WHITE PAPER

ixia

Packet Time Monitoring in Your Visibility Architecture

915-6953-01 Rev. A, December 2014

Table of Contents

Overview

More devices are connecting to more data from more sources. High user expectations for always-on access and immediate application response demands that IT organizations deploy a highly-scalable Visibility Architecture that eliminates network "blind spots"—where revenue-robbing and productivity-killing failures can grow unseen.

An increasing range of critical applications and services require timely delivery of network data such as:

Real-time interactive communication services:

- Voice over IP (VOIP)
  - Voice over LTE (VoLTE)
  - Videoconferencing and video chat
  - Virtual meetings
- Online gaming
- Connected financial services
  - Trading
  - Banking
- Industrial automation and industrial Ethernet
  - SmartGrid
- Time-sensitive networking (TSN)
  - Audio/Video Bridging (AVB)
- Mobile backhaul and synchronization
  - Precision Time Protocol (PTP)

High user expectations for always-on access and immediate application response demands that IT organizations deploy a highly-scalable Visibility Architecture that eliminates network "blind spots"—where revenue-robbing and productivity-killing failures can grow unseen.

In response to the need to monitor performance of timing-sensitive services and applications, innovative companies offer specialized tools for tracking this data. Many of these tools can be employed in a visibility architecture, but their success relies upon accurate timestamps of the monitored data. While timestamps may be added when packets arrive at the tool's network interface, this would require a tool to be deployed at every port where packets need to be monitored. Efficient visibility architectures may include access devices such as network taps, virtual taps, aggregators for traffic, and network packet brokers (NPBs) that feed data to the tools. In such a distributed architecture, it becomes critical for aggregation devices and NPBs to apply timestamps to enable effective use of timing analysis tools.

The Network Timing Domain

Before evaluating tools and methods to monitor timing-sensitive network resources, it is important to consider the vocabulary and concepts critical to the domain of network timing.

As a general principle, any timing analysis at the network packet layer begins with the introduction of timestamps associated with packets. Many methods of producing timestamps exist. Also, timestamps may be delivered with varying formats.

As a general principle, any timing analysis at the network packet layer begins with the introduction of timestamps associated with packets.

PCAP Timestamps

Timestamps may be included in an encapsulation of the packet, such as the common pcap format used to capture packets. In this case, the timestamp is included as part of the outer header encapsulating the packet, and the timestamp indicates the time at which that packet arrived at the encapsulating interface. Ordinarily this type of timestamp is indicated as Epoch Time, which is a count of seconds since a fixed point in time, such as the Unix Epoch, January 1st, 1970.

Protocol Timestamps

Timestamps may be inserted as an integral part of the packet, such as within fields in the protocol like Precision Time Protocol (PTP) or Network Time Protocol (NTP) which reserves fields for timestamp data. In this type of timestamp insertion, normally the timestamp indicates the time at which the packet was sent. These timestamps usually contain a seconds field and a subseconds field, offering a count of seconds since a fixed epoch such as the Unix Epoch January 1st 1970.

Proprietary Timestamps

Other timestamp methods are employed by various types of equipment, such as network switches and NPBs. Implementation may also vary, and given the proprietary nature, it is important to understand the method of timestamping that is used and the potential side effects of the method.

Key-Frame and Counter

In one set of proprietary timestamps formats, timestamps are appended to packets in a trailer, indicating a numeric value of a free-running counter. The rate of the counter varies by implementation, but for example it may run at 200MHz—which allows 5 nanoseconds for each count. A key frame is generated periodically when the counter rolls over, providing a real-time timestamp that can be used in conjunction with the trailer's counter value to determine the actual time of any packet.

In this type of implementation, it is typical for this trailer to be applied at the time the packet is transmitted from the switch or NPB, rather than on ingress. For this reason, the timestamp may not indicate the precise time that the packet arrived at the tap that was

used to gain access to the network. Some accommodation of the forwarding delay of the packet must be made. In some cases cut-through switches may be used that provide a relatively constant delay, and this delay can be compensated in the timestamp to give an indication of ingress time. In other cases of store-and-forward switches, changes in the packet length along with queuing effects can have a large impact on the accuracy of the timestamp.

One advantage of this method is that the timestamp trailer can be very small, since it only has to contain a count relative to periodic key frames. In the 5ns counter example, a 4-byte trailer will roll over and require a new key frame once every 24.5 seconds. Since these small numbers can be used, it is also less CPU-intensive to generate this type of timestamp compared with complete Unix-Epoch timestamps. One additional benefit is that given the small trailer requirement, little overhead is added to the data stream.

