Maintaining Network Visibility While Implementing Port Channel Architecture

GARLAND TECHNOLOGY

See every bit, byte, and packet®

The Enterprise Bandwidth Dilemma

The any-to-any connection of bandwidth-intensive applications such as HD video and real-time data backup solutions is forcing large organizations to find new ways to increase their bandwidth capacity efficiently. Network speeds are constantly evolving, but it's not always economically feasible for a company to migrate its infrastructure.

Gigabit Ethernet is essentially a necessity for enterprises in the financial services, retail and government sectors, but even that is quickly being overwhelmed by traffic demands. Moving to 10G network speeds might be the next move; but what if demands don't quite warrant a full shift from gigabit to 10G? Worse yet, what if the network is already at 10G speeds but the next move is a costly jump to 40G or 100G? Enterprise data centers need a more flexible solution for meeting bandwidth demands without actually having to migrate the entire network.

Port channel architecture is a stop-gap approach to meeting these bandwidth demands. Infrastructure migration is always the end-goal, but increasing network speeds is a decision that could take years to finalize. Instead, port channel architecture virtually widens the data pipes to help enterprises scale capacity.

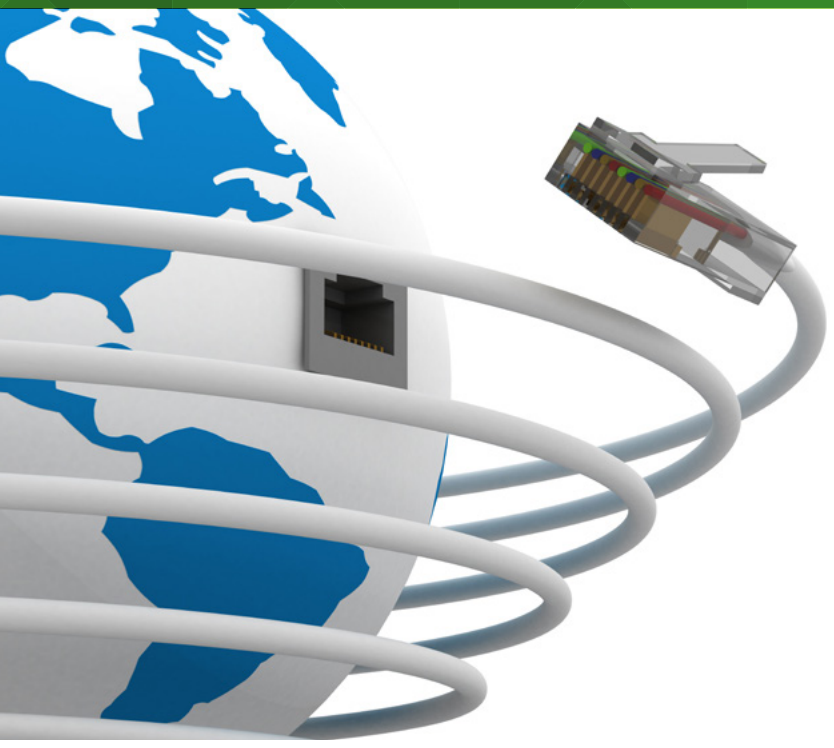

Implementing port channel architecture is a great performance solution, but virtually unifying separate physical links into one larger pipeline creates significant traffic monitoring challenges. To support port channel architecture without losing network visibility, architects must configure their networks to support aggregated and load balanced traffic well-above what the network can normally manage.

The Origin of Port Channel Architecture - Cisco's EtherChannel Technology

Port channel architecture is simply a vendor-agnostic name for Cisco's **EtherChannel®** technology. **According to Cisco**, "EtherChannel technology builds upon standards based on 802.3 full-duplex Fast Ethernet to provide network managers with a reliable high-speed solution for the campus network backbone."

Rather than limiting network element connectivity to one link, EtherChannel supports up to four links to scale up to 800Mbps (for four Fast Ethernet links), 8G (for four 1G links) or 80G (for four 10G links).

