

OPTIMIZING IP TRAFFIC OVER SATELLITE

SOLVING THE PERFORMANCE, CAPACITY AND COST CHALLENGES OF SATELLITE COMMUNICATIONS

WHITE PAPER

THE SATELLITE CHALLENGE

With the growing demand for remote access broadband services, satellite networks are becoming increasingly popular. In remote and rural areas, where there is little or no terrestrial communications infrastructure, satellite communications becomes the only viable option. Satellite links also provide a unique high-availability option as fail-over links to remote branches. Locations that would like to ensure connectivity in terrestrial network outages are utilizing satellite links to free themselves from those dependencies. Although the TCP/IP protocol is widely used to provide reliable data delivery it has challenges when running over wide area networks. When asked to run over satellite those challenges are exponentially increased. These stem from the inherent characteristics of satellite channels, such as large delays, increased error rates, and bandwidth asymmetry and availability.

Satellite communications solutions come with a high price tag, and often do not satisfy the performance needs of the business. Since most communications satellites are located in orbit at an altitude of 22,240 miles (35,786 km) above the equator, the distance for a packet to travel one way from the client to the server could easily approach 50,000 miles. Under even optimal conditions, this distance translates to considerable latency. Add chatty applications to the equation, and performance can often become completely unacceptable sometimes causing applications to fail entirely. Additionally, because they are dependent on atmospheric transport, satellites are especially prone to packet loss caused by environmental interference. Add all of these challenges up and satellite network reliance can prove frustrating, and at it's worst, a failed infrastructure investment.

The steadily rising demand from businesses, media companies and government entities has led to satellite bandwidth shortages and higher costs for the capacity that is available in the foreseeable future.

Complicating the picture, the available bandwidth from the current and near-term launching satellites has already been reserved. The steadily rising demand from businesses, media companies and government entities has led to satellite bandwidth shortages and higher costs for the capacity that is available in the foreseeable future. Expand Networks helps solve the performance, capacity, delivery, and cost challenges with solutions that dramatically increase the efficiency and scalability of satellite communications.

SATELLITE ENVIRONMENTS

Although enterprises use a variety of satellite solutions today, the two main environments used are Very Small Aperture Terminal (VSAT) satellite communication systems and Inmarsat Broadband Global Area Network (BGAN) service.

VSAT refers to receive/transmit terminals installed at dispersed sites connecting to a central hub via satellite using small diameter antenna dishes (0.6 to 3.8 meter). The network of VSATs can have a star or mesh topology. In the star topology, the hub station controls, monitors and communicates with a large number of dispersed VSATs. In a mesh topology a group of VSATs communicate directly with other VSAT in the network without going through a central hub. A hub station in a mesh network performs only the monitoring and control functions.

VSATs are flexible communications platform that are highly reliable and can be installed quickly. Typically, the hub antenna and the VSATs are owned, operated and maintained by a VSAT service provider. VSAT systems are attractive where the coverage area is large, where quick installation is required and where terrestrial alternatives are difficult to obtain.

By utilizing techniques such as bit-level compression, object-level caching and next generation quality of service, the amount of data being sent across these satellite environments can be reduced and the application performance can be enhanced significantly.

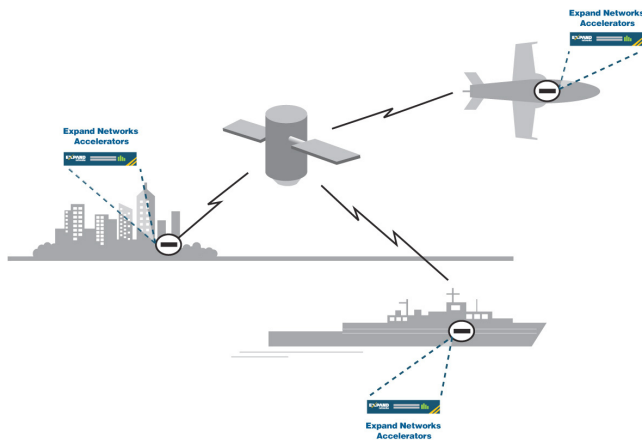
Inmarsat's BGAN is a mobile satellite offering with download speeds topping out at 492 Kbps and upload speeds in the 300 to 400 Kbps range. The main advantage of BGAN is its portability. Along with the benefits of the system there is one tradeoff; the price. The price for voice connectivity currently starts at approximately 99 cents per minute. The price of data connections ranges from \$5-\$10 a MB and is often capped with download per day limits.

