

Researched and written by:

A-TEAMGROUP

TRADING BEYOND THE HORIZON

Fragmentation Drives Multi-Market Execution

January 2010

An industry briefing prepared
by **A-Team Group** for



Introduction

In 2010, financial markets participants will continue to expand their trading activities as liquidity increasingly becomes fragmented, seeking alpha in new markets, best execution in dark pools, arbitrage opportunities across the order book and by implementing high frequency and complex, multi-leg, cross asset class strategies.

The successful operations – whether they be the proprietary desks of traditional broker/dealers, specialist high frequency and algorithmic traders, or quantitative hedge funds – will leverage a trading infrastructure that combines high performance analytical, algorithmic and order routing platforms with the lowest-latency access to multiple, geographically dispersed execution venues.

Multi-market trading – leveraging a fragmented market landscape – introduces new challenges, even for trading firms that have mastered the complexities of low-latency execution using approaches such as co-location and proximity hosting. Those mechanisms, while still relevant, provide a less complete solution when trading across markets that are geographically dispersed.

New entrants into the market for connectivity and proximity services include organizations that are themselves market participants, such as sell-side firms offering sponsored access and DMA, and liquidity venues, which are now providing global order routing networks, in some cases channelling order flow to their competitors.

Those service providers join traditional players including telcos, hosting companies and value-added extranet vendors, who often bundle trading applications with connectivity.

The bottom line: For multi-market trading, optimization of long-haul and metro communications links, combined with smart use of an optimized co-location strategy, is an imperative for achieving the lowest latency, and this requires an understanding of connectivity offerings at a deep, granular level.

This industry briefing explains the drivers for fragmentation and multi-market trading, the evolving landscape of market access, and explores connectivity and hosting approaches to minimize latency.

The Drivers for Multi-Market Execution

The financial markets landscape is continually evolving, providing new opportunities for participants, whether they be trading firms, execution venues, or providers of systems and services.

A fundamental driver for this evolution is regulation, generally set by government agencies with the goal of growing the local economy, establishing increased competition and a fair trading environment, or underpinning political initiatives. Essentially, regulation sets the rules by which participants leverage the new market landscape.

Regulation also encompasses the fiscal and taxation policies of different nations, which can impact the business climate within. As an example, a tax on large securities transactions has been mooted by the U.K. Government, and even by some lawmakers in the U.S. (although the current administration does not support such a tax).

Evolution is also a factor of innovation in technology, which allows participants to implement business approaches to take advantage of changes to the market landscape, by “playing” regulatory rules to their maximum.

Recent years have seen the introduction of several major regulatory initiatives in the U.S. equities markets:

- Reg ATS – for Alternative Trading System, allowed for the creation of new regulated execution venues to challenge traditional exchanges.
- Reg NMS – for National Market System, established the concept of a nationwide network of interlinked competing execution venues, including a rule that prevents the execution of a trade at one venue when a better price is being quoted at another.

The Attack of the ATS

As a result, the landscape for trading equities in the U.S. has evolved rapidly and significantly, with much increased competition. In a large part, the speed at which individual exchanges and ATS venues are able to execute transactions is fundamental to the share of order flow that they attract.

The New York Stock Exchange introduced an electronic counterpart to its trading floor through its acquisition of Archipelago, becoming a hybrid market. Still, its share of U.S. equities volume was just 28% as of November 2009.

Meanwhile, rival Nasdaq acquired electronic systems in the form of Brut and Inet (formerly Island) as well as smaller regional exchanges, such as the Boston Stock Exchange and Philadelphia Stock Exchange. Its share of U.S. equities volume in November 2009 was nearly 25%.

Market share battles continue, and not only between NYSE and Nasdaq. Newer entrants like Bats Trading and Direct Edge are steadily increasing their market share – around 10% for each in November. Others include Getco Execution Services and State Street’s Lattice.

In fact, more than 40 North American ATS venues – including trading firms’ internal crossing networks and lower-transparency dark pools - are listed in the 2009 edition of A-Team Group’s *Alternative Trading Systems Directory*.

