Practical Ultra-Low Latency Optical Transport Networks

Low Latency Networking Market Drivers

In today's highly sophisticated trading environment, market efficiency is rapidly entering a new era. Market information is now available in machine-readable format within microseconds, and high-end servers can process complex trading algorithms and execute large volumes of trades without any human intervention. Financial institutions constantly monitor the performance of their trading systems and are looking for ways to increase the speed of their mission-critical operations. To financial institutions, fractions of seconds carry significant monetary consequences—the speed of their network defines the success of their business. Optical networks provide the speed necessary to meet the responsiveness these users demand. Significant opportunities abound for technology vendors with the right combination of solutions.

Low-latency trading refers to the network connections that financial institutions use when connecting to stock exchanges and Electronic Communication Networks (ECNs) to execute financial transactions. With the spread of computerized trading, it is estimated that electronic trading now makes up 60 to 70 percent of the daily volume on the New York Stock Exchange (NYSE) and algorithmic trading almost half of that. American and European markets generally have a higher proportion of high-frequency algorithmic trading than other markets today.

Trading using high-performance computers has developed to where millisecond improvements in network speeds can offer a significant competitive advantage for financial institutions. Recent market research estimates that securities firms and hedge funds that specialize in rapid fire, high-frequency trading generated more than $20B in profits in 2009. A 1-millisecond advantage in trading applications can be worth $100 million a year. Although numerous factors influence trading-system latency, the latency introduced by a signal propagating through long optical fiber routes and optical transport equipment can be the most significant contributor. The only way to minimize fiber-related latency is to use the shortest fiber route between trading floors. The next area for latency optimization to focus is the physical optical transport equipment itself.

The total enterprise optical equipment market is worth an estimated $1B annually, with the financial vertical alone exceeding $500M annually. A significant market opportunity exists for network equipment suppliers to provide reliable low latency solutions that are simple to use. While the greatest latency efficiencies result from operating a private optical network, many firms lack the expertise or resources to build and operate their own private optical network alone. Increased adoption of outsourced services offers additional revenue opportunities for solution providers who can offer turn-key managed service solutions to these enterprises.

Network Design Approaches

Almost every optical vendor claims to have low-latency optical transport solutions, primarily because they all use the same network building blocks, including transponders, amplifiers, regenerators, and dispersion compensators.

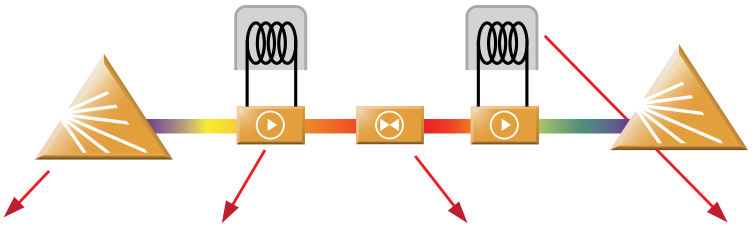

Transponder
- Converts grey signal to color channel
- FEC disabled

Amplifier
- Usually uses a combination of EFDA and Raman technology
- Raman provides lower latency but more complexity

Regeneration
- Electrically regenerates signal, may add jitter
- Segments the network and improves fault isolation

Dispersion Compensation
- Lower latency FBGs typically used
- FBG not a drop-in; network must be designed for it

Transponders adapt short reach optical 1 GigE or 10 GigE optical signals, converting them to specific wavelengths and launching them into a dark fiber network. Using passive multiplexers/demultiplexers to combine these signals with others allows for transporting them along the same dark fibers. In some implementations, transponders optically wrap signals using standard defined techniques such as optical transport networking (OTN). OTN encapsulation does not affect the data payload whatsoever, and the added overhead provides a management channel for maintenance and troubleshooting. For some data rates, such as 10 Gbps, OTN provides forward error correction (FEC) which helps compensate for bit errors that can occur along the optical transmission path. While FEC may recover lost bits resulting in better performance in the upper layers, for example by eliminating IP level re-transmission, the added overhead can result in additional nanosecond latency penalties.

Optical amplifiers boost the optical signal extending them over longer distances. While they amplify the signal itself, they also amplify any noise or distortion that gets accumulated along the optical transmission path. The typical amplifier technologies used in optical networks are either erbium-doped fiber amplifiers (EDFAs), which use coils of fiber within the amplifier unit itself to amplify signals. In low-latency networks, EDFAs contribute latency due to this internal fiber coil. Another amplifier is called Raman that typically is used in conjunction with conventional EDFA technology to enable applications such as long multi-span links, single-span links up to ~300km, and ultra-long haul links. Because

Raman amplification occurs within the transmission line itself, the quality of the line can dramatically affect the performance of the Raman amplifier and, therefore, the achievable signal gain. While Raman technology offers some advantages, it also causes many operational issues related to the deployment of the technology that require understanding, including:

- losses along the transmission line
- back reflection along the line
- connector quality
- fiber quality
- laser safety and automatic shut down and start up.

