Low-latency networks for storage infrastructure

shaping tomorrow with you

FUJITSU

A powerful and flexible Storage Area Network (SAN) based on low-cost-per-bit, with protocol-agnostic and secure transport over long fiber spans, can enable a host of valuable business capabilities for an enterprise. Today's most demanding applications—moving Virtual Machines (VMs), transporting video or executing transactions at high speed and in high volume via Ethernet, interconnecting high-performance computing (HPC) clusters via InfiniBand, or enabling synchronous storage services via Fibre Channel—all demand something more: ultra-low latency.

The gap in time between when an instruction is issued and when it is executed across a network has become a point of particular hypersensitivity among IT departments in a range of industries. Latency reduced to mere nanoseconds is emerging as simple table-stakes for the highest-value applications. It's a challenging standard that demands evaluation and optimization of end-to-end infrastructures and processes—spanning all connections across and beyond an enterprise—for any lapse that can be identified and eliminated.

The evaporation of tolerance for latency has influenced whole operational strategies in some cases. Whereas enterprises once simply contracted with carriers for site-to-site connectivity over shared networks, many now demand dedicated fiber-optic infrastructures for their most important routes and most time-sensitive applications. Even the electronics used to "light" the fiber networks now command scrutiny. Traditional approaches to common optical-transport functions from transmitter to receiver can yield levels of latency that were considered negligible only a couple of years ago but are untenable by today's standards.

Matching storage services to technology capabilities

SANs boomed in the first decade of the 2000s. A swarm of economic and regulatory pressures descended on enterprises, demanding that companies of almost every size and across almost every industry protect their operations by distributing information resources across geography, offsetting potential data loss and enabling rapid recovery. The technologies put in place to support the new disaster-recovery and business-continuity needs gave enterprises the opportunity to expand reliance on a range of compelling applications, such as clustering, grid and cloud computing. Whole new modes of dynamic operation were enabled, such as the ability to more cost-strategically deploy equipment and personnel across real-estate markets.

Wavelength Division Multiplexing (WDM) optical networking provides the high bandwidth and flexible support of traffic protocols (InfiniBand, Fibre Channel and various flavors of Ethernet, for example) that today's SANs require. But deploying the highest-capacity, screaming-fast router or multi-core-processor computer blade on a WDM-based network and contracting with a carrier for their highest-bandwidth connections won't necessarily satisfy the latency requirements of the enterprise's most time-senstive services.

More and more frequently, enterprises find that the most latency-intolerant SAN applications—moving VMs, synchronous disk mirroring, HPC clustering, video transport and processing of high-volume transactions, for example—demand the shortest-available, dedicated connections enhanced with the fastest electronics. Tremendous, unprecedented pressure is on IT managers to deliver even the most modest of reductions in delay anywhere in the ecosystem of equipment and processes that enable these latency-intolerant applications. This challenge demands uncovering and eliminating tiny traces of latency in the building-to-building fiber-optic transport connections themselves.

Uncovering latency in optical networks

Enterprises can undertake a number of technology improvements to improve latency: adoption of faster Central Processing Units (CPUs) and Network Interface Controllers (NICs), accelerated middleware appliances, and low-latency switches. Enterprises that require additional improvements in order to achieve ultra-low latency and optimize transaction rates can turn next to three common sources of delay in optical transport (Figure 1):

- Fiber delay: the length of the actual fiber cables
- Proximity delay: the physical distance to the fiber connection
- Equipment delay: the processing speed of the network equipment

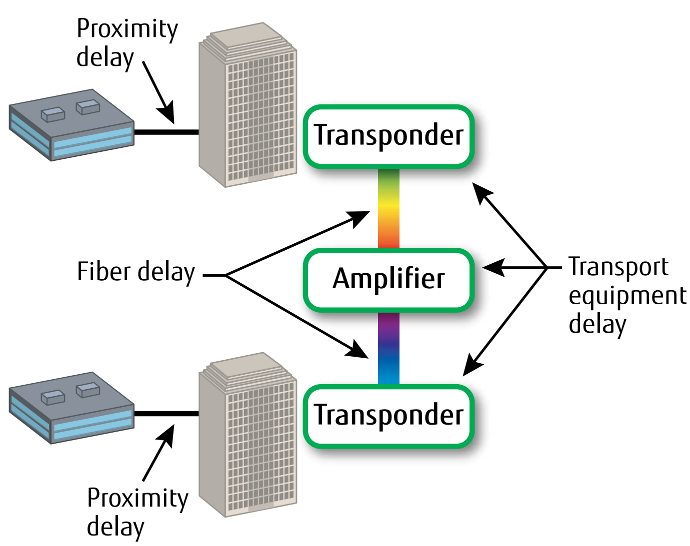

Figure 1: Sources of latency in optical transport networks

FUJITSU NETWORK COMMUNICATIONS INC.
2801 Telecom Parkway, Richardson, Texas 75082-3515
Telephone: (972) 690-6000
(800) 777-FAST (U.S.)
us.fujitsu.com/telecom

Fiber delay
It's a simple fact: the longer the physical route, the greater the time it takes for traffic to get from one end to the other. Cut roughly eight inches of optical fiber between two connected locations, and you cut a nanosecond of transport latency between the two sites. The competitive advantage of minimizing fiber delay can be significant for enterprises that do business in major cities, because the fiber routes within metropolitan areas often snake up and down manholes, across streets and along whatever available easements the fiber installer could secure. Enterprises requiring ultra-low latency must closely evaluate the directness or circuitousness of the routes available for dedicated infrastructures from dark-fiber providers.