MTFs Mushroom

In Europe, a regulatory overhaul known as MiFID – for Markets in Financial Instruments Directive – came into force in late 2007 and, like Reg NMS, has led to increasing fragmentation across the region, with new ATS venues (often referred to as Multilateral Trading Facilities, or MTFs) being created.

MTFs are fragmenting a marketplace that was once the battleground of national exchange powerhouses, such as the London Stock Exchange and Deutsche Borse. Those traditional exchanges have responded by setting up or acquiring their own MTF and dark pool facilities:

- The London Stock Exchange acquired a majority stake in the Turquoise MTF from its investment bank owners, and plans to merge it with its existing Baikal dark pool, drawing on low-latency trading technology supplied by another recent acquisition, of MilleniumIT.
- Major U.S. exchanges are setting up in Europe, through the establishment of the likes of NYSE Arca Europe (facilitated by NYSE's merger with the pan-European Euronext market) and Nasdaq OMX Europe (the result of Nasdaq's acquisition of the Nordic OMX Group).

Other new MTF entrants include Bats Europe, Chi-X Europe, Equiduct, Burgundy and Quote MTF, adding to the circa. 30 ATS venues operating in Europe.

Usefully, trading technology vendor Fidessa last year introduced the Fidessa Fragmentation index (FFI), which indicates the average number of venues a trading firm should visit in order to achieve regulatory-demanded best execution. All indications point to increasing fragmentation, for both major index and small cap stocks, with index values between 2 and 3 for FTSE-100 index participants.

The Rest of the World

Outside of North America and Europe, notable market landscape evolution – leading to fragmentation – includes:

- The introduction by the Tokyo Stock Exchange of the Arrowhead trading system, capable of supporting high frequency and algorithmic trading. As a result, local ATS ventures (known in Japan as Proprietary Trading Systems, or PTSs) are expected to mushroom. Chi-X is looking to set up in Japan, to join existing systems from BIX, CBX Asia and Liquidnet Japan.
- Regulatory initiatives in Australia that will see the granting of new licenses for exchanges, opening up competition and giving a boost to local ATS operators AXE and Chi-X Australia. Elsewhere in Asia/Pacific, the Singapore Exchange and Chi-X Global have joined to form an exchange-backed dark pool to compete with private initiatives.
- New trading venues opening up in Canada, to compete with the TMX Group (Toronto Stock Exchange), including the Canadian National Stock Exchange, Pure Trading, Chi-X Canada, Match Now and Omega.
- The creation of BM&FBOVESPA in Sao Paulo, Brazil, forming a securities exchange that in trading volume approaches NYSE, Nasdaq and Toronto. A trading link with the Chicago Mercantile Exchange (the exchanges own a small portion of one another's equity) allows order routing between them.

The Net Net

The establishment of new liquidity venues, usually based on the latest compute and networking technology has driven down central latencies and caused traditional exchanges to respond by upgrading their systems. Sub-millisecond latencies are now the norm for central matching of equities.

Globally, the financial markets story is one of increasing fragmentation, evolution of traditional exchanges by introducing low-latency electronic systems, emergence of competitors in the form of ATS, MTF and PTS systems, and continued globalization for all participants.

Market participants are looking to leverage the increasing fragmentation and opportunities provided by multiple low-latency trading systems to support algorithmic and high frequency trading strategies. In particular, they are investing in high performance computational and transactional systems, and also in smart connectivity options in order to maximize their potential for advantage.

The availability of 'off the shelf' algorithmic trading packages, often from vendors offering managed connectivity, has lowered the barrier to entry for trading firms looking to trade in fragmented markets, and so has encouraged further fragmentation.

