APPLICATION NOTE

Synthetic Latency Measurement with TradeVision

Measuring WAN latencies between your trading network sites

Introduction

When it comes to electronic trading in highly fluid markets, access to real-time and accurate market data is critical, and nanoseconds count. The speed and sequence with which market participants place and execute orders at the matching engine of a given exchange venue depend on many technical variables. One of these variables is the latency in accessing real-time market data. The latency of trading networks continues to be a key operational parameter that is tracked by many network operations and network architects in the finance sector. There are many methodologies and places to measure latency – between an exchange and a remote site, between two points within a trading network, and the latency of the wide area network between two or more geographically separated sites. This application note explains in detail how Keysight's TradeVision – the leading solution in the world for market data quality monitoring, can be configured to generate synthetic packets and accurately measure the latency and jitter values between this "mesh" sites at a single-digit nanosecond resolution. This capability is known as synthetic mesh latency (SML).

The latency of trading networks continues to be a key operational parameter that is tracked by many network operations and network architects in the finance sector. Nanosecond-level latency of the wide area network between two or more geographically separated sites can be accurately measured using Keysight's TradeVision.

Global Network Latency Footprint – Intercontinental Exchange (ICE)

Figure 1 below examines a latency footprint recently published by ICE. It shows latencies between ICE Global Network consolidated feed ticker plants across major sites. As a person responsible for global network operations or network infrastructure, how do you ensure your real, measured latencies do not substantially deviate from nominal averages? Note that latency and contributing factors to quality vary considerably depending on the transport mechanisms employed between network sites, where most hosts fall into the wireless, low-latency, or standard category.

KEYSIGHT TECHNOLOGIES

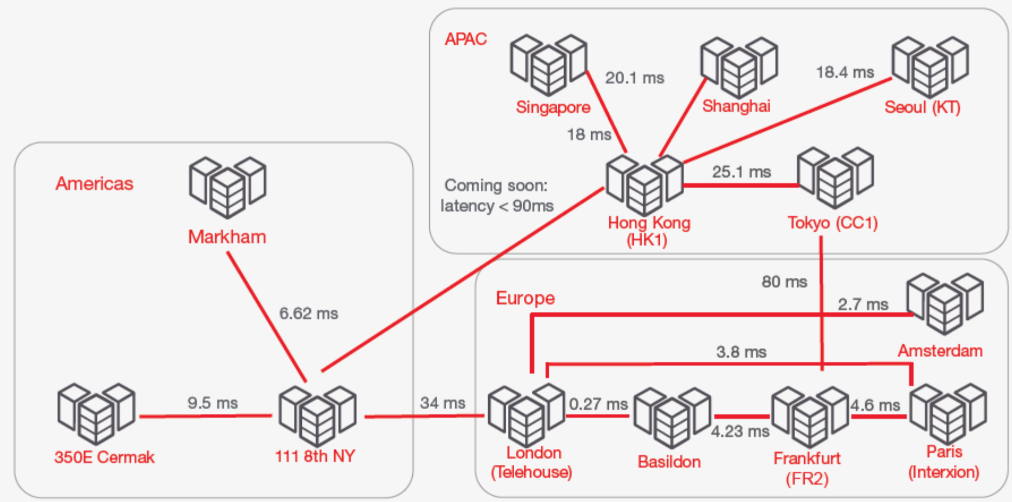

Figure 1. ICE global network latencies (source: Intercontinental Exchange, see https://www.theice.com/market-data/connectivity-and-feeds/network-topology-map)

Advanced Latency Analytics with TradeVision

Keysight's TradeVision is an advanced network visibility solution that serves as a market data analytics tool and a network packet broker. TradeVision allows IT teams to bring their market feed monitoring infrastructure together with network visibility management while allowing access to preprogrammed support for hundreds of trading venues. Easy to deploy, TradeVision detects sequence gaps and microbursts from more than 4,096 multicast channels in real time. It simultaneously monitors venue feed connectivity health continuously and provides visual dashboards and statistics with time stamping, accurate to the sub-microsecond. Most importantly, TradeVision supports advanced latency analytics, allowing the user to measure one-way application latency and jitter between the exchange venue and remote colocation sites throughout the network at various points of interest, as well as directly between colocation sites.

Getting into Details – Synthetic Mesh Latency

So when would you need to measure the latency footprint between consolidated feed ticker plants and shared colocation sites? Suppose the trading applications and analytics tools are deployed in the NY4 or Carteret sites, and the user is consuming MDP 3.0 market data from the Chicago Mercantile Exchange. The user would want the ability to continuously measure latency and jitter for all feeds and channels between these sites, across all transport paths: standard, low latency, and wireless.

