

High Performance Capture and Replay

WHITE PAPER

High Performance Capture & Replay

Regulatory compliance and the increased need for network monitoring and analysis at higher speeds are driving the requirements for high-performance capture and replay solutions. These must be capable of capturing and storing packet data in real time at up to 10 Gbps without losing packets.

To date, the only solutions capable of providing this performance have been expensive, proprietary hardware solutions.

However, off-the-shelf server platforms have recently seen enormous increases in computing power with the proliferation of multi-core CPU's, coupled with ever-faster bus and memory architectures. Matched with powerful, standards-based PCI Express packet capture and analysis off-load adapters such as Napatech's NT20E 2 x 10 Gbps network adapter, such systems can now meet or exceed the performance of proprietary monitoring solutions, and do so at much lower cost. A case in point is the new HammerHead 10 Gbps Capture and Replay system from nPulse Networks.

This white paper provides an overview of the demands of high-performance capture and replay, along with some details on the nPulse HammerHead system, and the underlying Napatech network adapter technology which it employs.

Need for High Performance Capture & Replay Solutions

IP networks and the Internet in particular are fast becoming the de-facto platforms for not only communication services, but IT services in general. IP networks now form the basis for almost all forms of communication be it data, voice or video and with the advent of LTE will also form the basis for mobile communications. The Internet is a cornerstone of "cloud" strategies across the IT world, which will extend the influence of IP networks beyond communication to almost all other aspects of our lives.

Clearly then, the focus must be on ensuring that IP networks and the Internet are functioning effectively and can be relied upon to be there when we need them.



Over the last several years, a significant trend has developed that has gone largely unnoticed – the significant investment in network appliances dedicated to monitoring, analyzing, testing, optimizing and securing IP networks.

These dedicated network appliances are deployed at critical locations in IP networks to not only detect problems, but to avert them. This entails not only examining each packet in real time, but also examining trends across packets or streams of related packets, which we can refer to as flows. Indeed it is the ability to monitor flows that can help to predict potential issues or potential illicit activity.

Monitoring of packets and flows can be done in real time, but often it is desirable to record packet information for thorough analysis at a later time. This requires a solution that is capable both of capturing high-speed data in real time and recording this information to disk in real time. This information should also be stored in a manner that allows fast retrieval for efficient analysis. As one can imagine, capturing 10 Gbps or up to 15 million packets per second over a period of hours or days requires a lot of disk space and an extremely efficient means of sifting through this data to find exactly what you are looking for!

Current High Performance Capture & Replay Solutions

Solutions do exist that can perform real-time capture, recording and replay at 10 Gbps. However, achieving this performance has required a proprietary hardware design, which leads to higher costs. This is also reflected in the price of these systems, which can range from \$100,000 to \$250,000. This has been accepted since no other alternative was apparent.

However, nPulse with the assistance of Napatech, has achieved the seemingly impossible – a high-performance 10 Gbps packet capture, record and replay solution based on a standard PC server platform. The nPulse HammerHead capture and replay system significantly redefines the playing field and offers a previously unattainable level of performance at a far more attractive price point than available solutions.

HammerHead's performance derives from three critical components:

- Napatech's intelligent adapters for real-time network packet capture and analysis
- nPulse's own high-performance software for packet capture and recording
- The careful tuning of off-the-shelf, multi-core servers and system components to maximize data throughput

Challenges of High Performance Capture & Replay

To understand the challenge of high-performance capture, record and replay solutions, one has to consider the sheer mass of data that can be produced and how this affects system performance.

In typical IP network operation, if packets are not received at the destination, they can simply be resent. This is one of the advantages of IP networks – the information will get through one way or another. Certain packet data can be prioritized to ensure that it has a higher likelihood of getting to destination first time. This is important for real-time applications, such as VOIP, Video over IP etc. There is therefore a mix of traffic sharing the same connections with different priorities.

This data will also have different sizes. While the goal is to use as large a packet payload size as possible, small packet sizes are sometimes advantageous for particular applications. For example, VOIP works better with smaller packets.

The worst case scenario is that all the packets in the 10 Gbps connection are 64 bytes long resulting in up to 15 million packets!

For network monitoring and analysis, each one of these packets needs to be inspected. It must be done in real time and none of the packets can be significantly delayed or lost.

This means:

- Capturing each packet in real time without packet loss
- Time stamping each packet with nanosecond precision
- Recognizing the source, destination, protocols and application used to send the packet
- Deduplicating and filtering so that only relevant packets are processed
- Forming flows and efficiently distributing packets and flows to multiple server CPUs for processing
- Recording each packet to disk

With 15 million packets, this means that these tasks need to be completed for a packet every 67 ns!

