LDA TECHNOLOGIES™

www.ldatech.com

NetRecap: Providing the Data to Answer "*Why?*"

WhitePaper

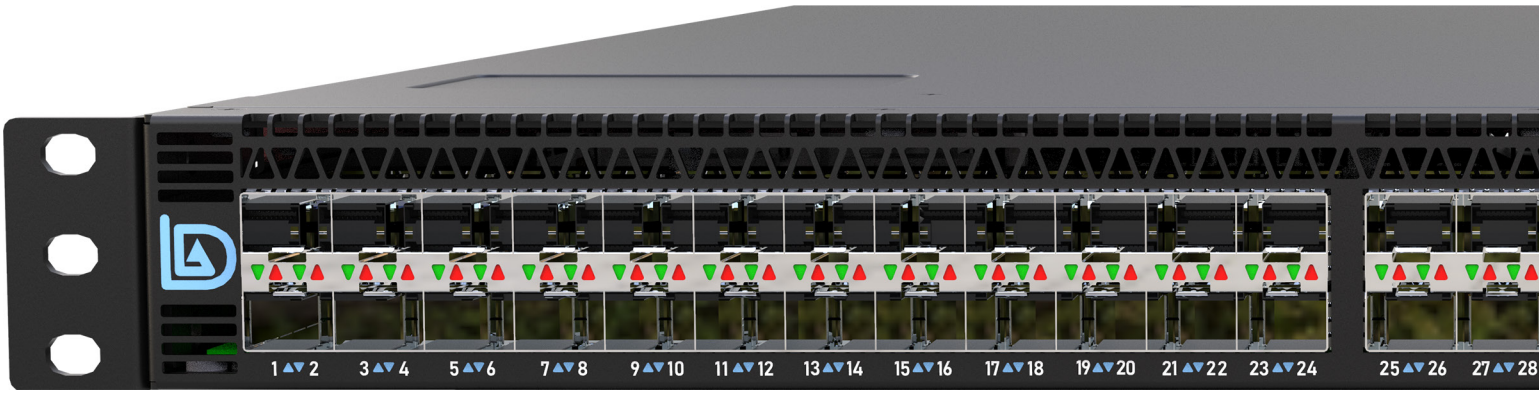

In “The Matrix” trilogy, the Merovingian, who is the all-knowing broker of information inside the Matrix, says that:

“’Why’ is what separates us from them, you from me.
‘Why’ is the only real social power, without it you are powerless.”

In networking, to know the “Why” requires that you capture every packet, with a highly precise timestamp, and its metadata in the form of VLAN tags. The LDA NEO with NetRecap can provide all of this information so your team can then determine the “Why” that will propel your business forward.

Introduction

At first glance, the LDA NEO looks like a network switch. Looks though, are deceiving; it is a fully programmable networking application platform executing on a server-class system with substantial resources. This enables NEO to not only route traffic like a switch, but it can also run multiple complex applications simultaneously. NEO performs many tasks in parallel on traffic such as Muxing, 40 GbE translation, bandwidth management, timestamping, VLAN tagging, packet cloning, and full lossless packet capture. In some use cases, NEO can compete exceptional well against existing dedicated network appliances while simultaneously outperforming them. It is not meant to replace these devices in all use cases as it has not been designed to address the many corner cases that might exist for these dedicated appliance products. LDA’s NetRecap is an application for the NEO reconfigurable network platform that enables timestamping, VLAN tagging, cloning, and in some configurations, full lossless packet capture. It collects all the data required for you to answer, “Why?”

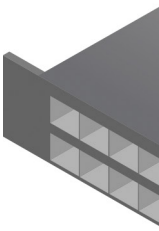

NEO, an Open Architecture Networking Platform

The LDA NEO Reconfigurable Network Platform is an open architecture built on an Intel Xeon server running Linux, with up to 128 GB of main memory, three U.2 NVMe internal drives, and a single PCIe x8 expansion slot. Being an open architecture allows LDA’s customers the flexibility to use whatever hardware capture card or generic high-performance Network Interface Card (NIC) the customer is comfortable with to perform network packet capture. Typically, this NIC is an Intel E810 that supports up to 100 Gbps of network capture. The LDA team has also tested and supports Napatech capture cards with the NEO platform for lossless packet capture. It is also possible to use other high-performance NICs in your NEO from Xilinx (formerly Solarflare) and NVIDIA (formerly Mellanox). Since NEO is an open platform, customers can run third-party advanced capture software from Napatech. The is also a software only capture platform utilizing DPDK from the leader in this space, Ntop. NEO’s ability to capture is limited only by the NIC’s capture capabilities and software selected by the customer. In Figure 2 below, we’ve provided a simplified architectural block diagram showing the layout of the NEO Reconfigurable Network Platform. All the colored components below are customer configurable.