Full Timestamp Trailer

In some other implementations, timestamps are appended as a trailer to each packet indicating the real time of that packet relative to some fixed time (such as the Unix Epoch, January 1st 1970). These trailers typically contain a count of seconds and some fractional portion of seconds, such as nanoseconds. Considering these counters, a 32-bit data field is required for the seconds count, and another 32-bit field is required for the nanoseconds. This method has a great advantage that each and every packet is provided with its own unique timestamp, so decoding and using these timestamps is much simpler.

In some other implementations, timestamps are appended as a trailer to each packet indicating the real time of that packet relative to some fixed time (such as the Unix Epoch, January 1st 1970).

Since larger numbers are required to be calculated for each packet, line-rate timestamping with full epoch timestamps in every trailer may require a more powerful CPU and overall a more powerful NPB.

Some implementations using a full timestamp trailer provide the most accurate timestamping by generating the timestamp on ingress of the packet. In this case, any processing or forwarding delay introduced between the NPB's network port and the tool port which is connected to the tool will not add error to the timestamp. This is the most accurate method of adding timestamps, used in high-performance packet brokers.

Terminology

Some terms used to describe packet timing have specific meaning that may differ from their common conversational meaning. This can add confusion when discussing or specifying timestamping equipment and technologies.

The terms "accuracy," "precision," and "resolution" may be easily interchanged in ordinary conversation, but in the realm of network timing they each have a specific definition. It is important to clarify these terms since they directly impact the applicability of any network timing monitoring solution.

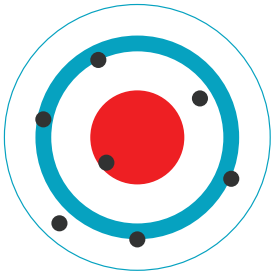

Low accuracy
Low precision

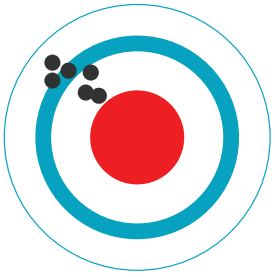

Low accuracy
High precision

High accuracy
High precision

Accuracy

Accuracy of timestamps produced by a NPB is the maximum difference between any timestamp and some traceable standard, such as Universal Coordinated Time (UTC). Measuring accuracy of timestamps produced by a NPB can be very difficult because it requires knowledge of the UTC transmission time of packets being measured. There is no practical way to evaluate the accuracy of timestamps in a deployed visibility network—they must be trusted wholesale. For this reason, the accuracy specification of a visibility tool is critical.

Accuracy of timestamps produced by a NPB is the maximum difference between any timestamp and some traceable standard, such as Universal Coordinated Time (UTC).

Prescision

Precision of timestamps produced by a NPB refers to the maximum difference in timestamp accuracy that is produced by the packet broker. So, for example, if no timestamp produced by the packet broker ever differs from UTC by more than 1µs, then it has 1µs of accuracy. But if all timestamps are within 200ns of one another, then it has 200ns of precision. To measure precision of the packet broker requires essentially the same challenging test procedure and specialized equipment as accuracy.

Resolution

Resolution is the minimum increment of time that the packet broker can measure between packets. Some clock or counter is usually used within the packet broker to mark time of a packet, and resolution is defined by this clock or counter. For example, if a 200MHz clock is used, then it counts every 5ns. The minimum resolution would therefore be 5ns. However, 10 Gigabit Ethernet (GbE) clocks operate at 156.25MHz, which means packets begin on boundaries of 6.4ns. The packet broker with a 5ns counter must round the 6.4ns cycle when a packet is received to the next available 5ns count.

Making Sense of Claims

Datasheets and marketing materials from NPB vendors often use oversimplified language when specifying the timestamp performance. It is easy to interpret the claims to mean something that they do not mean and, considering the difficulty of validating the claims, this has the potential for impacting the applicability to the solution. Some examples follow.

- 1 nanosecond precision – This may be interpreted to mean that the timestamp variability does not exceed 1ns. But in fact, what this normally means is that the timestamp itself can represent a 1ns numerical change. In other words, it usually means the timestamp contains a 9th decimal place.
- 100 nanosecond accuracy – This may be interpreted to mean that any timestamps produced are within 100ns of the UTC time at which that packet was received. However, the vendor may more typically be using the term "accuracy" to mean resolution, or in some cases they may mean precision. Often the number may be simply an indication of the potential accuracy of the system's time source, such as GPS, or the potential mathematical precision of the timestamp method, such as a 10MHz counter.