Another challenge is the sheer amount of data that needs to be stored and efficiently finding data again for replay. At 10 Gbps in both directions, this amounts to 20 Gbps of data to be stored or approximately 2.5 Gbytes of data every second!

However, consider how difficult it is to find data in this vast warehouse of packets? How can this be done quickly and efficiently?

The final challenge is achieving all of these high-performance requirements without overloading the server. In fact, as little CPU processing power should be used as possible.

What is Acceptable Performance?

From a packet capture perspective, the only performance that is acceptable is full line-rate packet capture, no matter the packet size with zero packet loss. All of the packets need to be examined and stored for effective analysis. Each packet should also be time stamped with nanosecond precision so that an accurate timeline can be established for later event analysis.

The recorder needs to be able to take this data and then record it to disk in real time. Any slower than real-time means that packets could be dropped making it difficult to perform meaningful analysis of the data later.

Retrieval and replay of the data stored does not have to be in real time, but should not present a significant delay. A smart indexing mechanism is needed to find data quickly.

If these performance requirements cannot be achieved, then the overall throughput of the solution is lowered and the connections cannot be fully utilized without resulting in loss of data.

Current solutions on the market fall into two basic categories:

- High-performance proprietary hardware solutions
- Low-end standard hardware solutions

The distinction between the two is that the high-performance solutions can meet the above requirements, but only by using a customized, proprietary hardware design that is tweaked for high performance. Needless to say, these are expensive solutions. The low-end solutions are based on standard off-the-shelf server platforms, but cannot meet the performance requirements above, since they are based on standard network interface cards, such as server Network Interface Cards (NICs).

A better alternative is the one deployed by nPulse using Napatech network adapters. The nPulse solution combines the best of both worlds, providing high-performance while leveraging the cost advantages of standard off-the-shelf server hardware. The difference is Napatech's intelligent adapters for real-time network analysis.

High Performance Packet Capture

Napatech intelligent adapters for real-time network analysis are designed for insertion into any standard PC server. They provide the same electrical and optical Ethernet interfaces as standard network interface cards, with the same form factor and compliant with the PCI interface standard. However, that is where the similarity ends.

Napatech network adapters are designed specifically for high-performance packet capture and unlike standard network interface cards, Napatech network adapters provide full line-rate packet capture of all packets, no matter the packet size with zero packet loss. This is achieved with less than 1% CPU load due to the extensive off-load of data handling features. In addition, Napatech network adapters provide a rich feature set for packet decoding, including recognition of tunneling protocols, deduplication, on-the-fly reconfigurable packet filtering and efficient distribution of data traffic to multiple CPUs based on flows or balanced loading.

The following graphs provide an indication of the performance advantages that Napatech network adapters can provide. In order to understand the significance of the throughput figures shown, the following table shows the theoretical throughput limit for a 10 Gbps port:

Note that these measurements are presented for Ethernet frames, which carry IP packets in their payload.

As can be seen, as the size of the frame (and thereby the size of the packet) decreases, the throughput performance decreases. This is due to the increased overhead that the preamble and inter-frame gap between each Ethernet frame introduces at low Ethernet frame sizes. The preamble is used for synchronization while the inter-frame gap is a minimum idle period between transmission of frames. (defined as 9.6 nanoseconds for 10 Gbps Ethernet).

| | Maximum load at zero loss for different frame sizes (in bytes) | | | | | | |
|---|---|------------|------------|------------|------------|------------|------------|
| | Measured | 64 | 128 | 256 | 512 | 1024 | 1518 |
| 10G Ethernet (theoretical max throughput) | Gbps | 7.6 | 8.6 | 9.3 | 9.6 | 9.8 | 9.9 |

Theoretical maximal throughput for a 10 Mbps Ethernet port

For example, for 64 byte Ethernet frames transmitted at 10 Gbps, the preamble is 8 bytes and the inter-frame gap is 12 bytes. This means a total of 84 bytes needs to be transmitted for each 64 byte frame. In other words, only 76% of the transmitted data is Ethernet frame data.

The following graphs show the performance of Napatech network adapters compared to standard network interface cards in a standard PC server:

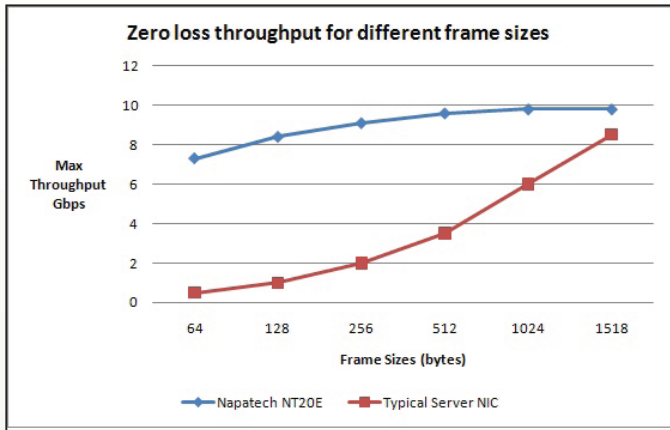


Figure 1a: Napatech NT20E throughput performance

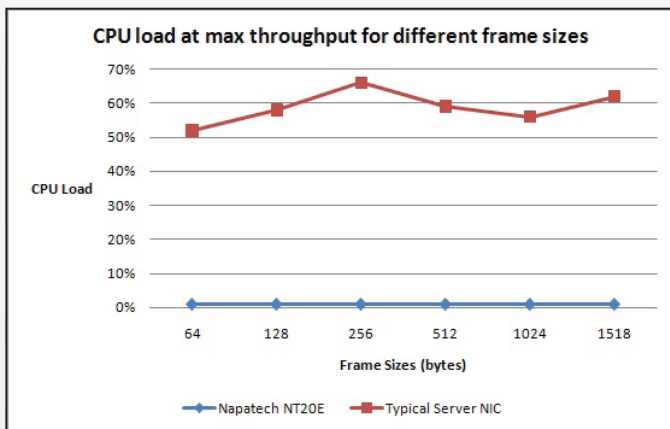


Figure 1b: Napatech NT20E CPU load performance

The Napatech network adapter can provide throughput performance at the theoretical limit, whereas a standard network adapter performance drops dramatically with lower frame sizes until less than 1 Gbps is supported without dropping packets.

At the same time, CPU load is almost zero as more of the data handling is performed on the adapter than by the CPU. As can be seen, at 10 Gbps, this reclaims almost 2/3 of the available CPU processing that can now be used by the network application analyzing the traffic.

The poor performance of standard network adapters is due to the fact that they are not designed for analyzing real-time data traffic. They are designed for communication where the

majority of the packets are large and where only packets that are addressed to the server interface are processed. The packets are processed directly by the CPU, which is notified every time a packet arrives. Since this is not a continuous activity in normal communication, this is acceptable and performance is not seriously impacted.

However, in real-time network monitoring, all traffic needs to be captured and processed. Neither the CPU nor the network interface card can handle this much traffic at one time resulting in very poor throughput performance as the number of packets increases.

High Performance Packet Capture and Recording

nPulse's HammerHead uses the capture abilities of the NT20E adapter to ensure that 100% of the packets from a 10 Gbps link being monitored are transferred into internal memory, with minimal loading on its CPU's. To store this traffic for later replay, and to make room for further captures, these packets must then be "recorded" to HammerHead's 16 TByte internal disk space.

Recording to disk at a sustained rate of 10 Gbps is a substantial technical challenge. Even the fastest server class disk drives can typically only sustain about 1 Gbps of read/write performance. HammerHead achieves a full 10 Gbps performance through sophisticated, parallel management of many physical drives operating as a single, highly-tuned file system.

Writing packets to disk storage, and reading them back again, at a sustained rate of 10 Gbps is a key part of HammerHead's value proposition, and is unique amongst systems built from off-the-shelf components. It is a critical enabler for accurate regeneration of captured traffic.

High Performance Packet Replay

Accurate regeneration of captured traffic requires packets to be retransmitted with the same time spacing as the original packet stream. This preserves the layer 2/3 traffic profile on the target network or device, and reproduces the original dependencies and impacts on applications at the higher layers.

Napatech's NT20E adds extra information, which is referred to as an extended descriptor, to each packet as it is captured, and before it is transferred to host memory. This includes a highly accurate time stamp (10 nanosecond resolution) as well as substantial meta-information about the packet such as flow identifying parameters (IP addresses, ports, and traffic type), protocol type, VLAN tags, MPLS tags and more.

When the HammerHead product replays captured traffic, it uses the capture time stamp to accurately reproduce the inter-frame gap of the original stream. The host application passes a packet stream to the adapter for transmission, and the adapter computes the delta in the time stamps of sequential packets, then waits that long before releasing each one, removing the extended descriptor in the process. In this way, the regenerated packet stream is a high-fidelity reproduction of the original real-world traffic.

Note that the other meta-data in the extended descriptor can also be used for advanced functions such as redistributing packets from a single merged capture stream across multiple transmit ports, again with the goal of preserving the exact flow and timing structure of an original multi-port capture. The flow and timing information of the extended descriptor can also be used to build a comprehensive index of large volumes of traffic data. For example, other nPulse appliance products built with the Napatech adapters use such an index to allow selective replay of only certain time windows or certain designated flows from an otherwise overwhelming capture of many Terabytes of traffic.