TradeVision extends advanced latency analytics to accomplish continuous measurements by generating synthetic UDP packets sent between source and destination "mesh pairs" deployed throughout global colocation sites, as shown in Figure 2.

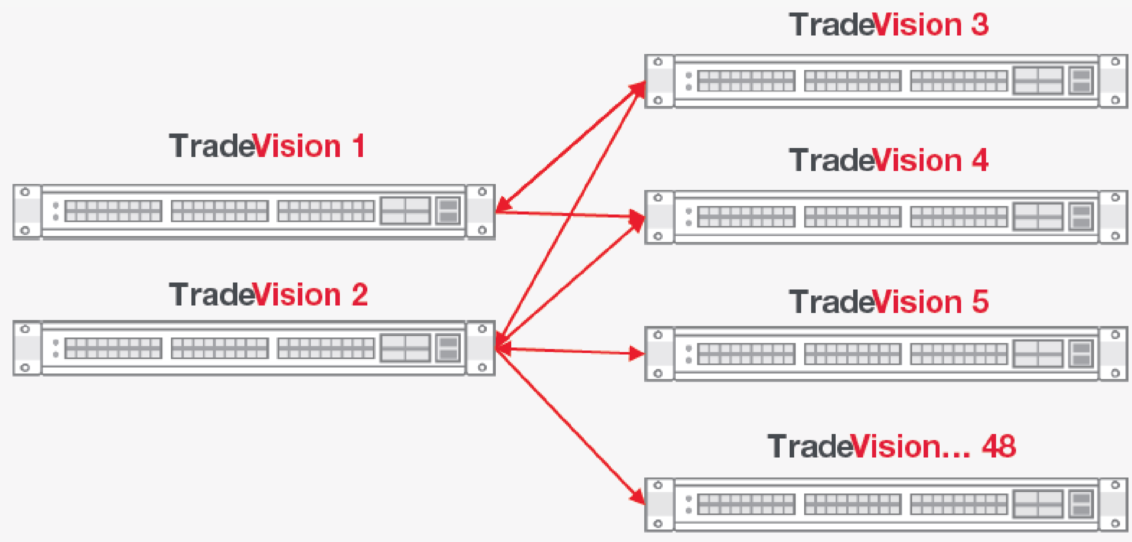

Figure 2. TradeVisions in a Synthetic Mesh Network

Each TradeVision appliance in such a deployment configuration generates synthetic packets with a unique nanosecond resolution time stamp inserted as a 15-byte trailer before the packet's frame check sequence. A single source appliance may be configured to send to as many as 48 destinations, leveraging the existing Layer 3 transport network. Since both devices in a given synthetic latency mesh pair sync via PTP to a GPS-connected grandmaster, you can make latency and jitter measurements on the destination TradeVision with the same visual dashboards employed for continuous network operations.

Figure 3 illustrates how the TradeVision source appliance connects to respective sites.

