PEP Deployment and Bandwidth Management Issues

Charles Younghusband, Peter Slade, and Jeff Weaver

XipLink, Inc.
Suite 800, 3981 boul. St-Laurent, Montréal, QC, Canada H2W 1Y5
{cyounghusband,pslade,jweaver}@xiplink.com

Abstract. This paper will discuss current deployment scenarios for Performance Enhancement Proxies (PEP) technologies in broadband satellite access systems from the perspective of one PEP technology provider. Recent improvements such as DVB-S2 can provide substantial gains at the link layer. In order to achieve further efficiency gains, the satellite industry is now forced to look elsewhere – namely other layers in the data communications network stack. Satellite terminal manufacturers are now moving beyond basic TCP acceleration techniques to more comprehensive optimization techniques that incorporate advances in data compression and flexibility for more deployment scenarios. Some of the advances for PEP technology are in part due to CPU and memory technology advances, resulting in increasingly affordable access to computing power, allowing PEP manufacturers deliver substantial performance and bandwidth savings gains.

1 Introduction

Performance Enhancing Proxies (PEP) are TCP split-connection solutions to address the properties and limitations of wireless communications using various techniques, transparently or non-transparently to the end-hosts, and have been previously addressed by the IETF in RFC 3135 [1]. PEPs can also perform other functions such as compression and application specific optimization. This is of particular relevance to satellite communications given the high cost of bandwidth. They can also be combined with other intermediate-node network technologies such as virtual private networks (VPNs).

This paper will discuss current deployment scenarios for Performance Enhancement Proxies (PEP) technologies in broadband satellite access systems from the perspective of one PEP technology provider. Recent improvements such as DVB-S2 can provide substantial gains at the link layer, and the satellite industry is now forced to look elsewhere – namely other layers in the data communications network stack -- for efficiency gains. Satellite terminal manufacturers are now moving beyond basic TCP acceleration techniques to more comprehensive optimization techniques that incorporate advances in data compression and flexibility for more deployment scenarios. In part this is due to improved PEP technology, but some of the advances for PEP technology are due to CPU and memory technology advances, results in an increasingly affordable access to computing power, allowing PEP manufacturers deliver substantial performance and bandwidth savings gains.

K. Sithamparanathan (Ed.): Psats 2009, LNICST 15, pp. 77–84, 2009.

This allows PEP manufacturers to reduce cost and deliver substantial gains. For instance, very small (pocket-sized) solutions are now available to users to further broaden the deployment availability of advanced PEP and link optimization technology.

A PEP technology provider looking beyond a narrow technology application is also required to examine many different deployment scenarios, not just in the context of a particular satellite networking technology but the topology and impacts of encryption and the possibly use of multiple paths or meshed and hybrid networks. This paper will briefly discuss some issues with deployment scenarios for VSATS and Mobile Satellite Services and bandwidth management issues that arise from the combination of the deployment scenarios with optimization techniques.

This paper assumes basic familiarity with PEPs and other optimization techniques used for satellite and wireless data systems.

2 PEP Technology

Besides the well known TCP Acceleration to address the limitations of standard TCP/IP when used over satellite [2], some of the techniques now being applied to reduce bandwidth demands include compression, filtering of rogue Internet traffic, shaping of traffic, data suppression, optimization of particular application protocols, header compression as well as lossy, application oriented compression techniques. They generally have the net effective of maximizing the available capacity based on the goals of the service provider. These benefits vary by application but are substantial enough that virtually any data application over satellite should consider the use of PEP technology, with the one exception being live streaming video applications which typically are pre-compressed and are delivered via UDP.

PEPs can be proprietary or standards-based. The dominant standards-based solution is SCPS-TP [3], or the associated I-PEP [4] standard developed by the SatLabs Group. Modern PEPs are generally user transparent. SCPS-TP and I-PEP also are by default network transparent; they use an enhanced TCP approach that retains the IP address and port numbers during and through acceleration. The SCPS-TP approach of using an enhanced TCP (with more options) also benefits from being able to incorporate ongoing enhancements, drawn from the research community, directly to the PEP.

3 Satellite PEP Deployment Scenarios

There are many factors to be considered for the deployment of PEP technology into networks incorporating satellite-based technology as a primary communications basis. This section will review and comment on the PEP technology aspect.