It's important to view this stop-gap approach to bandwidth flexibility not as a means to cut costs, but rather as a means to bridge the gap between infrastructure migrations without succumbing to dropped packets.

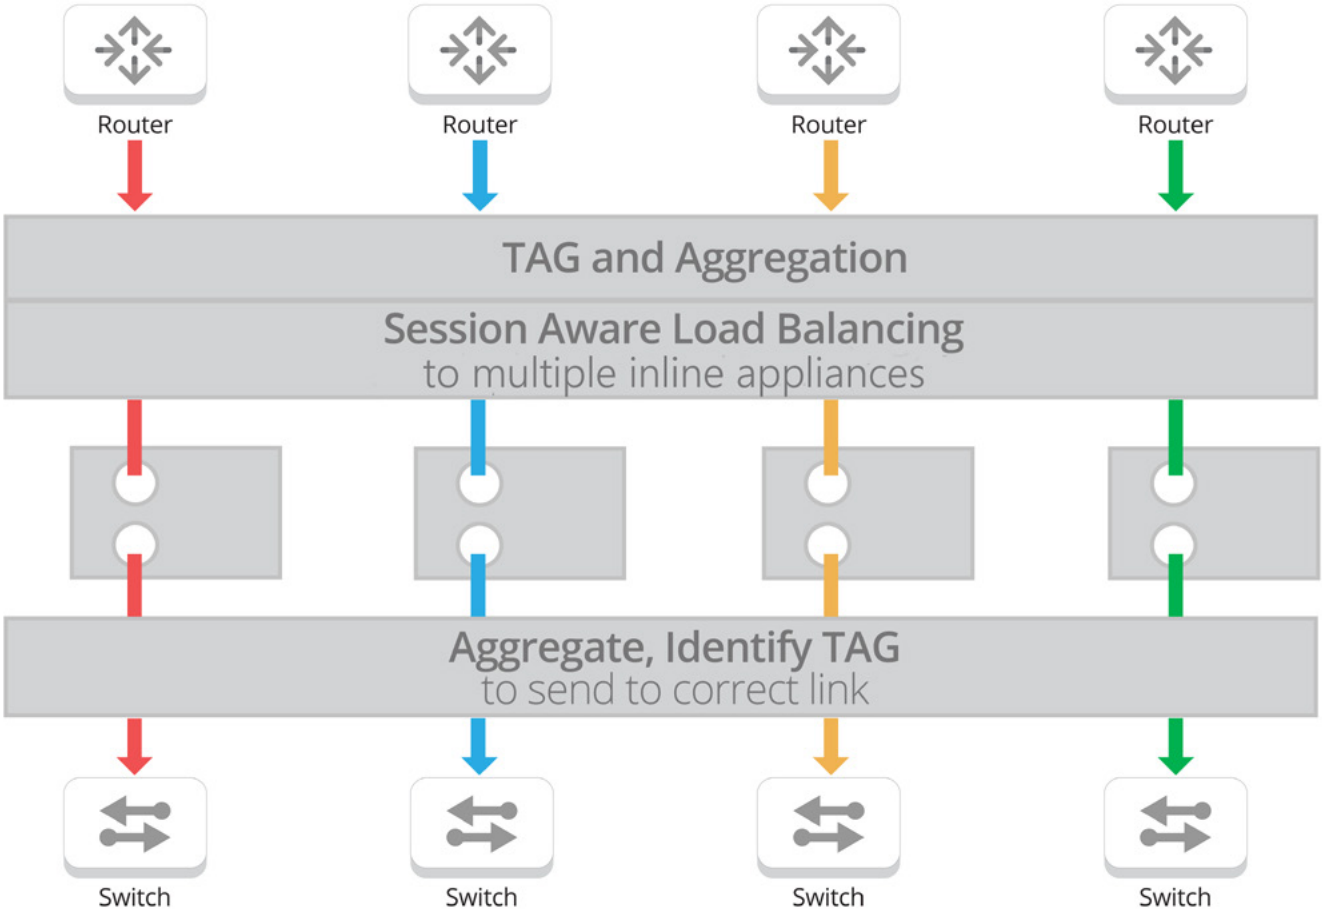

Why Port Channel Architecture is the Right Stop-Gap for Meeting Bandwidth Demands

Consider a network environment with a 1G link between a router and a switch. Many companies attempt to scale bandwidth by simply adding more 1G links to the network element connection. However, there is a fundamental issue with this approach - routing and switching protocols such as the spanning tree protocol (STP).

