Packet Deduplication

Eliminate Duplicate Packets to Drive Network Monitoring Efficiency

Network engineers can reduce both CAPEX and OPEX by reducing or eliminating duplicate packets. Did you know that:

- Duplicate packets arriving at network monitoring tools generally waste both storage and processing capacity, reducing the window of network visibility
- Duplicates can overload a tool's input port, resulting in dropped packets
- Tools such as FireEye and TeaLeaf that track interactions at the session layer cannot manage duplicate packets
- Duplicate packets can make up as much as 55% of total monitoring traffic volume

The APCON INTELLAFLEX Packet Controller blade eliminates duplicate packets from 10G data streams at line rate, reducing the volume of packets to avoid tool overload and errors on session-based appliances. This saves both operating expense on monitoring activities and capital expense on redundant tools.

Find out how duplicate packets are generated, and specifies the challenges and inefficiencies that duplicate packets bring to a network monitoring program. Further, see how the APCON INTELLAFLEX Packet Controller blade eliminates duplicate packets to streamline the monitoring process.

Comprehensive network monitoring is a requirement in modern data centers. Application performance requirements, resource allocation and planning, and the ever-present threat of security breach combine to create a compelling case for complete network visibility at all times.

Network monitoring data sources such as SPAN (Switched Port Analyzer), Mirror ports, or TAP (Test Access Port) create exact copies of data traffic without interference. The result is that monitoring tools often receive multiple copies of the same packet. In some cases duplicate packets can represent as much as 55% of the traffic delivered to monitoring tools.

How Duplicate Packets are Generated

There are many ways that duplicate copies of packets are generated by a network monitoring system. The vast majority of duplicates are generated by SPAN or Mirror ports and Ethernet Taps that feed the network monitoring system. These duplicates are a natural consequence of capturing all data, and eliminating this duplication by design would be technically challenging.

For example, consider Figure 1. In this case, the SPAN port configured on this switch will report two copies of certain packets – once when the packet is received into VLAN 100, and again when the packet is sent out from the same VLAN. This occurs because the SPAN configuration is set to report both receipts and transmissions on this VLANso that all traffic is reported and none is missed.

Figure 1 – Configuring a VLAN SPAN using the "Both" parameter to include Tx and Rx traffic on a Layer 2 switch generates duplicate packets.

SPAN Configuration: VLAN 100 Both (Tx/Rx)

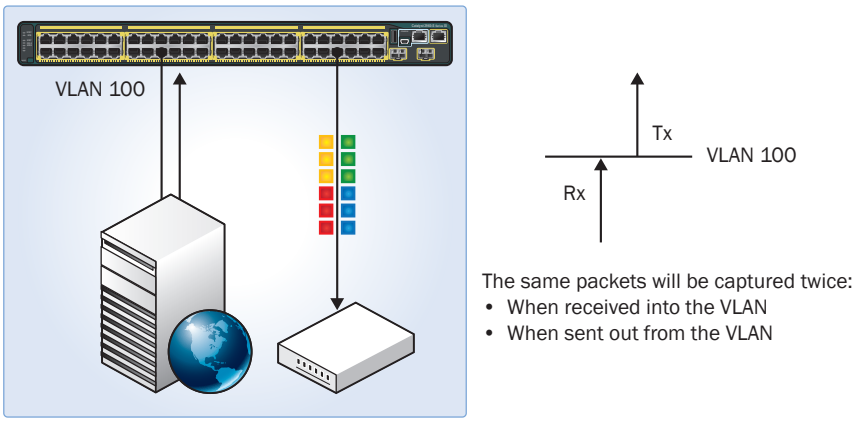

The same packets will be captured twice:
- When received into the VLAN
- When sent out from the VLAN

Figure 2 shows another SPAN configuration that generates duplicate packets. This SPAN port is set to report packet receipts from VLAN 100 and VLAN 200. Any packets that are received by VLAN 100 and sent on to VLAN 200 will be reported twice through this SPAN.

Figure 2 – Configuring a multiple-VLAN SPAN where a server communicates with other servers on the same switch generates duplicate packets.

