HPE aruba networking

Business white paper

Data center microsegmentation

Deploying distributed, stateful east–west firewall with the HPE Aruba Networking CX 10000 with AMD Pensando

HPE GreenLake

The HPE Aruba Networking CX 10000 with AMD Pensando™ has brought to market a new category of enterprise switch—a distributed services switch—with unique characteristics particularly suitable to address a diverse set of demanding use cases across data centers, edge, and cloud. This includes functionality typically delivered by enterprise-class services appliances and foundational, best-in-class networking enabled through HPE Aruba Networking AOS-CX data center switches.

Historically, IT organizations have focused on the fabric for connectivity—but that is only half of the challenge. Infrastructure services need to support disaggregated application scale. These have all been different appliances or VMs in the fabric—but historically have not been part of the fabric. This leads to complexity, different vendors to manage, traffic tromboning across the fabric, and complexity between the networking and services teams.

It's time for a better way

The HPE Aruba Networking CX 10000 Series Switch

The multifunctional CX 10000 performs functions typically delivered by expensive services and network appliances—such as firewalling and security policy enforcement, VPN tunnel termination, and network address translation—at a reduced cost, in a variety of network roles for enterprise data center, cloud, and edge deployments.

Integrating stateful service capabilities within a data center switch, the CX 10000 moves security and visibility closer to where applications and workloads are processed without changing an organization's existing network architecture or software configurations. This creates significant opportunities for any organization to improve its data center security posture and visibility while reducing the cost of acquisition and simplifying operations.

In this paper, we will unpack the technical and TCO advantages of deploying distributed network and security services for organizations who are considering the CX 10000 to provide stateful east–west firewall within the network fabric—at the top-of-the-server rack (leaf switch).

Top-security use case—isolation / segmentation / Zero Trust

Distributed stateful east–west firewall segmentation in data centers

The cybersecurity threat landscape has changed dramatically in recent years. Today, adversaries are more motivated than ever to penetrate enterprise data centers and steal valuable information.

In response, adopting the concept of Zero Trust is the number one trend in enterprise security practice today. For the data center, this means trusting no entity on the network by default and distrusting all traffic unless a security policy explicitly allows it. According to NIST SP 800-207,

"Zero Trust security models assume that an attacker is present in the environment" and that a Zero Trust architecture is "designed to prevent data breaches and limit internal lateral movement."

Segmentation is key to preventing unwanted lateral movement, by statefully inspecting all east–west traffic in the data center and applying policies that stop bad actors from moving through the internal network.

Several approaches have been tried to achieve segmentation in the past, but with minimal success:

- Hardware **traditional firewall** appliance-based segmentation
- **Virtualized firewall appliance-based** segmentation
- **Software agent-based** segmentation
- **Network switch** (stateless ACL)-based segmentation

Distributed firewall services vs. traditional firewall

A solution based on traditional (next-generation) firewalls would appear to be ideally suited to stopping east–west lateral movement, even though they are originally designed to inspect north–south traffic.

Whether physical or virtual appliances, they were repurposed by many organizations to serve as internal east–west firewalls to segment the network.

However, there are several key problems with this approach.