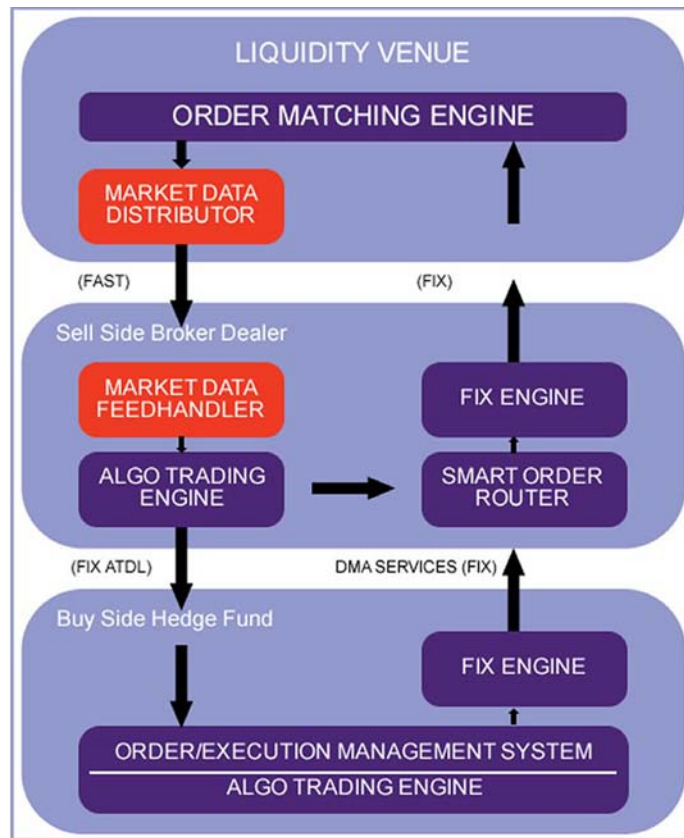
A-Team's Take:

Fragmentation is here and here to stay. Indeed, it is accelerating and becoming an increasingly global phenomenon. Leveraging fragmentation is a challenge but offers significant opportunities for those trading firms able to master the landscape of multi-market liquidity and execution.

Market Connectivity: The Changing Landscape

Execution of trading strategies by definition involves more than one market participant. Most commonly, trading involves an actual customer, which might be an investment manager or hedge fund (the buy side), an exchange broker/dealer member (the sell side) and an exchange (or other liquidity venue, like an ATS).

Traditionally, these three types of market participant would be physically and geographically separate, with their respective trading systems linked by various telecommunications services. A simplified overview is:



Co-location Catches On

As trading markets evolved, with new types of market participants, market fragmentation and the introduction of algorithmic and high frequency trading, so too have the operational and technical approaches that underpin those markets. Some developments have been:

- Liquidity venues moved their electronic systems to be hosted in highly secure and scalable data centers, with access to cheap power and high bandwidth connectivity.

As an example, in the NYC area, New Jersey has become a popular region for such hosting, where hosting companies including Telx, Equinix and Switch and Data operate data centers.

In particular, Nasdaq operates a primary data center in Carteret, NJ, while NYSE is constructing a data center in Mahwah, NJ, scheduled to go live in August. Bats operates from the Savvis data center in

Weehawken, NJ, while Direct Edge is currently located at Knight Capital’s data center in Jersey City, NJ and plans to move to Equinix’s Secaucus, NJ data center this year.

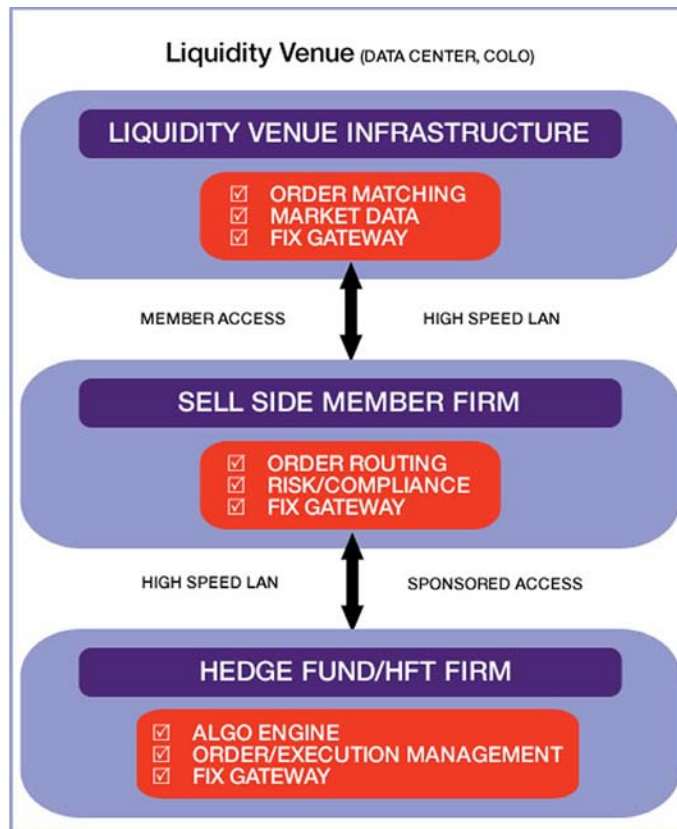
- Those liquidity venues, in partnership with the hosting companies that they leverage, began offering co-location services to their sell-side member firms, allowing them to site their trading systems in the same data center, connected via cross connects or risers (cable going from their switch across the room to the member firm switch).