By utilizing techniques such as bit-level compression, object-level caching and next generation quality of service, the amount of data being sent across these satellite environments can be reduced and the application performance can be enhanced significantly. Also, as explained, BGAN and VSAT services can get quite expensive. WAN optimization appliances, such as Expand Networks Accelerators, can alleviate some of the performance and cost concerns by mitigating these challenging networks.

OPTIMIZATION OVER SATELLITE

Performance Enhancement Proxy (PEP) technology can mitigate the effects of latency, help fill the link with data and improve network performance. Installing a pair of PEPs at either end of a satellite link can trick each local network into believing the remote, satellite-linked network is right next door. However, not all PEPs are alike, and bundled satellite modem PEPs are constrained in their capabilities and deliver limited results.

Unlike simple PEPs, Expand Accelerators apply a mix of TCP acceleration, link conditioning, compression, and application-specific acceleration techniques to increase the performance of applications despite the degraded conditions. They offer extensive caching, compression, and QoS capabilities to overcome congestion and latency on the WAN to provide the most effective use of the available bandwidth. Expand also offers advanced technologies such as packet fragmentation, to reduce the effect of large file transfers and similar applications (e.g., FTP) on sensitive real-time traffic such as VoIP and server based computing (Citrix/MS Terminal Services/VDI).



Protocol Acceleration

TCP is subject to a number of limitations on a WAN that severely effect it's performance. The Expand Accelerator's TCP acceleration (PEP) overcomes these limitations to increase performance for all TCP applications. It's based on the Space Communications Protocol Standard (SCPS) developed by NASA and the U.S. Air Force. SCPS is a transparent, highly reliable set of TCP extensions that interoperates with other SCPS based devices . With SCPS, the Accelerator acts as a transparent TCP proxy to all TCP traffic to overcome latency by using a variety of techniques, such as enlarging transmission windows for higher throughput, overcoming TCP slow-start and advanced congestion-avoidance mechanisms.

Expand's SCPS implementation is the highest performing in the industry. Along with it, a number of other features help deliver significantly better application performance for satellite users in a number of ways.

Expand's SCPS implementation is the highest performing in the industry. Along with it, a number of other features help deliver significantly better application performance for satellite users in a number of ways:

- **Window Scaling.** Standard TCP stacks on workstations typically support a maximum transmission window of 64 KB. This means that, at any one time, only 64 KB of data may be transmitted without receiving an acknowledgement (ACK). When transferring data over a long-haul WAN connection, the maximum threshold must be increased to avoid severe under-utilization of the expensive link.

In order to keep the pipe full, Expand Networks' TCP acceleration creates a much larger window to allow the link to be fully utilized. When the window size is enlarged, the destination does not send as many ACK packets, resulting in further improvement in bandwidth utilization

With TCP acceleration, a file transfer that used only a fraction of the available bandwidth will now fully utilize the link and complete the transfer faster. Sessions will scale faster and fill the pipe immediately, effectively avoiding the "saw-tooth" performance curve in standard TCP sessions.

- **Error Detection and Proactive Resolution.** Standard TCP assumes that all packet loss is caused by network congestion. But, on a dedicated and controlled WAN link packet loss could also be due to a bit error that can occur from noise in the channel or other environmental conditions. TCP Vegas and TCP Reno congestion avoidance mechanisms have been designed to address these network issues and adjust the transmission speed accordingly. The TCP Vegas algorithm emphasizes packet delay, rather than packet loss, as a signal to help determine the rate at which to send packets. TCP Vegas detects congestion at an early stage based on increasing Round Trip Time (RTT) values of the packets in the connection. TCP Reno, on the other hand, detects congestion only after it has actually happened via packet drops.

By implementing both TCP Vegas and TCP Reno network congestion avoidance mechanisms, Expand Accelerator's TCP acceleration feature can perform the appropriate corrective measure both when there is network congestion and when there is bit-error packet loss. Expand Accelerators can adapt to any satellite environments since it allows the users enable one or disable both of these congestion avoidance mechanisms.