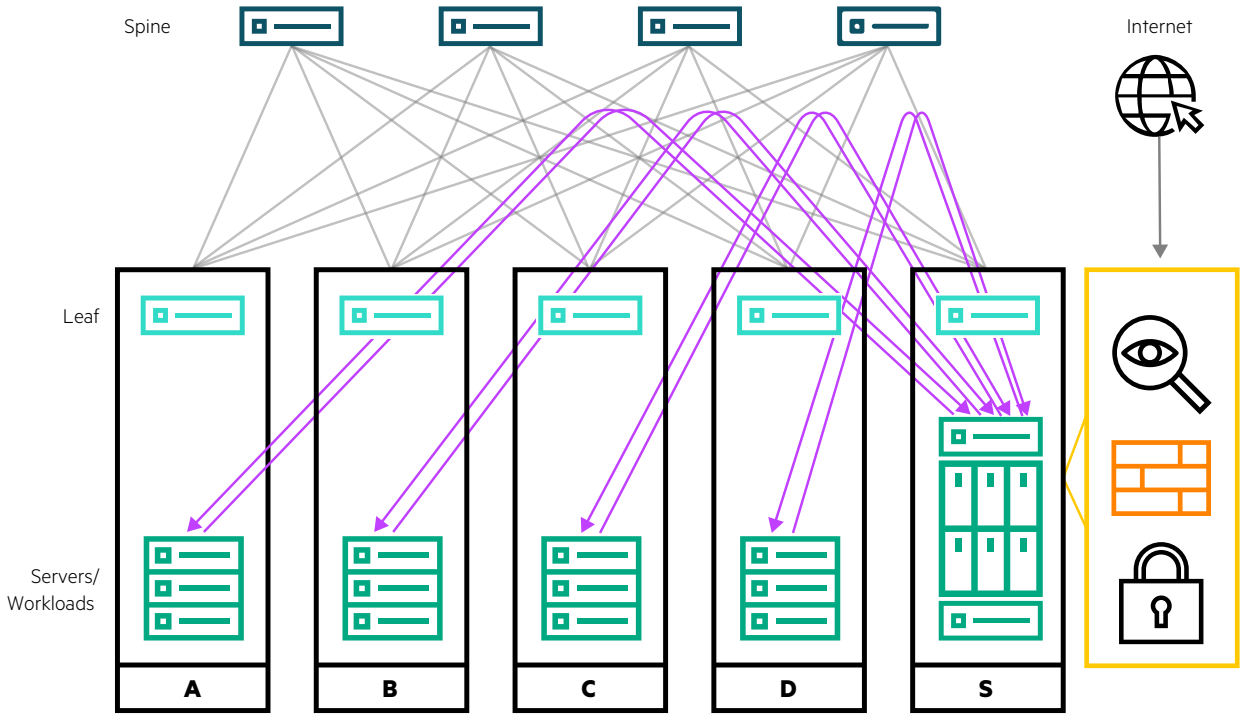

Figure 1. Centralized services architecture. Leaf-spine fabric with services leaf

- **Complexity:** As shown in Figure 1, traffic tromboning is a big issue when inserting appliances in modern data center fabrics. Getting traffic to and from them symmetrically is challenging, requiring techniques like policy-based routing (PBR) or VLAN stitching.
- **Capacity and cost-prohibitive:** Using next-generation firewalls for east–west traffic can be cost-prohibitive. Such firewalls quickly run into capacity problems if trying to inspect all internal data center traffic, creating the need for multiple firewalls that must be periodically upgraded to address traffic increases.

- **Static policies:** Next-gen firewalls don't explicitly consider data center application architecture in their design. They remain blind to the relationship between workloads and applications. Today an application may comprise multiple workload types, microservices, and containers that run in the data center or the public cloud. In a virtual environment, workloads can be spun on and off at any minute and moved around dynamically within the data center. The same IP address can be reused as the workloads come and go. All these make traditional firewalls very ineffective as firewalls to enforce stateful policies for east–west traffic.

Distributed firewall services vs. software agents

The typical software agent-based solution consists of two components: an agent and an orchestrator. The agent resides inside the virtual machine, the pod, or the physical server and enforces security policies distributed by the orchestrator. While this approach allows for very granular control over the interaction between workloads inside a data center, it has several challenges.

- **Manageability and scalability:** This approach requires the installation of the software agent in every workload. In a large enterprise data center environment with more than 10,000 workloads, maintaining this vast number of agents will soon put significant pressure on the support organization, affecting the solution's manageability and scaling.
- **Server resources:** All traffic enforcement and filtering are done in the software. No matter how small the percentage of CPU and memory a vendor claims to consume for its solution, in the real world, host-based agents sap compute resources and impede performance, wasting precious resources that otherwise could be used for business applications.
- **Cost:** Customers usually purchase a subscription license on a per-agent per-year basis. In a large deployment, these annual subscription costs become exorbitant at enterprise scale.
- **Vulnerability:** The proliferation of agents is also becoming a security issue as they are user-level processes that run inside the host operating system, if the host is compromised, these agents can be stopped or bypassed, effectively stopping segmentation.
- **Protection coverage:** Any component in the data center that cannot support the installation of a software agent will remain unprotected. This may include the hypervisor control plane, storage servers, appliances, unsupported operating systems, or a memory copy of VMware vSphere® Storage vMotion® traffic.