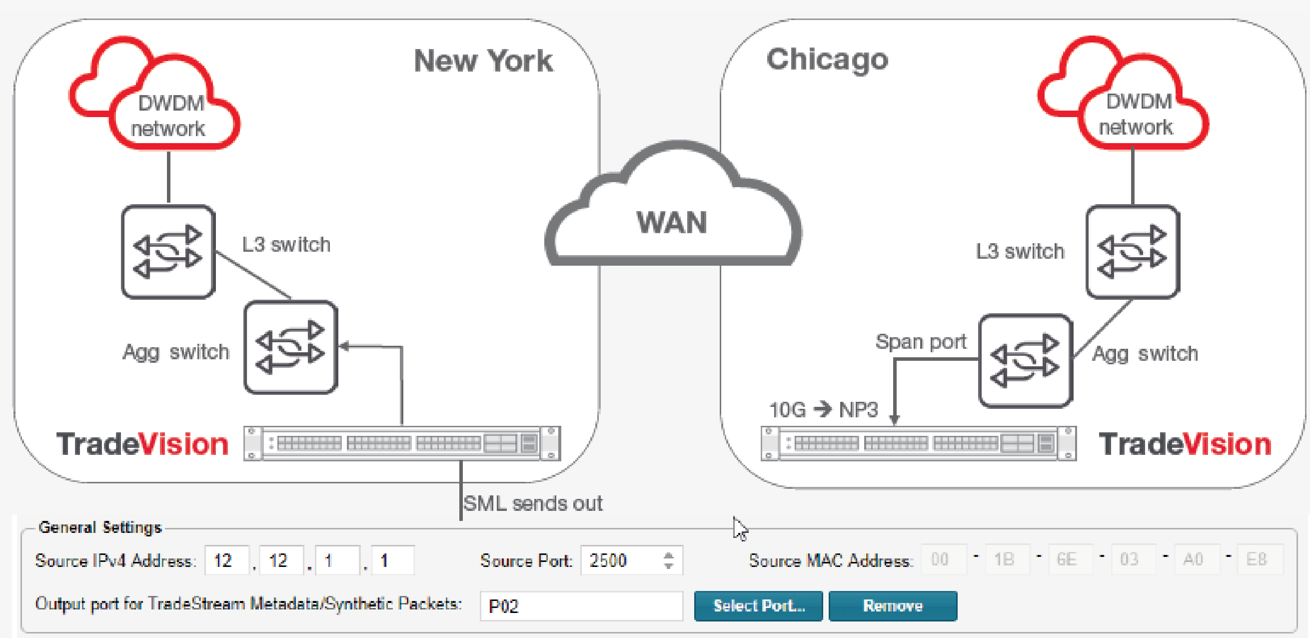

Figure 3. Synthetic Mesh Pair Deployment Configuration

The scenario above presents a high-level setup in two different locations (A-Side, B-Side) and how the latency is measured using synthetic mesh latency (SML) UDP packets in TradeVision.

On the source TradeVision unit, represented by the A-Side setup on the left hand side, you can configure up to 12 destinations (Ipv4: Port and Description) which are connected to an L3 switch You can configure the SML packet size, the number of packets, and the frequency to send packets. Each SML packet includes an Ixia PTP-synchronized timestamp trailer. It is important to note that PTP must be configured and synchronized in order to get accurate latency and jitter SML measurements.

The TradeStream output port is used to send the synthetic routable packets from the source TradeVision unit. The SML packet is routed through the network to a destination L3 switch, based on the configured SML destination IP address. At the destination L3 switch, represented by the B-Side setup on the right-hand side, there are two configuration options available:

- Static ARP for the IPv4 address matching the SML destination IP address on the L3 switch and Static MAC matching the SML destination MAC address on the L2 switch cabled directly to the TradeVision network port;
- Forward all IPv4 address:Port packets matching the SML destination IPv4:port pair to a mirror (span) port which is cabled directly to the destination TradeVision network port. The L4 port is the one that differentiates the SML packets originating from different TradeVision units.

On the destination TradeVision unit, you can configure the channel (IPv4 address:Port) that is specified for the source TradeVision SML message, and use the two-point latency (TPL) feature to measure latency and jitter for these SML packets from the source TradeVision unit to the destination TradeVision unit.