- **Dynamic Bandwidth Adjust.** In addition to TCP Vegas and TCP Reno algorithms discussed above, Expand Accelerators have also implemented the Dynamic Bandwidth Adjust feature. Through a real-time feedback mechanism, the Accelerator can automatically adjust the bandwidth it sends to the WAN when congestion in the network occurs. This feature also provides effective traffic optimization when multiple paths with different bandwidths or delay characteristics exist. This mechanism ensures the most reliable optimization of all IP traffic at all times, and can also help in environments with multiple satellite links and backup links for disaster recovery.
- **Fast Start.** The proxy operation of the Accelerator means that TCP sessions are terminated locally, so that establishing and tearing down TCP connections takes place at LAN speeds. This increases the link utilization and overcomes TCP's slow start and congestion avoidance on both sides of the network, resulting in better and more immediate response for users.

Application Acceleration

Expand also offers several “plug-ins” that add advanced caching techniques and packet aggregation capabilities to optimize CIFS and accelerate specific applications such as HTTP and FTP; highly interactive applications such as Citrix, Terminal Services, and Telnet; and virtual desktop solutions using Virtual Desktop Infrastructure (VDI).

The Accelerator goes far beyond simple queuing with dynamic QoS that is tightly integrated with application acceleration and adapts to network conditions to maintain chosen priorities even during periods of high congestion.

- **HTTP and FTP.** Web-enabled applications are characterized by many small objects (logos, graphics, etc.) that require multiple round-trips over the WAN, resulting in poor performance. Using caching techniques, Expand Accelerators serve the objects locally at LAN speeds. Eliminating repetitive content transfers over the WAN speeds delivery and saves bandwidth for both HTTP and FTP. Utilizing local termination of the HTTP session also speeds application response. The Accelerator’s byte-level caching and compression work in combination with its Layer-7 QoS to enhance response times; it seamlessly compresses web services (HTML, XHTML, Javascript, J2EE, JSP, etc.) and works on most MIME-types (user-configurable).
- **Interactive Applications.** Virtual desktop solutions such as Virtual Desktop Infrastructure (VDI) and highly interactive applications, such as Citrix Presentation Server (XenAPP), Terminal Services, and Telnet, are characterized by a request-reply protocol with relatively small packets. Although Citrix ICA and RDP protocols in particular have been optimized to deal with WAN latency, they still pose a challenge to application acceleration solutions that use block caching. As noted above, because blocks are larger than the size of interactions that these applications generate, it is difficult to get a cache hit without imposing additional latency required by the buffering needed to build a sample of sufficient size to generate a hit. However, Expand Accelerator’s byte-level caching and compression has the granularity necessary to optimize highly interactive applications, and is capable of increasing throughput by an average of 300% and peaks of more than 1,000%.

In addition, for interactive applications and solutions such as Citrix, Terminal Services, Telnet, and VDI, Expand offers a plug-in using packet aggregation that optimizes the bandwidth utilization further, increasing the number of user sessions by an average of two to three times and peaks of more than 10 times. It does all this with superior network, server, and user performance and all on the same infrastructure. More sophisticated than either a data reduction or compression technique, this application-specific optimization actually multiplexes sessions together temporarily for transport over the WAN.

Caching

Expand Accelerators offer caching at multiple levels: bit, byte, object, and file, so that, unlike other solutions, it can deliver benefits not only for “the usual suspects” (standard TCP-

based applications), but for non-TCP and interactive applications as well. It also incorporates advanced predictive algorithms to accelerate traffic the first time it's seen. The Accelerator's integration with Microsoft Domains enables it to offer synchronous file access with SMB Signing, offering full security and protection of data integrity.

The Accelerator offers a DNS acceleration service that can overcome the poor performance associated with a centralized DNS server deployment over a satellite connection. The Accelerator provides a transparent DNS caching solution that can process both UDP and TCP DNS requests, answering them locally. This feature does not require any special configuration on PCs or servers. It is simple to deploy and use. By enabling DNS acceleration on the remote Accelerator, the DNS records are cached and available locally across the entire branch office. This, in turn, significantly reduces latency, bandwidth consumption, and DNS server load, and improves the user experience.

