

High Performance VPN Solutions Over Satellite Networks



January 2004 © Encore Networks, Inc.

Specifications are subject to change without notice.

Encore Networks, Inc. • 45472 Holiday Drive • Dulles • Virginia • 20166

Tel: 703-318-7750 • Fax: 703-787-4625 • Email: info@encorenetworks.com • Web:

Enhanced Packet Handling Both Accelerates And Encrypts High-Delay Satellite Circuits

Characteristics of Satellite Networks?

Satellite Networks have inherit round-trip delay characteristics of about 1/2 second for typical geo-stationary satellite circuits. Such delay, which pauses some performance issues for voice applications, might lead to problems such as "conversation collisions" where both parties can be talking at the same time.




Fig. 1: Satellite circuit introduces a round trip propagation delay of 1/2 second.

Data protocols also face problems with long delays. Two problems in particular handicap satellite links for data applications:

- 1) **Throughput limitation.** TCP senders cannot exceed the rate at which the receiver can acknowledge receipt of packets--satellite latency effectively caps standard TCP throughput per session (RFC-793).
- 2) **Security.** Transmissions from satellites are available to anyone with a suitable receiver.

Separate solutions for each problem have been available for years. Manipulation of TCP header fields with a Performance Enhancing Proxy (PEP) server can fool the end points into increasing the throughput on a satellite connection. Encryption of the data defeats interceptions off the air. However, standard IPsec encryption of the IP packet hides the fields in the TCP header and so prevents an acceleration proxy from changing them.

In the past, an end user had to make a choice and compromise on what is more important.. Today, via improved handling of packets, integrated solutions are available in a single device.

TCP Throughput Over satellite:

Throughput is not a problem with one-way transmissions, such as media broadcasting, for which the only limit is the capacity of the transponder or relay equipment in the satellite. The sender can transmit at the maximum circuit capacity 100% of the time because:

- the circuit is dedicated and fully available,
- there is no need to acknowledge data or correct errors, and
- the receiver applies no "back pressure" to limit throughput.

Transmission Control Protocol (TCP), was designed for entirely different conditions:

- low network latency,
- maximize bandwidth utilization,
- error detection and correction and
- congestion avoidance and recovery via flow control mechanism.

TCP is the transport protocol that operates end-to-end to ensure accurate delivery of data. The TCP header follows the IP header (Fig. 2). Higher layer protocols ride on TCP (that is, their headers follow the TCP header in the IP packet). In general, any protocol that desires guaranteed delivery will rely on TCP, including HTTP, telnet, and FTP.

A Complete Data Packet:

MAC header	IP header	TCP header	Data in Higher Level Protocols
------------	-----------	------------	--------------------------------

TCP header portion:

00	01	02	03	04	05	06	07	08	09	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
Source Port																Destination Port															
Sequence Number																															
Acknowledgment Number																															
Data Offset				reserved				ECN				Control Bits				Window															
Checksum																Urgent Pointer															
Options and padding																															
Data																															

Fig. 2: Packet headers for TCP/IP to deliver a reliable service to higher layer protocols.

Note in Fig. 2 the TCP header fields called **Sequence Number**, **Acknowledgment Number**, and **Window**. These fields control data flow between two TCP devices, A and B. Also, they are used to implement sliding window scheme that ensures an effective flow control between the end points.

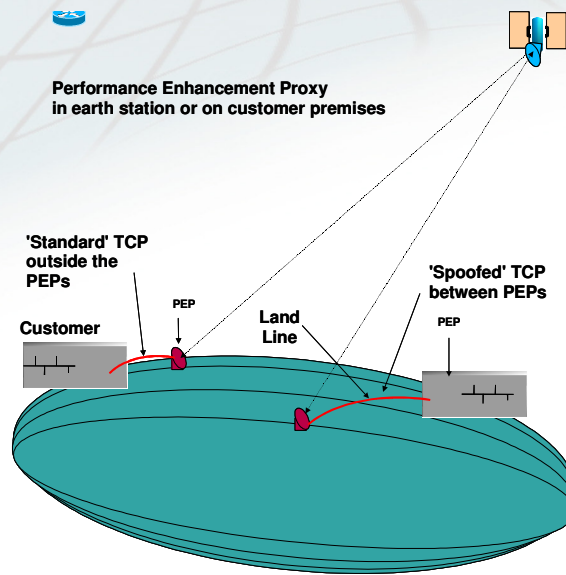
TCP senders may transmit all the bytes in the window value. This allows packets to be burst onto the satellite channel in groups. When the window is exhausted, the sender must stop until it receives another **Acknowledgment Number** (with a Window size). In this way, each side of a TCP connection controls the rate at which it receives data; it applies backpressure by reducing its window.