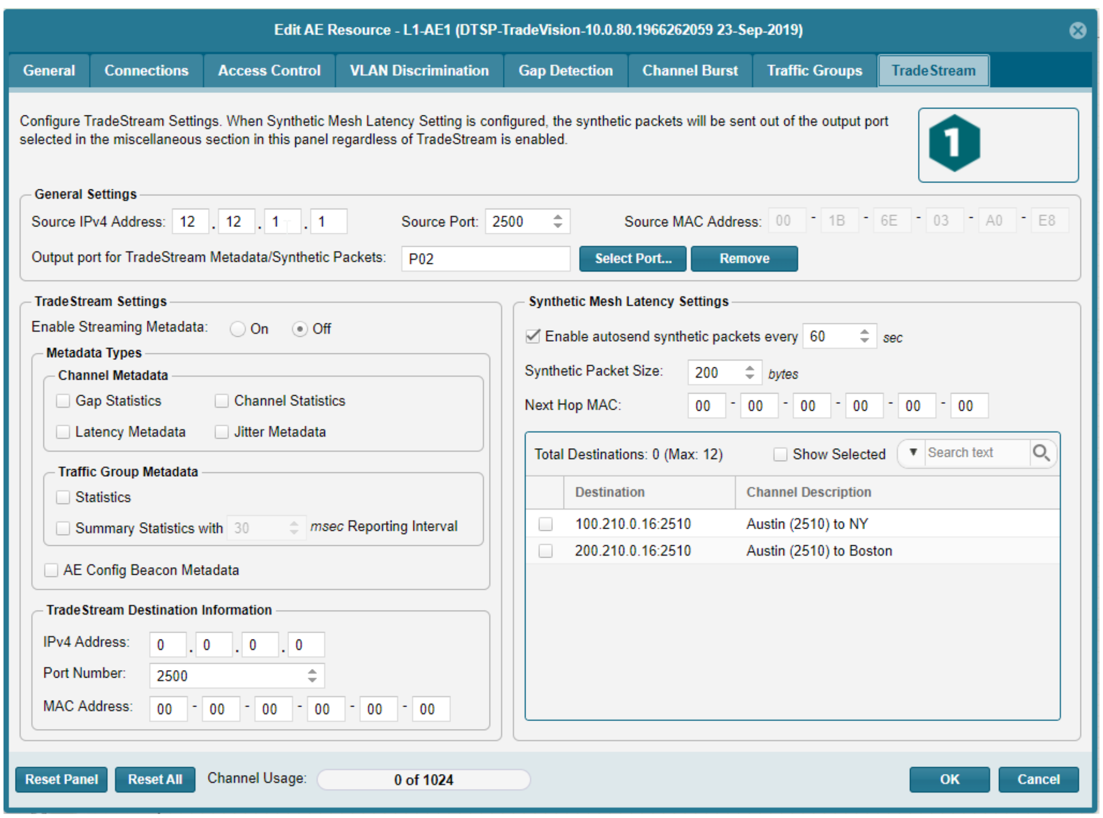

To generate synthetic mesh packets, configure TradeStream settings for each Analysis Engine (AE) of the source TradeVision unit as follows:

- Enable autosend synthetic packets every n sec: When enabled, one synthetic packet is sent out of the TradeStream output port at a certain interval. The valid range is between 1 and 60 seconds. The default value is 60 seconds.
- Note: You should not send synthetic packets at a rate slower than the Latency Sample Interval.
- Synthetic Packet Size: Size of the synthetic packet in bytes. The valid range is between 1 and 1,500 bytes. The default value is 200 bytes.
- l Next Hop MAC: Enter the MAC address of the first gateway on the path towards the remote TradeVision.
- In the Destinations section, you can select up to 12 customized mesh channels which you have already defined in the Custom Channel Configuration window (48 total per TradeVision appliance). At least one destination needs to be defined when Send Synthetic Packet is enabled. Each selected channel uses one channel from the 1024 channels available for each Analysis Engine (AE).

Custom Mesh Channels

Synthetic Mesh Latency uses custom mesh channels with unique IP:Port configuration per channel to send timestamped synthetic UDP packets between SML mesh pairs. A source TradeVision unit can send synthetic mesh UDP packets up to other 12 destination TradeVision units, therefore a maximum of 132 mesh channels can be added for each TradeVision unit.

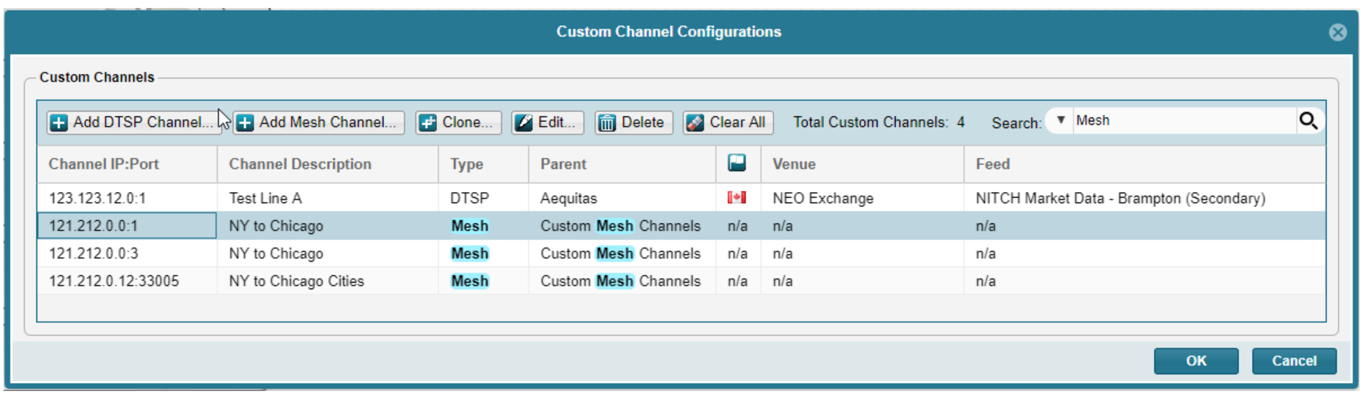

When measuring latency, it is recommended to use a description that would easily help identify the source TradeVision unit on the destination side. This description is post pended to the Channel IP:Port in the tree view for Latency and Traffic Groups.

