NETRONOME

WHITE PAPER

Enabling Efficient and Scalable Zero-Trust Security

FOR CLOUD DATA CENTERS WITH AGILIO® SMARTNICS

THE NEED FOR ZERO-TRUST SECURITY

The rapid evolution of cloud-based data centers to support virtualized services and applications has created significant challenges for data center security architectures. Traditional data center security strategies have typically relied on perimeter-based security appliances, while assuming that the interior of the data center could be trusted. But in today's multi-tenant environments, tenants and applications cannot be trusted, and the potential for threat injection inside the data center is significant. To properly protect tenants and application workloads in this "Zero-Trust" environment, security functions must be distributed and associated with the tenants and workloads directly with fine-grained control. This problem is exacerbated by the dramatic increase of east-west traffic in the data center, driven by the virtualization and distribution of applications driven by centralized control and automation

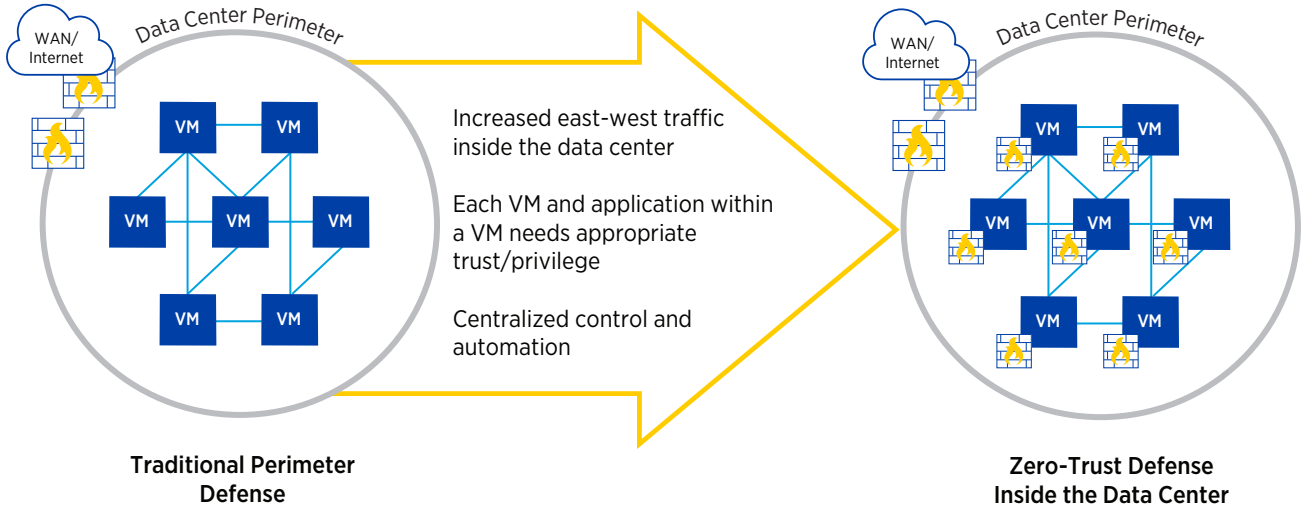

Traditional Perimeter Defense

Zero-Trust Defense Inside the Data Center

SECURING THE WORKLOADS, NOT JUST THE NETWORK

Traditional security approaches have involved protecting resources at the network level, such as specific ports, VLANs, and network addresses. This could typically be accomplished with simple static rules. But in today's multi-tenant cloud environments, security policies must be much more tenant and application-specific, resulting in more complex and fine-grained policy rules. These fine-grained rules may include virtual overlay network identifiers and application

related information, for example. This type of fine-grained security control that is associated with specific applications or workloads is often referred to as zero-trust security.

THE NEED FOR STATEFULNESS

When talking about workload or application-level security, an important requirement that comes up is the ability to apply stateful filtering on a per-connection basis. The key here is to consider the state of a flow in combination with matching on the header fields to determine if a packet should be permitted or denied. Stateful inspection requires every connection passing through the firewall to be tracked so the state information can be used for policy enforcement. A typical implementation would use a connection-tracking table to report flows as NEW, ESTABLISHED, RELATED, or INVALID to the firewall. The policy applied to the system then dictates how packets are evaluated within the stateful firewall. This type of bidirectional connection tracking is critical to adequately protect many server applications.

AGILITY AND ELASTICITY IS KEY

A critical advantage of cloud networking is to have the ability to quickly and easily migrate workloads across servers and data centers to maximize resource utilization, provide resiliency, and support rapid scale-out. This means that security policies that are associated with tenant and application workloads must also be able to migrate seamlessly along with them. In modern SDN data center environments, this implies tight integration of security policy distribution within the control and orchestration infrastructure, such as OpenStack. The server networking datapath is the typical enforcement point for these fine-grained security policies. The server networking datapath is typically implemented with a virtual switches or routers for forwarding and overlay control, while Linux netfilter iptables functions are typically used for stateful security policies. The figure below shows a typical OpenStack deployment with Security Groups that are programmed into the server networking datapath to provide agile security control while working in concert with the overall orchestration and management infrastructure.

