

Port Aggregator Tap Guarantees Correct Packet Ordering

Solution Brief



Port aggregation is an important technology for increasing network monitoring efficiency. It enables a tool to see both sides of a full-duplex conversation through a single network interface controller (NIC). Accessing network traffic with Port Aggregator Taps doubles the number of links that can be monitored by a tool with a fixed number of NICs.




Figure 1: Port Aggregator Taps double monitoring capacity (per NIC)

Unfortunately, port aggregation has the potential to deliver traffic with packets from the two sides of the link out of sequence relative to their arrival times, and this can upset protocol and forensic analysis applications. This paper explains why maintaining correct packet order is important, why most Port Aggregator Taps cannot guarantee correct packet ordering, and how Net Optics solves this problem in a new generation of Port Aggregator Taps with internal timestamping.

Port Aggregator Tap Guarantees Correct Packet Ordering

Solution Brief

The Packet Ordering Problem

An inevitable consequence of combining two traffic streams onto a single set of wires is that the timing of the original traffic streams cannot be preserved perfectly. For example, if packets arrive from both directions at exactly the same time, one packet must be sent first on the aggregation port, and the other packet must be delayed.

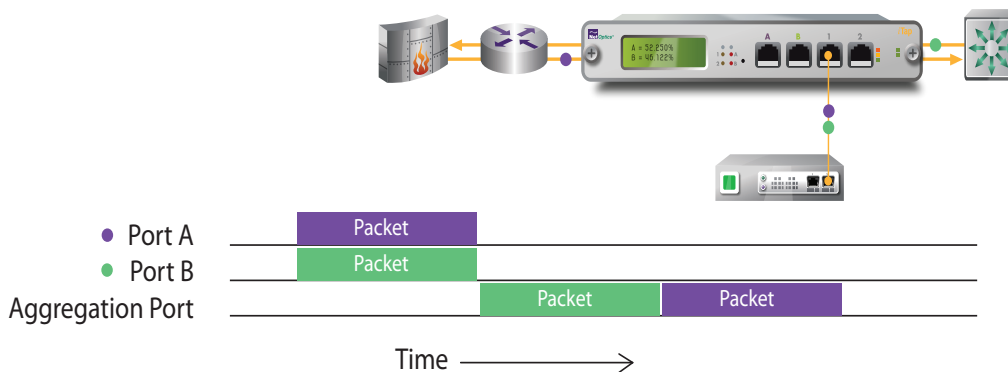


Figure 2: Port Aggregation affects relative packet timing

Fortunately, for many monitoring and analysis tasks, the absolute timing of packet arrival is not important. However, what is often important is the packet order — that packets appear on the wire in the same sequence that they occurred in the two traffic streams. For example, consider a typical handshake sequence, and what can happen if the packet order is not preserved.

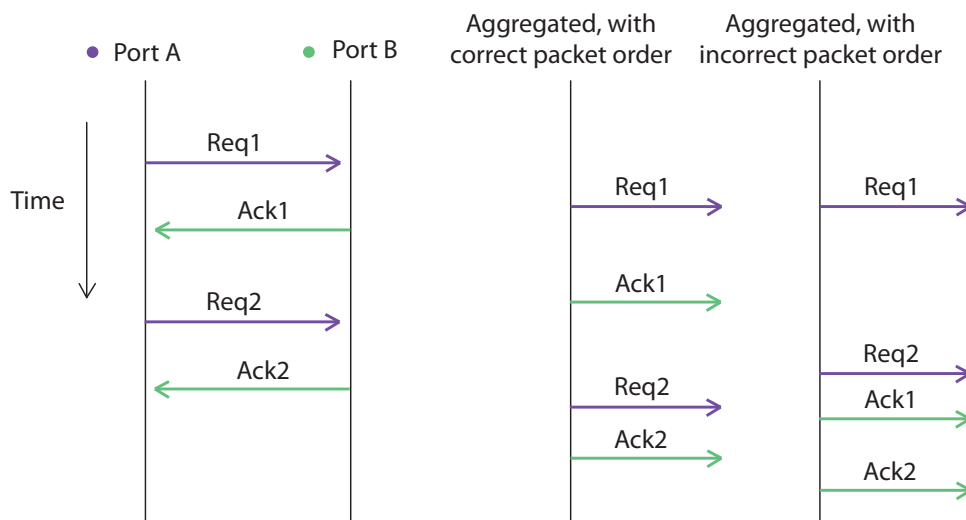


Figure 3: Handshake sequence shows the importance of packet ordering

As can be seen in Figure 3, a program or person looking only at the aggregated traffic stream with incorrect packet order as shown on the right would observe that the first request had not been

Port Aggregator Tap Guarantees Correct Packet Ordering

Solution Brief

acknowledged before the second request was issued, and that might constitute a protocol violation. But that would be a false error because the handshakes actually occurred in the proper sequence, so the analysis would be faulty. An investigator could waste a lot of time chasing the false error, and it may be impossible to understand the true issue at all in the presence of these false errors.

