How to Architect a More Secure Enterprise Through Zero Trust

Zero-trust architecture has been central to government efforts to improve cybersecurity. In simple terms, zero trust assumes that any user (or device) on your network could be a malicious actor. In that case, users' identity and network privileges should be verified each time they attempt to access network resources, not just at the perimeter.

Think of zero trust like security in an office building. There should be two main priorities:

1. **Least privilege:** Every employee's key pass might get them through the front door, but that doesn't mean they should have access to every office or storage room in the facility. They should only get access to the resources needed to do their job.

2. **Network segmentation:** For least privilege to be effective, there must be barriers in place — just like an employee would need to use their key card to enter a different floor or hallway of the building. This segmentation stops what is known in zero trust as lateral movement.

But while the concept is straightforward, the implementation can be tricky, especially as agencies dig deeper into the technical details. That is the case with network segmentation.

Stumped by Segmentation

According to a 2023 survey ...

93% of IT and security decision-makers say segmentation is critical to preventing attacks such as ransomware.

30% of organizations have deployed segmentations in more than two mission areas.

WRONG WAY:

How Not to Do Segmentation

The problem is that the most common approaches to network segmentation are riddled with shortcomings. That includes:

Software Agents

- **Manageability/scalability:** Requires software agent in every workload
- **Server resources:** Agents draw on compute resources, impede performance
- **Cost:** Annual subscription costs become exorbitant at enterprise scale
- **Vulnerability:** If host systems are compromised, agents can be bypassed

Traditional Network Switches (Stateless Access Control Lists, or ACLs)

- **Stateless:** Provides no session tracking for traffic flows
- **Lack of Transmission Control Protocol (TCP) security checks:** Leaves servers vulnerable to denial-of-service attacks
- **Scale limitation:** A switch can only support a limited number of ACL entries, with no support for group-based policy
- **Static rules:** ACLs unable to adapt as workloads are spun up or down or migrated within the data center

Traditional Next-Generation Firewall Appliance

- **Complexity:** Appliance becomes a traffic choke point
- **Capacity/cost:** Multiple firewalls needed to keep up with traffic
- **Static policies:** Next-gen firewalls not made for dynamic environments
- **Limited segmentation:** No option for granular application segmentation

Those approaches fall short because they repurpose existing solutions to serve a new requirement. And none of them are a good fit, forcing agencies to compromise on security, performance or manageability.

Stateless vs. Stateful

Stateless

A stateless system does not track traffic and therefore can only follow static policy rules around address information.

vs.

Stateful

A stateful firewall recognizes flows in a network and keeps track of the state of sessions, making it possible to dynamically block traffic that deviates from unexpected patterns — a key tool, for example, in defending against distributed denial of service attacks.

Zero Trust-Ready:

The Distributed Services Switch

Fortunately, there's another way: a switch with integrated stateful software services.

Key capabilities include:

- A distributed stateful firewall for monitoring traffic within the data center
- Zero-trust segmentation
- Pervasive telemetry
- Stateful Network Address Translation (NAT)
- Encryption
- Load balancer

Benefits of this approach include:

- **Distributed and scalable architecture:** It reduces network congestion and complexity.
- **Increased cost-effectiveness:** With all services built into the switch, this provides similar total cost of ownership to a traditional data center platform.
- **Context-aware policies:** It accommodates the dynamic nature of data center virtualization.
- **Increased visibility:** It sees all workload-to-workload communication, allowing the platform solution to provide more granular segmentation.
- **Zero impact on server resources:** All traffic inspection happens on the switch with hardware acceleration, not on the rack servers, with minimal latency.
- **Increased resiliency:** With security enforcement through the switch, segmentation remains effective even if servers are compromised.
- **Comprehensive protection:** This covers any connected data center asset, such as bare metal servers, hypervisors, virtual machines, storage servers, etc.
- **TCP security checking:** TCP handshake validation, half-open session tracking, DDoS protection, etc., are included as part of the stateful firewall feature set.
- **Highly scalable rules:** Security policy has practically unlimited scaling, up to millions of rules.

Unlocking the Potential of Zero Trust

Zero trust is a federal mandate because it addresses the most pressing security needs of the modern agency enterprise. But to achieve its full potential, agencies need to get the technical details right, and that begins with network segmentation.

The Aruba CX 10000 series switch with Pensando, available from Iron Bow, is tailored to support the transition to network segmentation. It delivers stateful software-defined services inline, at scale, with wire-rate performance and orders of magnitude scale and performance improvements over traditional data center switches, at a fraction of the total cost of ownership. This is how zero trust becomes a reality.

To learn more visit here.

govloop
IRON BOW TECHNOLOGIES
HPE aruba networking