This co-location, which has become a significant business for liquidity venues, essentially cuts propagation latency – that latency due to the distance data needs to travel – to almost zero.

- Sell-side firms introduced sponsored access services to their buy-side customers, allowing those customers to leverage the sell-side firm’s membership status to co-locate their trading systems at liquidity venues.

The overall result of these developments have been to move liquidity venues’ data centers from playing a supporting role, to a central one. “We really see our data centers as the future of our market,” said Stanley Young, CEO of NYSE Technologies and co-global CIO of NYSE Euronext, at a recent conference.

The establishment of co-location facilities, and their availability to all via sponsored access, has introduced an IT landscape that enables, even encourages, latency sensitive trading applications to flourish.



The Network Effect

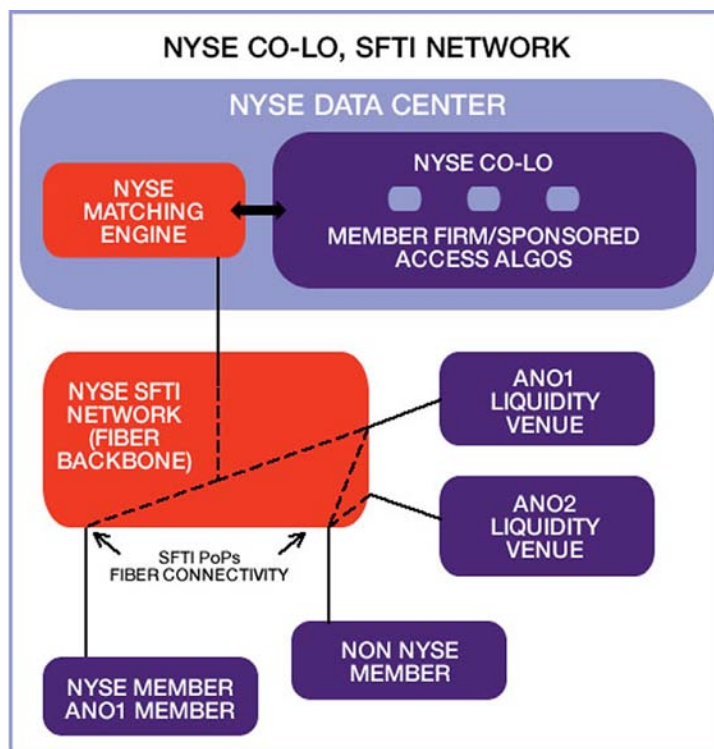
While co-location offers a clear benefit for trading firms, in many cases it is not practical or even desirable as a complete market access solution. In particular:

- Many trading firms require access to more than one liquidity venue to support trading in different securities and across different asset classes, in order to execute customer business, perform complex multi-leg transactions, or for hedging purposes. This tends to favor hosting at a data center that offers the best compromise in terms of latency to several venues.
- There is not always a business justification for investing in co-location to achieve the lowest latency for every liquidity venue. In general, choices are made with regard to where co-location is required, and where it would not be cost-effective. For example, a trading firm's high frequency trading activities in certain stocks would almost certainly call for co-location, while execution of many algorithmic strategies, and customer orders, is likely to be less latency sensitive.
- The costs of co-location can be significant, not only for initial deployment but also for ongoing operations, system and software upgrades and even physical moves, since liquidity venues are continually assessing their data center strategies.