The Packet Ordering Hazard Under Aggregation

Most Port Aggregator Taps currently on the market do not preserve packet order, and therefore certain types of traffic analysis may not be possible using these Taps for monitoring access. The reason packet dis-ordering occurs can be seen in the basic Aggregator Tap design.

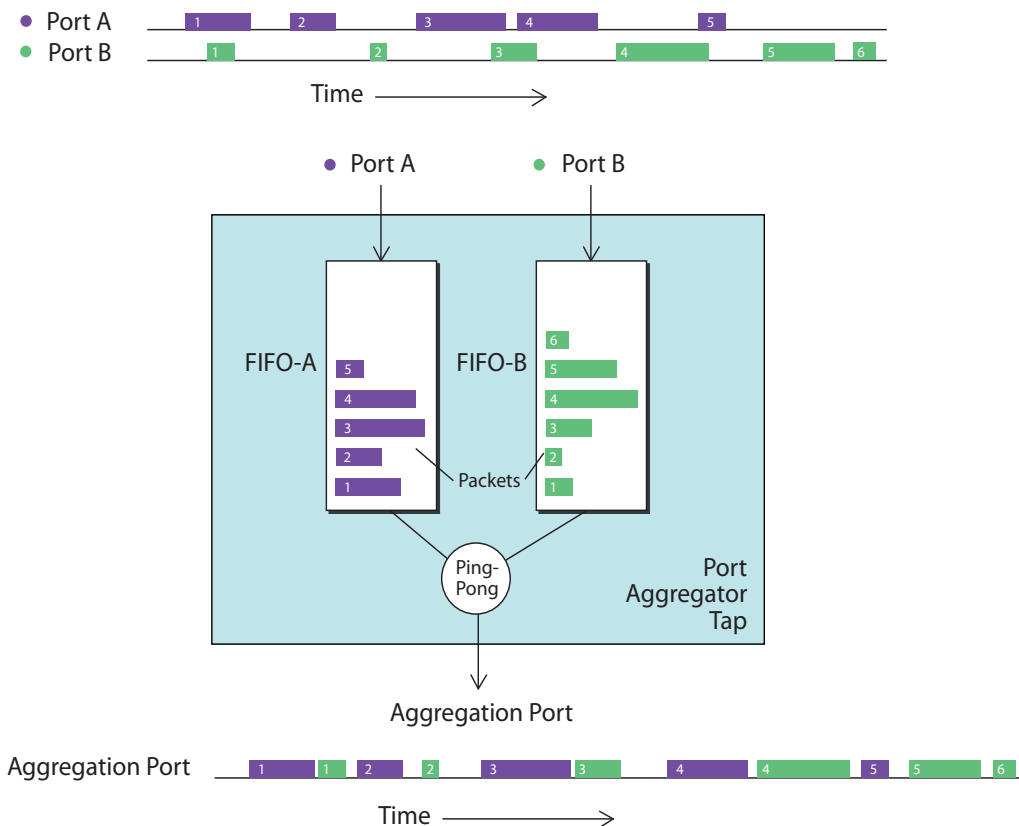


Figure 4: Typical Port Aggregator Tap design

In this typical design, packets received at Port A are stored in a first-in-first-out (FIFO) buffer, and packets received at Port B are stored in a second FIFO. As packets are pulled out of the FIFOs and sent to the aggregation port, the nature of a FIFO buffer guarantees that packet order is preserved within the Port A packets by themselves, and within the Port B packets by themselves. However, the order between a packet from Port A and a packet from Port B depends on when the Tap pulls the packets from the FIFOs, not on the original arrival order of the packets.

In fact, the information about the actual arrival order of the packets is not available at the outputs of the FIFOs, so the Tap has no way to know what order they should be in. Therefore, the Tap uses a simple

Port Aggregator Tap Guarantees Correct Packet Ordering

Solution Brief

round-robin or ping-pong algorithm for pulling packets from the two FIFOs, pulling one packet from FIFO-A (if one is available), then one from FIFO-B (if available), then A, then B, and so on.

Most of the time, the simple ping-pong algorithm results in maintaining the correct packet order. However, cases arise where the ping-pong algorithm results in out-of-order packets. Consider the case illustrated in Figure 5, where a long (maybe a jumbo) packet arrives on Port A while small packets arrive on Port B. By the time the long packet A-1 is transmitted on the aggregation port, three small packets have already arrived in Port B's FIFO. To maintain correct packet order, all three of these packets should be transmitted before the next packet on Port A. The Tap ping-pongs to FIFO-B and transmits packet B-1; but by the time B-1 has been transmitted, packet A-2 is available in FIFO-A. Therefore the Tap ping-pongs to FIFO-A and transmits A-2 before sending B-2 and B-3. The result is out-of-order packets on the aggregation port.