SPAN Configuration: VLAN 100, 200 Rx Only

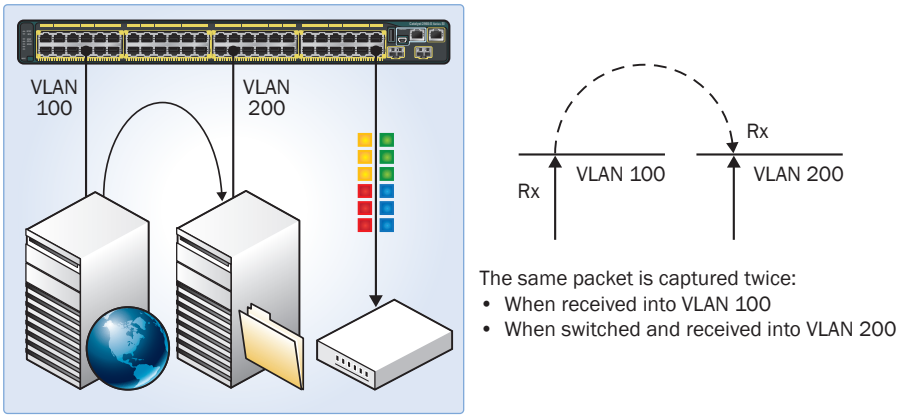

The same packet is captured twice:
- When received into VLAN 100
- When switched and received into VLAN 200

Figure 3 shows a SPAN port configured to report traffic being received and transmitted by two particular ports on the switch. Any packets traversing these two ports will be reported twice.

Figure 3 – Configuring a multiple server port SPAN using the "Both" parameter to include Tx and Rx traffic, and where the server communicates with other servers on the same switch generates duplicate packets.

SPAN Configuration: Fa1/18, Fa1/40 both (Tx/Rx)

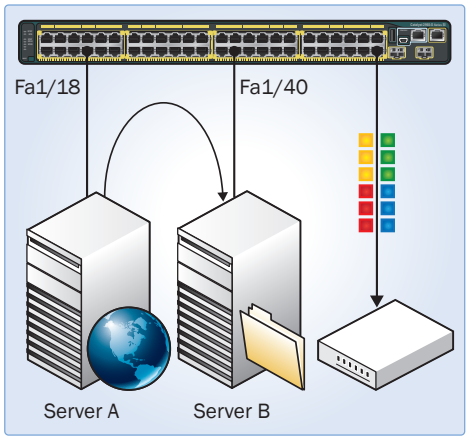

The same packet is captured twice:
- When entering port Fa1/18
- When leaving port Fa1/40

Aggregating packets to a monitoring system will often result in duplicate packets being reported. For example, in Figure 4 the monitoring system aggregates data that has been delivered from Capture Points 1 and 2. Capture Point 1 reports packets from both Switch A and Switch B. However, Switch B also reports its traffic through Capture Point 2, so all packets coming from Server 2 (and any other servers attached to Switch B) will duplicate some of the data seen at Capture Point 1.