In order to make timestamps, a packet broker must be supplied with some time source.

Time Sources

In order to make timestamps, a packet broker must be supplied with some time source. The time source can define the baseline for accuracy of timestamping in that system.

Note: The time source used for packet timestamps in a NPB can define the maximum accuracy that the NPB could achieve, but achieving that accuracy depends on the remainder of the system design. Actual performance can vary considerably. The precision of timestamping is always determined by the NPB's system design, independent of the time source.

Time in networking equipment can be aligned in two ways. Synchronization means the system has real-time alignment with some traceable source such as UTC. Protocols such as NTP or PTP can provide synchronization, and GPS provides a source of synchronization around the world over the air. Syntonization is when the frequency or rate of a clock or counter is aligned with a reference. When comparing timestamps between sources or performing relative time measurements, such as a latency or delay measurement, syntonization is all that is required. In order to provide precise and accurate timestamps, both synchronization and syntonization are necessary.

There are a variety of time sources which can be supplied to a packet broker to facilitate timestamps.

Local Real-Time Clock and Oscillator

Most modern networking equipment includes a hardware real-time clock that can be set either manually by the user or is preset at the factory, and then time is kept from that point forward by a local oscillator. The method used for setting the local time indicates the quality of synchronization, and the oscillator type and quality defines the accuracy of syntonization.

For example, one particular device may contain a real time clock (RTC) that is set at the factory using NTP to within 10ms of UTC, and it may be equipped with a Stratum 3 oscillator with 4.6ppm of accuracy, which means it may accumulate up to 4.6µs of error

each second. Timestamps created using this time source may have good precision but the accuracy is unreliable since it begins to deviate at a rate of about 3 seconds per week from the very moment it is built.

Pulse Per Second (1PPS/PPS)

One method of providing time synchronization across networked equipment is 1 Pulse Per Second (1PPS or PPS), which sends an electrical pulse at the "top" of every second synchronized with UTC. If the RTC is set within +/- 0.5 seconds of UTC and then maintained by synchronizing with a 1PPS pulse, then very good time synchronization, or accuracy, is possible. During the time between 1PPS pulses, sub-second time values are derived based on a high-frequency counter or oscillator, which will define the limit of precision. Combining 1PPS with a high-precision oscillator is therefore critical.

Global Positioning System (GPS / GNSS)

GPS satellites contain atomic clocks and transmit very accurate time, which can be easily received by GPS receiver hardware that is embedded into networking equipment. Coupled with the right oscillator and quality design, GPS can provide 100-200ns accuracy to UTC as a time source.

Precision Time Protocol (PTP)

PTP, also known as IEEE1588 or simply "1588", is a network protocol specifically designed to synchronize networking devices across wide areas. PTP is growing in popularity in datacenter equipment because it provides superior accuracy to NTP, while not requiring specialized cabling or antennae as would be necessitated by GPS or 1PPS. PTP can operate over ordinary network infrastructure. The potential accuracy of PTP as a time source depends almost entirely upon the network and traffic over which it is deployed. In certain conditions, PTP can provide +/- 1µs of accuracy to UTC (about 10x that of GPS). The accuracy of PTP, however, is variable depending on network load.

Network Time Protocol (NTP)

NTP is the classic method of synchronizing time over computing devices and networks, and it is ubiquitous throughout the entirety of the Internet. NTP can provide time accuracy on the order of 10s of milliseconds, which may be sufficient for event notification such as SMTP or syslog, but ordinarily is not sufficient for monitoring time-sensitive network services or applications.

NTP is the classic method of synchronizing time over computing devices and networks, and it is ubiquitous throughout the entirety of the Internet.

What To Monitor

Tools that monitor performance of timing-sensitive applications and services on the network can use timestamps to calculate different metrics or Key Performance Indicators (KPIs).

Latency

Latency is a classic network performance metric, which at the basic level requires the evaluation of timestamps applied to the same packet as it passes through two locations in the network. By comparing the timestamps, the latency of the network segment can be monitored.

