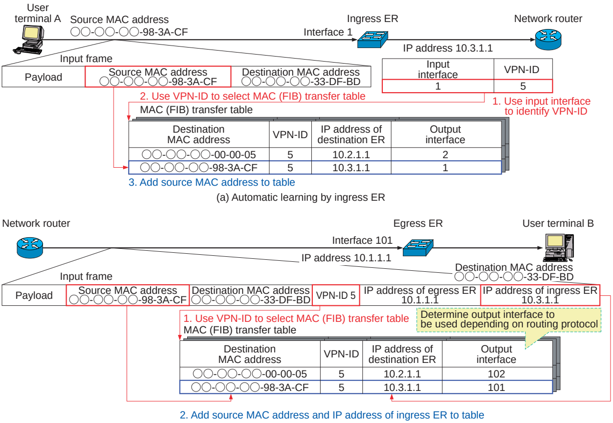

Input interface	VPN-ID
1	5

Destination MAC address	VPN-ID	IP address of destination ER	Output interface
○○-○○-○○-00-00-05	5	10.2.1.1	2
○○-○○-○○-98-3A-CF	5	10.3.1.1	1

(a) Automatic learning by ingress ER

Destination MAC address	VPN-ID	IP address of destination ER	Output interface
○○-○○-○○-00-00-05	5	10.2.1.1	102
○○-○○-○○-98-3A-CF	5	10.3.1.1	101

(b) Automatic learning by egress ER

FIB: forwarding information base

Fig. 2. Procedure for automatically learning MAC addresses.

the packet is then transferred along the WAN.

Once these transfer tables have been automatically built, data transmissions between ERs are performed via unicast transfers*2. However, if the IP address of the destination terminal is known but no information is available to identify the ER to which it belongs, then an ARP (address resolution protocol) packet used to resolve the MAC address from the IP address is delivered to the ER from the terminal. In such cases, the ER creates a number of copies of the packet equivalent to the number of remote nodes existing within the same VPN, and broadcasts the copies to all user nodes within the VPN.

To allow for the movement of terminals between ERs, existing entries within the transfer tables may be overwritten and updated using the header information contained in newly arrived packets. In addition, to allow for the fact that there are restrictions on the number of entries that can be maintained in an actual implementation, an aging time*3 parameter has been provided to allow for the deletion of entries for which there have been no communications over an extended period.

7. Types of services

The structure of the networks needed when an L2-VPN system is used to provide services over an existing FLET'S network is shown in **Fig. 3**. **Figure 3(a)** shows the type of network that will be used when service is to be provided by an NTT service company in conjunction with a VPN service provider, and **Fig. 3(b)** shows the type of network that will be used when service is to be provided by an NTT service company on its own.

In cases where service is to be provided by a VPN service provider in conjunction with an NTT service company, the VPN service provider will supply settings servers used to manage ERs, and all the other equipment and facilities supplied by any ordinary ISP to provide connectivity to the FLET'S network, and will earn revenue by providing VPN connections between users. The NTT service company will obtain its revenues from fees for using the FLET'S network.

In cases where service is to be provided by an NTT service company on its own, the company will provide inter-LAN connectivity services in an Ethernet layer while ensuring security between different VPNs using FLET'S group access. The NTT service com-

*2 Unicast transfer: A type of communication targeted at a single specified node. Opposite of multicasting.

*3 Aging time: Rather than being designed simply to store the learned MAC addresses permanently, the system is designed to return to the pre-learning state (i.e., to forget) after a specified amount of time has passed since learning. Aging time is used to refer to how long MAC addresses are stored in this way.

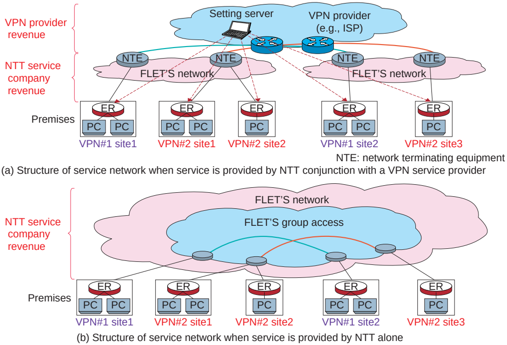

(a) Structure of service network when service is provided by NTT conjunction with a VPN service provider

(b) Structure of service network when service is provided by NTT alone

Fig. 3. Structure of service network when using an L2-VPN system over a FLET'S network.

