



Whitepaper

Inter-Party

Latency

Management

Contents

1. Who's Responsible for Latency?	4
2. Inter-Party Latency Management	5
3. Secure, Scalable Latency Measurement	5
4. Cross-Party SLA Verification	6
5. Collaborative Trouble-Shooting	7
6. CorvilClear™ – A Solution for Inter-Party Latency Management	9
7. Isn't it time you were able to look outside your own four walls? Now you can!	10

1. Who's Responsible for Latency?

No single party controls the full end-to-end path for electronic trading or market data delivery. Many members of the trading industry have invested substantially to reduce latency within their own infrastructure; but ultimately everyone remains dependent on partners to


Traders worry that important investment decisions about where to co-locate and which systems to upgrade might be misguided

ensure adequate end-to-end performance. Participants today have little visibility into sources of latency outside their own four walls. Traders can monitor order-entry round-trip latency, but finding out where this latency is incurred is much harder – in your own infrastructure, across a service provider's network, or in the market center itself? The latency experienced by unidirectional market data traffic is a particularly frustrating blind-spot for all parties.

This situation leads to nervousness, and sometimes to finger-pointing. In recent polls by low-latency.com, 55% of respondents felt that their market data providers should do more to reduce latency, and 87.5% felt that market centers and exchanges are not pulling their weight when it comes to latency. These polls were informal but are corroborated by our own conversations with industry participants. Traders worry that important investment decisions about where to co-locate and which systems to upgrade might be misguided, or could be nullified by a lack of corresponding commitment from their partners. And they also worry that inequitable treatment might leave them at the mercy of other market players.

These perceptions follow inexorably from the current lack of end-to-end transparency. Without doubt, many industry participants recognize the importance of latency. Some companies measure and advertise latency figures for their own trading platforms. But these measurements are not standardized or made in a consistent fashion that allows

Figure 1
No single party controls the full end-to-end trading loop. How can partners work together to achieve superior latency performance?



different platforms to be compared. They are rarely viewed as impartial, and they do not reflect the full end-to-end latency that traders experience.

2. Inter-Party Latency Management

Confronting these challenges leads us to propose an inter-party latency management (IPLM) solution that market centers, providers and participants can jointly use to monitor latency and performance across their respective systems. By adopting a joint approach to latency management, all parties can overcome the technical obstacles that currently hinder end-to-end visibility. Additionally, this approach facilitates transparency between parties and promotes collaborative problem-solving.

Market providers enhance the value of their services by adding trader-to-venue performance visibility

These benefits translate into strategic advantages for each segment of the trading loop. Traders achieve superior execution quality and increased trading success, through better insight into trading latencies and their impact on algorithm performance. Market centers attract greater liquidity from latency-sensitive clients by increasing end-user confidence. Market providers enhance the value of their services by adding trader-to-venue performance visibility, and remove any risk of being seen as slow. All parties gain by reducing the time and effort spent pin-pointing and fixing inter-party latency performance problems.

3. Secure, Scaleable Latency Measurement

Measuring latency for messages travelling between two parties involves generating information about when and where each message is seen, and then exchanging this information so that transit times can be determined. The solution we propose is entirely passive – we do not modify or delay messages in any way, and the exchange of measured information occurs out-of-band. Our solution operates as a stand-alone system attached to each party's network, and requires no integration with existing systems or protocols.

Clearly the exchange of measurement data between parties must be secure. We use a policy-based approach to security, giving each party full control over who is allowed to access data, how they are required to authenticate, and which data they have access to. A party can securely peer with multiple opposite parties while ensuring that each of them has access only to data relating to their own messages.

Exchanging latency data will obviously consume network bandwidth. Since trading parties are often connected via expensive WAN links, we've designed our data protocols to be very lightweight. The added overhead is less than 4% of the traffic being measured. This ensures that latency measurement can run alongside trading and market data over the network infrastructure already in place, with no requirement for new bandwidth or special treatment.

