HIGH-SPEED DEDUPLICATION HELPS FINANCIAL PROVIDER OPTIMIZE ITS CYBERSECURITY MEASURES

Actual Concern

Zero-Day Attacks are a nightmare for organizations and often can cause the institution to pay a ransom in the millions of dollars, in addition to causing long-lasting damage to its reputation. The Log4J vulnerability is the latest of those nightmares where an issue was exploited and allowed remote access and execution of the exposed system. When an attacker has control of a system, they can move laterally to other systems and attempt to exfiltrate data. The challenge is when they get to the point of exfiltration, the attacker is very smart to utilize encryption and protocol weaknesses to exfiltrate the data without security tools raising an alarm. The key is to catch them moving laterally and mitigate their ability to exploit your valuable assets, which include corporate and customer data.

The Customer

One of APCON's long-time financial banking institutions — with multiple corporate offices and billions in revenues — recently noted that it cannot cover all vulnerabilities in an expansive network environment. Better visibility and deduplication capabilities would allow security tools to perform their functions at optimal levels.

Solution Needed

The customer needed to upgrade their infrastructure to support growth and identify better ways to secure the financial institution. They needed to support many 40Gbps links and, going forward, were centralizing their security and networking tools to a tool farm. They needed the ability to deduplicate the data stream before sending packets to their tool farm. As they were exploring the Network Packet Broker market, one of the key considerations was to be able to support aggregation of traffic averaging 125Gbps or more with bursty conditions. This requires the deduplication engine to actually support traffic above 150Gbps to ensure every packet is inspected for copies before sending data to the tool farm.

The customer chose APCON for having the ability to not only deduplicate the traffic with bigger incremental pool sizes, but also to support line rate multifunction capabilities like Port Tagging, Protocol Stripping, and Load Balancing.

The Deduplication Process

When working with Network Packet Brokers and collecting traffic from your infrastructure, consider how you handle duplicate traffic. There are vendor tools that have intelligence built into the system so that if it receives duplicate packets from the Network Packet Broker, it will build an incorrect model. This shortcoming will send you on a wild goose chase which can cause the firm to utilize precious IT resources and waste money. Here are some things to consider on this subject:

- Network Packet Brokers provide deduplication in increments, meaning there is a certain amount of bandwidth or pool of resources it can handle for each deduplication service.
- The more ingress traffic ports that need to be deduplicated, the bigger the incremental pool sizes need to be. When you have several 40G/100G connections coming in where the packets need to be deduplicated, deduplication pool resources can fill up and not properly deduplicate the traffic, whereby your tools will process and provide erroneous results.
- Oversubscription of the deduplication pool must be managed to ensure the deduplication service functions properly. This goes into the previous point above but also to the next point.
- The order of services, or when deduplication occurs, is also important. If you assign different services like Protocol Stripping or Packet Slicing or both, when does that service get processed? When does the Deduplication service get processed?