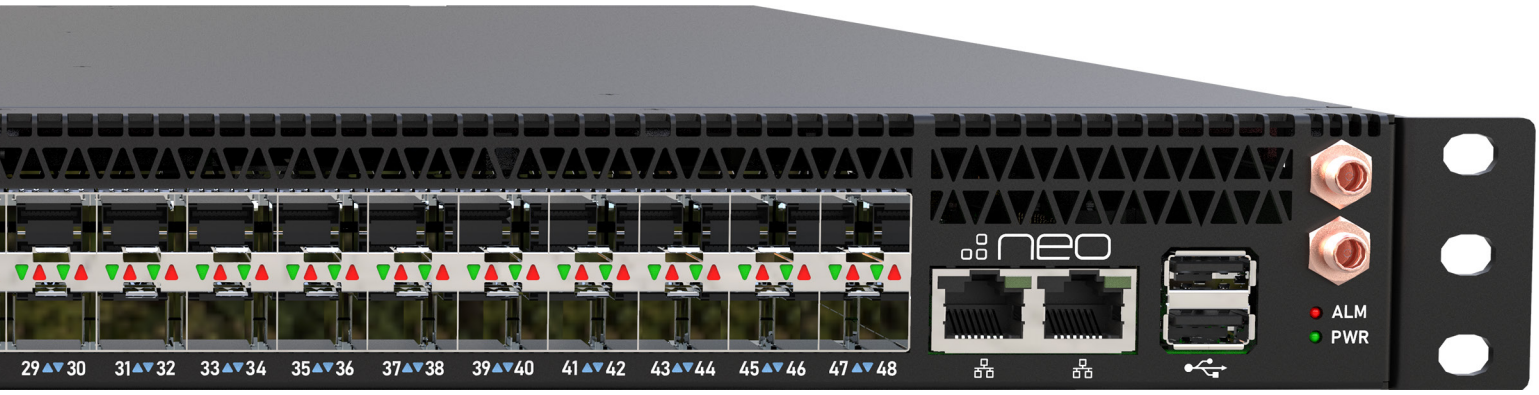

Figure 1 Front View - LDA Technologies NEO Reconfigurable Network Platform

Many other appliance vendors in this market provide only closed architecture solutions that require additional licenses to enable key features. These vendors often charge premium prices for additional "qualified" memory, storage, or networking modules. Often users are even barred from upgrading the hardware or installing their applications on these appliances. LDA expects customers to use the NEO as they see fit; they bought the platform; therefore, they should be allowed to install whatever hardware or software makes them comfortable. This includes currently unsupported capture or networking cards and third party memory and NVMe drives that have not yet been qualified by LDA. The flexibility and cost savings resulting from the NEO being an open platform is one of the main features that has attracted customers to the NEO platform.