As a result of this multi-market access requirement, there has been increased interest in establishing Central Proximity Hosting – where a trading firm's data center(s) is located in an optimal or best compromise location to connect to multiple liquidity venues – and, in particular, a number of data center operators have established facilities in order to satisfy multi-market connectivity in the NYC Metro area, in locations such as Clifton, North Bergen and Weehawken in New Jersey.

In addition to central proximity, trading firms often leverage order-routing networks, or extranets, providing connectivity to a number of liquidity venues. Such networks are often tied to software applications for trade execution, from the likes of SunGard/GL Trade and Fidessa, while others are from specialist connectivity providers, including BT Radianz, Fixnetix, Savvis and Transaction Network Services (some of whom also provide data center co-location services).

A notable entrant into the network space is the New York Stock Exchange, which is combining its Secure Financial Transaction Infrastructure – aka SFTI – with its new co-location facilities to offer access to multiple liquidity venues, including those that compete with the exchange's (once) core business.



To that end, the NYSE is actually refocusing its core business, as explained by Stanley Young: “We’ve actually changed our model. We’re moving from being a place where transactions occur to being what we call a fabric player. We allow

industry participants to meet virtually and we get paid for that. That model is now driving all of our thinking.” Such a shift in strategy from what was once the world’s leading equities exchange underlines some key market dynamics for the next decade:

- Fragmentation of trading markets will continue, and is a trend that even the biggest of the traditional exchanges cannot reverse.
- The primary business model emerging with regards to liquidity venues is one of who can provide the requisite high performance and low-latency access to it, along with value-added facilities, such as trading strategies and risk management.
- The actual business of providing liquidity in markets is becoming a commodity, with an increasing low barrier to entry.
- The successful trading firms will be those that can leverage the landscape of fragmented liquidity and make optimized choices for co-location/proximity hosting and connectivity based on their requirement for minimizing latency, time to market and price in support of their business strategy.

A-Team’s Take:

Connectivity – and its quality as measured by latency, flexibility, reliability, scalability and performance – is the key differentiator for any winning business trading in the world of fragmented liquidity.

Optimizing the Latency out of Market Connectivity

From the perspective of a trading firm, especially those engaged in algorithmic and high frequency trading, the imperative is speed. And, specifically, not how fast one's market connectivity is in absolute terms, but being faster than that of their competitors. (NOTE: they can be only 100ms faster and that is good enough - they just have to get there first)

It is also a case of being as fast as one needs to be. No two trading firms are alike. Each has their own business model, or models, and each will have their own low latency and connectivity requirements. Consider the latency profile for a (fairly typical) sell-side firm that is engaged in a just a few activities:

- For a select basket of U.S. equities, it engages in high frequency trading strategies. Execution latencies of single digit microseconds are required for this. The firm might also offer sponsored access to its customers, also seeking that low-latency access.
- The firm also runs proprietary algorithms designed to highlight and trade arbitrage opportunities in U.S. equities cash markets (in New York City) and futures index markets (in Chicago). Latencies in the order of a few hundred microseconds up to a few milliseconds are the current norm.
- For its institutional and hedge fund clients, who are engaged in fairly traditional executions to manage their portfolios, latencies of 10s of milliseconds are acceptable, though likely across a broader range of instruments and asset classes.

For such a trading firm, there is no single trading architecture to support its diverse businesses. Rather, the firm will need to blend a mix of co-location with one or more liquidity venues, NYC Metro connectivity to several venues, and long haul connectivity to the Chicago markets, either directly or via a third party data aggregator.

The Optimization and Flexibility Challenge

The challenge faced by many trading firms in 2010 will be to optimize their IT and connectivity infrastructures to support diverse business models within a fragmented, multi-market execution landscape.

Optimization is key to ensuring the best overall latency profile, so that requisite latency is achieved to each liquidity venue. Beating the latency goal at one venue, while missing it at one or more others is a net fail for a trading firm.

As well as optimization with respect to overall latency, flexibility to adapt to change needs to be built in. It is said that change is a constant in the financial markets, and in the world of liquidity venues, that is certainly true. Within even the confines of the U.S. equities markets, recent history has seen the emergence of new trading venues and new electronic markets, mergers and consolidations of traditional markets, physical movement of data centers, and the introduction of new co-location facilities, new order-routing and extranet services. This change is set to continue into the new decade.