We've also designed our latency measurements to be precise. Trading parties who are located in the same city can achieve sub-millisecond latencies while those who are co-located can reach the microsecond level. Our solution will measure latency in each direction across any

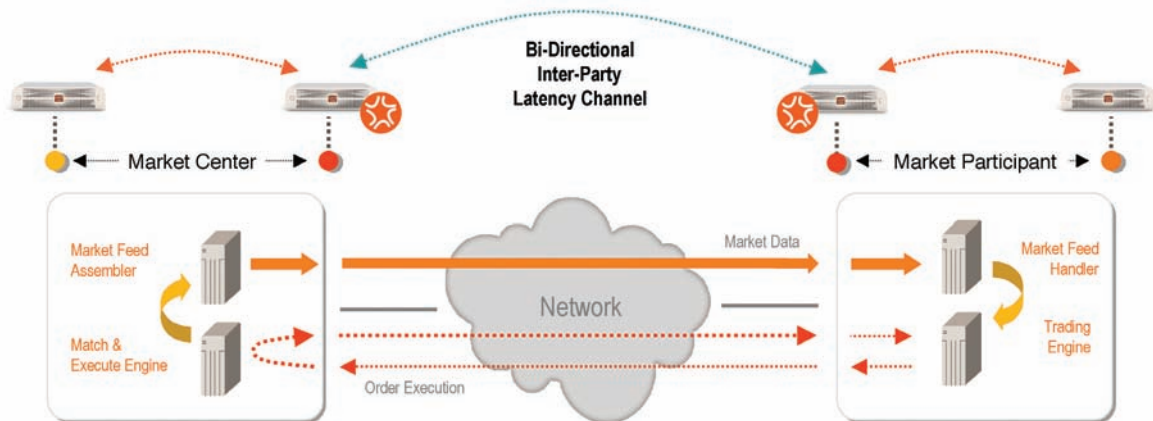


Figure 2
Corvil's inter-party latency management solution operates as a stand-alone system attached to each party's network, providing secure exchange of latency information over a lightweight protocol that runs alongside existing applications.

domain with microsecond precision. Uniquely, we do not require any form of external clock synchronization to achieve this – thus avoiding the deployment headaches associated with accessing GPS or CDMA signals for the purpose of clock sync.

4. Cross-Party SLA Verification

When a performance or connectivity problem crops up between two parties, it's not unusual for a debate to ensue about its nature, location and cause. Sometimes it's hard to agree about whether there even is a problem. Inter-party latency management can remove the need for this debate, because it provides each party with a complete set of underlying performance metrics and a common framework for deciding what is and is not acceptable performance.

Our IPLM solution lets partners define and verify comprehensive SLA targets for every packet or message travelling across each segment of the trading infrastructure. For example, two trading parties can easily determine whether or not a third intermediary party is achieving acceptable performance. Equally, a service or platform provider can use the technology to demonstrate attainment of agreed SLA targets to clients.

IPLM also facilitates an 'apples-to-apples' basis for comparing performance across multiple partners or service providers. To support such comparisons we provide a configurable centralized dashboard where key data from multiple inter-party measurement channels can be reviewed at a glance, and where channels can be ranked based on performance versus SLA targets.

Two trading parties can easily determine whether or not a third intermediary party is achieving acceptable performance