Lastly, after configuring the customized mesh channels on the destination TradeVision unit, you have to select the Ixia Timestamp (15 byte trailer pre-FCS) option in order to correctly measure the latency on the mesh channels using TPL, because the synthetic packets from the source TradeVision unit have the Ixia timestamp inserted in the trailer.

Monitoring Mesh Network Latency with Live Dashboards

After your source and destination appliances are configured and you are generating synthetic packets, you can use TradeVision's dashboard to monitor the average, minimum, and maximum latency and jitter measurements in real time, and generate threshold-based alerts if desired. Measurements below were obtained in our lab environment using L3 production switches, showing min, max, average latency and jitter values, with three different switch configurations (direct fiber, static ARP, port mirroring).

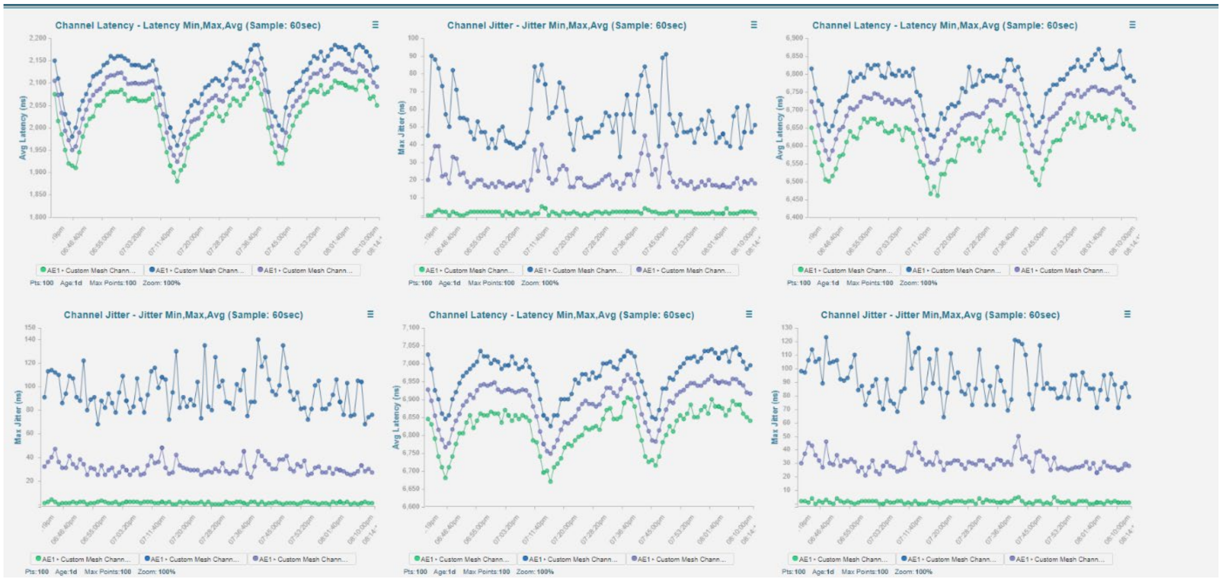

Beyond SML: Which Latency Measurement to Use Where

This application note has provided the background and solution details for measuring latency throughout the global network infrastructure using TradeVision's Synthetic Mesh Latency.

Looking beyond SML, it's important to at least have a high-level understanding of other advanced latency analytics capabilities and how to leverage these to meet your unique network operational monitoring needs. The following two tables summarize latency objectives, associate them with one of the three latency measurement options offered, and provide additional configuration insight for each.

Table 1. TradeVision latency analytics based on objective

Objective	One-way latency	Two-point latency	Synthetic mesh latency
One-way latency between exchange and a remote site	✓		
Latency between two points within the network		✓	
Inter-data center / colocation latency over WAN			✓

Table 2. Configuration details for OWL, TPL, and SML latency methodologies

Configuration	One-way latency	Two-point latency	Synthetic mesh latency
Support for external switch time stamps		✓	
Min # of TradeVisions needed	1	1 (when using external time stamps)	2
Maximum supported	N	2	48
Latency and jitter	✓	✓	✓

Learn more at: www.keysight.com

For more information on Keysight Technologies' products, applications or services, please contact your local Keysight office. The complete list is available at: www.keysight.com/find/contactus