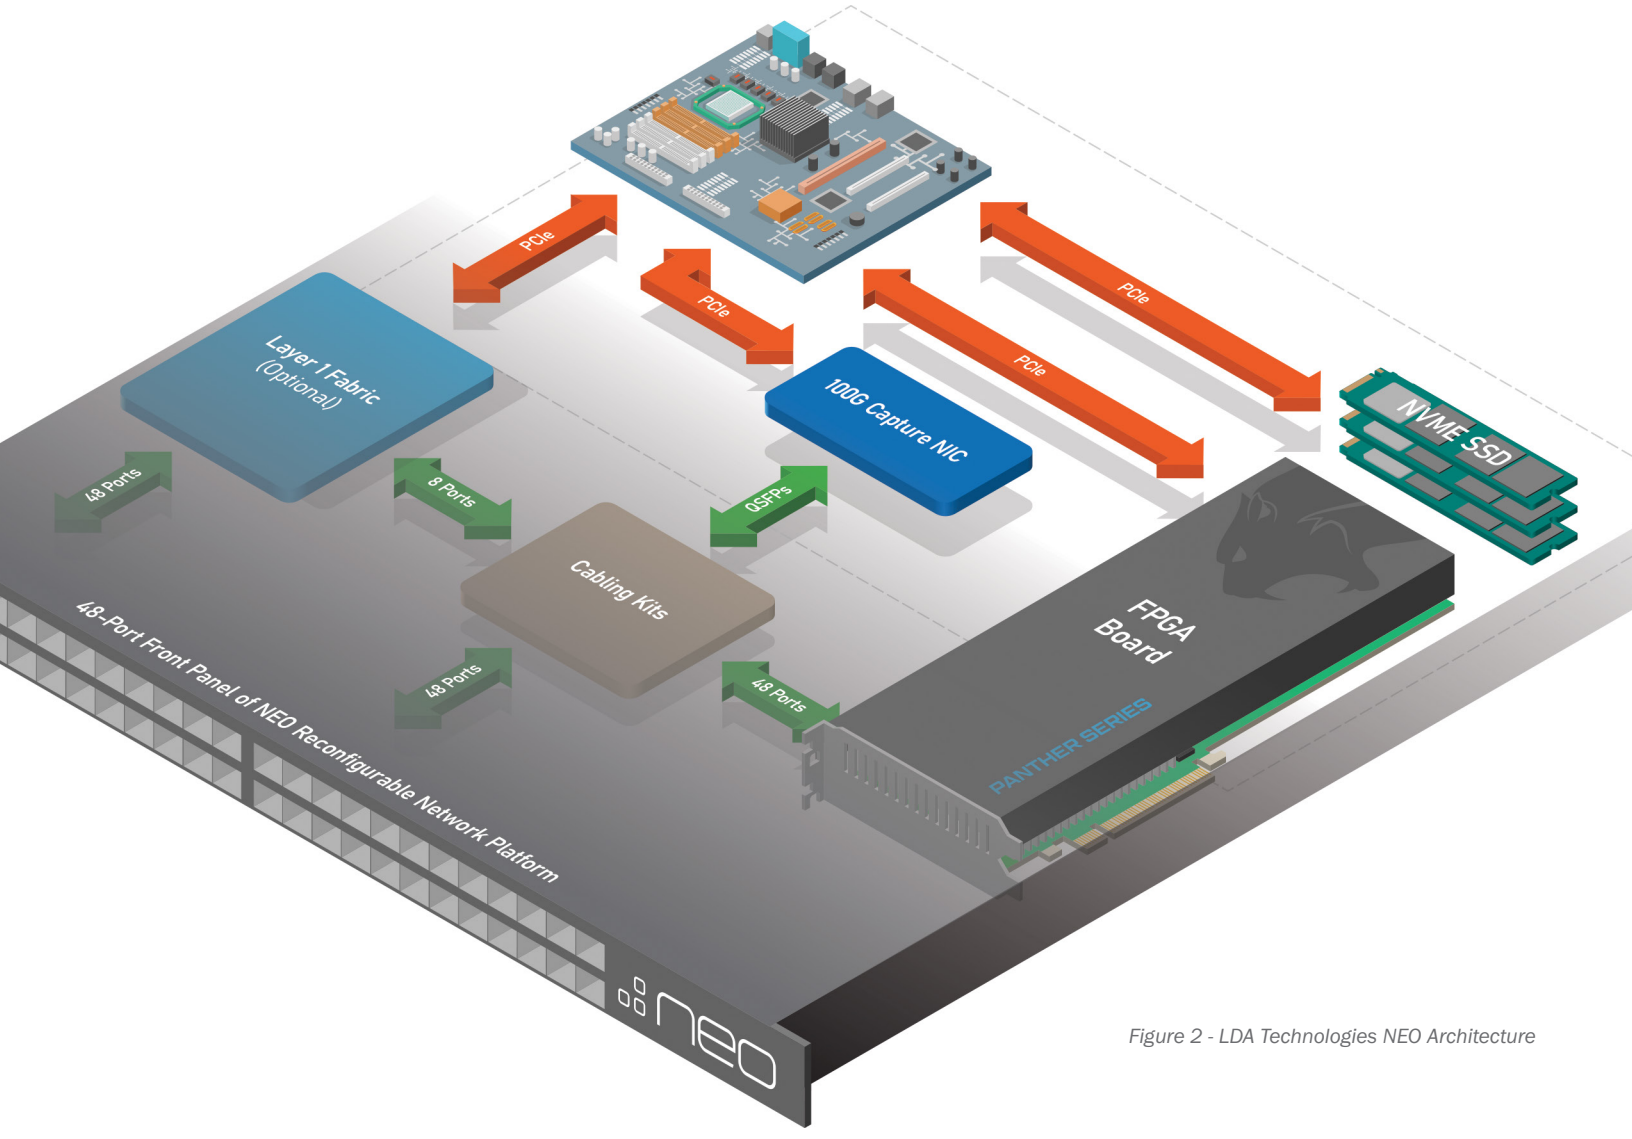

Figure 2 - LDA Technologies NEO Architecture

NetRecap Overview

NetRecap was designed for both 25 GbE and 10 GbE traffic analysis. By default, all traffic coming into each of the forty-eight ports on NEO receives a standard precision timestamp and a unique VLAN tag. Eight of these forty-eight ports can then be reconfigured to receive high-precision timestamps with an accuracy of 96 picoseconds for 10 GbE or 40 picoseconds for 25 GbE. Time is perhaps one of the most important metrics to track when doing network packet capture.
As the Merovingian said:

"Yes, of course. Who has time? Who has time?
*But then if we never *take* time, how can we have time?"*

Unlike most other capture solutions, NEO's accuracy is measured in 10 s of picoseconds; this is like a carpenter moving from measuring wood using a tape with 1/8th inch marks to one with millimeter marks (about 1/25th of an inch). Precision is essential; actually, knowing the time between when two specific packets arrive is sometimes the whole purpose of a packet capture project.

NEO VLAN tags each packet with a port-specific value, which makes it easier to correlate traffic as it is aggregated together. NEO can then steer all ports directly to the network capture card. These cards often work with the operating system to collect packets and pass them to a capture application via various kernel bypass techniques. This allows these applications to cache network packets into the main memory until they can be inspected or permanently stored on NVMe drives. NEO uses NVMe drives with four lanes of PCIe Gen4 with an effective writing speed of 4.4 GB/sec. If the storage subsystem is properly configured, all three drives can support writing 100 Gbps of capture to disk. A second version of the NEO is available that supports up to four NVMe drives, which would push this writing capability well beyond 100 Gbps.

The NEO provides a single eight-lane PCIe slot for a NIC for capture; almost any current 10/25/100 GbE card should work. This makes it easy to fit NEO with NetRecap into your existing capture infrastructure or build a whole new, highly performant one. If you've been capturing network packets for years, then you may already have a robust infrastructure using Napatech, Solarflare, or Intel with Ntop already installed. NEO with NetRecap can be configured to fit right in using the capture NIC and software package you're most familiar. The performance of your NEO when doing packet capture is only limited by the NIC and software capture solution you select.

Basic Packet Search and Analytics

NEO comes with CentOS Linux preinstalled. This enables one of the most straightforward use cases, which leverages tcpdump for basic diagnostics. Tcpdump uses the familiar libpcap kernel API to gain access to the packet flow. By default, all packets accessible via libpcap flow through the kernel, limiting the peak packet rate through libpcap to about three million packets per second. A more performant version of libpcap called xdpcap utilizes eXpress Data Path (XDP) technology. Xdpcap should lift many of the kernel limitations that gate the performance of libpcap. Either way, one could simply use tcpdump or xdpdump to gather packet data for your analysis.

Advanced API Capture and Replay

For the past several decades ntop has been focused on packet capture performance, features, and interfaces. They bring the following tools to the LDA NEO platform in their suite of network capture products, here are some of those tools and a brief summary of their capabilities:

- **PFRing™ ZC (Zero Copy)**: This is the network driver that provides kernel bypass and works through the DPDK interface on Linux. It provides a simple and clean Application Programming Interface (API) that can then be used by other tools in the ntop suite.
- **N2Disk**: An advanced network traffic recorder that captures network packets and stores them to disk. With N2Disk you can capture packets over a long period of time or define a rolling temporal window for capture, perhaps only saving packets for a day as available storage becomes scarce.
- **nProbe™**: Goes one level deeper moving from packets to flows by utilizes NetFlow v5/v9/IPFIX allowing you to deploy both a probe and collector designed to replay NetFlow flows. The flows from the probe can then be forwarded to ntopng for analysis. Flows can also be exported to tools like Kafka and ElasticSearch (using a plugin).
- **ntopng**: Is the next generation version of the original ntop. It is an advance libpcap analysis engine enabling you to see exactly what is going on within a series of captured packets or real time as data is flowing through the network. With it you can quickly identify the top talkers (send and receive), round trip times, TCP statistics and much more.
- **Disk2N™**: Was designed to play traffic that has been previously recorded back into the network respecting inter-packet time. Disk2N also handles packet reforging, meaning that it will reforge source/destination MAC/IP/Port address on the fly, recomputing the destination MAC in case of multicast and recalculating the checksum.

There are a number of other ntop tools available, but these are the initial ones we expect NEO users may be interested in utilizing.

N2Disk for replay is perhaps the best way to understand how an algorithm might respond. With N2Disk market data could be replayed back into a port with the inter-packet times respected and key packet data reforged. This would enable a developer to fully exercise their algorithms by running a series of simulations against it. To quote the Merovingian:

“You see there is only one constant. One universal. It is the only real truth. Causality. Action, reaction. Cause and effect.”

