Managing The Edge

Security Starts at the Edge of Your Network

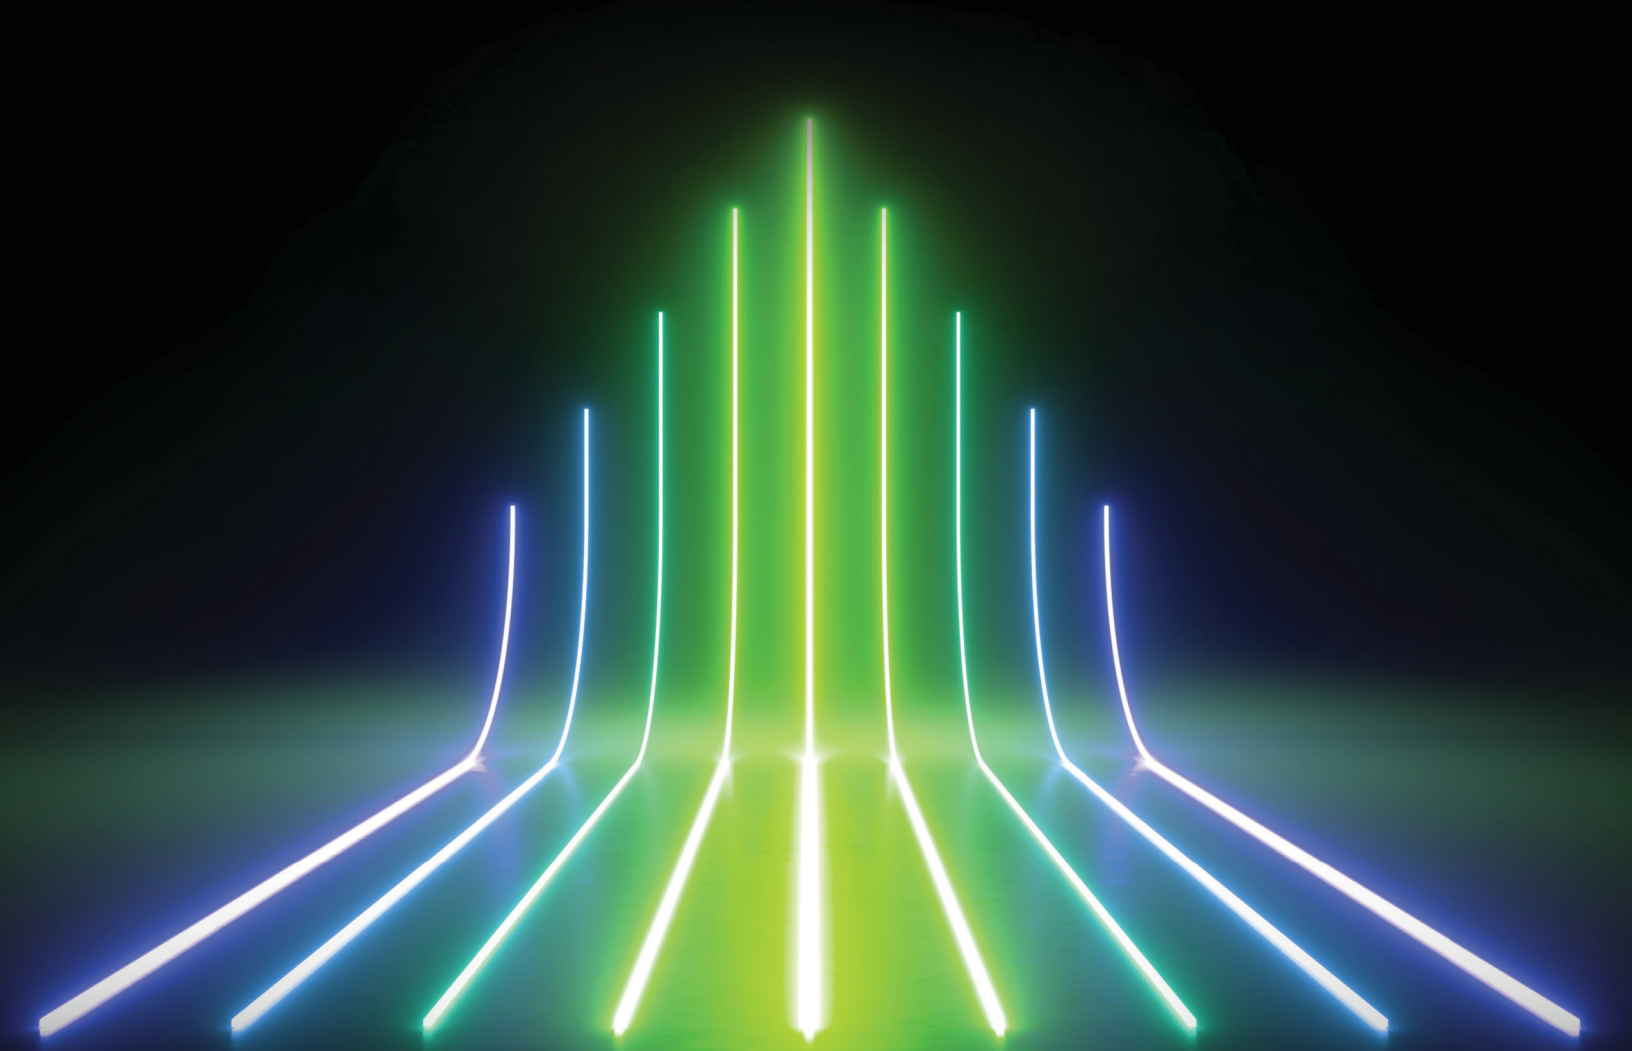

OTD BİLİŞİM
GLOBAL VAD

GARLAND TECHNOLOGY

The Evolution of Enterprise Network Security

For networking professionals who have been around since the early 1990s when firewalls first came into play, there's an understanding that there were certainly simpler days of cyber security. In fact, back then there were very few network security engineers, let alone the security architects that exist today.

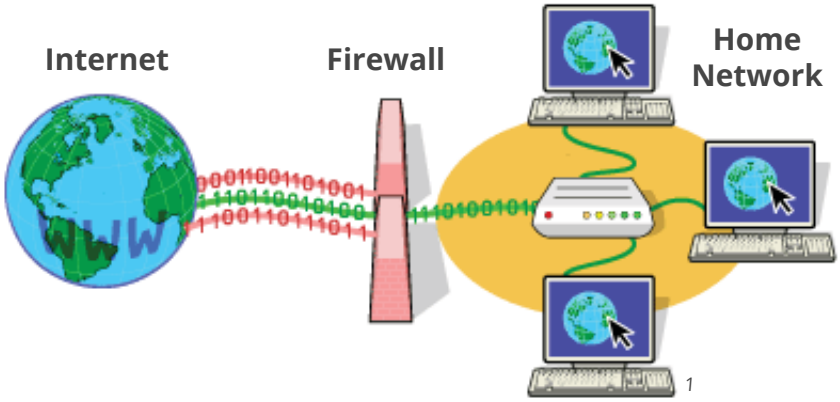

For network security engineers in the 1990s, threats were weak enough and traffic demands were low enough that they could simply pull the plug on the network, taking it down entirely to troubleshoot. At that time, networking professionals were focused on building up the network core to improve packet forwarding — managing the edge of the network wasn't even a thought.

> The new networking reality is a hybrid network supporting internal business apps and cloud-based solutions — all of which function at the network's edge.

With cyber attackers becoming increasingly sophisticated, security architects are tasked with bolstering network edge defenses with inline appliances.

Image by How Stuff Works

Early iterations of firewalls were simply placed on the live wire to capture traffic data; but networks no longer consist of just a primary link with a secondary link positioned for any necessary support. Instead, security architects must contend with multiple links and load balance accordingly. Having so many links connected at the edge of the network has caused the cyber security challenges that exist today, and has left security architects with multiple pressing questions:

- **How do I connect more than one tool at the edge?**
- **How do I access data at the edge?**
- **How do I distinguish between inline and out-of-band connectivity?**
- **How do I manage my day-to-day traffic?**
- **How do I install multiple devices and manage them individually?**

In addition to these barriers to modern network security, security architects must ensure protection without sacrificing any uptime. According to recent research from Ponemon Institute, the average cost of an unplanned data center outage is almost $9,000 per minute. To meet the demand of 100% network uptime, security architects are turning to bypass network TAPs to manage their inline device(s), which allows for simple, flip-of-a-switch ability to take an active inline device and make it out-of-band for maintenance and troubleshooting, while still providing 100% network uptime.