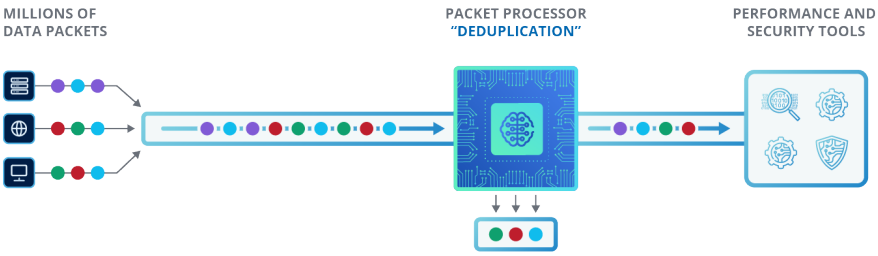

FIGURE 1

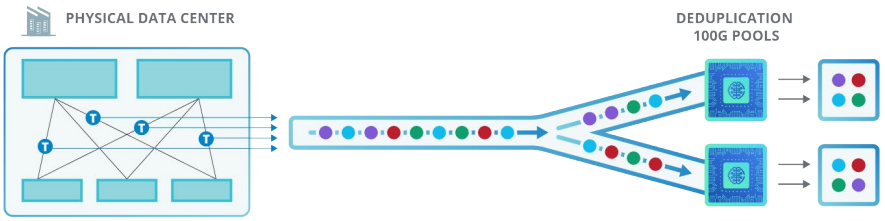

FIGURE 2

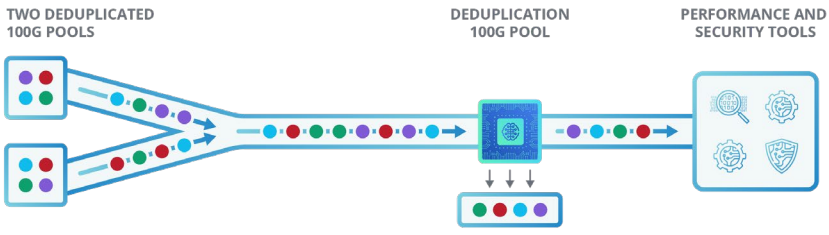

FIGURE 3

Deduplication Use Cases

It is important to understand that there are advanced processing modules used to hold a packet in cache, then explore all packets in the pool for any duplicates. As shown in Figure 1, millions of packets come into the advanced processing module. The module has a packet that it uses to identify any packets matching the criteria of a copy. Once there is a copy identified, the system removes the packet/s and sends the deduplicated packets to the tool. This process is happening constantly and very quickly.

See the diagram in Figure 2 that illustrates a very complex challenge. A large deployment in which HR and Accounting share applications and resources are spread throughout the infrastructure. The traffic could be collected via TAPs or SPAN, but the volume of traffic may exceed 100G, which exceeds the capacity of one deduplication pool and requires that it be broken into two pools.

Once the deduplication process is complete, you're not done! Since there is a possibility that duplicates can exist from conversations across the infrastructure, we must send that deduplicated traffic from separate pools into one aggregate pool to be deduplicated again before sending the data to the specified tool. This will ensure no duplicate traffic exists. Most vendors only have increments of 100Gbps of deduplication pools.

The APCON E36 Line Card has an Advanced Processing Module and handles (2) 200Gbps per service engine, which means it can handle more ports to deduplicate traffic per pool. On the left in Figure 3 is an example of the E36 Line Card 200G pool where other vendors are often limited to 100Gbps increments.

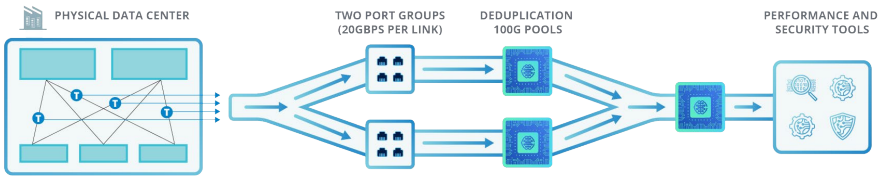

FIGURE 4

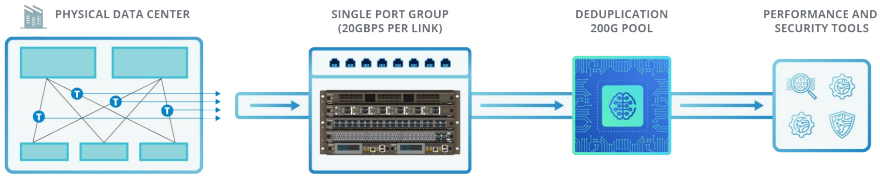

FIGURE 5

Consider a use case where the traffic from those links will be at or above 20Gbps and the deduplication pool has a maximum capacity of 100Gbps. The system should only group 4 ports each to ensure the pool is not oversubscribed as shown in Figure 4. If the traffic being inspected and deduplicated in that pool does not include all traffic, technicians need to create yet another pool for aggregation and deduplication before forwarding data to the tool of choice.

What happens if the traffic from deduplication pool 1, which went through its first phase of deduplication, is more than 100Gbps and needs to be deduplicated again to ensure there are no copies? You would need to create a new service for the deduplication process to take Pool 1 and combine it with Pool 2 to ensure no duplicate traffic gets to the tools.

As in Figure 5, when utilizing the APCON's IntellaView E36 line card you can see the difference between 100Gbps increments of deduplication pools to 200Gbps. All eight ports can be grouped into one deduplication pool.

This is easier to manage and prevents the potential further deduplication process needed to ensure all traffic is deduplicated before going to other tools.

Aside from addressing the deduplication pool sizes and working to control oversubscription, APCON handles the following with the E36 line card:

- The IntellaView E36 line card provides 36 ports that can go up to 144 ports in breakout mode to cover 10G/25G/40G/100G speeds.
- Each port operates at full line rate to process services like Protocol Stripping, Packet Slicing, Tunnel Management, and Port Tagging.
- IntellaView E36 provides two Deduplication pools of 200Gbps increments on the line card. This deduplication service can also service other line cards in the chassis.
- Conducts deduplication last when other advanced services are applied, which means if slicing or stripping is done, more packets can fit into a deduplication pool.

We hope the technical information provided above helps illustrate the challenges related to managing deduplication and what to look out for. We hope you find this helpful and hope to hear from you.

APCON Solutions for Networks

APCON, Inc. • 9255 SW Pioneer Court, Wilsonville, Oregon 97070
+1 503–682–4050 • 1–800–624–6808 • **apcon.com**

22003-0222