3.1 Optimization of VSATs

The dominant scenario, and best understood, is of course to deploy PEP technology closely bracketing the satellite link on VSAT networks. TCP acceleration can be integrated and deployed within a terminal itself. As the TCP protocol particularly suffers on networks with variable latency and bandwidth, TCP acceleration technology has been commonly integrated for some time. Integrated TCP acceleration in the terminal unit is also becoming common on more traditional SCPC links as well. With the historical price sensitivity satellite modem manufacturers face, these modems did not make use of more advanced processors and thus often provide little benefit beyond TCP acceleration and enforcing QoS on the available bandwidth. Meanwhile, there is demonstrated benefit to compression and other techniques to optimize the bandwidth that has not been deployed significantly as of yet. However with continued increase in the CPU power and memory availability compared to the relatively modest increases in terminal bandwidths, this can be expected to change. Embedded optimization in the terminal will provide substantial benefit both to the user in terms of response time and to the operator in terms of capacity. However, the embedded PEP technique will only act on TCP/IP data that has not been encrypted so it is generally applied to Internet access and various commercial services.

PEP technology can be deployed on a stand-alone, dedicated purpose basis (as a network appliance) or incorporated with other network appliance technologies such as VPN.

Network layer encryption such as IPSec completely hides the TCP layer headers from acceleration devices, as such it cannot benefit from acceleration. Proper encryption also has the property of randomizing the data, rendering it uncompressible and inaccessible to application layer optimization. SSL based VPNs can have the SSL element accelerated, but not the encapsulated TCP/IP data so the problem persists.

Notably, IPSec would normally be preferred over SSL-based VPNs, as SSL VPNs use TCP thus doubly complicating the issues with the TCP protocol. While the actual TCP, or the SSL-TCP may be accelerated, there are actually two levels of TCP to be accelerated (TCP-in-TCP) and so the solution will still suffer from standard TCP's limitations and poor performance.

3.2 Mobile Satellite Services Optimization

Mobile Satellite Services (MSS) further adds a new set of challenges for a PEP manufacturer; including more demand on bandwidth, more susceptibility to link outages, routing flexibility for multiple communication paths (possibly both satellite and terrestrial) and the resulting adaptation required for links with dramatically different characteristics including jitter, latency, peak and minimum bandwidth and asymmetry. In addition to effectively serve the MSS market one must consider in particular physical characteristics that properly consider the deployment environment.

In the transportation sector, multiple access techniques may be made available. These are often configured independently from a PEP: a tunnel through multiple alternate network paths may be setup to isolate the remote network, yet they can lead to substantial changes in the characteristics in link. This is common of ships and aircraft which may need to employ alternate technologies or frequencies based on their location, as well as access technique. For instance, a ship may use VSAT technology at sea, switch to BGAN as they enter a country's waters and may no longer be licensed to use the VSAT technology, and change yet again to a wired or wireless as they enter a harbor.

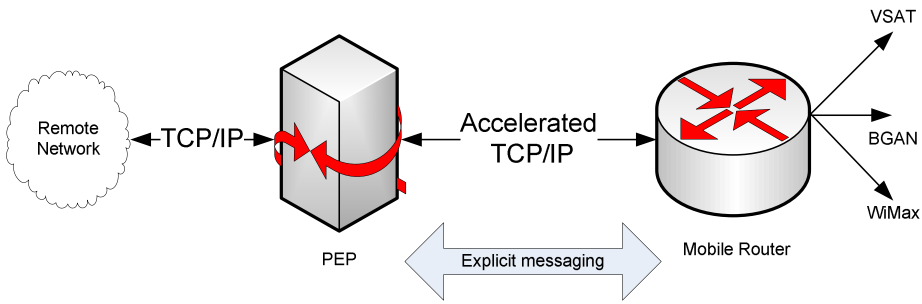

Fig. 1. Explicit messaging between elements to signal link changes

The use of a network transparent optimization protocol, where the accelerated connections carry the same IP addresses, port numbers and appropriate DSCP tags can be easier to use. The caveat for new PEP installations is to be wary of intermediate shaping devices already deployed which may attempt to shape the accelerated traffic downwards and work against the TCP acceleration.

A single-configuration "best fit" TCP acceleration scenario using a preset congestion control strategy (e.g. TCP NewReno, Vegas or Westwood or some other proprietary congestion control algorithm), while worthy of further research, currently cannot provide the effective solution for all deployment., Hence this approach lowers the delivered bandwidth and ultimately the competitiveness of a network solution incorporating satellites. However, in cases where a VPN exists between the PEP and the mobile router, a best fit solution is the only option however explicit probing or artificial intelligence techniques could be used to determine the type of link in actual use and reconfigure accordingly.