STP exists on the LAN side of an Ethernet network to act as the logical center of traffic with the primary job of blocking redundant links. When links are simply added to a router/switch connection, STP blocks see them as redundant and render them inactive unless in the case of failover. STP blocks help avoid loops in the network, but are not ideal for architects searching for ways to increase bandwidth quickly.

The Transparent Interconnection of Lots of Links (TRILL) protocol has largely replaced STP because of its support for seemingly redundant links while still protecting against loops. However, both TRILL and STP exist in Layer 2 (the Data Link Layer) of the network which limit their ability to help enterprises expand capacity.

Port channel architecture utilizes Layer 3 of the OSI stack model (the network layer) to enable greater flexibility in network design when planning for greater bandwidth demands. The approach aggregates physical links between network elements to virtually scale bandwidth capacity without introducing concerns about redundancy.

Cisco uses a proprietary port aggregation protocol (PAgP) to enable their EtherChannel technology. However, this technology can be accessed by any enterprise - not just Cisco shops - with the right use of aggregation, load balancing and traffic filtering.

The key to successfully implementing port channel architecture is to understand the differences between configuring the technology to monitor traffic with out-of-band appliances and ensure visibility for in-line security appliances.

Out-of-Band Monitoring Appliance Visibility for Port Channel Architecture

With port channel architecture, network managers can take a standard 10G connection and turn it into two or four physical links that are recognized as one virtual pipeline. This yields 20G or 40G of bandwidth in both the eastbound and westbound directions - and 40G or 80G of traffic overall.

The specific monitoring problem with boosting bandwidth so much is that port channel architecture allows the traffic to flow down one virtual lane despite the fact that there are up to four separate physical links. Although the the traffic can all flow to the connected network TAP, the wrong approach allows packets to travel down any of the four links. Consider the following example scenario:

- In a conversation over VoIP, packets are sent between a router and a switch.
- With port channel architecture with four links between the router and switch, packets from one party are sent partially on Link 1 and partially on Link 2.
- The second party in the VoIP conversation sends packets over Link 3 and Link 4.