This is a successful strategy when tapping 1 or 2 links, however there are multiple tapping scenarios that security architects must understand to properly meet the demands of their environments.

Challenges When Deploying Many Inline Security Appliances

Network security has evolved to a point that companies have a somewhat standard set of inline security appliances that go beyond basic firewalls. These appliances include: Next-gen firewalls (NGFW), Distributed denial of service protection (DDoS), data leakage prevention (DLP), intrusion prevention systems (IPS), compliance systems, web application firewalls (WAF), SSL Decryption, network monitoring, forensics and more.

While these inline security appliances have become all but necessities for protection, there is still reluctance on the part of security architects to deploy the full stack. Some of the more debilitating challenges that keep security architects from deploying so many inline appliances include:

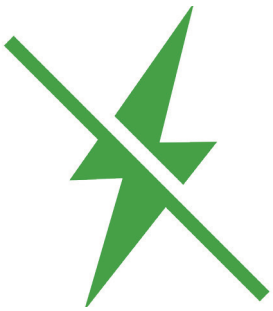

The Possibility of Network Outages

Security architects face challenges from two perspectives. On the one hand, not deploying the necessary set of inline security appliances can result in network outages due to data breaches. However, without proper deployment, inline security appliances can quickly become a single point of failure (SPOF) in the network. With such focus placed on network uptime, network outage concerns can often seem insurmountable.

Deployment Complexity

When security architects need to deploy inline security appliances into an enterprise network that was designed years ago, it can be difficult to determine proper and efficient placement. Legacy infrastructure may also cause incompatibilities with modern tools.

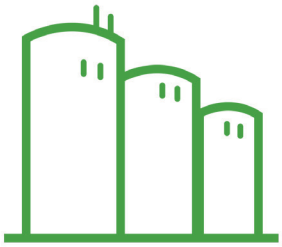

The Silo Issue

Today's network trends are all about breaking down silos for greater efficiency, but as inline security appliances are deployed they can become silos themselves. As links are routed to specific appliances, improper deployment can lead to segmentation that leads to a disconnected security stack.

Inline Security Appliances Require Many Ports

Many security architects turn to SPAN ports to connect out-of-band security appliances. However, there are only so many SPAN ports for connectivity and they can easily be overwhelmed by traffic demands. Without a support system of network TAPs, it's nearly impossible to meet the visibility needs of all the security appliances.

All of these challenges exist within the controlled environment of a single enterprise network. However, the new reality for many large enterprises is an increasing number of branch offices and remote sites.

> With the number of cyber attacks increasing exponentially each year, security architects must find an effective way to deploy these inline security appliances and manage the edge of the network — both locally and remotely.

Benefits of Using Inline Bypass TAPs

In today's security landscape, tools that sit inline with production traffic will typically have the ability to fail-to-wire if the appliance goes down. But this doesn't eliminate the need for bypass TAPs. When bypass TAPs are coupled with the ability to remotely manage the bypass functionality of the TAP, network administrators gain significant control over their network. Bypass benefits include:

1 Quickly rule out causes of network issues by enabling “forced bypass mode” on the bypass TAP. This will effectively remove the inline appliance from being inline with the production network. If the network problem still persists, then you have immediately ruled out that security device from being the cause of the issue. All throughout this process the bypassed security appliance will still be receiving packets, so it can be brought back inline without missing a beat.

2 Eventually a security device will need a firmware update. The process of updating an appliance generally requires a maintenance window where the network will be affected as changes are made. By making use of the “forced bypass mode” functionality, the network can remain up while the inline appliance is taken down for maintenance. If there are redundant security devices, forcing bypass mode can trigger the high availability (HA) appliance to take over. Once maintenance is complete on one device, that device can be brought back inline and the redundant appliance can then be taken down for its maintenance.