Compression

The Accelerator's compression option can compress almost any type of traffic across the WAN, saving valuable bandwidth and allowing more applications and users over existing links. Expand Accelerators compress at the byte level as well as object and file level and can provide the benefits of compression and caching all IP traffic, including interactive traffic and real-time transmissions. Expand's compression is also flexible, offering both out-of-band transparency and tunneling modes. Expand's true transparency preserves existing infrastructure investments and maintains full visibility through any management solutions, while tunneling is available for those network configurations where WAN visibility isn't desired and full packet compression is delivered.

Quality of Service

The Accelerator's advanced QoS and traffic shaping mechanisms and intuitive management interface makes it easy to prioritize applications and guarantee bandwidth to the applications critical to a business. The Accelerator goes far beyond simple queuing with dynamic QoS that is tightly integrated with application acceleration and adapts to network conditions to maintain chosen priorities even during periods of high congestion. Not only does the Accelerator QoS share bandwidth not needed by higher-priority applications, its unique application-aware technology prevents higher-priority or highly accelerated applications from choking out lower-priority ones.

Layer-7 QoS and traffic discovery capabilities can be used to classify, monitor and prioritize network applications according to business objectives. Layer 7 QoS can guarantee optimal application performance regardless of WAN conditions by assigning a priority or bandwidth guarantees and limits for each application with full granularity. The Expand Accelerator QoS not only understands the difference between applications in terms of their importance to the business, it understands the difference between different sites in terms of their application usage. For example, the Accelerator can ensure that an ERP application gets higher priority than a CRM application for manufacturing sites, while the opposite is true for call centers.

Furthermore, unlike many competing solutions, the Accelerator provides both outbound and inbound QoS. Even if a remote site does not have an Expand appliance, the datacenter

appliance can throttle application behavior at the remote end to prevent link and server overload.

Packet Fragmentation

Even with proper QoS priorities, applications that transfer large amounts of information, such as CIFS, FTP, and backup systems, can effectively starve real-time applications such as VoIP and video over the low-speed links that are generally used for smaller remote sites. The problem is that even though the real-time application may have priority, the bulky nature of large-transfer applications takes too long to clear the link even when queuing and traffic shaping are enabled. The added latency that results can make VoIP, for instance, impossible for many branch offices. The solution applied by the Expand Accelerator is to reduce the size of data packets and intelligently fragment packets depending on the effective link speed and VoIP traffic profiles. Packet fragmentation also stabilizes network jitter and latency, maintaining optimal Accelerator performance for VoIP.

Centralized Management

The Accelerators and Expand Network's centralized monitoring and management system, ExpandView, provide powerful monitoring and graphical reporting for full application-level visibility and cost-effective end-to-end network management. Expand Accelerators are full NetFlow compliant and replace the need for costly probes. The Accelerators can be monitored and managed through the Web User Interface, CLI(SSH/Telnet), SNMP, or as entire environments through ExpandView. Expand's SCPS PEP and broad acceleration technology configurations, management, and reporting are all consolidated and exportable for reporting.