The Ability to Multi-task

One of the distinguishing features of Napatech network adapters is the ability to define flows and distribute traffic efficiently to multiple server CPUs. With multi-core CPU servers now available, the ability to fully utilize these architectures and perform more processing in parallel greatly enhances the overall performance of the network appliance.

Napatech network adapters have the ability to decode packets and recognize the key layer 2 to 4 protocols used. This provides information on the source, destination and application used to send the data. In addition, Napatech network adapters can recognize tunneling mechanisms used, such as VLAN, MPLS, GRE, GTP and SCTP.

Based on this information, the adapter is able to build hash keys that can identify related packets as being part of the same "flow". Multiple types of flows can be defined based on a combination of the header information and the types of protocol used. 17 different configurations are supported by Napatech, which can be mapped to specific protocols or types of packets allowing automatic selection of the right hash key type when a given packet arrives.

Based on the hash key, packets can then be directed to a specific CPU or set of CPU cores for further processing using a zero-copy Direct Memory Access (DMA) process. This places packets directly into the host memory of the CPU without the need to copy packets. This increases performance, lowers CPU utilization and lowers latency or delay.

How the CPUs are used and the traffic load is distributed is completely in the control of the user as the hash key mapping mechanisms and assignment to CPUs can be changed on the fly. This helps to ensure that all available processing power is used optimally.

Introducing the HammerHead

nPulse's HammerHead product is a line-rate 10 Gbps Capture and Replay system, designed for highly accurate reproduction of captured, real-world traffic. It can also optionally introduce malicious traffic into the regenerated stream.

With rapidly growing deployment of 10 Gigabit networks, and the proliferation of sophisticated network devices designed to work at these high speeds, realistic traffic testing is more important than ever. At 10 Gbps, even small errors or configuration issues can have major performance, security and cost impacts. But capturing, recording and generating 10 Gbps of traffic on a sustained basis are all difficult technical challenges, and previous approaches have relied on proprietary hardware and software. This has resulted in high-cost solutions beyond the reach of many of the developers and network managers who need them.

HammerHead solves this problem by using high-performance, standards-based, off-the-shelf hardware components, such as multi-core CPU servers and advanced disk subsystems, to dramatically reduce cost without sacrificing performance. nPulse also relied on Napatech's PCI-based NT20E network adapter to ensure 100% accurate packet capture and replay performance, further off-loading the system CPU to allow even higher application performance in areas such as the storage, manipulation and retrieval of multi-Terabyte traffic files.

The result is an affordable but powerful solution for ensuring maximum network, device and application performance through the capture and replay of real-world network traffic at full 10 Gbps line rates. Stress-testing networks in this way exposes potential weak points and bottlenecks that could starve network performance, introduce security vulnerabilities, or increase operational costs.



Figure 2: Hammerhead 10 Gbps Packet Capture & Replay

nPulse Networks' HammerHead provides state-of-the-art 10 Gbps test performance at a fraction of the cost of proprietary alternatives by carefully matching Napatech's PCI standards-based network adapter technology with powerful, off-the-shelf, multi-core server hardware and an advanced multi-threaded software architecture. Target applications include:

- Ensuring functionality of 10 Gbps networks and devices under full load
- Stress-testing network devices and services to identify performance limits
- Confirming that new network deployments deliver the performance and reliability required
- Verifying that security devices will block internal and external threats

For more information, visit www.npulsenetworks.com.



Figure 3: Napatech NT20E 2 x 10 Gbps network adapter

Napatech intelligent real-time network adapters address a broad range of network appliance applications ranging from network monitoring & analysis to network test & measurement, network security and network optimization. Napatech network adapters can be used to improve performance where there is a need to capture all packets in real time at high speeds without losing packets.

Napatech provides two families of network adapters:

- A capture family for packet capture applications where the network appliance is sitting off-line (e.g. network performance monitoring or intrusion detection)
- An in-line family for in-line applications where the network appliance needs to both receive and transmit packets (e.g. intrusion prevention systems)

For both families, 4 x 1 Gbps and 2 x 10 Gbps Ethernet network adapters are available.

For more information, see <http://www.napatech.com>

Europe, Africa and Asia

Napatech A/S
Tobaksvejen 23 A, 1
DK-2860 Soeborg
Denmark

Tel. +45 4596 1500
Fax. +45 6980 2970
www.napatech.com
nteusales@napatech.com

US West Coast and Americas

Napatech Inc.
650 Castro St., Ste. 250
Mountain View, CA 94041
US

Tel. +1 888 318 8288
Fax +1 650 618 1401
www.napatech.com
ntussales@napatech.com

US East Coast

Napatech Inc.
10 N E Business Ctr. Dr., Ste. 115
Andover, MA 01810
US

Tel. +1 888 318 8288
Fax. +1 978 824 9414
www.napatech.com
ntussales@napatech.com