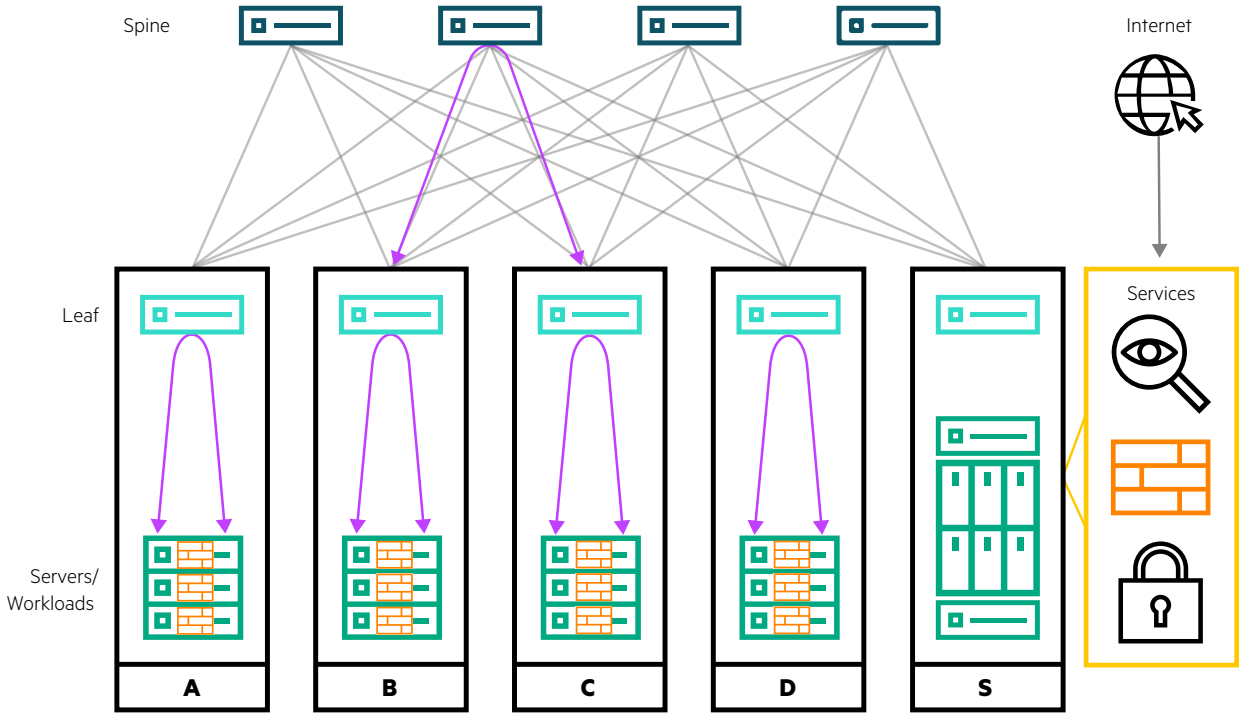

Figure 2. Centralized (distributed firewall agent) services architecture. Leaf-spine fabric with software agents

Distributed firewall services vs. network switches (stateless ACLs)

This approach leverages the standard stateless ACL feature supported in most network switching platforms to filter east–west traffic within the data center. Stateless ACL-based filtering merely leverages generic access lists and checks if a packet matches the source, destination address, and port numbers in the header to make a permit/deny decision. Whether they receive a single packet or thousands, each packet is treated individually. The switch doesn't track the ordering of packets in a TCP session or maintain or understand the connection state.

This approach will quickly run into many restrictions for a large data center deployment.