Proximity delay
The space near a fiber access point is a finite quantity, and delay is shaved with each inch closer to the optical gear. Given that real estate near fiber junction points is so in demand, enterprises with the most challenging latency requirements might reach out to lease space from co-location providers.

Equipment delay
Finally, there's the actual speed of the optical equipment across the end-to-end transport network. When an enterprise's acceptable level of latency is measured in full seconds, the delays injected by the gear that traffic encounters along a fiber network are probably not of consequence. But it is incorrect to believe that data in a glass network actually moves at the speed of light. Every device introduces some amount of delay, and, therefore, any non-optimized transport equipment across the information path can thwart business models that are based on the most challenging network latency limits.

Eliminating these sources of equipment delay is trickier because different devices carry out key, common optical network functions in very different ways, yielding very different degrees of latency.

A deeper look at transport latency
Color conversion, amplification, dispersion compensation and regeneration are all common optical-networking functions that are carried out among different optical-networking solutions. The enterprise that fails to understand how its Wavelength Division Multiplexing (WDM) gear accomplishes these four basic functions risks adding significantly to the time it takes for network instructions to be executed after being issued. And, in latency-sensitive algorithmic trading, a firm's whole business strategy could be jeopardized by ill-fitting implementations of necessary optical-transport functions across the fiber strands linking key locations.

Color conversion
In color conversion, traffic signals are converted, or "transponded," from gray to a specific light color (wavelength). Frequently, multiple colors are aggregated, or "muxponded" into a single high-speed channel across a fiber-optic network. Electrical Time Division Multiplexing (TDM) techniques might incur milliseconds of latency. Examples are:

- Optical Channel Data Unit (ODU) encapsulation and thin film filters
- Forward Error Correction (FEC) algorithms
- Performance monitoring
- Protocol conversion
- Clock recovery

Best-in-class, purpose-built transponders and muxponders have emerged that produce latency in only the high nanoseconds to low microseconds.

Amplification
Traffic signals naturally weaken as they travel mile after mile along an optical fiber network. Optical amplifiers are designed to offset this weakening effect by boosting the signal. The Erbium-Doped Fiber Amplifier (EDFA) is one of the most commonly deployed types of amplifier device. But, because delay of hundreds of nanoseconds is introduced by each EDFA on an optical route (even microseconds in the case of some high-gain, dual-stage EDFAs), the most latency-sensitive enterprises must instead turn to alternative, optimized architectures that do not rely on integrated doped fiber spools or that add significant "noise," for example. High amounts of noise and low gain figures both require higher numbers of amplifiers, which increases latency. Recently introduced innovative Raman amplifier technologies do not require the addition of supplementary fiber, leading to a shorter optical path and more efficient transmission. Raman amplifiers also deliver low-latency inline amplification from within the fiber itself.

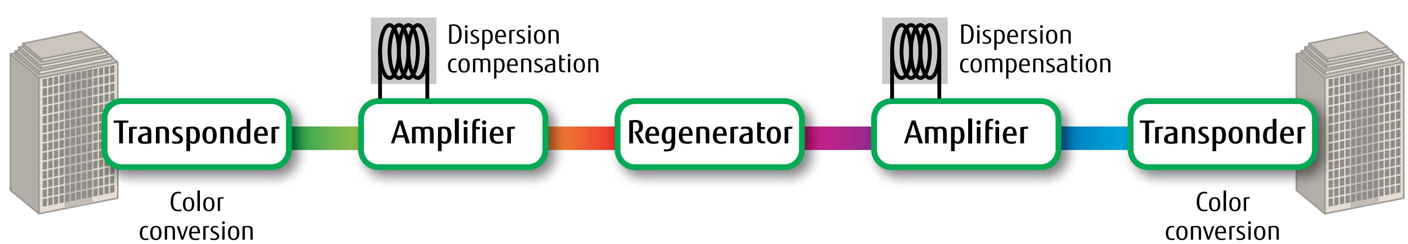

Figure 2: Optical networking functions that can increase latency

Dispersion compensation

Traffic signals also degrade because of "chromatic dispersion," a phenomenon in which usually high-speed optical channels smear into a rainbow of colors and bleed into neighboring traffic streams. Service issues result. Dispersion Compensating Fiber (DCF) can successfully reverse that effect, but the long spools of the special fiber that are required to accomplish the objective results in unwanted latency. And when improperly located, DCFs demand additional amplifiers in a fiber network—producing even more delay. Alternative remedies that leverage Fiber Bragg Gratings (FBGs) can deliver dispersion compensation while delivering negligible amounts of latency.