5. Collaborative Trouble-Shooting

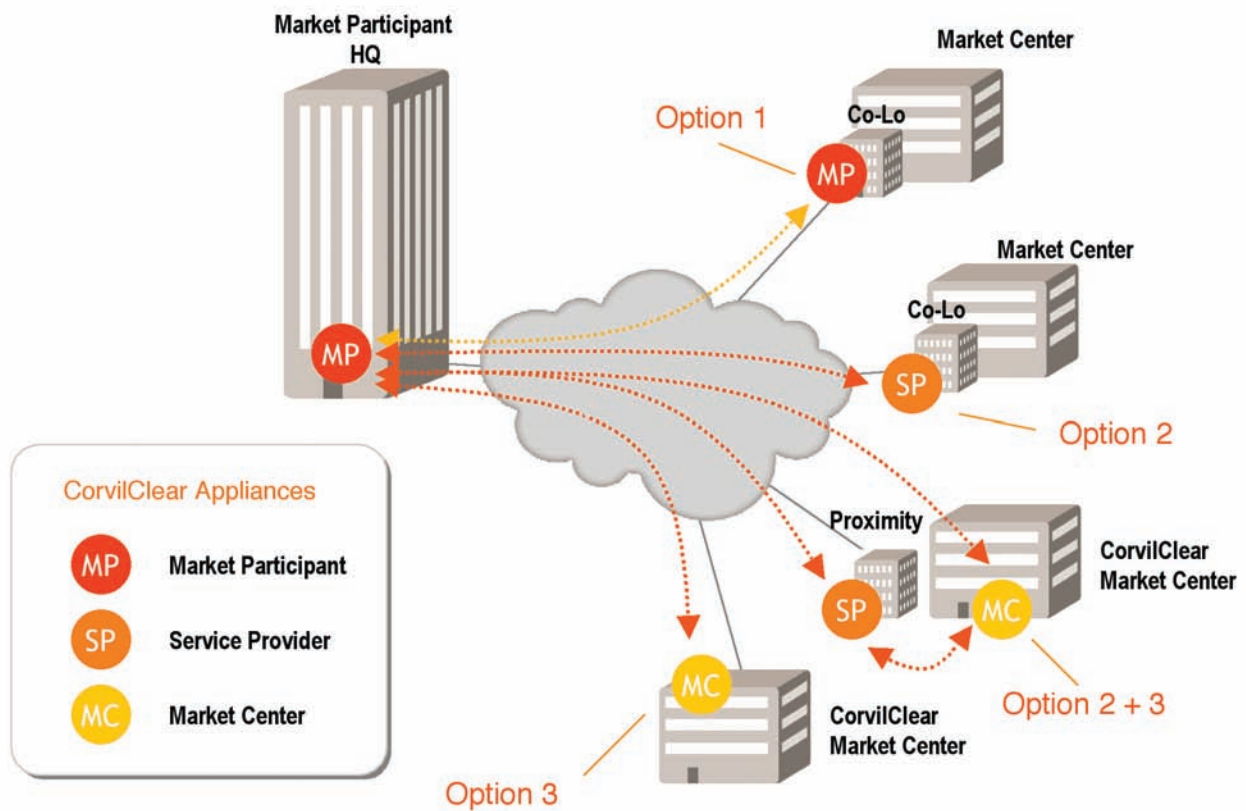
How much easier would it be to fix an inter-party performance or connectivity problem if both sides had access to a common set of advanced data capture and trouble-shooting analytics? As well as delivering end-to-end visibility to both parties, inter-party latency management can also provide access to equivalent sets of detailed data from the opposite sides of a connection, making it much easier for staff working together on each side to see where a problem is occurring and how it can be resolved.

To help users exploit this opportunity for faster trouble-shooting, we've included a range of data capture and analysis functions in our IPLM solution

Figure 3

Key benefits of inter-party latency management.

Inter-Party Latency Monitoring	<ul style="list-style-type: none">• Microsecond precision latency measurement using patented auto-clock synchronization algorithms suitable for multi-party environments where it is challenging to have common external clock synchronization.• Flexible lightweight method for exchanging latency and loss performance data over existing network infrastructure.• Policy-based security rules for safe sharing of performance data between one or many inter-connected parties.• Centralized and real-time configurable view of end-to-end performance for all peering parties.
Cross-Party SLA Verification	<ul style="list-style-type: none">• Verify SLA compliance of external services and infrastructure against latency and loss objectives for every packet or message.• Common language and rules for 'apples-to-apples' specification and comparison of performance across all parties.• Policy driven alarming automatically detects SLA violations along end-to-end market data and execution path.
Collaborative Trouble-Shooting	<ul style="list-style-type: none">• Common performance data set including packet captures with embedded per-packet end-to-end latency data.• Peering parties have access to a common set of advanced tools and analytics facilitating a common interpretation of diagnostic findings.