Duplicate Packets from SPAN and Aggregation

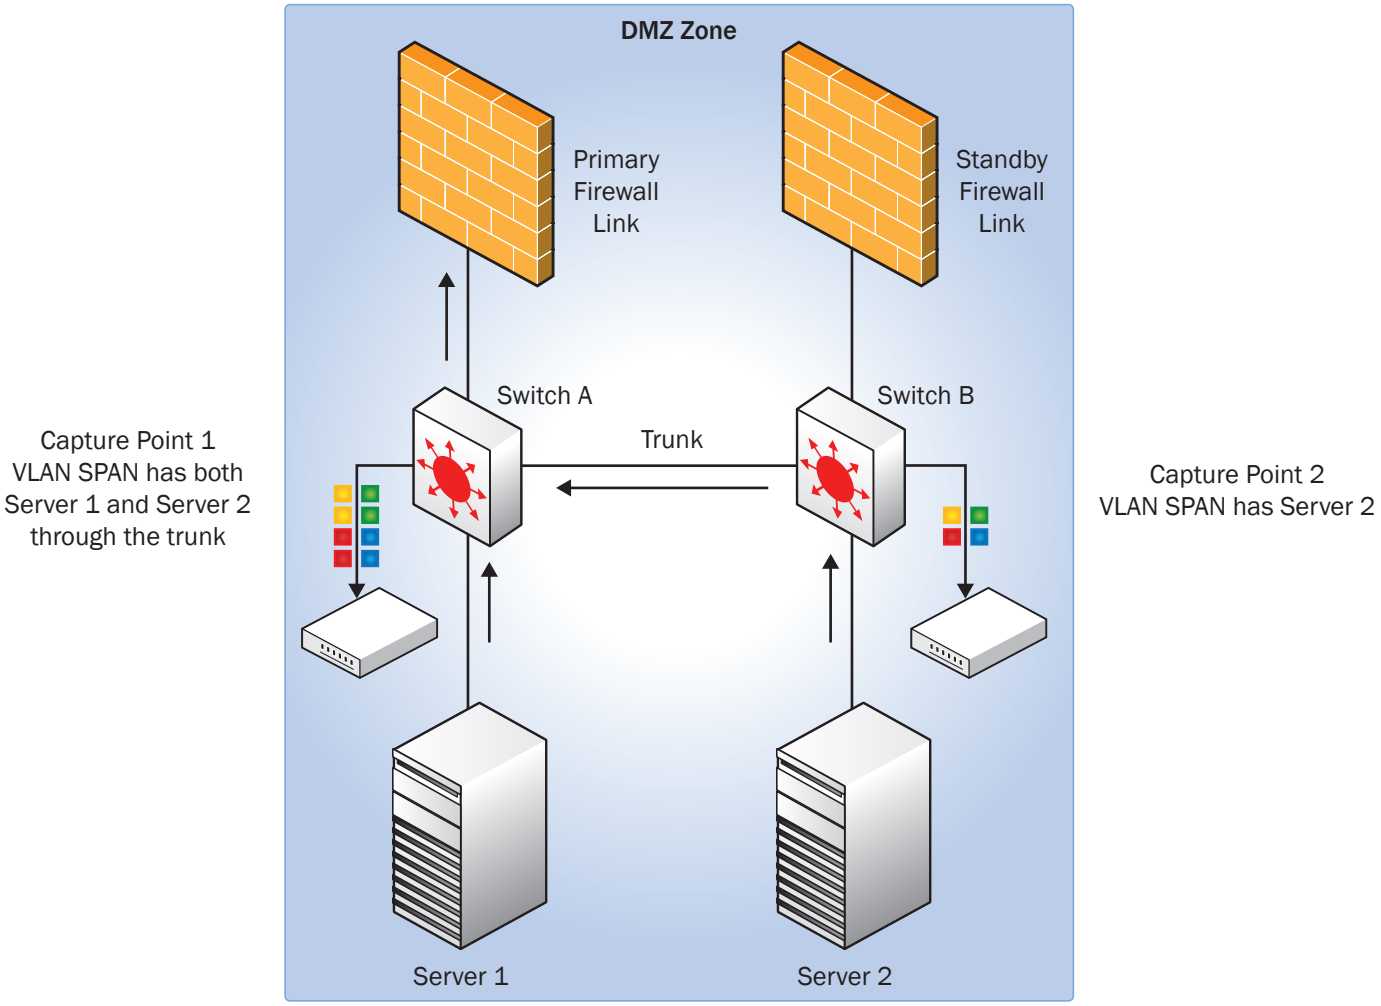

Figure 4 – Aggregated SPAN traffic from both capture points 1 and 2 generates duplicate packets for Server 2 as well as all servers connected to Switch B. Network engineers cannot rely on capture point 1 during failover, as it depends on capture point 2 for data.

Duplicate Packets Hinder Network Monitoring

With duplicate packets arriving at network monitoring tools regularly, certain conditions arise that reduce the efficiency and effectiveness of your monitoring strategy. For example, many monitoring tools are at or near capacity when simply handling regular production network traffic. In the worst case, oversubscription from duplicate packets may cause packet loss of required traffic. In many cases, data centers have invested in additional equipment to handle the load, raising both capital expense and operating expense to administer and use the tools.