Raman amplification is the process whereby pump energy propagating along the transmission line causes signal amplification through stimulated Raman scattering. The actual transmission cable itself is used as the amplification medium requiring many safety considerations when using Raman. In addition, Raman amplifiers typically contain multiple pump laser diodes within an amplifier unit.

Optical performance limitations, such as transmitter optical launch power, optical receiver sensitivity, optical signal attenuation from the fiber itself, and optical dispersion that occurs as a signal travels down a fiber path require the eventual regeneration of signals. Regeneration refers to the process where the optical signal is converted back into the electrical domain for retiming, reshaping, and retransmitting, known as 3R. This conversion process may result in added latency and, depending on the vendor's specific implementation regenerators, may or may not result in higher latency than optical applications. Regenerators, however, provide an important intermediate maintenance access point for monitoring and troubleshooting networks, which is even more important considering that FEC is often disabled for 10 Gbps transmission for low-latency implementations.

In some cases, to offset the distortion that occurs in optical transmission, some vendors propose dispersion compensation modules (DCM) to effectively reverse the accumulated optical dispersion effects. Traditional DCMs often use long coils of internal dispersion compensating fiber that add latency and are unsuitable for low-latency network designs. Many vendors propose an alternative type of DCM called fiber Bragg grating (FBG), which does not use these fiber coils and therefore provides lower latency. Regardless, determining where to place a DCM is important and requires considerable network planning and a thorough understanding of the physical fiber characteristics, because no two networks are the same and network design involving DCMs is not trivial.

Low Latency Network Implementation Considerations

As already highlighted, all optical transport equipment vendors virtually use the same network building blocks, so one may ask what are the key differences between vendors? Many factors require consideration when implementing a low latency network.

First, users must determine the most appropriate interface rate for their given application. For instance, understanding the tradeoffs between 1 GigE and 10 GigE when it comes to throughput and serialization delay to determine which offers the lowest latency for their needs.

The next step is selecting the appropriate transponder. In many cases, network designers choose to disable FEC when using 10G transport to reduce latency, which is acceptable but requires using a more robust optical line design. Specifically, without FEC, data becomes more susceptible to degradation effects resulting from optically amplified transmission paths. In many cases, users require a combination of 1 GigE, 10 GigE, and in some cases other protocols to be transported. Therefore, it is important to have a flexible network design that can be easily evolved to support more or varying traffic types and data rates.

Network designs are often optimized for either a specific type or specific number of channels to achieve the lowest latencies possible, which typically involves reducing the number of regenerators used in the network. Unfortunately, future additions, new wavelengths, or changing from one data rate to another often results in a network redesign, including adding amplifiers and physically repositioning and redeploying DCMs. This redesign obviously results in service disruptions, increases costs and latency, and decreases user productivity, satisfaction, and overall business profitability.

The biggest contributors to network design are the characteristics of the physical fiber itself. Network designs must allow for effects from distance and fiber splices from deployment throughout the life cycle of the network. Fiber plants age and evolve, and new splices get added. A key implication from these considerations applies to the selection of amplifier technology.

As discussed earlier, Raman amplifiers are often recommended for use in low latency networks, because they use the transmission fiber itself as the amplification medium and, therefore, offer lower latency than traditional EDFA. However, systems using Raman amplification differ from conventional EDFA-based systems. For example, high loss along the transmission line, and in particular at loss points occurring close to the Raman amplifier, can severely decrease the available pump power and the achievable Raman gain. Loss can occur because of dirty or faulty connectors or sharp bends and other stress points along the fiber.

High back-reflection can also occur at discrete loss points along the transmission line causing part of the pump-energy propagating along the line to be back-reflected and returned to the pump laser diode where it originated. A high level of back-reflection can degrade the performance and decrease the available pump power. Dirty or faulty connectors may cause loss and/or back-reflection, resulting in reduced available Raman pump power.

An additional risk is that the high Raman pump power passing through the connectors can harm or degrade the connectors over time. Minimize the number of connectors close to the Raman amplifier, and those that remain must be specially designed to accommodate high power (for example, E2000 connectors). Directly splicing the output of the Raman amplifier to the outside cable plant requires specialized technicians and more time, both of which could impact the ability to rapidly turn up or restore services or to quickly accommodate network changes. However, many fiber providers do not allow direct splicing into third-party fiber plants. Special precautions are required when connecting Raman amplifiers compared to EDFA, as well as higher technical complexity and risk.