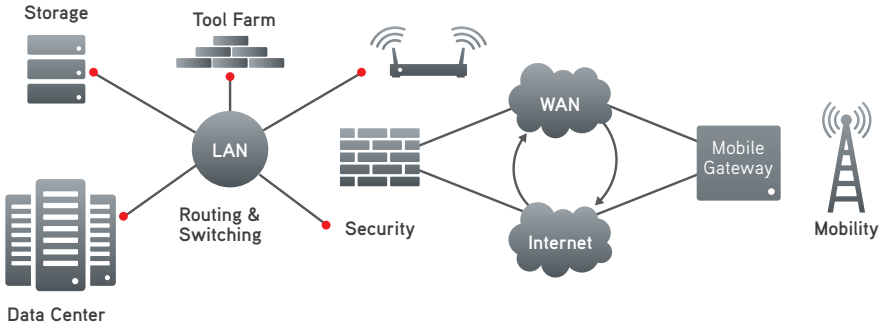

Latency is a classic network performance metric, which at the basic level requires the evaluation of timestamps applied to the same packet as it passes through two locations in the network.

Many networked applications and services rely on low latency in order to function correctly. Monitoring latency on these services and applications is important to ensure QoE and support SLAs, as well as monitor for trends that may lead to degradation and failure of these networks.

Since two locations must be monitored in order to measure latency, then potentially more than one NPB must be used to create a timestamp at each of the multiple locations. Time synchronization is required in order to make such a measurement, and the accuracy of these timestamps to UTC is important. The Full-Timestamp Trailer method of packet timestamping might be preferable for latency monitoring, so timestamps on each packet can contain all of the information required to make the evaluation without having to refer back to a key-frame.

The accuracy and precision required for monitoring latency in a network depends on the application requirements, as well as the requirement of the tool.

- Voice over IP (VoIP), Voice over LTE (VoLTE) – In order to facilitate a two-way interactive conversation, voice applications over the network require low latency. For example, ITU-T G.114 recommends a maximum of a 150ms latency for each direction. Segments of the network could require monitoring for their contribution to the overall latency—especially considering that this limit applies across the entire data path including the Internet or a mobile backhaul network that lies between the endpoints of the conversation.
- Videoconferencing, Video Chat and Virtual Meetings – As companies continue to increase collaboration of globally dispersed teams, and end users increasingly move to videoconferencing and video chat services for their personal communications, the quantity of traffic on networks that requires low latency will continue to rise. To keep

productivity high, it is important to monitor the latency of these services. Latencies should be kept to under 300ms end to end, and monitoring segments of the network for latency introduced in this type of traffic can give early indication of problems that may arise.

- Real-Time Market Data – Many automated banking and financial systems rely on timely delivery of market data as well as other information that impacts market conditions. In this environment, even microseconds count. Specialized tools are adapted to monitor all aspects of timing-related effects on networked trading systems, and it is critical to deliver accurate sub-microsecond timestamp data to these tools in order to maximize their effectiveness.
- High-Frequency Trading – One other hot topic over the past few years in the time-critical network domain has been High-Frequency Trading. Latency is critical in this application, and equally critical is time synchronization. Packet brokers offering sub-microsecond timestamps for the continual monitoring of not only trading data latency but also system time synchronization are necessary.

Packet / Frame / Bit Rate

Certain applications require bit rate or packet/frame rate to be constant in order for them to operate properly. For example, network providers, IT infrastructures and content providers are becoming more and more reliant on high QoE associated with streaming media services such as video on demand to retain customers and revenue. Monitoring the bit rate or packet/frame rate is made possible by the addition of precise timestamps from the visibility architecture elements such as NPBs.

Certain applications require bit rate or packet/frame rate to be constant in order for them to operate properly.

Time Synchronization

Time Synchronization is important for certain networks and applications, and monitoring of time synchronization requires accurate real-time timestamps to be applied to the monitored data by the NPBs.