If a given network monitoring system uses data storage for record-keeping and event reconstruction, duplicate packets can reduce available storage by up to 55%, leading to unnecessary expenditures on additional storage. The same issue applies to streaming network event recorders, reducing the available recording time window by up to 55%.

Less quantifiable but just as important is the fact that analyzers are less effective when analysis must be performed around duplicate packets. Excessive duplication leads to a lack of precision in analysis and skewed network statistics. Further, some functions such as session re-establishment cannot be performed with duplicate packets in the system.

Solution: IntellaFlex Packet Deduplication

The solution to packet duplication is to use an intelligent data traffic management switch to examine each data stream arriving from the production network. When fitted with an IntellaFlex Packet Controller blade, an APCON IntellaPatch Series 3000 switch identifies and removes duplicate packets.

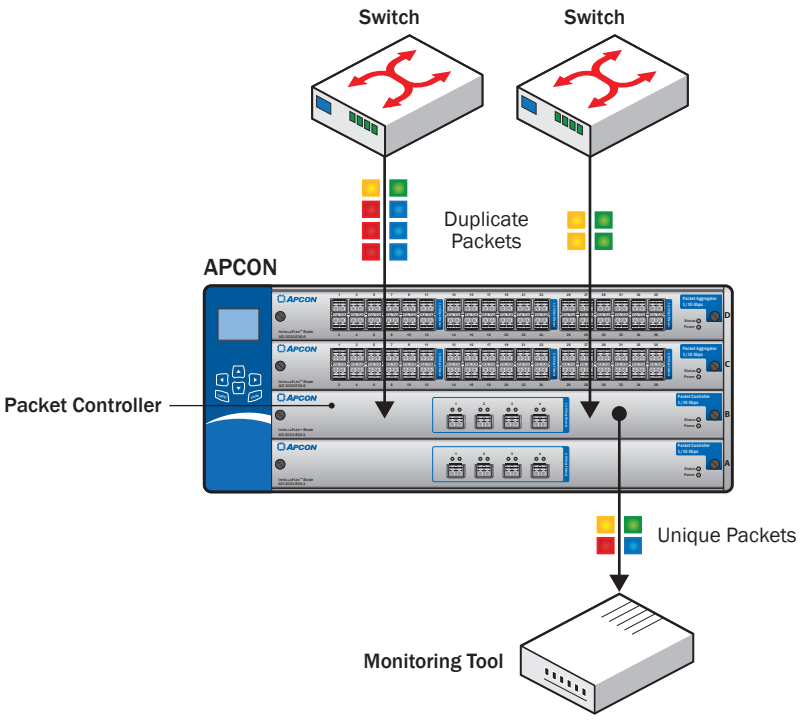

Figure 5 – In this example, duplicate packets are removed by the APCON switch before passing the data stream to monitoring tools.

Standard Deduplication Scenario

For a representative deduplication scenario, see Figure 6. In this configuration, four data streams from network data capture points arrive at the APCON switch through an IntellaPatch Layer 1 port blade. The data streams are passed through the switch backplane to the IntellaFlex Packet Controller blade for deduplication. This is simply one of many configurations for deduplication.