One constant is that fast time to market in leveraging new venues and trading systems is a key goal, since capturing market share and leveraging early high-margin trading operations remains as fundamental to profitability as it's always been. Flexibility is key in reducing time to market.

The formidable challenge for trading firms, then, is to address these fundamental optimization and flexibility requirements.

The CFN, Low Latency Leader, Technology Response

CFN Services was founded in 2005 as a spinoff from CSX Fiber Networks to provide customized and turnkey network design, planning, deployment and managed services, drawing on extensive experience and knowledge of fiber network deployment since 1983.

As well as operating its own low-latency global exchange turnkey network, offering rapid connectivity – which currently connects Chicago, the New York/New Jersey Metro area, Washington, DC, London, Paris, Frankfurt, Sao Paulo, Tokyo, Hong Kong and Singapore – CFN is focused on providing customized connectivity solutions, designed to meet individual customer's requirements for latency, time to market, price, reliability and support for already existing network infrastructure.

Drawing on a wealth of expertise and experience, a custom CFN solution typically encompasses co-location/ proximity facilities, combined with optimized connectivity that leverages the most appropriate and lowest-latency lit and dark fiber optic spans available.

If one had access to a tool like CFN's FiberSource® (see below) and were to review the multiple fiber paths available from various providers, one would see that few are direct. In some cases, a single fiber provider may have a direct path between points A-B, but their solution between locations C-D is far from efficient. This is where CFN offers a unique advantage.

In addition to recognizing and utilizing the most direct routes from various fiber providers, CFN also integrates the best segments from providers and creates new, more direct paths between two locations. The resulting hybrid managed service is often over a lower distance and a corresponding lower latency. Indeed, a key differentiator for CFN is that its solutions are carrier agnostic, utilizing the best cable spans available to create ultra-low latency reliable networks.

CFN Managed Service offerings provide an evolutionary path so that bandwidth is available when required without a sizeable up-front expense, thus providing a trading firm with the flexibility to respond to market changes.

Management of optical networking and the equipment necessary to deliver the inherent advantages is a complex venture. Granted, given enough time and manpower, any firm can develop some expertise to design and deploy a fiber network.

However, the key issue is that, like trading strategies, one really needs to understand all the nuances, which are not obvious. There are many variables, such as optical loss due to Polarization Mode Dispersion (PMD) compensation, and chromatic dispersion, different hardware effects on latency, not to mention that a successful low-latency solution requires consistent monitoring, management and improvement plans.

Aside from available optical paths and technologies that CFN continually evaluates and optimizes, and the extensive network testing and compensation performed to deliver peak performance, the experience it has developed over 25+ years gives it the background knowledge and expertise necessary for the design and management of the networking components and the network itself.

Dual Network Operations Centers (NOCs) with 24/7/365 availability and the in-house optical engineering experts at CFN take the bulk of the network management off a trading firm's shoulders; enabling them to concentrate on the more important tasks of matching trading strategies to the most effective optical networking being provided by CFN.

By leveraging the CFN relationship, trading firms ensure that a coordinated optical networking and proximity hosting approach is combined with the most effective trading tools; and this comes with no additional OpEx for hiring additional resources to manage the optical networking component.

Now, let's review some basics of latency within fiber networks:

- Latency is largely a function of the speed of light, which is 299,792,458 metres/second in scientifically controlled environment; which would equate to a latency of 4.8 microseconds for every kilometre. But when measuring latency of data we need to account for the fiber optic cable, though it seems pure, it is not a vacuum so refraction of light needs to be accounted for. For measuring latency in long haul networks the calculation is actually 4.9 microseconds for every kilometre. In shorter metro networks, the latency performance rises a bit more due to building risers and cross-connects and can be as high as 5 microseconds per kilometre.
- It follows that to calculate latency of a connection, one has to know the distance travelled by the fiber, which is rarely a straight line, since it has to traverse geographic contours and obstacles, such as roads and railway tracks, as well as other rights-of-way.
- Due to imperfections in the fiber, light degrades as it is transmitted through it. For distances of greater than 100 kilometres, either amplifiers or regenerators need to be deployed. Accepted wisdom has it that amplifiers add less latency than regenerators, though in both cases it can be highly variable, and so needs to be taken into account. In particular, legacy spans are more likely to make use of higher latency regenerators.