In addition, the output power of Raman pumps is much higher than EDFAs and well above the designated safe level of radiation. Sometimes the pump power propagating along the fiber reaches levels that are hazardous to eyes and skin. As a result, Raman amplifiers have many built-in safety features that increase the complexity of their operation and deployment.

Optical transmission systems are typically required to comply with Class 1M hazard requirements according to IEC standard 60825 Part 2. In the case of an accidental connector opening or a fiber break, all lasers and transmitters must reduce power to a safe level, typically within 1 second of detection. Raman amplifiers must support safety mechanisms allowing automatic shut down of the pump. If a disruption occurs along the transmission line that causes the Raman amplifier to shut down, an automatic restart procedure initiates. Safety mechanisms increase the complexity and risk during installation, as the optical line could be up/down during installation. Autonomous Raman amplifier safety mechanisms may also increase the time to restore service if an issue arises during normal operation.

Another factor that providers must consider when implementing a low latency optical transport network is that low latency is worthless in unreliable or malfunctioning networks. Providers must choose their building blocks carefully. The telecom industry has defined specific methods for calculating and predicting hardware reliability. During the design stage, it is important to develop products with the fewest single points of failure possible. Industry calculations measuring predicted failure rates are based on detailed statistical analysis defined by the Telcordia standards body in SR-332-CORE.I2 to derive the predicted failures in time (FIT) of each product.

Network designers must consider maintenance, operation, and reliability. A network is only as fast as it is reliable; likewise, it must also be easy to troubleshoot. More importantly, rapid fault isolation and restoration is vital when failures occur. Optical network designs involving more regeneration nodes provide several advantages, including easier fault isolation and rapid restoration per wavelength as well as intermediate per channel statistics. Regeneration nodes further segment the network into smaller sections which is especially beneficial for fault isolation in long-distance networks, simplifying network turn up, and increasing overall system reliability. Another benefit is the elimination of DCMs, which saves money and reduces complexity.

The up time of each wavelength increases when using a network design that implements more regeneration nodes. Figures 1a and 1b compare line design options using different combinations of optical amplifiers and Regens. Because implementation of regeneration nodes typically have lower FIT values than amplifiers, a network design involving more of them will increase its uptime; likewise it decreases the probability that all wavelengths on a multi-wavelength network will experience down time.

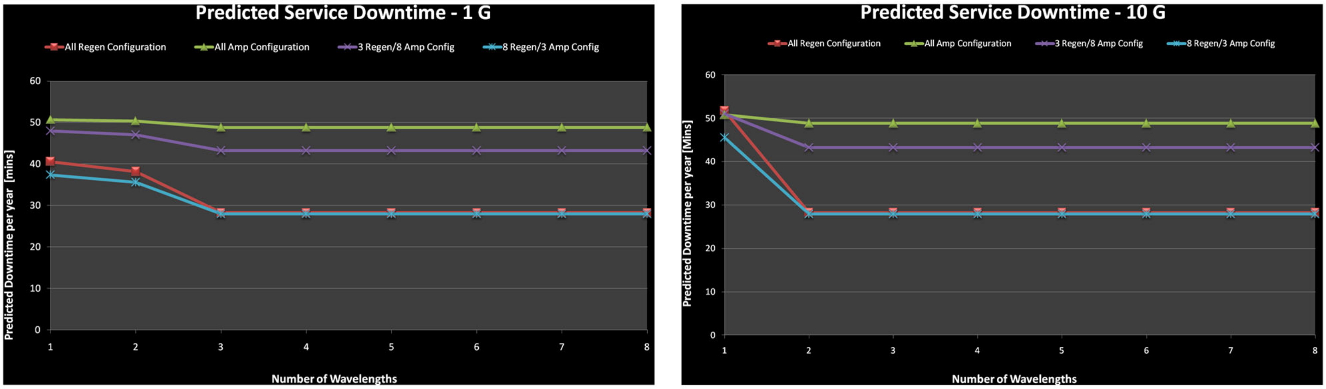

Figure 1a and 1b. Comparison of line design options using combinations of optical amplifiers and regeneration nodes

The effect is even more noticeable when using amplifiers with higher FIT values. For example, Raman amplifiers have higher FIT values than EDFA, resulting from the use of more pump diodes, to an estimated extra 1200 FIT values per amplifier. Quite simply, Raman-optimized network designs have lower up times than alternative approaches. Figure 2a and 2b compare the predicted system down time for a 1 GigE and 10 GigE wavelength of an optical network design using an all-Raman/EDFA design compared to one using all-Regen. The Regen-based network design has a considerably lower predicted annual down time.