A PEP manufacturer can provide an interface to a middleware component which could be written to adapt the TCP acceleration and possibly other elements to changes in the link characteristics if they are known. Such a configuration means that an adaptive congestion control can be avoided, and a more aggressive congestion control approach can be used to maximize throughput.

There is no standard at this time for such an interface: in part this is due to the different information that may be available. The information in the message could take the form of an explicit rate, a profile referencing a specific link technology, a value representing available buffer space in bytes (possibly in a matrix referenced by different queues), the latency, bit error rate or jitter. Depending on the PEP, it could provide benefit to communicate these bandwidth properties to the other PEPs in the system. This could be done with an extended I-PEP option for instance.

This market also often demands more bandwidth that can be readily supplied by current MSS services such as Inmarsat's BGAN service. While very bandwidth scalable VSAT terminals now exist, the same can not be said of most MSS services without a dramatic increase in cost and physical deployment size. As these networks use shared contention strategies of some kind and make use of multiple hops, they can also deliver less bandwidth that expected; for instance an FTP transfer through BGAN can be expected to go between 150kbps and 300kbps *[inmarsat reference]* despite the maximum terminal capacity of 492 kbps. TCP acceleration can double the link utilization. Then other improvements further multiply the benefit: compression can improve the performance anywhere from zero for pre-compressed video, a typical 2:1 to higher compression ratios for text and raw sensor data, as well as minimize the data consumption bill for per-per-bit services. PEP technology can then be considered to help bridge modern broadband MSS services to VSAT technologies.

The deployment of such PEP technology into these MSS must be accessible and affordable. The demand for something portable has driven XipLink, Inc. to release a 'XipStick'; a network appliance smaller than a typical compact mobile phone. It provides two Ethernet interfaces, a USB which can be used for power as well as a USB-based network interface for a laptop. This makes it an ideal deployment platform for typical MSS terminals and be carried in one's pocket. The standalone nature permits government level encryption devices to be used in line, or it can also provide secure IPSec-based VPN services. Unlike PC software solutions, it can also act on all data from a network intelligently, and there is no PC software to maintain. Remarkable for this discussion is that this compactness and computing power (sufficient for intelligent compression strategies) at a reasonable price level simply was not available until recently.

Communications systems for aircraft obviously have their own challenges for the certification of equipment installed on an aircraft. The importance for PEP technology to operate on general purpose hardware is further reinforced by this market and many similar certification-heavy environments. A software embedded PEP can be used with pre-certified hardware equipment in order to shorten the time to deploy the system.

The configuration of the hub-based system tends to be different on MSS systems as well. Typical MSS only has a small number of terminals connected and transmitting at any one time. As well, they are less likely to share a receive channel and instead behave more like a dedicated link attempting to traverse a shared network, such that

Fig. 2. A Portable PEP Appliance

they have an upper boundary on their maximum bandwidth. In the case of BGAN networks, a hub PEP system needs to treat each remote BGAN terminal as an individual basis with a fixed maximum. This is different than the VSAT scenario with many terminals sharing a DVB receive path, allowing unused bandwidth can be redistributed and shared amongst the active terminals.

3.3 PEP Client Software Based Solutions

A PC-based software solution can be an option for the deployment of PEP technology for both VSAT and MSS markets. Although this obviously eases additional hardware issues, there continues to be considerable hesitancy for service providers to put themselves in a position to be maintaining PC software and the impacts of different operating systems, service packs and other applications which may interfere. Furthermore, if there is more users on a network accessing a terminal a PC based solution will not be aware of the network utilization of the other nodes and therefore must use less effective adaptive congestion control techniques: there is benefit to a funneled solution where one node will see all of the traffic on the network. However a software-based solution has the benefit of end-to-end VPN possibilities if installed in the operating system in such a way that it is applied before the VPN tunnel is created. This approach is particularly attractive for employees who find themselves only occasionally accessing the network over satellite and a corporate IT department has mandated a particular PC-based VPN client.

The satellite-based service provider has a similar issue with IT infrastructure maintenance; most providers position themselves to provide access. The use of application-specific optimization creates a grey area; when provided by a service provider it is more likely to confuse service provision with IT services to properly maintain application-specific optimization as the application specific optimization must be updated with new versions and changes to the IT infrastructure. This is of particular importance to satellite networks given the wide distribution of nodes and the difficulties of on-site servicing. This encourages the simplification of PEP technology. However, application-specific optimization can provide substantial benefits for business users when their applications send highly repeatable data or the applications are unnecessarily highly interactive.