A complicating issue with TCP is its congestion-avoidance algorithm, which misinterprets the long delay and high bit error rate from long trips through the atmosphere. TCP assumes that all loss and delay are caused by network congestion. As a result, it cuts back the transmission rate by reducing the Window size, and then slowly ramps the send rate back up. This algorithm further reduces the efficiency of the channel.

Performance Enhancing Proxy (PEP):

To improve throughput, the TCP session is broken into segments (Fig. 3). The total circuit avoids the throughput cap by modifying the behavior of TCP between the Performance Enhancement Proxies (PEPs) at each side of the satellite link. PEP is not yet an official standard, but is used to describe the set of proprietary mechanisms satellite vendors and systems integrators created to mimic or spoof TCP over a satellite connection.

Fig. 3: TCP proxies manipulate Sequence and Acknowledgment fields to raise throughput.



In a TCP proxy implementation, rather than one session being established, there are three: one TCP session from each user end point to the local PEP, and another session between the proxies. A proxy on the ground deals with its adjacent end point using standard TCP, so the customer device needs no modification. Prompt ACKs from the PEP give the sender permission to send more, even before the previous window of data reaches the far-end earth station. The customer can send at the same rate available on the LAN, up to the capacity of the satellite circuit. At the far end, the PEP appears to be a normal router port.

Over the satellite hop, proxies may work with a larger window or with a protocol that doesn't use windows and acknowledgments (for example, User Data Protocol, UDP). Either way, the proxies take full advantage of the available bandwidth to achieve higher throughput.

Encryption for Accelerated TCP:

Serious concerns for confidentiality often translate into a requirement for encryption. To ensure privacy, enterprises prefer to encrypt end-to-end. Encryption can be added to PEP, but only between ground stations where the proxies are located. The preferred method of encryption on IP Virtual Private Networks (VPNs) is "IP Security" (IPsec), which comes in two formats, both of which interfere with TCP acceleration.

1) Encapsulating Security Payload (ESP) encrypts each user IP packet, including the TCP header, and places it inside a new IP packet generated by the customer's VPN router. Encryption prevents the PEP proxy from seeing or modifying the TCP ACK and Window fields, so these sessions can't be accelerated

2) IPsec Authentication Header (AH) doesn't encrypt the payload, and AH leaves the TCP header visible. However, the strong authentication process rejects a packet in which PEP modifies a header field, which also prevents acceleration by PEP.

Internet service providers commonly offer VPN based on "secure tunnels" created with ESP. ESP is the most secure approach if carried end to end among routers on customer sites. To allow PEP over satellite, however, the ESP session may be decrypted at the ground station, just before data reaches the PEP proxy. The existence of an "unencrypted six inches" on the cable between IPsec router and the PEP is unacceptable to many enterprises and government agencies as they mandate to have a complete end-to-end secure solutions.

The result has been a trade-off between high performance and high security using IPsec VPNs. Standard IPsec can't be accelerated by PEP. The additional processing time to encrypt/decrypt (and compress/decompress) further lengthens the ACK cycle, cutting throughput.

Best of Both

A mechanism has been developed to allow current PEP acceleration techniques to access the TCP header, yet still provide IPsec VPN standards based security between two endpoints across a satellite connection. Encore Networks (www.encorenetworks.com) has developed a technique that preserves the authentication and encryption integrity of the IPsec VPN standards, yet allows the TCP to be spoofed over the satellite connection. The technique is called Selective Layer Encryption (SLE).

Encore Networks has implemented various flavors of SLE in its BANDIT product line of VPN customer premises routers. The flavors differ in where in the packet the SLE control information rides across the Internet. The SLE may use:

- ✓ the TCP options field (SLE-TCP),
- ✓ the IP options field (SLE-IP),
- ✓ the TCP payload (SLE-Transparent mode),
- ✓ an encrypted packet tunneled within another IP packet (SLE-VPN).

SLE uses the TCP options field to announce and control the handling of the encrypted information. Test results have demonstrated interoperability with different satellite modem vendors that preserve the integrity of TCP fields across the satellite link.

However, not all vendors of satellite modems support TCP options (although they have been defined for many years). Some satellite systems have PEPs that ignore the additional header length occupied by the options. On these systems, the receiver never sees the TCP option fields--and therefore can't recognize encrypted data. To deliver acceleration over satellite modems that don't preserve TCP options, Encore Networks enhanced SLE with a second mode of operation (SLE-IP) that doesn't rely on TCP option fields, but on similar fields in the IP header.

A third mode of operations (SLE-Transparent) is independent of both IP or TCP options field. The TCP Data field is used instead of the IP or TCP Options fields. Everyone passes the Data field. This mode of SLE opens the door to interoperate with other satellite modems that do not support any options fields, essentially all PEP processes and all satellite modems.

The last mode of operations (SLE-VPN) is similar to SLE transparent mode with the exception of Network Address Translation (NAT). All SLE implementations except for SLE-VPN require "NATing."