- **Stateless:** Its stateless nature provides no session tracking for traffic flows. Lack of session tracking requires the policy to explicitly include rules for the return traffic of each flow, which effectively opens all ports to the server and provides zero protection for clients.
- **No application layer gateway (ALG) support:** Applications that use multichannel protocols, such as FTP or MSRPC, require data channels to be created dynamically on the fly. These dynamic channels have known significant difficulties in association with stateless ACLs.
- **Lack of security:** A standard ACL solution does not provide any security checks for TCP sessions, such as handshake validation, or half-open session checking. This means application servers are vulnerable to DDoS attacks.
- **Scale limitation:** Standard ACLs in network switches are usually implemented in the switch's TCAM table for faster speed. However, TCAM table size is limited in most switch products, which results in a limited number of ACL entries that can be supported. In addition, the lack of support for group-based policy makes it even less scalable.
- **Manageability:** Standard ACLs are difficult to manage at scale. Implementing the right set of ACL entries at the correct port or SVI can quickly become a nightmare for network administrators.
- **Static rules:** By design, standard ACL rules are implemented by simply using layer 3/4 IP addresses and ports. They remain blind to the relationship between workloads and applications. In a virtual environment, workloads can be spun up or down at any point and dynamically migrated within the data center, making standard ACLs very limited and cumbersome when enforcing policies for east–west traffic.

The HPE Aruba Networking CX 10000 overcomes the limitations and challenges of all these legacy alternatives, providing a more robust, scalable, and flexible solution.

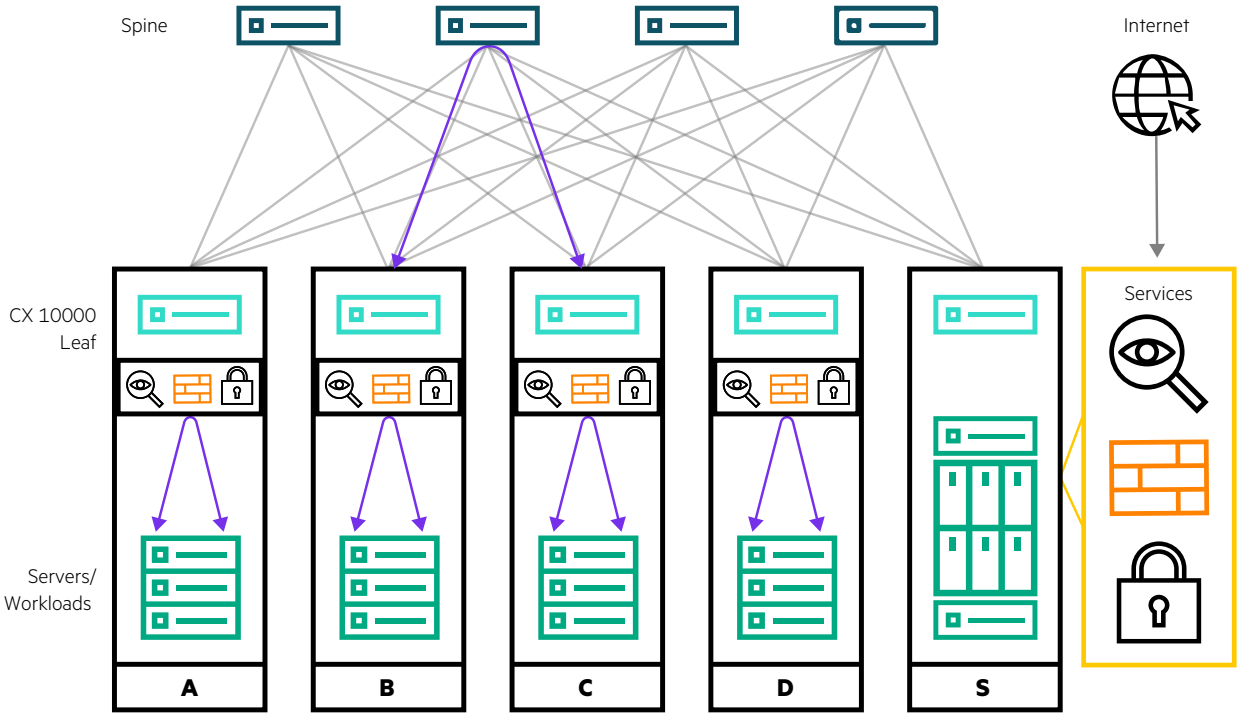

Figure 3. Distributed services architecture. HPE Aruba Networking CX10000 leaf switch / fabric accelerated network and security services