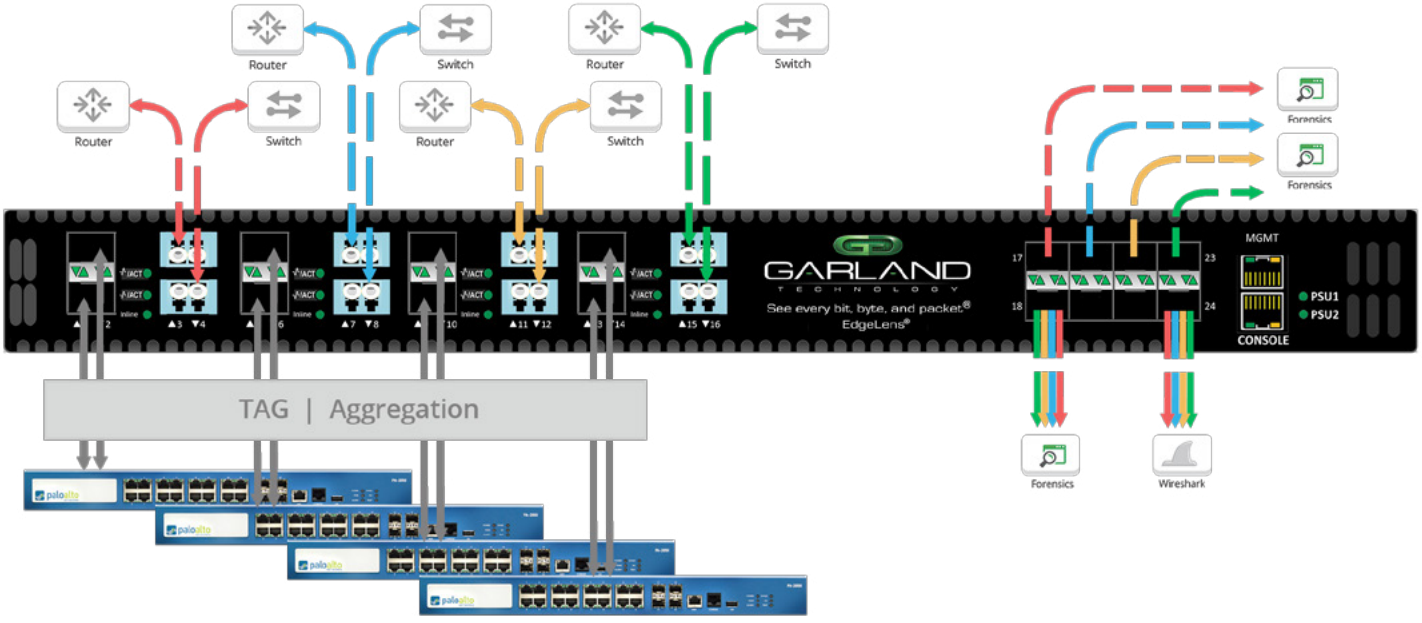

For network managers who understand the advantages of network TAPs for total visibility, it might seem reasonable to tap all four of these VoIP links to ensure every bit, byte and packet® is captured. However, tapping these segments of the conversation can't provide a total picture of the traffic for monitoring purposes. For out-of-band monitoring appliance visibility in port channel architecture environments, network managers must make use of aggregation and load balancing features of network TAPs such as **Garland Technology's FAB.**

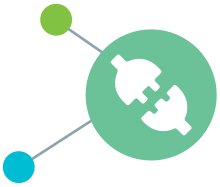

Aggregation: Network managers plug all four of the physical links between the router and the switch into the FAB. The FAB takes all of the traffic across these links and combines them into one set of packets. When the packets have been aggregated, the FAB can then send traffic to monitoring devices.

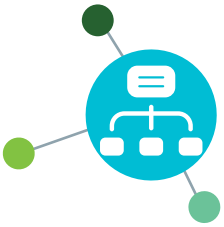

Load Balancing: Because port channel architecture enables such high bandwidth in a network that doesn't fundamentally support such high speeds, load balancing is necessary to avoid any oversubscriptions. The load balancing function of the FAB takes the, say, 40G of port channel architecture traffic and distributes it evenly amongst out-of-band monitoring appliances. Network managers can break the traffic up into 1G or 10G segments and send it to an equivalent number of monitoring appliances to ensure the network can handle the traffic that port channel architecture allows.

Because the FAB doesn't have any proprietary protocols, network managers are free to configure the network TAP and its aggregating and load balancing functions in accordance with specific network demands.

Garland's FAB can ensure 100% network visibility for out-of-band appliances in port channel architecture environments, but supporting active, in-line security appliances is another matter. The FAB can still be used, but it must be accompanied by external network TAPs.

In-Line Security Appliance Visibility for Port Channel Architecture

The aggregation and load balancing processes are the same when delivering network visibility to both passive out-of-band and active in-line appliances. However, active security solutions require an extra step because they must send traffic back through the network TAPs.

Because the traffic is being sent to active in-line appliances, this traffic will be returned back into the network after being examined by the in-line appliances. This traffic must be tagged by the FAB so that when the in-line appliance returns the traffic back to the FAB the FAB will know where the traffic came from so it can be returned to the proper link. The FAB will remove the VLan tag before sending the traffic back into the network. When traffic from the four physical links is received by the network TAP, it is tagged to specify which link it came from. The traffic is then aggregated and load balanced to security appliances. When the appliances have scrubbed the packets for any malicious information, the traffic is sent back to the network TAPs, the tags are removed, and the packets are sent westbound on the same links that sent them eastbound.