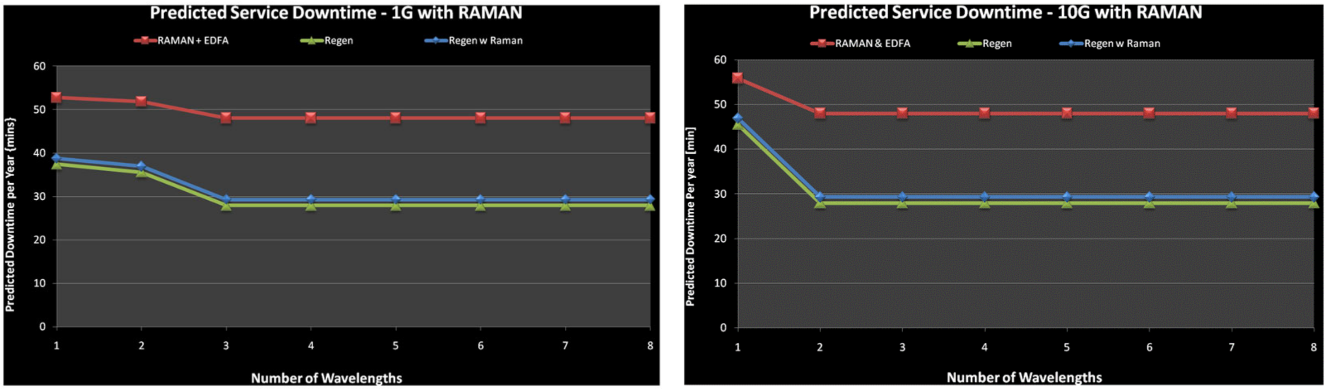

Figure 2a and 2b. Comparison of down time for a 1 GigE and 10 GigE wavelength using all-Raman vs all regeneration node designs

It is important to use the right combination of building blocks with lower FIT values when implementing a low latency transport network, particularly one carrying financial transactions from algorithmic and high-frequency trading.

Use Raman amplifiers carefully and only when latency improvement requirements warrant rather than as a general implementation practice.

Low Latency Network Implementation Considerations

Providers must make a number of key considerations when designing and implementing low latency optical transport networks. While latency is an important criterion, other factors are equally important, including implementing a solution that is simple to turn up without requiring an in-depth understanding of photonic theory or requiring special expertise. Similarly, while a network must be fast, it must also be reliable. When issues do arise, they must be easily and quickly isolated and remedied. Finally, the network itself must be capable of evolving to accommodate changes in the number wavelengths transported and varying data rates. Far too often network designers optimize only a single variable, such as latency, resulting in a network design with limited flexibility, such as:

- Physically placing DCMs specific to the wavelength count or mix.
- Placing amplifiers in locations that only accommodates a specific type or maximum number of wavelengths.
- Networks that do not include intermediate locations for fault monitoring.
- Networks designed with only optical amplifiers complicates containing a fault or issue affecting one wavelength from affecting others.
- Using various amplifier technologies can increase system complexity such that simple network changes are no longer possible without extensive expertise, expense, and user disruption.

In general, low latency network designs must fully consider all requirements, including latency. Reliability, maintenance, and the flexibility to evolve are as vital as network implementations that combine the solution building blocks appropriately for balance. Specifically, select network equipment with low latency and low failure probability and determine the appropriate combination of Regen and amplifier technologies.

JDSU Ultra-Low Latency Tool Kit

The JDSU Ultra-Low Latency network architecture combines the optimal mix of amplifier and regenerator technology while factoring in deployability for:

- simplified per channel fault isolation
- simplified turn-up with minimal complexity and risks
- high reliability and with rapid time to restore
- flexibility regarding system adds and changes
- latency reduction options through the sensible use of Raman amplification
- low latency.

For more information please visit: **www.jdsu.com/ultralowlatency**.

Test & Measurement Regional Sales

NORTH AMERICA	**LATIN AMERICA**	**ASIA PACIFIC**	**EMEA**	**WEBSITE: www.jdsu.com/test**
TEL: 1 866 228 3762 FAX: +1 301 353 9216	TEL: +1 954 688 5660 FAX: +1 954 345 4668	TEL: +852 2892 0990 FAX: +852 2892 0770	TEL: +49 7121 86 2222 FAX: +49 7121 86 1222	