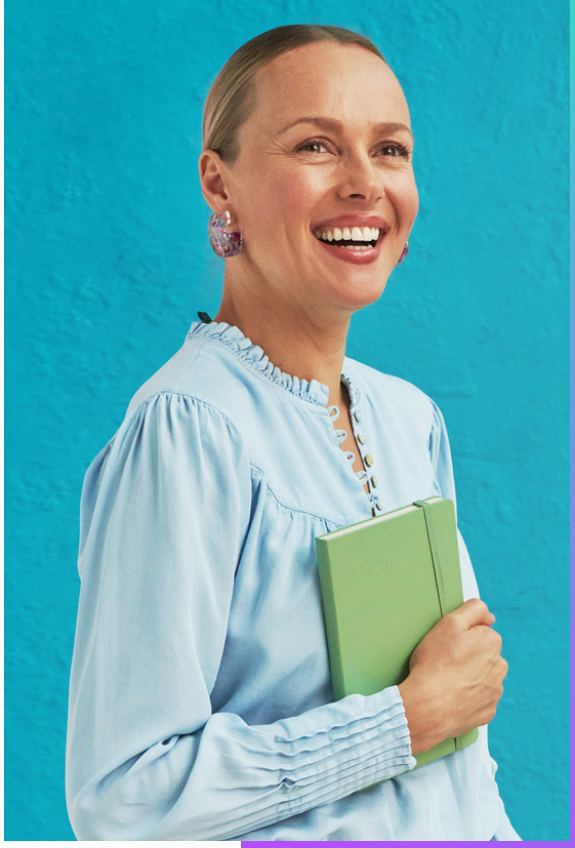

Advantages of using the CX 10000 solution include:

- **Distributed and scalable architecture:** The architecture inspects traffic right at top-of-rack (ToR) switches, removing the need to hair-pinning traffic to traditional centralized appliances and reducing network congestion and complexity. Aggregate east–west inspection capacity scales as additional server racks are added to the pod. active/active stateful firewalling with state-sync via VSX, and the firewall state follows if a workload moves from one rack to another.
- **Cost-effective:** All stateful services, including firewall, come built into each and every leaf switch in the data center, priced optimally within range of the TCO of a typical standard data center incumbent platform.
- **Context-aware policies:** Context-aware segmentation policies accommodate the dynamic nature of data center virtualization. Segmentation rules can be configured based on the virtual workloads' tags or names. As virtual workloads are spun up, deactivated, or relocated; their associated policies remain effective without manual reconfiguration.
- **Visibility and granular segmentation:** Because all server traffic traverses the ToR, the CX 10000 has visibility to all workload-to-workload communications, no matter if they are inside the same server, on different servers, within the same rack, or across different racks. This deeper visibility allows the CX 10000 to provide granular segmentation across the entire data center fabric.
- **No impact on server resources:** All traffic inspection happens on the CX 10000 switch with hardware acceleration outside the rack servers, with minimal latency.
- **Stateful connection and ALG support:** The CX 10000 solution intrinsically supports stateful connection tracking. Return traffic of existing flows is automatically allowed. ALG support is also built in, including FTP, TFTP, and MSRPC compared to a network switch (ACL).
- **TCP security checking:** Comprehensive TCP session security checking—such as TCP handshake validation, half-open session tracking, and DDoS protection—is included as part of the stateful firewall feature set, compared to a network switch (ACL).
- **Highly scalable rules:** The policy rules in the CX 10000 solution are implemented in data plane DRAM, instead of TCAM, with the support of a highly efficient lookup algorithm. This method removes the TCAM size limitations while maintaining line-speed traffic filtering. The CX 10000 security policy has practically unlimited scaling, up to a million rules.
- **Simplified policy and management:** By design, CX 10000 security policy is configured at the entire cluster level. Sites need only consider who is allowed to talk to whom inside the data center when configuring the security policy, without worrying about how policies will be deployed and where; which CX 10000 switch and which port is automatically determined by the management plane, which will apply these policies for optimized deployment.

The HPE Aruba Networking CX 10000 offers superior security and lower TCO

Prior to the availability of the CX 10000, DC network and security operators were having to deploy a combination of L2/3 switching (Cisco, Arista with VXLAN), providing only a course-level segmentation (with no stateful services or Zero Trust).