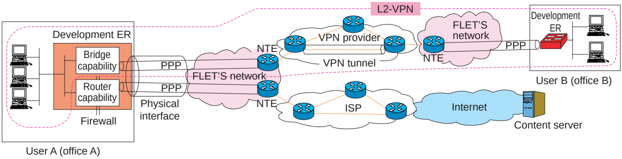

Fig. 4. Example of a system providing simultaneous access to the Internet and to an L2-VPN network.

pany earns revenues from fees for using the FLET'S network, FLET'S group access, and L2-VPN service.

8. Use for simultaneous Internet access

If two separate PPPoE (point-to-point protocol over Ethernet) sessions are used, then the ERs may be used to provide simultaneous access to a VPN and the Internet. An example of such a system is shown in **Fig. 4**. In a system such as this, it is possible to provide Internet access for multiple users over a single VPN by obtaining an Internet account for only a single node within the VPN and allowing users at other nodes to access the Internet through the VPN. Note that to allow for Internet connectivity the system has been designed with a router, and to provide security against illegal access to VPNs from the Internet it has been given firewall capabilities.

9. Edge routers and settings servers

An edge router is shown in **Fig. 5**. It is more compact than an ordinary broadband router, making it suitable for use in the home. Settings servers may be created by installing the applications software we have developed on a commercially available PC. ER settings may be specified either directly or from a server. In the direct method, ER settings are specified over a telnet connection from a LAN terminal. When specifying settings from a server, ER settings are first recorded in the settings server and the settings are then downloaded from the server when the ER is first started up. It is also possible to force the settings to be downloaded whenever it is necessary to do so. An example of a system using server-specified settings is shown in **Fig. 6**. Note that when the server settings

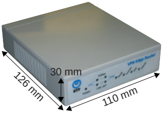

Fig. 5. Appearance of L2-VPN edge router.

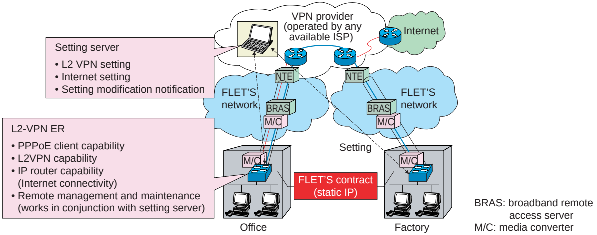
VPN provider (operated by any available ISP)
Internet
Setting server
• L2 VPN setting
• Internet setting
• Setting modification notification
NTE
NTE
FLET'S network
FLET'S network
BRAS
BRAS
M/C
M/C
L2-VPN ER
• PPPoE client capability
• L2VPN capability
• IP router capability (Internet connectivity)
• Remote management and maintenance (works in conjunction with setting server)
M/C
M/C
Setting
FLET'S contract (static IP)
Office
Factory
BRAS: broadband remote access server
M/C: media converter

Fig. 6. Handling of edge routers by settings server.

method is used it also becomes possible to specify user settings or perform isolation testing centrally from a single remote location.

10. Future issues

In this article, we described how we developed i) settings servers that can be used to manage edge router settings remotely from a single location and ii) compact edge routers designed for use in the home, which can be used to provide Ethernet access over an IP-VPN network. They make it possible to provide a low-cost best-effort wide-area Ethernet service to small and medium-sized businesses and SOHO users. While this completes the work necessary to demonstrate that this is a commercially viable system, several issues remain. These include improving the ease of operability of the system and improving the user interface.

Akira Saito
Chief, Research and Development Center, NTT East Corporation.

In 1997, he joined NTT Access Network Service Systems Laboratories, Makuhari, Japan. From 1991 to 2003, he was engaged in R&D of high-speed downloading using multiple TCP connections and system engineering of carrier access networks. He moved to his present position in November 2003.

Masakazu Miyamoto
Access Service and Network Architecture Project, NTT Access Network Service Systems Laboratories.

He received the B.E. degree in electrical engineering and the M.S. degree in computer science from the University of Yamanashi, Kofu, Yamanashi in 1997 and 1999, respectively. In 1999, he joined the NTT Access Network Service Systems Laboratories, Makuhari, Japan. From 1999 to 2003, he was engaged in development of ATM access network systems and researching dynamic traffic engineering on IP/MPLS networks. Since 2004, he has been improving maintenance and operation after development of L2-VPN system.

Kouichi Suto
Senior Research Engineer, NTT Access Network Service Systems Laboratories.

He received the B.E. degree in electrical engineering from Iwate University, Morioka, Iwate in 1977. Since joining NTT Laboratories in 1977, he has been active in developmental research on optical fiber trunk transmission systems and optical fiber subscriber transmission systems. From 1999 to 2002, he was engaged in development of an access control server system for IP-VPN service. Since 2002, he has been engaged in developmental research on access network services using ADSL and/or FTTH systems. He received the Electronics Letters Premium from IEE in 1978. He is a member of IEEE.