3 When setting up a new remote site or introducing a new inline security device to the network, a bypass TAP can expedite the process with minimum down-time. Instead of scheduling time, disconnecting cables to plug in the new device, and then completing the base configurations, having a bypass TAP already in place can allow a new appliance to be configured and then seamlessly brought into the network once the configuration has been validated. This can again be accomplished through “forced bypass mode” as copies of live network traffic will still be passed to the new appliance, but the new appliance will have no way to affect the production network. This way you can configure the device with actual traffic without worry.

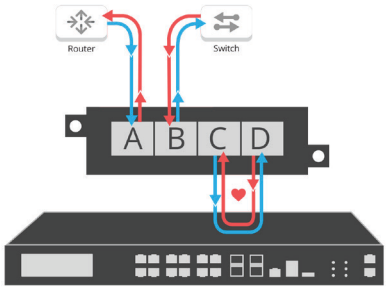

Diagram 1. Bypass mode, active inline

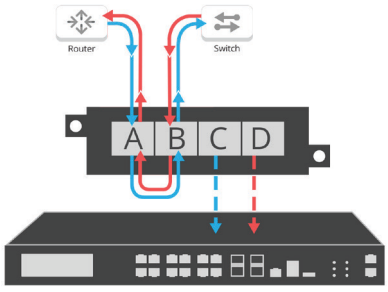

Diagram 2. Bypass mode, off-line

Chaining the Edge
What Today's Network Visibility Should Look Like

Previous white papers and articles have discussed the inefficiencies of SPAN ports for network visibility and security appliance connectivity; but security architects need a more concrete answer to network edge management questions.

To overcome the glaring inline security appliance and network edge management challenges, security architects must turn to "chaining" for proper deployment. Chaining (aka daisy-chaining) is a simple wiring scheme in which multiple devices are wired together in a sequence. Having inline security devices chained in se- quence at the network edge forces traffic to be inspected and sanitized before it can enter or exit the network. While network security is increased through chaining, a traffic flow bottleneck is created: if any of the security appliances goes down, the entire network will go down.

When building a Network Edge visibility fabric, Network Packet Brokers (NPB) must be considered, as they provide the means to chain security devices without adding points of failure the network. In addition, NPBs will create a single, unified layer of visibility across all the chained security appliances.

NPBs are equipped with various specifications, including network speed, number of ports, and supported optics. What defines a NPB is the ability to connect multiple inline or out-of-band devices with the functionality to aggregate, filter, regenerate, and load balance to the appliances it serves. When it comes to managing inline appliances at the edge, NPBs must be able to monitor the health of the active inline devices and prioritize the integrity of the network link should the inline appliances fail.

With this layer mediating between the actual inline appliances and the flow of network traffic, security architects can achieve the necessary level of edge protection and man- agement without worrying about single points of failure or port misman- agement. In a modern hybrid network supporting internal business apps and cloud solutions, this is how network visibility with chaining the edge should work:

1. Traffic flows from internet to the internal network where it is copied by the Network Packet Broker (NBP).
2. The NPB then routes packets through the chain of inline security appliances. The security devices can see all the data - every bit, byte, and packet.
3. When all inline security tools have played their parts, traffic flows back to the NPB and continues on into the network core.
4. Copies of the network traffic can be obtained by the NPB at any point in the forwarding process. This copied traffic is then sent to connected out- of-band monitoring solutions for analysis and forensics. Rule-based filters can also be applied for specific tools to see only the packets they require.
5. Visibility of traffic from both before and after the inline appliances allows security architects to compare packets at both ends to spot any deviations from baseline expectations.

Having a layer of visibility separating inline security appliances from the live wire means security architects can spot any performance issues and troubleshoot without disrupting the flow of traffic. Because uptime is so important for modern enterprises, the chaining approach to security architecture is essential to efficiency and appliance effectiveness.

Security architects in companies of all sizes face the same inline security appliance chal- lenges; but that doesn't mean there is a one-size-fits-all solution for implementing the NPBs that are necessary for visibility. Depending on individual network specifications, there are different solutions for managing the edge of the network.