Using a FAB and external network TAPs for in-line appliance visibility in port channel architecture environments is just one option. **Garland Technology** also offers an integrated solution for port channel architecture - the **EdgeLens® hybrid bypass TAP** with built-in packet broker.

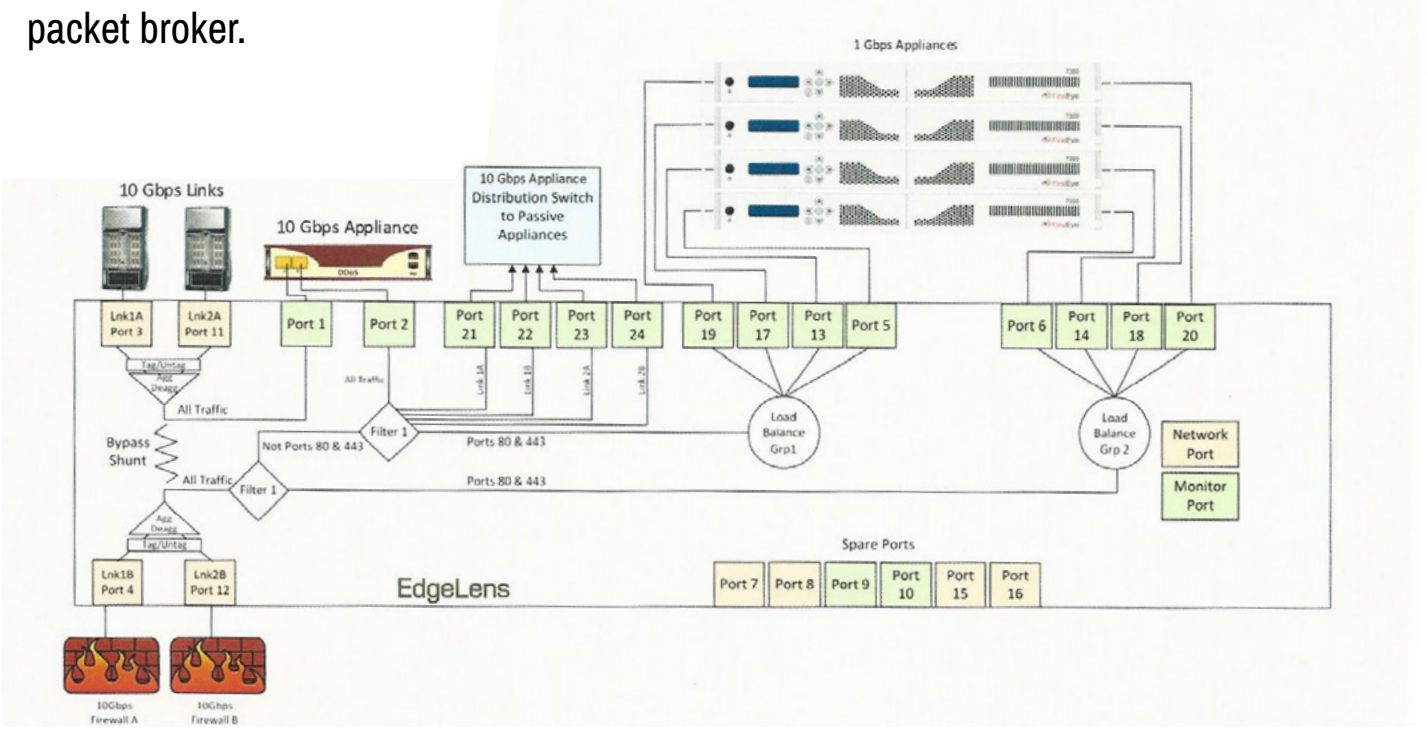

Understanding the EdgeLens® Integrated Port Channel Architecture Solution

With the EdgeLens® integrated network TAP solution, network architects can consolidate their designs by supporting multiple active and passive appliances from one box. The EdgeLens® can monitor ports to become "session aware," load balancing aggregated traffic either to passive distribution switches or active, in-line security appliances.