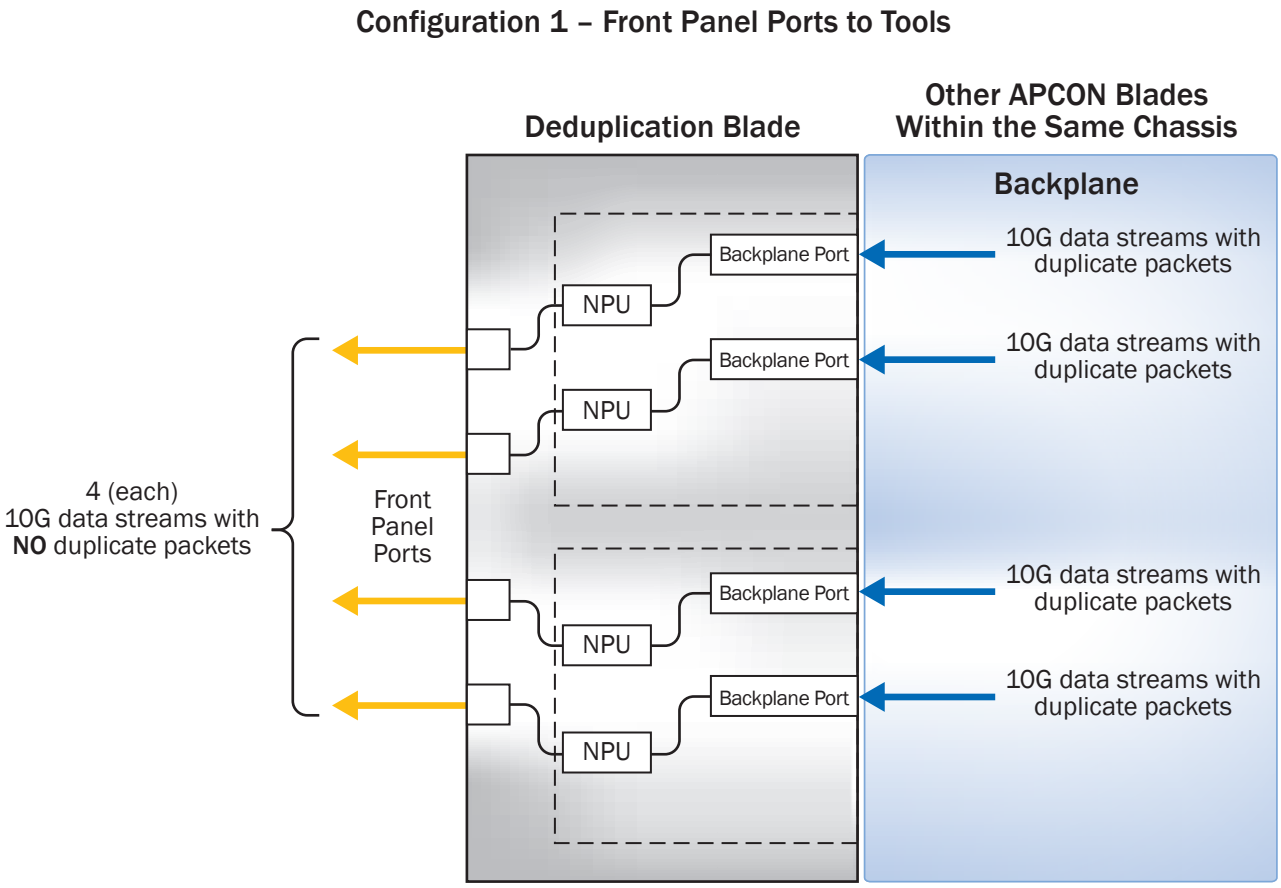

Figure 6 – The IntellaFlex Packet Controller allows you to deduplicate up to four 10G data streams. In the diagram above, four streams enter the blade through the 10G RXAUI backplane ports, pass through the Network Processing Units for deduplication, and then pass out through the front panel ports.

Switch Backplane Deduplication

Figure 7 shows an alternate configuration where two data streams are deduplicated using only the backplane ports and sent back out through other blades in the IntellaPatch switch chassis. In this configuration the front panel ports are not used. However, to meet future needs, data traffic could be brought in through one of the front ports, deduplicated by one of the unused Network Processing Units, and passed back out another front port. At any given time, all four 10G front ports, four NPUs, and four 10G backplane ports may be configured for use.