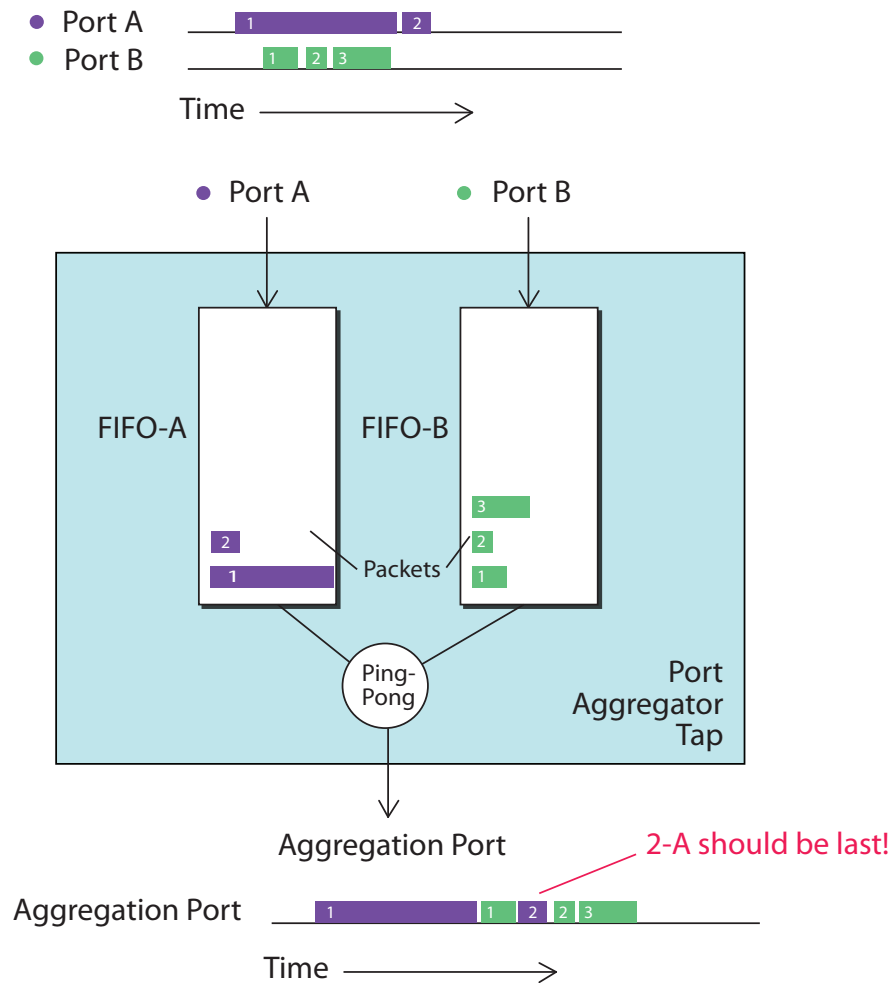


Figure 5: Packet order can be incorrect with ping-pong algorithm

This problem case occurs often enough in real-life traffic that packet order errors can seriously impact analysis of aggregated traffic streams.

Port Aggregator Tap Guarantees Correct Packet Ordering

Solution Brief

The solution

Net Optics offers a feature in some of our Port Aggregator Taps that guarantees correct packet order in the aggregated traffic stream. This is accomplished by internally timestamping the arrival of each packet, and storing this timestamp with the packet in the FIFO buffers. When the Tap pulls packets from the FIFO buffers, instead of using a ping-pong algorithm, it examines the timestamps of the available packets and always chooses the packet with the earliest timestamp. (Of course the timestamp is discarded and not sent out the aggregation port. However, an additional Timestamping feature may be enabled to substitute the timestamp for the CRC bytes in the packet for debugging or timing analysis purposes.)