To help users exploit this opportunity for faster trouble-shooting, we've included a range of data capture and analysis functions in our IPLM solution. First of all, we provide an advanced form of packet capture that records every packet travelling across the network, and also records precisely how

We've linked our packet capture and SLA monitoring functions so that parties at either end can quickly identify SLA violations within the captured data

Figure 4

Deploying CorvilClear™. Traders can take the initiative themselves by installing CorvilNet™ appliances at co-location points (Option 1). Market service providers can use co-location to enable CorvilClear™ for multiple clients simultaneously (Option 2). Market centers can provide latency visibility to the core of their trading platform (Option 3). The last two options can be combined, enabling transparency to both the edge and the core of the trading venue and between service provider and market center (Option 2+3).

long it took to transit (and whether it got dropped along the way – measurements are provided for both directions). We've linked our packet capture and SLA monitoring functions so that parties at either end can quickly identify SLA violations within the captured data. Parties can examine their own

captured data in detail using our built-in event drilldown GUI, or they can filter and export the data in truncated or encrypted form, for use with other tools.

Our solution also incorporates a broad range of advanced analytics that both parties can use to assess network or application behaviour

Our solution also incorporates a broad range of advanced analytics that both parties can use to assess network or application behaviour. Examples include TCP performance measurements specially adapted for long-lived trading connections, message sequence number analysis for market data traffic, configurable microburst detection and bandwidth/resource requirements analysis, top talkers/listeners and conversations. We plan to keep on extending this range by working directly with our users to identify the analytics that they find most useful in the field.

6. CorvilClear™ – A Solution for Inter-Party Latency Management

Corvil's solution for inter-party latency management – CorvilClear™ – is implemented using the CorvilNet™ family of appliances and is already in use in the field amongst our customers. Figure 4 illustrates the different ways in which the system can be deployed. Market participants can take the initiative themselves by installing CorvilNet™ appliances at co-location points, and using them to support

Market centers have the option to enable transparency all the way to their matching engines and feed assemblers, by installing CorvilNet™ appliances in the core of their trading platform

CorvilClear™ latency management between the different market centers where they trade, and their other corporate sites. Market service providers can use a similar style of co-location deployment to provide latency management to multiple clients simultaneously. Market centers have the option to enable transparency all the way to their matching engines and feed assemblers, by installing CorvilNet™ appliances in the core of their trading platform. The ultimate form of visibility is possibly a combination the last two – CorvilClear™ to both the edge and the core of the trading venue.

7. Isn't it time you were able to look outside your own 4 walls? Now you can!

CorvilClear™ allows for the first time latency management outside your own four walls. You can quickly identify and remove latency blind-spots that currently afflict trading performance, while helping you to work more effectively with your partners.

You can quickly identify and remove latency blind-spots that currently afflict trading performance

This newfound transparency leads to higher confidence, lower risk and better control of costs. Already, traders are finding creative ways to exploit the latency information even further. For example, by leveraging detailed latency measurements to drive smart order-routing systems; or in algorithm back-testing for latency sensitivity using data sets that include real-world latency measurements.

We invite you to explore the new set of possibilities enabled by CorvilClear™.



Corvil Ltd, 6 George's Dock, IFSC, Dublin 1, Ireland

T +353 1 859 1000 **Tech Support** +353 1 859 1010 E info@corvil.com W www.corvil.com

Copyright © 2009 Corvil Ltd. Corvil is a registered trademark of Corvil Ltd.
All other brand or product names are trademarks of their respective holders.