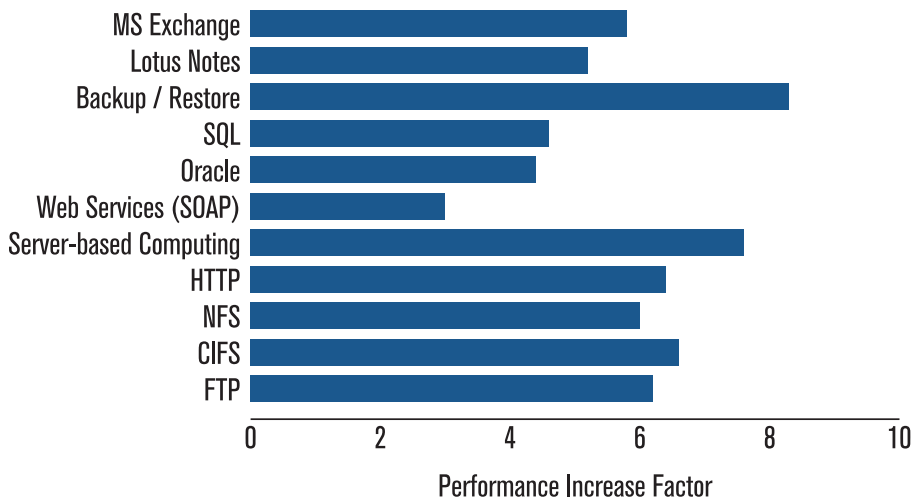
PERFORMANCE OVER SATELLITE

By deploying Expand Accelerators, enterprises can enhance application throughput significantly. A series of laboratory tests were conducted to characterize the Accelerators' performance using various applications protocols over satellite links. This section summarizes the test results.

Application Performance

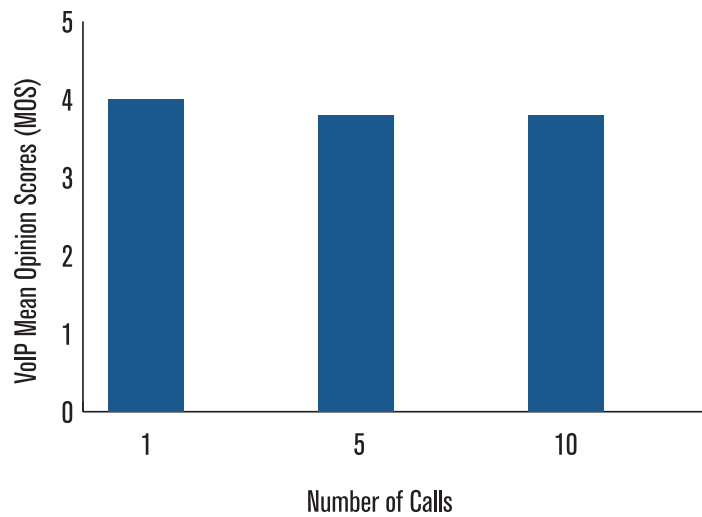
Network Testing Labs, an independent technology research and product evaluations firm, recently tested the performance of Expand Networks Accelerators . In these tests, several different application protocol streams were funneled through Expand appliances. The results were obtained across both real and simulated Satellite links with latency of 650 msec and data transmission rates of 64 Kbps, 512 Kbps and 1.5 Mbps.

The following chart summarizes the relative performance benefits of the WAN link optimizers for different protocols. . The performance increase factors denote by how much the optimization devices improved the performance of the WAN link. For example, a WAN acceleration factor of 2 would mean the data traversed the link in half the time it did in un-optimized form. In another words, a performance increase factor of 2 means the optimization device made the link carry data twice as fast.



VoIP Performance

Network Testing Labs also measured the Mean Opinion Scores (MOS) with the WAN links carrying Voice over IP (VoIP) and Expand appliances optimizing this traffic. VoIP MOS provides a numerical indication of the perceived quality of VoIP calls. The MOS is expressed as a single number in the range 1 to 5, where 1 is lowest and 5 is the highest perceived quality. In these tests, MOS of 2.9 was measured when Expand Accelerator was in bypass mode and no optimization took place. The following chart shows the VoIP MOS test results for each set of tests.