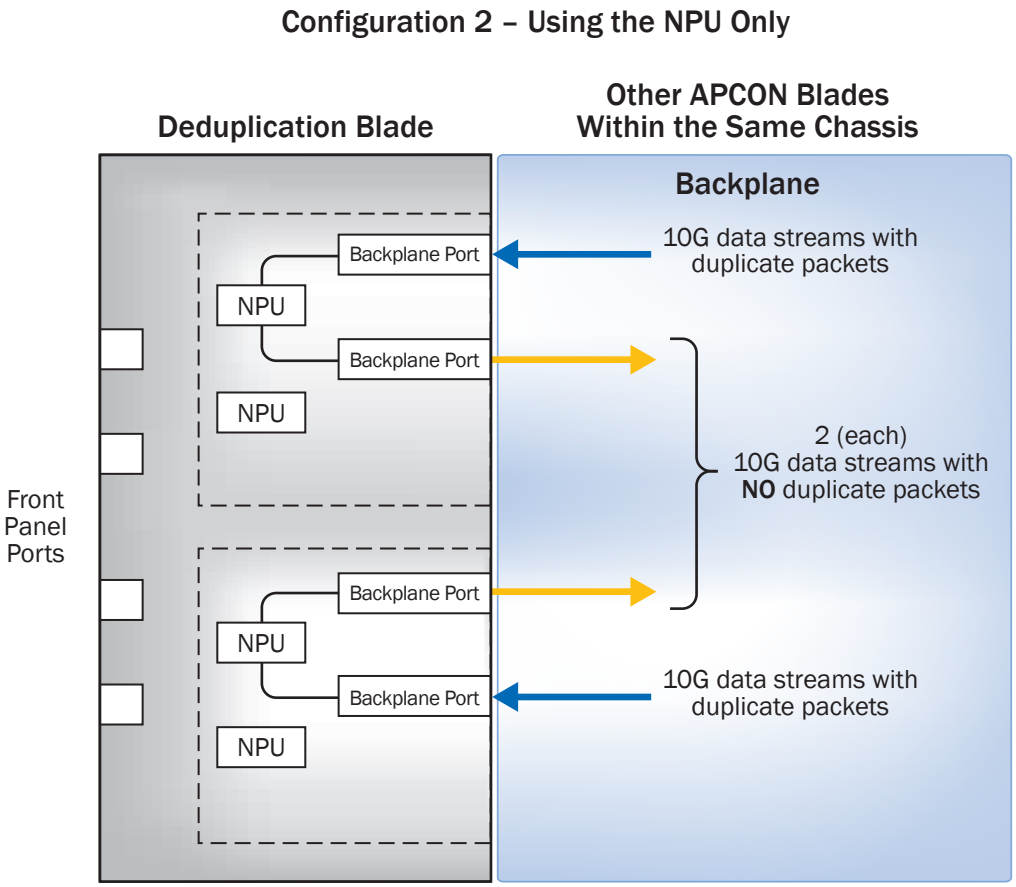

Figure 7 – While not as efficient as the implementation shown in Figure 6, it is allowable to use backplane lanes exclusively. However, only two streams can be deduplicated when using all four backplane lanes. Two 10G streams of data can be directed from the backplane through the Network Processing Unit for deduplication and then back out the backplane to another blade in the switch.

Packet Aggregation and Deduplication

In Figure 8, a total of eight data streams containing duplicate packets enter the switch through an IntellaFlex Ethernet packet aggregator blade. The IntellaFlex blade offers the ability to merge several data streams into a single unified and filtered stream. In this configuration, the eight input streams are aggregated into four streams passed through the switch backplane to the Packet Controller blade. After deduplication in the NPUs, the deduplicated data streams efficiently pass out through the front ports to the monitoring tools. If the aggregated data streams exceed 10G bandwidth, ingress filtering on the IntellaFlex blade will be required to bring the aggregated stream down to 10G.