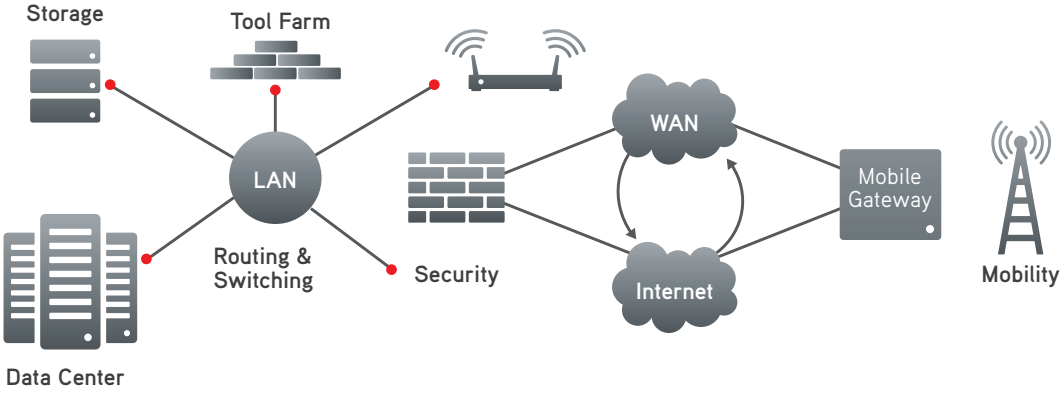

- Mobile Backhaul networks depend on time synchronization—especially for TDM networks and VoLTE—to facilitate call handoff, provide efficient use of spectrum, and maximize bandwidth. All of these add up to customer retention and QoE. Monitoring of real-time timestamps at the nanosecond level is therefore important to ensure QoE in mobile backhaul networks.
- Industrial Automation requires time synchronization provided ordinarily over GPS, PTP, and IRIG-B. Monitoring of this time synchronization at the microsecond level is necessary to ensure reliable operation of Ethernet synchronized automation networks.

- Smart Grid and Power Delivery is the next generation of management of power grids and power systems. Power operators use combinations of GPS and PTP to precisely synchronize power generation plants and user loads in order to efficiently and effectively deploy available power resources to where they are needed most. Sub-microsecond real-time timestamps on monitored packets are necessary to monitor synchronization of power networks.
- Audio/Video Bridging (AVB) uses PTP to synchronize audio and video devices across a network with tolerances in nanoseconds. Accurate, sub-microsecond timestamps must be provided by the NPB in order to monitor the effectiveness of AVB synchronization.
- The Internet of Things brings about networking of many different types of devices, many of which require real-time synchronization. In the growing world of connected devices, monitoring of time synchronization in the Internet of Things will become an important part of network visibility, including sub-microsecond real-time timestamps provided by the Network Packet Broker.

Bandwidth

Most current network visibility architectures provide some bandwidth statistics, but without time synchronization and accurate timestamping of packets, fractional errors in calculation of bandwidth may accumulate. When precise allocation of bandwidth is required, such as in mobile backhaul networks or service providers' SDN deployments, these errors in calculation can lead to reduced efficiency of network resource deployment. Sub-microsecond timestamps from NPBs allow for precise calculations of bandwidth utilization in networks.

Most current network visibility architectures provide some bandwidth statistics, but without time synchronization and accurate timestamping of packets, fractional errors in calculation of bandwidth may accumulate.

Events

Many events that occur on a network have some time dependency, and monitoring equipment should be equipped with the facility to report accurately on the timing of such events. This also provides captures or monitored packets to tools that can be correlated to collected events such as syslog or SMTP. Typically event monitoring provides millisecond resolution and real-time accuracy synchronized by NTP (UTC). Therefore timestamps on packets linked to these events provided by NPBd must be accurate within milliseconds of UTC and provide sub-millisecond precision.

Packet Order and Bursts

Certain applications require packets to be received in order, and monitoring or capturing packets to analyze for packet order is advised. Likewise, particularly in an environment where oversubscription of tool ports is in use or N:M aggregation is in place at the network layer, bursts can be present that are not readily identified at the tool due to re-queuing from aggregation. Accurate, sub-microsecond timestamps made at the ingress from the NPB can provide visibility into burst behavior and detailed packet order analysis in packet captures and tool port analysis.

Conclusion

As this brief overview shows, timing is a critical component of network visibility where time-sensitive applications are in use. While there are many ways to gain this visibility (many of them outlined here), working with a partner who has a holistic view of your needs is the best way to match the right technology to meet your goals and budget.

Ixia and our partners are the only providers today who can design, test, and monitor time-critical networks and applications no matter the scale of your network. Let us know how we can help you.

As this brief overview shows, timing is a critical component of network visibility where time-sensitive applications are in use.

Ixia Worldwide Headquarters
26601 Agoura Rd.
Calabasas, CA 91302

(Toll Free North America)
1.877.367.4942

(Outside North America)
+1.818.871.1800
(Fax) 818.871.1805
www.ixiacom.com

Ixia European Headquarters
Ixia Technologies Europe Ltd
Clarion House, Norreys Drive
Maidenhead SL6 4FL
United Kingdom

Sales +44 1628 408750
(Fax) +44 1628 639916

Ixia Asia Pacific Headquarters
21 Serangoon North Avenue 5
#04-01
Singapore 554864

Sales +65.6332.0125
Fax +65.6332.0127

915-6953-01 Rev. A, December 2014