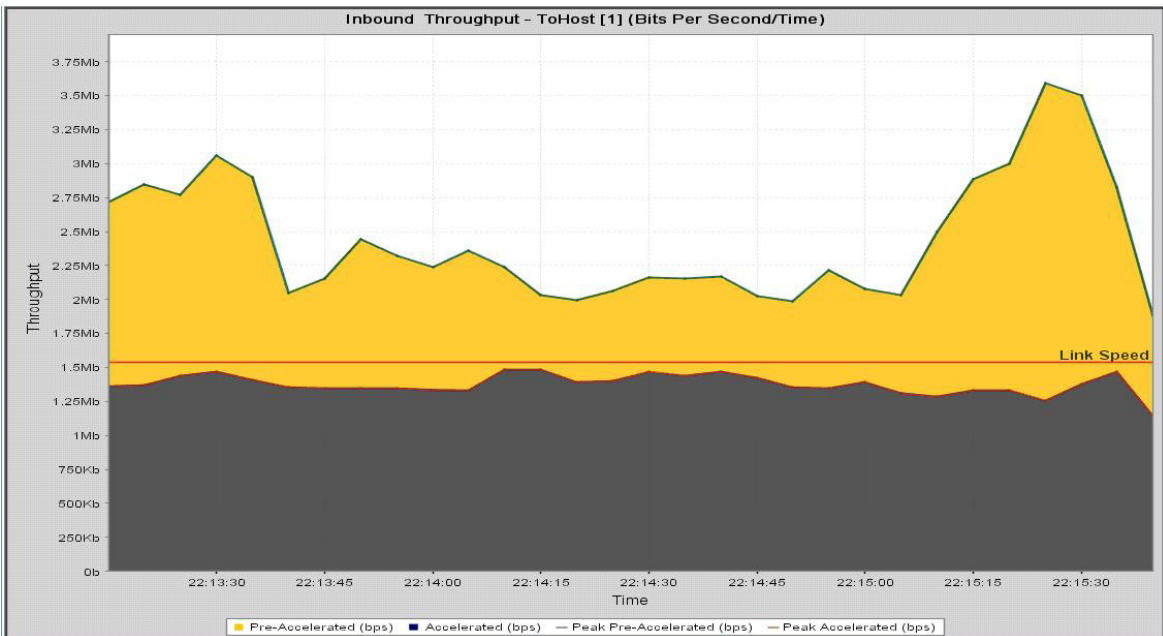
Citrix Performance

In another laboratory setup, Expand tested application performance over simulated satellite links with 600 ms of round-trip delay. Three scenarios were tested: baseline Citrix, Citrix with TCP acceleration and fully optimized Citrix. For all these cases, Citrix's SpeedScreen Latency

Reduction Manager was enabled. SpeedScreen is a latency reduction mechanism Citrix uses to provide local text echo and mouse click feedback. The configuration for each of these scenarios and results of these tests are shown in the table below

Scenario	Configuration	Performance
Baseline Citrix	<p>Citrix Configuration: SpeedScreen & Native Compression enabled</p> <p>Expand Configuration: Bypass</p>	User experience was unusable. It took almost 10 seconds for the login window to appear. Flash video would not play at all. Navigation through Windows was very difficult. ~2-3 Sec delay when typing in MS Office.
Citrix with TCP Acceleration	<p>Citrix Configuration: SpeedScreen & Native Compression enabled</p> <p>Expand Configuration: TCP Acceleration enabled, No Compression</p>	User experience was poor, but usable. TCP Acceleration had trouble keeping the link full. Flash video was very jerky until it locked up the entire session. Navigation through Windows resulted in approximately 1-2 seconds screen response time. ~1-2 Sec delay when typing in MS Office.
Fully Optimized Citrix by Expand	<p>Citrix Configuration: SpeedScreen enabled & Native Compression disabled</p> <p>Expand Configuration: TCP Acceleration & Compression enabled</p>	User experience was good. Flash kept the session up & operational. Navigation through Windows was less than 1sec delay. Less than 1sec delay when typing in MS Office.

The diagram below shows the effective throughput over the simulated 1.5 Mbps satellite link. The yellow area indicates the effective Mbps transferred over the link to the remote site and the blue area indicates the actual bandwidth consumption on the link after compression. The Expand solution compressed up to 4 Mbps on a 1.5 Mbps link, which translated into improved response time for end-users, full utilization of the investment in bandwidth, and significantly faster and more reliable throughputs.



SUMMARY

Satellite communications, while fulfilling unique business requirements, comes with significant network challenges. Just as the demand increases the limited satellite service availability is severely strained. Expand's solutions for optimizing satellite communications work holistically to give users and providers much greater control over, and utilization out of, their networks. The technology's standards-based approach combined with next generation compression, caching, and conditioning techniques provide satellite environments the ability to squeeze the maximum efficiency out of their challenging network conditions. Expand's solution achieves all this using broad reaching technology. It's impacts span all IP traffic, all traffic types (including interactive and real-time applications), using an easily managed set of advanced technologies. The Expand Accelerator is a tightly integrated, multi-service platform that delivers a complete application acceleration solution. As proven by their 10 years of accelerating the world's satellite environments, Expand has developed the world's best satellite acceleration solution.