4 Bandwidth Management

4.1 Remote Terminals and Fully Meshed Networks

PEPs embedded within a terminal can work in tandem with the bandwidth allocation mechanisms. They can be made to substantially increase the performance of networks, particularly if the satellite network is configured primarily or entirely in a demand-assigned bandwidth allocation for maximum efficiency. Coordination between the terminal and PEP can have the effect of making congestion control and flow control equivalent: the goal being to right size the data queued to send within the terminal. This section revisits the generalized explicit messaging scenario as depicted by Figure 1, but focuses specifically on a dedicated link to be optimized.

The queue size in the output buffer of the TDMA modem should not be so much as to increase latency unnecessarily, and not so little that the terminal cannot assert to the bandwidth controller that it is demanding resource of the network. This can allow the system to minimize buffering of bulk (TCP) data. On most TDMA networks, this queue length can vary dramatically on a millisecond basis, as they typically use a burst time plan between 30ms to 120ms, so such a solution must be highly reactive.

However the coordination between the terminal and the PEP is not possible when VPN is enabled between the PEP and the satellite modem. Indeed this is a major driver of sales of standalone units to networks that already have TCP acceleration embedded in the terminals. At least in many government level encryption solutions, the VPN prevents devices within the network from communicating directly the external entity outside of the network. However intelligent algorithms best-fit have been implemented that permit 100% utilization of the TDMA link return channel. However unlike some of the mobility applications described above, this is specific to one link technology. The proper implementation of DSCP (which is promoted through the encryption process) will allow simultaneous real-time data such as VOIP to flow properly into the right queues.

4.2 Hub PEP Issues

Adaptive Coding and Modulation (ACM) adds further challenges to the hub as the effective bandwidth to any terminal may change; thus requiring flexibility to shift bandwidth allocations and in order to maintain Service Level Agreements (SLAs). The combination of ACM, QoS/shaping/SLA enforcement, TCP acceleration and advanced compression creates many challenges. Namely, they all may act at different layers of the network stack: link layer, IP layer, TCP layer respectively and the compression which may exist in some form at any layer. Thus cross-layer and likely cross-device interaction between the devices/functions can be very significant in support of developing optimal hub configuration.

For instance, if the PEP is installed downstream of the shaping device, then the traffic will be shaped prior to compression and the bandwidth possibly under-utilized. If the PEP is installed prior to the shaping device, then the PEP may be forced to run in an adaptive congestion control to allow the shaper to shape, or it must mirror the shaper configuration. Feedback systems, such as the one described earlier, can be developed, or a completely integrated solution can be deployed. QoS within the PEP system can allow it mirror the network properties and use a more aggressive bandwidth control mechanism instead of an adaptive congestion control algorithm.

System scalability is a key issue for the acceleration of Internet traffic. For public networks, it can be seen that the number of TCP connections transiting the hub can be very high, particularly in areas where PCs are poorly maintained the network is rife with viruses. Peer to peer and other applications further load the network. It has been seen that the use of Satellite for backhaul network applications may use two orders of magnitude greater TCP connections than private or government networks.

4.3 Comments

Application specific PEPs fail to yield the important benefits from a generalized technology. General PEP technology providers need to be flexible to fully address the specific needs of particular TCP/IP applications performed over different satellite-based technologies, as well as terrestrial wireless and even wired configurations. An important aspect of this is flexibility for PEP TCP acceleration, independently in each direction.

The availability of affordable but reasonably advanced processors yields new opportunities for inline streaming compression. The compelling economics of advanced optimization, while beyond the scope of this paper, must be more thoroughly considered by the satellite industry. The use of various bandwidth management techniques, compression, shared caches or data suppression techniques must also be considered in view of the network topology and scale and the PEPs must be adaptable to new requirements.

References

1. Border, J., et al.: Performance Enhancing Proxies Intended to Mitigate Link-Related Degradations, RFC 3135 (June 2001)
2. Consultative Committee for Space Data Systems, Space Communications Protocol Specification (SCPS): Transport Protocol, CCSDS-714.0-B-1, `http://www.ccsds.org`
3. The SatLabs Group, `http://www.satlabs.org`
4. RFC 2474 – Definition of the Differentiated Services Field (DS) in the IPv4 and IPv6 Headers