CFN of course goes beyond those basics as part of its network design and planning, looking at all the variables of the network latency, optimizing them in a manner that compliments a firm's trading strategy and back-end infrastructure.

- It is a popular misconception that dark fiber (cable that has been deployed by a telco or other network/extranet provider, but which has not had networking equipment attached) always has the lowest latency between two points. This can be true in some instances, but there are many instances when a lit circuit is the better choice when low latency is an absolute priority, due to a shorter circuit path dark fiber availability and/or better equipment configuration.

In practice, CFN has found that in a few cases, routing of available dark fiber added considerable distance compared to lit alternatives, and so may exhibit overall higher latency, even taking into account latency due to networking equipment on lit spans.

Another factor is that for long-haul distances, dark fibre is less available than it is in metro areas, which might limit its use.

- For lit network spans over long haul distances, it is common for network operators to deploy Points of Presence (PoPs) to allow network access for communities along the route, hence improving the economics of operating the cable. The downside, as CFN highlights, is that such PoPs add latency, and so are not desirable for those seeking the lowest end-to-end latency.
- The network topologies available for connectivity between endpoints – for example, a trading firm's data center and that of a liquidity venue – are also important considerations for CFN when evaluating overall distances and latency due to networking equipment.

In general, mesh networks, where endpoints are directly connected, are most desirable as they offer the shortest routing, but they may not be available for particular connectivity requirements.

Leveraging ring networks, which connect several endpoints might be a cost-effective, available approach, but here it is important to understand routings, network protocols used, and the potential for bottlenecks in shared spans.

In many instances, a hybrid metro network is the best solution when examining the variables of cost,

latency and time to market.

- CFN understands the impact of different network protocols in use over lit spans, in terms of both economics and latency. Layer 1 protocols, such as SONET, exhibit the lowest latency but do not support shared use of spans. As such, costly bandwidth might not be utilized.

On the other hand, layer 3 packet protocols allow sharing of spans, but this inherently increases latency and can lead to bottlenecks if peak bandwidth requirements are not built into the planning.

A CFN solution allows traffic to remain at layer 1 whether a fully meshed point-point or a ring topology is used, so removing unnecessary latency.

- CFN also embraces the concept of Central Proximity Hosting, where a data centre location and associated trading applications are optimized for both the termination of long haul circuits and for establishing the most effective metro networking solution. This allows latencies of individual connections to liquidity venues can be matched to trading priorities and strategy.
- CFN's expert staff leverage FiberSource®, a unique knowledge base of telecommunications infrastructure, data centers, proximity and co-location sites, using it in the rapid identification, design, costing, and optional delivery of customized multi-carrier fiber-based networks and services. In particular, FiberSource® contains information on:

- Access to 550 Global Carrier Fiber Networks
- 350,000+ miles of fiber routes
- Longhaul and Metro Networks covering 100's of metropolitan market areas
- Access to 25,000 Data Centers and Lit Buildings
- Fiber Proximity data on more than 200,000 wireless sites
- 80,000 Tower and Rooftop Collocations Sites
- Visibility into existing Fiber served Sites
- Details of the wireless and wireline carriers that serve each site or collocate in the facility
- Central Office (CO) location and wire center information across the US covering 20,000+ LEC CO
- Fiber conduit information

With FiberSource® Financial Advisor, CFN can offer professional services, helping firms: plan network strategy and design lowest-latency solutions for specific routes, combining dark and lit fiber, interconnects and co-locations; network designs and cost estimates for connections between locations; all with an understanding of exchange, liquidity venues and market data provider locations.

Conclusions and Summary

Regulation and technology innovation continue to be the underlying drivers of change in market structure and the business models and execution strategies adopted by trading firms.

As a result, fragmentation of liquidity in the financial markets is increasing, becoming a global phenomenon and expanding beyond equities markets to other asset classes, such as futures, options, fixed income and foreign exchange.

Liquidity providers and those who provide access to it are engaged in continual competition to provide lower latency offerings, a phenomenon often referred to as the “low latency arms race.”