To address the requirement for stateful data center security services (east–west firewall, microsegmentation), operators were/are faced with having to deploy traditional hardware-based NG firewalls (Fortinet, Juniper), costing hundreds of thousands of dollars each while only providing limited scale

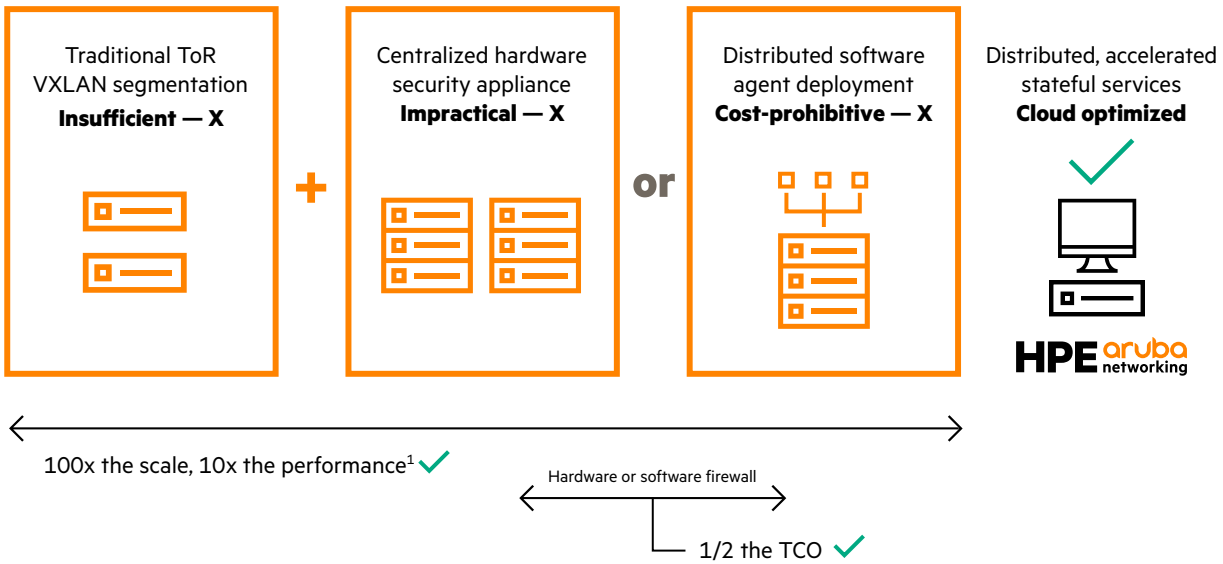

Figure 4. DC network segmentation, east–west firewall deployment options. Stateful services, 800G, east–west firewall segmentation across 500 servers/32 racks

[1] Based on the internal HPE analysis.

Deploying software-based firewall agents on all their servers (for example, Illumio, NSX) is another option to provide east–west firewall/microsegmentation—but this option is also very disruptive, difficult to manage, and very costly.

Let's take a closer look at the TCO advantages that the CX 10000 offers over each of these options.

In this example, the HPE Aruba Networking CX 10000 provides a lower TCO of $1.069M (or 53%[2] savings) over three years, versus a traditional design configured with next-gen firewalls and standard L2/3 top-of-rack Ethernet switches.