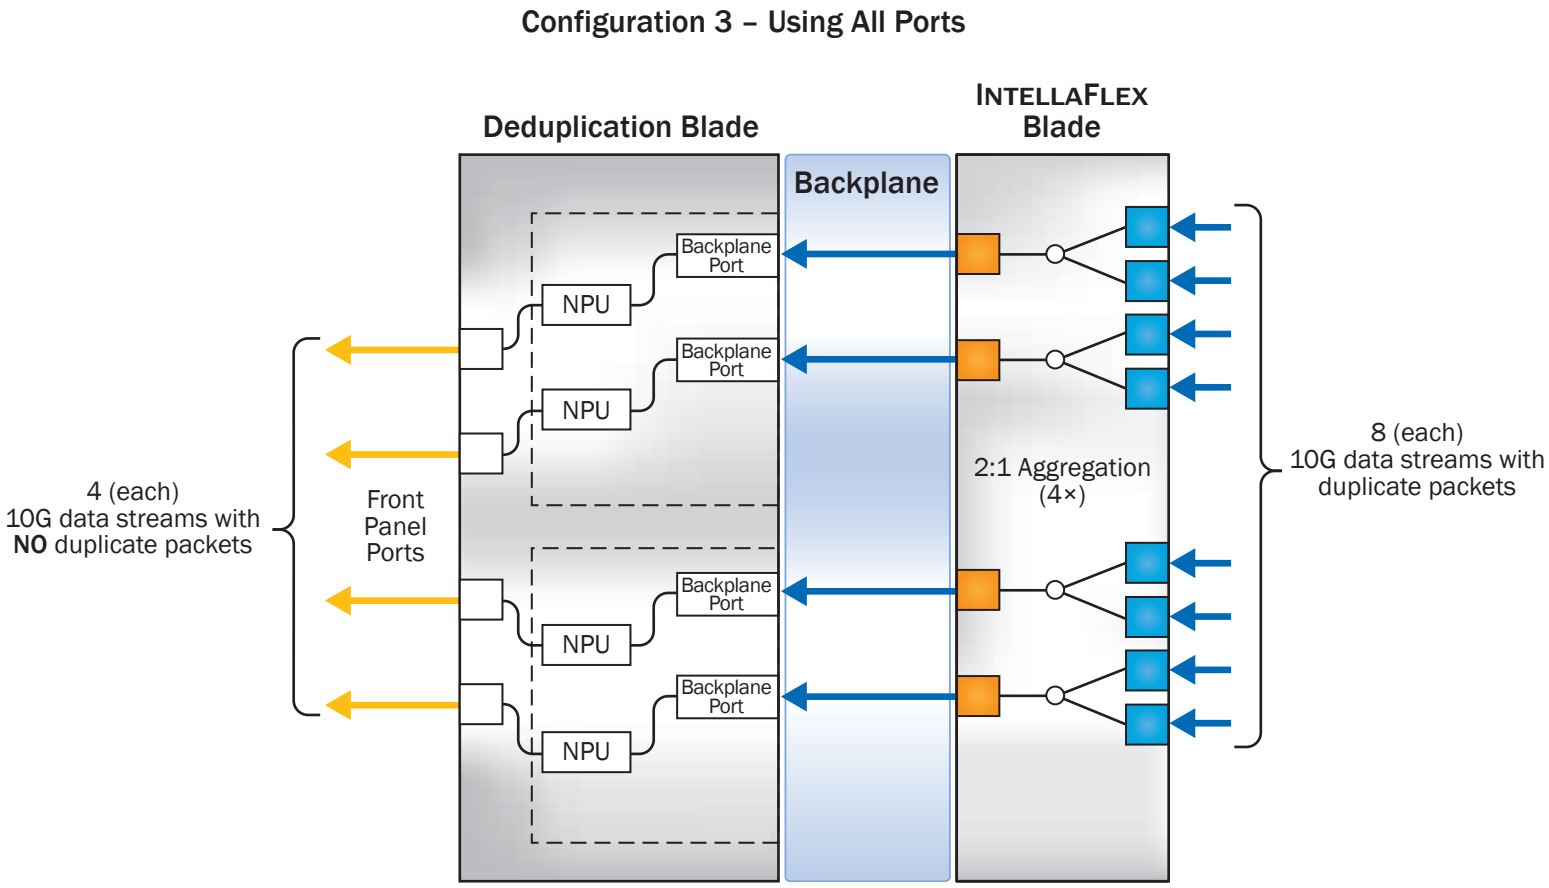

Figure 8 – In this configuration, 8 x 10G data streams are deduplicated by aggregating two streams into one on a separate IntellaFlex blade, then routing the resulting 4 data streams through the backplane lanes of the Packet Controller, through the Network Processing Units, and then out through the front ports of the Packet Controller blade.