Regeneration

Signal regeneration is another common optical network function. Variants of the feature offset signal degradation over longer links, but any of the traditional implementations of electrical regeneration can introduce delay. Techniques that entail FEC termination and setup, non-optimized routing and cabling among filters and cards, for example, can inject hundreds of microseconds of latency. Purpose-built, low-latency approaches to regeneration slash this delay into the nanoseconds.

Fujitsu low-latency network solutions

Constructed expressly for the job of low-latency networking, The Fujitsu FLASHWAVE 7420 Metro/Enterprise WDM platform delivers breakthrough performance for synchronous SAN applications such as business continuity, disaster recovery, moving VMs, HPC clustering, high-volume transaction processing and video transport. Only the Fujitsu platform—employing transparent wavelength conversion and purpose-engineered techniques for inline amplification, dispersion compensation and signal regeneration—enables IT managers to eliminate so many sources of optical-transport latency end-to-end across infrastructures and ensures enterprises the fastest-possible connectivity along key routes.

Furthermore, the FLASHWAVE 7420 system delivers flexible, cost-effective, protocol-agnostic support for all of the leading services in play across enterprise networks, such as Ethernet, InfiniBand, Fibre Channel, Fibre Channel over Ethernet (FCoE)/Data Center Bridging (DCB) and Fiber Connection (FICON). The Fujitsu solution also offers flexible scalability (fast expansion of up to 80 transmission channels of service over two strands of fiber), cost-effective power and space efficiency, unmatched distance support of fiber spans of up to 2,000 kilometers, physical-layer encryption and optical-line monitoring, and a safety net of best-in-breed technology, personalized network monitoring and application support to ensure highest reliability. Beyond the industry-leading FLASHWAVE 7420 platform, Fujitsu offers IT managers valuable consultancy in designing the optimal, lowest-latency solution for their needs; an array of latency-optimized cards, and ease-of-use support (interoperability tests, implementation and setup, surveillance and operation, and key partnerships). Together, these convey to enterprises a differentiating low-latency networking capability for the most innovative SAN capabilities—available uniquely from Fujitsu.

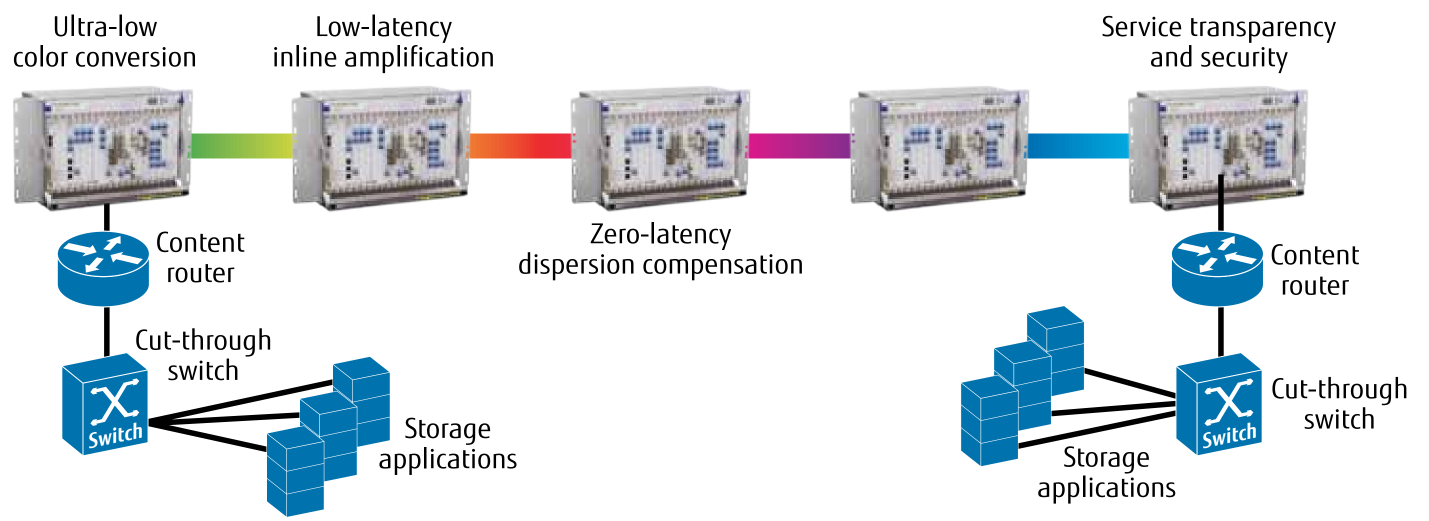

Figure 3: End-to-end low-latency advantage

FUJITSU NETWORK COMMUNICATIONS INC.
2801 Telecom Parkway, Richardson, Texas 75082-3515
Telephone: (972) 690-6000
(800) 777-FAST (U.S.)
us.fujitsu.com/telecom