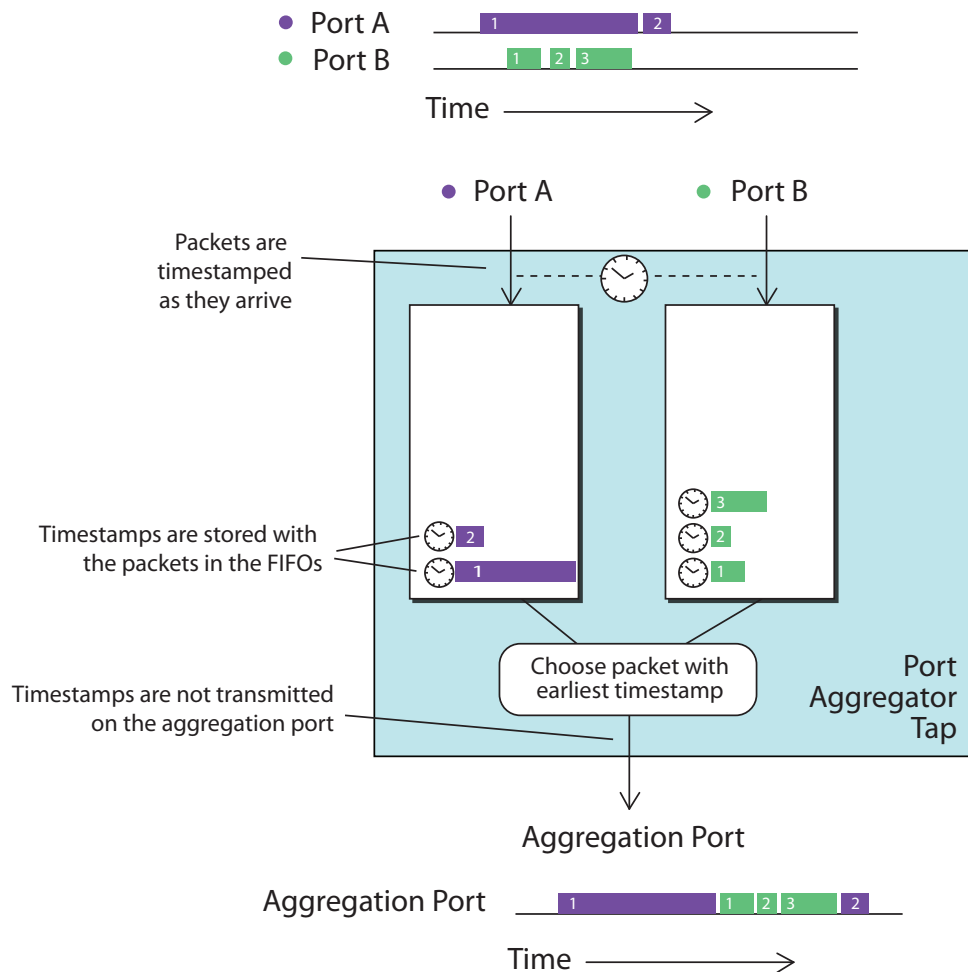


Figure 6: Timestamping solves the packet ordering problem

With this Packet Ordering feature enabled, protocol, performance, and forensic analysis applications can depend on the Port Aggregator Tap to accurately preserve the sequence of packets in time, removing the hazard of false errors due to packet dis-ordering by the Tap.

Port Aggregator Tap Guarantees Correct Packet Ordering

Solution Brief

Half-Duplex mode

An additional limitation remains when using port aggregation: When the aggregated bandwidth of the full-duplex traffic exceeds the bandwidth of the aggregation port, the aggregation port is not physically capable of transmitting all the traffic and packets must be dropped. In other words, as long as bandwidth utilization on the link is less than 50 percent in both directions, all of the traffic can be passed to the aggregation port. But if the combined traffic is more than 100% of the available bandwidth, packets are dropped.

While it is impossible to overcome the limitation of the aggregation port's bandwidth, the new Net Optics Port Aggregation Taps offer an alternative approach to the problem. By engaging the Tap in Half-duplex mode, the link traffic arriving at Port A is mirrored to monitor Port 1 and traffic arriving at Port B is mirrored to monitor Port 2, performing the function of a standard half-duplex network Tap. In Half-duplex mode, all of the traffic from both directions is passed to the monitoring tool, even if the bandwidth utilization is 100 percent in both directions (200 percent aggregated). Of course, now the monitoring tool needs two NICs to access all of the traffic simultaneously, since aggregation is not happening. However, Half-duplex mode can be a valuable time-saver because you don't have to wait for a maintenance window or get a configuration change approved to swap your Port Aggregator Tap for a standard half-duplex Tap when a high-utilization situation appears.



Figure 7: Half-duplex mode supports 100% utilization in both directions

Conclusion

The Packet Ordering feature now available in Net Optics Port Aggregator Taps brings a new level of usability by guaranteeing that the aggregation operation does not change the sequence of packet ordering on a link. Monitoring and analysis tools are assured that full-duplex conversations can be observed in an aggregated traffic stream without false errors being introduced by packet dis-ordering. Moreover, Net Optics offers additional value-add features in our new Port Aggregator Taps including Timestamping for precise timing analysis and Half-duplex mode for high bandwidth utilization environments. Combined with other Net Optics monitoring access devices, Port Aggregator Taps create a versatile monitoring access platform (MAP) for improved visibility and security-threat management across the entire network.

For further information on Tap technology:

<http://www.netoptics.com>

Net Optics, Inc.

5303 Betsy Ross Drive

Santa Clara, CA 95054

(408) 737-7777

info@netoptics.com

Customer First!