Port 80 and Port 443 are where Internet traffic enters the enterprise network. These ports must be monitored by active security appliances. The EdgeLens® tags this traffic and load balances it to an appropriate distribution of web application firewalls (WAFs). While the EdgeLens® ensures traffic begins and ends on the same link after being scrubbed by WAFs, it also sends 100% of the traffic to a distribution switch for passive monitoring appliances.

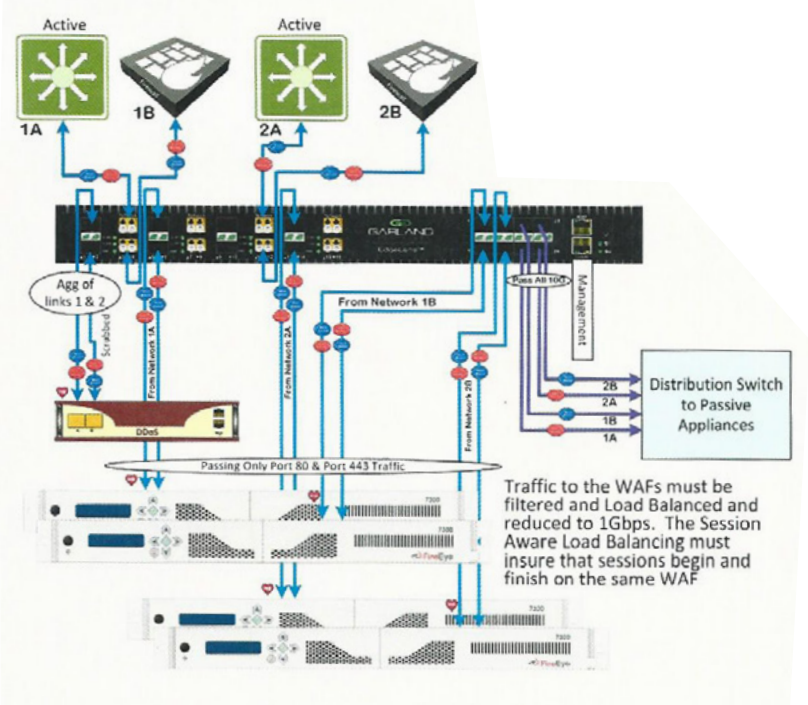

By combining the aggregation and load balancing with the tagging abilities of packet brokers all in one box, network architects can meet bandwidth demands without having to deploy extra equipment and making the network design more complex.

The Bottom Line about Port Channel Architecture

While the technical details regarding the tagging, aggregating and load balancing processes necessary for port channel architecture are complex, the reasons to adopt the technology are simple - bandwidth demands are growing out of control and enterprises can't always migrate their infrastructures accordingly.

Whether you're planning to make the jump from 1G to 10G or 10G to 40G/100G, port channel architecture can act as a quick stop-gap for meeting immediate bandwidth needs - as long as network managers ensure visibility with proper connectivity.

The EdgeLens® hybrid bypass TAP isn't made specifically for port channel architecture. However, it's combination of filtering, aggregating and load balancing features, as well as it's use of Garland's tagging process, make it an ideal solution for companies looking to implement Cisco's EtherChannel technology in a more vendor-agnostic environment.

If you want to learn more about implementing port channel architecture without losing sight of every bit, byte and packet® within your network, **contact Garland Technology** for a closer look at the EdgeLens® integrated solution.

Garland Technology is all about connections – connecting your network to your appliance, connecting your data to your IT team, and reconnecting you to your core business. It's all about better network design. Choose from a full line of access products: a network TAP that supports aggregation, regeneration, bypass and breakout modes; packet brokering products; and cables and pluggables. We want to help you avoid introducing additional software, points of failure and bulk into your network. Garland's hardware solutions let you **see every bit, byte, and packet®** in your network.

Contact

Sales, quotations, product inquiries:
sales@garlandtechnology.com

Garland Technology, LLC.
New York | Texas | Germany