Four Common Inline Security Appliance Tapping Scenarios

There are NPB solutions available for a wide range of inline security appliance connectivity needs. Whether the environment is a small business with one or two inline appliances or a massive enterprise operating at advanced network speeds, security architects must choose an approach appropriate for their network needs both today and tomorrow.

Four common scenarios include:

- The 1G Data Center Solution with Remote Site Management
- A High Availability (HA) Solution
- A 10G Chaining the Edge with Media Conversion
- The EdgeLens® - Advanced Edge Management

1G Data Center
with Remote Site Management Solution

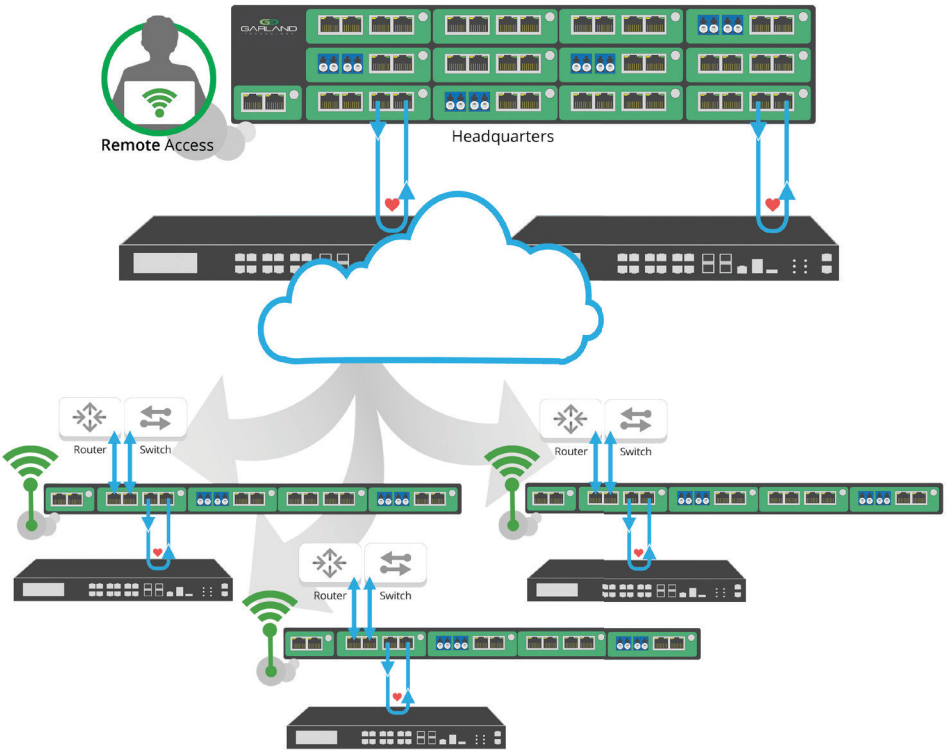

This is a common solution for companies with multiple remote locations that may not have on-site, local administrators (retail, banking branches, etc.). In these situations, the main site may have a robust security architecture with multiple inline devices. The remote sites may have a significantly scaled down security presence with only a few inline security devices. Because of the lack of on-site IT resources, ensuring the inline appliance is both active and functioning can become a difficult task. Using inline bypass TAPs with remote management helps to both ensure network connectivity and quickly troubleshoot devices without having to be on-site at the remote office.

Benefits of this solution include:

- Scalability - Add modular TAPs when required
- 1G Media Conversion - fiber to copper
- Remote management via GUI or CLI