For example, consider an HTTP client (Fig. 4) at a Remote Office trying to access an HTTP server at Headquarters. The client's IP address is private, 10.21.222.100, so it must be NATed (to 65.220.69.80) to reach the Internet. The server also has a private address, 10.9.231.30, but is known on the Internet by its (NATed) public IP address, 68.153.166.93. This connection normally requires two NAT operations.




Fig. 4: SLE-VPN feature of the BANDIT tunnels connections between private IP addresses across a VPN on a public network.

The disadvantage of applying NAT in this way is that every WEB (HTTP) server shown at "Headquarters" in the diagram (on the 10.9.231.0 private network) requires a public IP address. New public IP addresses are sometimes difficult to find, which is making them harder to get and more costly.

SLE-VPN resolves this issue by "tunneling" the SLE connection in a VPN session. That is, the SLE session is encrypted with IPsec as the payload in another IP packet which carries public IP addresses.

The SLE-VPN connection is encrypted, hiding its IP addresses. Therefore the HTTP client can now use private addresses, its own and that of the server, to establish the HTTP session. For example, HTTP client 10.21.222.100 can access HTTP servers 10.9.231.30 or .20 or .10 by going directly to the private address instead of using the public address.

Servers containing confidential information that do not have a public IP address are invisible from the Internet. Only participants in the VPN have the ability (not necessarily the permissions) to reach these internal servers.

All-in-one Security:

An encrypted PEP session, maintained between two Encore Networks proxies at ground stations, operates at high speed, to take full advantage of the capacity of the satellite circuit. From the proxies out to the user sites, normal IPsec packets behave according to the standards. Ground transmission may be leased lines or an ISP. To link the two encryption formats, Encore switches packets between the PEP encryption session and the IPsec session within the same enclosure. (Fig. 5)

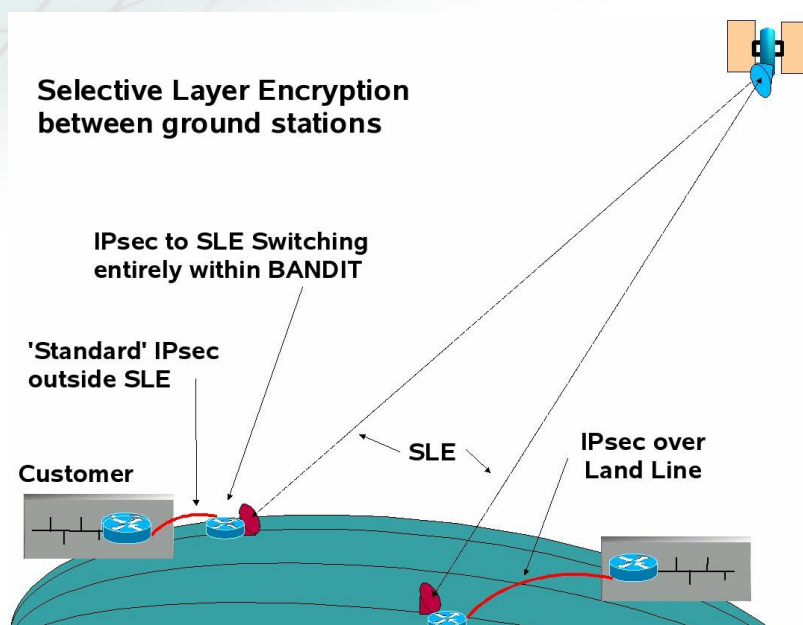


Fig. 5: Encore switches IPsec encrypted packets onto SLE link over satellite.

Encore BANDIT Enables End-To-End IP VPNs over Satellite

The high latency of satellite networks creates issues for the channel-efficient transport of TCP/IP over satellite connections. The introduction of Performance Enhancing Proxy (PEP) accelerators to spoof TCP and improve the performance of IP connections over satellite links has subsequently prevented the extension of standards based IPsec VPN connections out to satellite endpoints.

To serve this need, Encore Networks (www.encorenetworks.com) developed Selective Layer Encryption (SLE) that allows an IPsec standards-based VPN tunnel to be accelerated over a PEP satellite connection. Selective Layer Encryption creates satellite VPN solutions with IPsec that are both secure and channel-efficient.

About Encore Networks

Encore Networks, Inc (www.encorenetworks.com) is a leading developer of integrated IP+Legacy VPNs and security applications, converged signaling, and broadband voice and data solutions for both carriers and enterprises. Encore provides advanced security solutions for wireless, satellite, and wireline networks that include encryption, stateful inspection firewall, IP VPNs, support of legacy data protocols, and built-in dial backup and fail-over capabilities. Encore's innovative signaling and data solutions include signaling conversion for voice migration from circuit to packet, a broad line of IAD, VPN router, CPE, and host products to transform legacy data networks for the broadband IP infrastructure.