With ntop’s Disk2N (Disk to Network) application, the NEO can easily become a platform that exposes you to both the “Why” and the “effect” of network traffic.

Lossless Packet Capture, Analysis and Storage

Line-rate lossless packet capture is the holy grail of network capture. NEO with a Napatech Link™ NT series SmartNIC and Napatech's suite of capture tools is that holy grail. With a Link NT200A02 in the NEO, you can easily capture up to 100 Gbps of network packets, without loss, to NEO's internal NVMe drives. Much like SolarCapture Pro, Napatech's tools are API driven and provide kernel bypass capture directly into memory. The NT200A02 is perhaps Napatech's most advanced card supporting multiple 10/25/40/100 Gbps Ethernet interfaces. Napatech's Link Capture software is unmatched in the industry and affords you everything mentioned in the previous use cases. Napatech's Link Capture has been benchmarked with all the critical applications that consume captured packets, and these are Suricata, Snort, Zeke, Wireshark, TRex, and N2Disk (ntop's sister application of Disk2N).

If high fidelity packet capture is your objective, then Napatech's Link Capture platform is your only real choice. As mentioned, they are the only platform that can claim lossless line-rate capture at speed up to 100 GbE. Also, they're the only environment that supports the complete set of the following advanced features:

- Stateful flow management
- Multi-CPU distribution
- CPU-Socket load balancer
- Multi-port packet sequencing (using NEO's VLAN tags)
- Correlation key
- Deduplication
- Line-rate performance

Packet capture is Napatech's core business; they've listened to their customers and, as a result, have developed some features unique to Napatech's Link Capture platform. Stateful flow management directly within the SmartNIC is perhaps Napatech's newest and most advanced feature yet. With stateful flow management, you can push down into the NIC the flows of interest, and it can shunt or drop those that aren't required, thereby relieving the host of having to process the packets from those flows. Furthermore, the SmartNIC can then apply match action technology to the selected flows performing packet transformations within the NIC before the packets being captured are placed in system memory by the Link Capture platform.

Another critical feature of Link Captures is distributing packets between cores on the same CPU or another CPU in the system up to a total of 128 cores. While SolarCapture can distribute packets to other cores, they need to be on the same physical CPU. Furthermore, SolarCapture cards do not support a bifurcated PCIe hardware bus, allowing an adapter to steer packets to more than one destination CPU socket directly. The above two distinctions are critical if you want to handle line-rate traffic at speeds above 10 GbE because the sheer volume of packets can easily overcome a small finite set of CPU cores. Unleashing as many CPU sockets and cores as possible to capture is the best way to ensure everything performs flawlessly as packet rates ramp up.

Finally, and perhaps the Napatech Link Capture platform's best feature is its highly precise replay functionality. Link Capture replay allows for precise inter-frame gap control and nanosecond replay precision, facilitating exact session replication, which is a requirement when determining exactly "why" an algorithm is not performing correctly.

Summary

LDA Technologies NEO Reconfigurable Network Platform is a proven production environment for networking. By enabling LDA's NetRecap capability, then adding another vendor's capture card and software, you're preserving that production environment while bolting on a high-performance capture system that runs alongside your production environment, in the same chassis. It does NOT change the flow of your production traffic; there is no need for optical splitters or another "U" of rack space for a capture appliance from yet another vendor. All of these issues would open production up to more potential points of failure. The Merovingian said that:

"Choice is an illusion created between those with power and those without."

NEO controls all the data flows. It gives you the power to effectively manage your production environment while also providing you with the choice to select a capture platform, so perhaps, in this case, the Merovingian was wrong. NetRecap enables the capability to capture network packets, and as NEO said in The Matrix:

"Choice. The problem is choice."

With NetRecap enabled on the LDA NEO Reconfigurable Network Platform, the problem then becomes which capture platform to select, the choice. Perhaps you're already familiar with capture, but you've been instructed by management to keep it within a tight budget. In the past, you've been using Intel cards with ntop; this capture solution utilizing Intel's newest E810 could be installed into a NEO with ntop at a very affordable price point. Perhaps you require lossless packet capture, with extensive fine-grained replay capabilities so you can find the "Why" no matter where it is, then there is Napatech.

NEO with NetRecap enables the capability to do non-disruptive network packet capture within a production environment, at several different price points, and with varying abilities. LDA provides you with the freedom to choose which capture solution best fits the needs of your business.

LDATech.com
1 (800) 738-8163
info@ldatech.com

WP-NPTDTAW | Whitepaper | 201104

LDA TECHNOLOGIES™