Configuring the Deduplication Engine

One of the technical challenges inherent in packet deduplication is determining whether a given pair of packets are in fact duplicates. For example, duplicate packets arriving from different VLANs may have slightly different Ethernet header information, but the packet payload and the balance of the header will be identical.

APCON's advanced web-based graphical switch configuration software allows network engineers to specify the IPV4 and IPV6 packet header fields to be compared for deduplication purposes. The packet payload is always checked for duplication. These deduplication settings also allow the engineer to specify a time window for which to retain packets to be checked for duplicates. Figure 9 shows the deduplication settings for one NPU on the IntellaFlex Packet Controller blade.

When activated, the IntellaFlex Packet Controller blade examines each packet as it enters the NPU and creates a unique digital signature for that packet based on the selected header fields and a comparison of the payload information. As more packets arrive at the NPU, exact duplicates are discarded. It is important to note that all deduplication comparisons are made within a given NPU. Duplicate packets arriving on a data stream directed at a different NPU will not be detected.

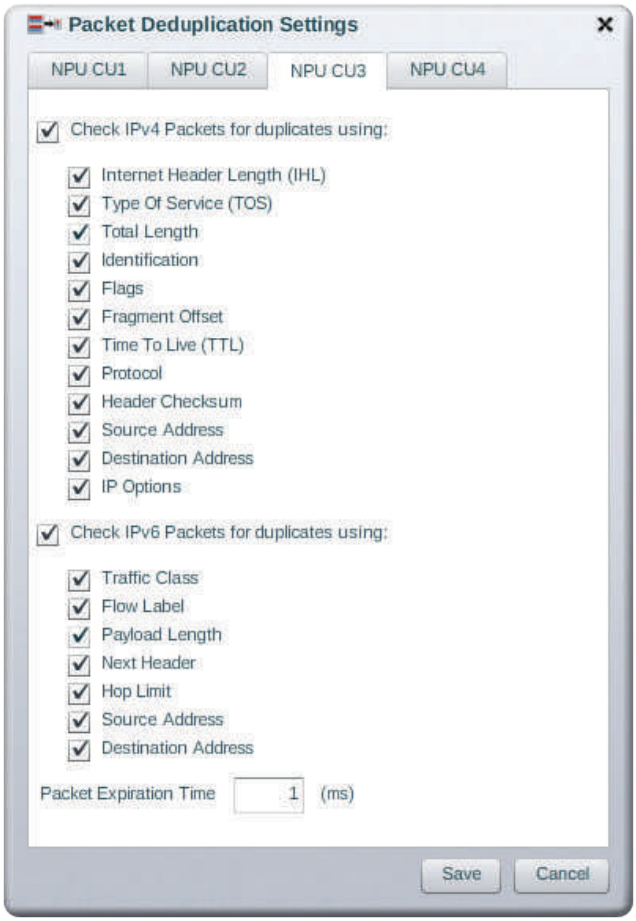

Figure 9 – Packet Controller Deduplication Settings.

Conclusion: Packet Deduplication for Network Monitoring Efficiency

Duplicate packets are a natural consequence of a comprehensive network monitoring program. However, with APCON's IntellaFlex packet deduplication technology, duplicate packets need not interfere with effective network monitoring. APCON's IntellaPatch Series 3000 switches can be equipped with a range of APCON IntellaPatch and IntellaFlex blades to meet the needs of any size network. APCON technology saves unnecessary capital expenditures on redundant monitoring equipment by delivering only the required data to network monitoring tools – making every packet count.

Contact Us

Please email sales@apcon.com if you have any questions

ABOUT APCON

APCON develops innovative, scalable technology solutions to enhance network monitoring, support IT traffic analysis, and streamline IT network management and security. APCON is the industry leader for state-of-the-art IT data aggregation, filtering, and network switching products, as well as leading-edge management-software support. Organizations in over 50 countries depend on APCON network infrastructure solutions. Customers include Global Fortune 500 companies, banks and financial services institutions, telecommunication service providers, government and military, and computer equipment manufacturers.

12100-R4-0613

APCON Solutions for Networks