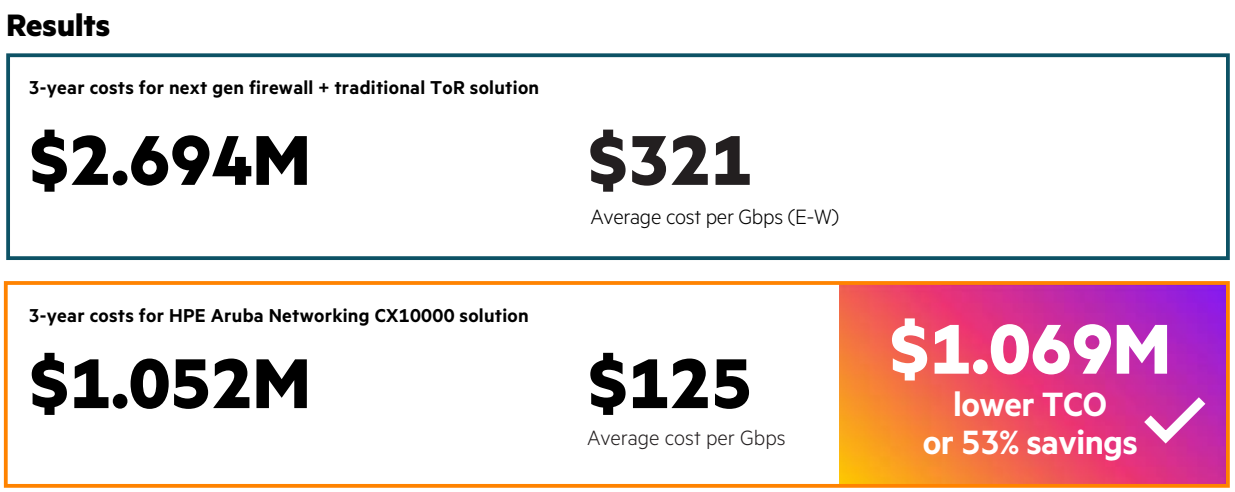

Figure 5. Three-year TCO analysis of new DC deployment with 500 servers implementing east–west firewalling

Three-year cost for next-gen firewalls and standard L2/3 ToR Ethernet switches
Assumes 500 servers, 24 server racks, 400 Gbps/rack, 8400 Gbps total bandwidth. Assumes 75% and east–west traffic needed inspection. Requiring eight traditional NG firewalls and two standard L2/3 ToR switches per server rack. Compared with an HPE Aruba Networking CX 10000 with AMD Pensando configuration providing accelerated stateful firewall services embedded into the L2/3 ToR switch.

In this next example, the HPE Aruba Networking CX 10000 provides a lower TCO of $1.269M (or 57%[3] saving) over three years, versus a traditional design configured with software-based firewall agents and standard L2/3 top-of-rack Ethernet switches.

Figure 6. Three-year TCO analysis of new DC deployment with 500 Servers implementing east–west firewalling

Three-year cost for software-based firewall agents and standard L2/3 ToR Ethernet switches
Assumes 500 servers, 24 server racks, 10 workloads/server, 5000 agent licenses and two standard L2/3 ToR switches per server rack. Compared with an HPE Aruba Networking CX 10000 with AMD Pensando configuration providing accelerated stateful firewall services embedded into the L2/3 ToR switch.

[2, 3] The TCO analyses are based on hypothetical examples, using specific industry assumptions. Individual customer configurations will vary based on specific designs and configurations.

Conclusion

While data center networking has evolved over the past decade, providing higher performing 25/100/400G leaf-spine topologies to address the volume and velocity of emerging application architectures, security and services architectures have not.

With the explosive growth of east–west traffic in the data center, centralized security appliances are proving inefficient, expensive, and difficult to manage. Simply put, hair-pinning traffic to an appliance sitting at the data center edge introduces heavy performance, cost, and operational penalties.

This problem is further exacerbated by microservices-based applications, where traffic may not even need to leave a physical host to go from one service to another. This means some application traffic may never be inspected by a hardware firewall, IPS, or other security device—leaving enterprises vulnerable to attack from within the enterprise itself.

The HPE Aruba Networking CX 10000 Series Switch with AMD Pensando provides an entirely new class of switching solution to overcome the limitations of legacy architectures. The CX 10000 will allow operators to extend industry standard leaf-spine networking with stateful distributed microsegmentation, east–west firewalling, NAT, encryption, and telemetry services—all delivered inline, all the time, on every access port, closer to where critical enterprise applications run.

Our HPE Aruba Networking distributed services architecture expands Zero Trust deeper into the data center, to the network-server edge, delivering fine-grain microsegmentation, dramatically scaling and strengthening the security of mission-critical workloads—delivering 100x the scale and 10x the performance at 1/3 the TCO of traditional solutions.

Learn more at

ArubaNetworks.com/products/switches/distributed-services-switches/

Visit **ArubaNetworks.com**

Make the right purchase decision. Contact our presales specialists.

Contact us

Hewlett Packard Enterprise

a50010559ENW