Providing access to this fragmented liquidity has become as important – and potentially more profitable – as providing liquidity itself.

Trading firms are increasingly seeking low-latency access to liquidity from multiple venues, to support both proprietary and agency business, but face challenges balancing their business goals with the investment required in low-latency technologies.

Intelligent selection of data center, proximity, co-location and connectivity services is vital in the construction of trading infrastructure to support profitable trading in fragmented markets.

Such selection requires deep technical understanding of: connectivity technologies and network protocols; a knowledge base of available dark and lit fiber, co-location and proximity centers and locations of exchange and liquidity center matching engines, market data provider and extranet PoPs; as well as broad experience of the practical pitfalls that can impact latency.

CFN contact

lowlatency@cfnservices.com
www.cfnservices.com



About CFN Services

CFN Services, the Low Latency Leader, is a managed telecom infrastructure services company providing network services for the Enterprise, Public Sector and Carrier Markets, specializing in network design, planning, deployment, and managed services, including: Low Latency Global Exchange Connectivity, Global, Regional and metro network design and cost optimization and mobile backhaul optimization.

CFN Services specializes in Data Center optimization ensuring the long haul network enhances the Enterprise distributed network strategy. CFN Services leverages FiberSource®, a global knowledge-based platform that can view all available dark and lit fiber, collocation, and lit buildings; providing the ability to quickly identify and design ultra low latency solutions. Learn how CFN Services can ensure you are Optimizing the Power of your Network www.cfnservices.com.

A-TEAM GROUP

About A-Team Group

A-Team Group provides a range of global online news and in-depth research reports focused on the business of financial information technology.

A-Team Group serves its global client base of IT and data professionals within financial institutions, technology and information suppliers, consultants and industry utilities with insight into the business of electronic trading, market data, reference data, risk management and the impact of regulation upon these industry segments.

A-Team Group's publishing division delivers more than just the news. A-Team knows how difficult it is to keep on top of all the news and developments in the industry so they monitor it for you and, leveraging the knowledge of their team of industry experts and extensive network of contacts, they bring you insight into the stories, trends and topics that matter.

A-Team's online news service, A-Team Insight, features industry segment Channels providing regular updates with analytical Insight articles that explore trending stories, as well as a comprehensive newswire covering all relevant announcements. In addition, A-Team produces a monthly PDF digest of the key news and features and a quarterly flagship magazine A-Team IQ. Coverage reaches across North America, Europe and Asia/Pacific. Subscriptions are required for full access to Insight articles, but free trials are available. Find out if you qualify for a complimentary subscription and sign up for a free 30-day trial at: www.A-TeamGroup.com/complimentary-access.

A-Team Group's research division provides industry professionals with focused and in-depth research offerings to better understand the specific uses of data and technology in today's trading and investment processes across the financial enterprise from front to back office. These include a series of topical white papers, survey-based research reports and focused directories (eg: algorithmic trading, valuations and alternative trading systems directories). Many of A-Team's research publications are available for free at: www.A-TeamGroup.com/site/research.

A-Team offers custom research solutions, commissioned by clients seeking answers to specific questions for in-house product development or marketing, or looking to support their marketing activities, promote thought leadership and generate sales leads. Find out how our custom research solutions can boost your marketing campaigns by contacting A-Team Group.

A-Team Group's events division produces a series of Insight Exchange events annually. These events combine A-Team's expertise in financial markets IT with thought leadership from world-class technology innovators and practical experience from financial market practitioners. For a schedule and more information, visit: www.A-TeamGroup.com/InsightExchange.

A-Team also partners with customers to produce custom physical and webinar events.

For more information about A-Team Group, visit www.A-TeamGroup.com.

Researched and written by:

A-TEAMGROUP

For more information about CFN Services Low Latency Global Exchange Connectivity Solutions
Contact:

CFN Services
The Low Latency Leader

Lowlatency@cfnservices.com
Americas Phone: +1-703-788-6633
<http://www.cfnservices.com>

Ask about our solutions for NY/NJ, Chicago, Toronto, Sao Paulo, London, Paris, Frankfurt, Tokyo,
Korea, Hong Kong and Singapore.

An industry briefing prepared
by **A-Team Group** for