High Availability (HA) Solutions

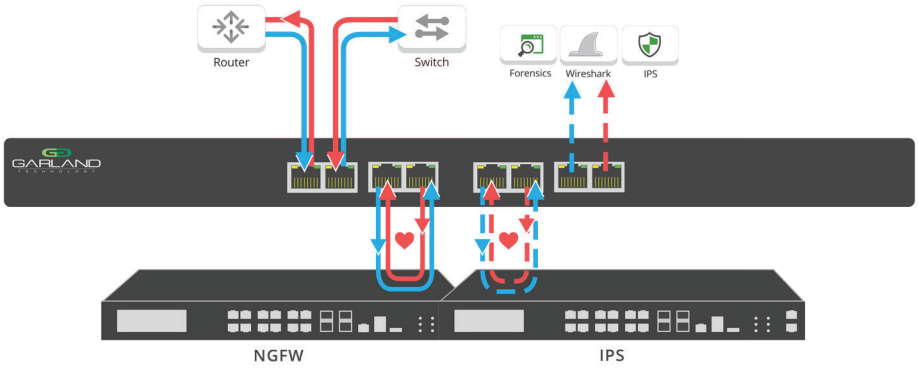

Diagram 2. High Availability (HA) solution for Active/Passive, provides failover from primary device to backup appliance

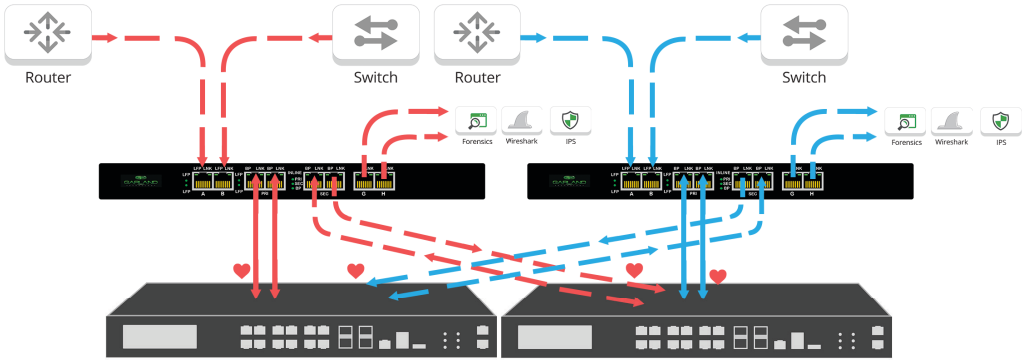

Diagram 3. High Availability (HA) solution for Active/Active, provides failover if either active device fails.

In this solution, a purpose-built inline TAP is capable of tapping a single network link and distributing the traffic out to two different inline appliances. In addition to the inline connectivity, copies of the traffic can also be sent to out-of-band tools as well.
This type of TAP was originally developed in a time when Internet Service Provider (ISP) connections were expensive and providers were limited. This type of TAP enabled network security to be highly available by using two security devices on a single link. If one security device failed, a second one was already ready and waiting to take over.

Today, the cost of internet connections has dropped considerably and having redundant links from multiple ISPs for failover has become the standard. Security devices also have the ability to work in tandem and gracefully failover in the event of a failure on one device. By adapting the deployment design of this type of inline TAP to add in a second appliance, environments with redundant internet connections and security devices in HA can still benefit from the additional layer of resiliency provided by the inline TAP.

Chaining the Edge
Support Up to 4 Active, 10G Inline Devices

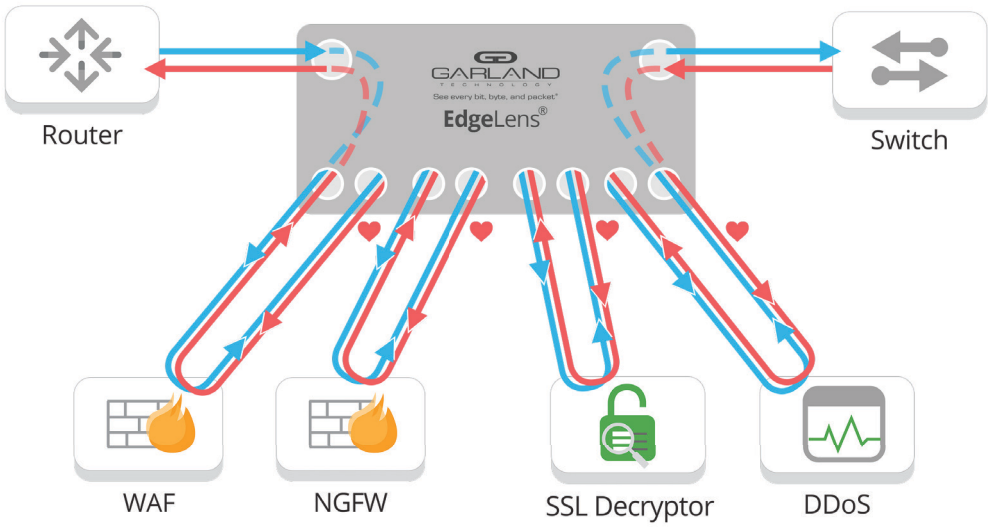

This solution is geared more toward security architects that must support high-traffic organizations (for example, retail or financial services companies). The solution monitors up to four inline security appliances and maintains the failover assurance that large enterprises require.

With a flexible 1U design, security architects can mix and match fiber and deploy up to four bypass TAPs.

While these solutions offer the functionality and flexibility that security architects require, there is still demand for a more efficient way to direct traffic from an ever-growing chain of inline security appliances. This is where the EdgeLens®, a hybrid bypass TAP with packet broker capabilities, comes into play for enterprise security architects.

How Security Architects Can Meet Advanced Edge Management Needs

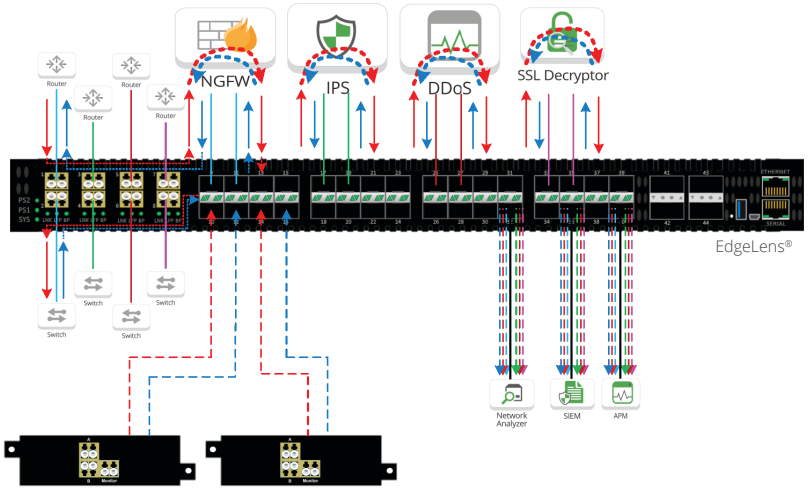

The EdgeLens® is the solution for security architects tasked with managing multiple inline security appliances and load balancing to support increasing bandwidth demands. As a hybrid bypass TAP and packet broker, the EdgeLens offers four integrated 10G TAPs with 8 built-in failsafe network ports and 32 SFP+ and 4 QSFP+ monitoring ports.
This enables security architects to filter, aggregate and load balance the inline data stream of security appliances and monitoring solutions.

With EdgeLens, security architects are free to tap a 10G circuit and filter it for separate 1G appliances. This is a cost effective way to attach multiple devices through a single network TAP all within a 1U chassis for data center efficiency.

In this example, we highlight management of up to four (4) inline appliances, while replicating traffic for use with all out-of-band tools. The EdgeLens has four 1G/10G TAPs that can provide inline bypass capabilities for up to four 1G/10G inline security devices.

Benefits of chaining with the EdgeLens:

- Gain intelligent and optimized packet visibility
- Access to both inline and out-of-band security/monitoring tools
- Enable real-time security proof-of-concept evaluations without impacting the network
- Shield monitoring devices from cyber attackers
- Increase efficiency of inline and out-of-band tools

Conclusion

Security architects are under tremendous pressure to manage the edge of the network, defending the valuable core from increased cyber threats. The more typical tapping scenarios offer flexibility, but the EdgeLens is the all-in-one solution for security architects looking to achieve an effective chaining approach.

If you want to learn more about how you can manage the edge of the network with the EdgeLens® Inline Security Packet Broker or EdgeSafeTM Bypass TAPs, contact Garland Technology today for a Design-IT consultation to discuss a security design tailored specifically to your needs.

Setting Yourself Up for Security Success

Want to learn how we can help you implement network security best practices?
Book a free Design-IT consultation and one of our engineers will work directly with you on designing your network connectivity strategy.