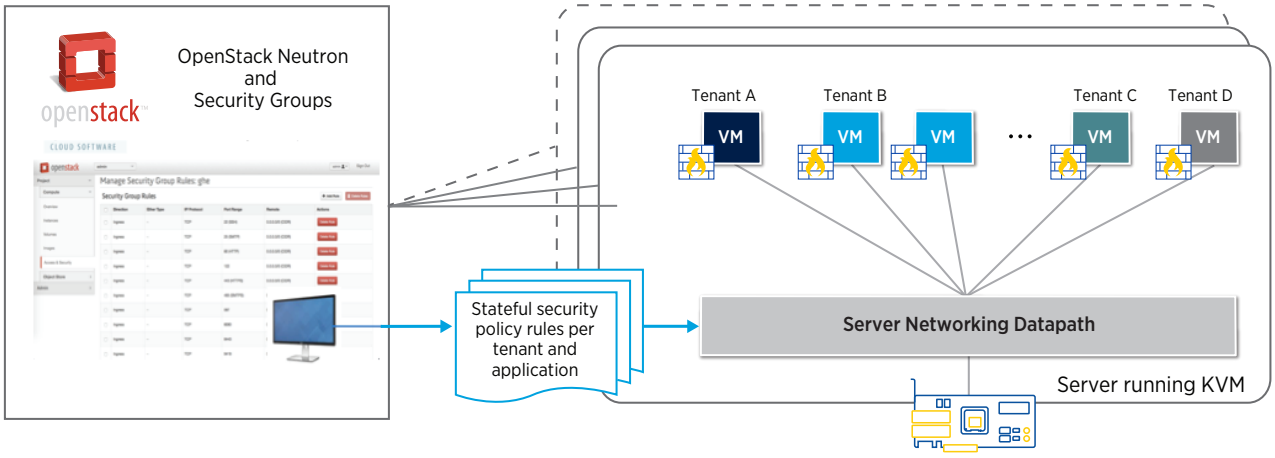

THE EVOLUTION OF SECURITY SUPPORT WITH OPENSTACK

OpenStack has a feature called Security Groups that provides a mechanism to implement stateful security in the form of connection tracking on a per VM basis. The connection tracking function itself, also referred to as Conntrack, is implemented in the Linux IPTables module. But

since most cloud data center deployments use OVS for network overlay processing and other tasks, this resulted in a very cumbersome implementation, because a separate Linux Ethernet bridge needed to be instantiated and connected to the OVS bridge, in order to get the stateful connection tracking function. In addition to the increased configuration complexity, the extra processing steps increase the CPU load, resulting in poor performance. More recently, the connection tracking functionality has been integrated into OVS 2.5, eliminating the need for the extra Linux Ethernet bridge and resulting in some improvement, as shown in the figure below:

Connection Tracking with OVS and Linux Bridges

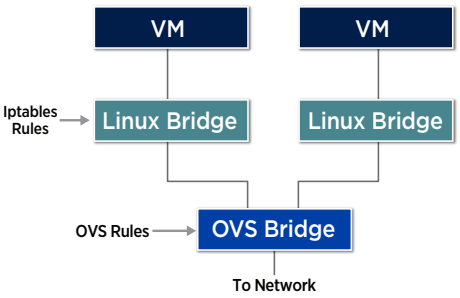

Integrated Connection Tracking with OVS 2.5

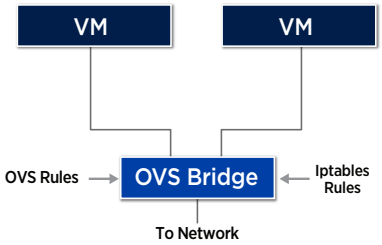

But even with the improvement in OVS 2.5, the Conntrack function itself is still quite CPU intensive and can degrade the server networking datapath performance significantly when enabled. For this reason, Conntrack is an excellent candidate for offload, together with the virtual switching function, to an acceleration device such as a SmartNIC.

PERFORMANCE AND RESOURCE IMPACTS OF STATEFUL SECURITY ON SERVERS

In modern data centers the goal is to achieve efficient scalability by leveraging the computing power of many thousands of servers and breaking data center tasks up into smaller, more manageable workloads that can be flexibly allocated to available resources. This means that a data center's service capacity will grow at the same pace (linearly) as the server footprint. An often overlooked assumption with this technique is that most, if not all, of the CPU resources in a given server are available to run the application workloads. But deploying security on servers creates a scenario where the security services are competing for the same server resources as the applications, which is counter-productive to the revenue goals of the data center. Let's examine some of the specific problems that result from server-based virtual switching with security processing:

Overlay Network Support and Tunneling: Network tunneling is an intense process from a compute perspective requiring lookups on several headers (inner and outer headers) as well as inserting new headers in the case of encapsulation. Network tunneling implemented on the server CPU becomes a significant bottleneck at high network speeds due to lack of compute parallelism. This is because the packet arrival time exceeds the packet processing time on the server, which cannot support the core or thread capacity to sufficiently parallelize these network tasks to keep up with arrival rate.

Access Control with Tens of Thousands of Rules: Using a host-based ACL or filtering table for access control exposes a major weakness when using a server CPU for networking. Server architectures rely heavily on various levels of CPU caches to achieve their greatest possible

performance. The required ACL rule capacity to support tens to hundreds of VMs per server can lead to lookup table sizes that are much larger than the cache sizes in typical servers, leading to a state of continuous thrash, which significantly reduces performance.

Stateful Tracking of Millions of Sessions: Keeping state on millions of flows exacerbates the cache-thrash problem even further. Per-connection state tracking is an additional data structure that must be accessed with every packet received, putting greater pressure on caching and memory subsystems. Server CPUs rely on high cache hit rates for networking performance, but stateful firewalling significantly reduces, or eliminates, any effectiveness of an x86 cache.

When considering security posture and data center scale-out, Zero-Trust Stateful Security is an effective strategy in theory. But as we have noted, there are severe penalties on each server in the form of bottlenecks and increased resource consumption, leaving few resources left for revenue-generating VMs, and erasing many of the benefits of cloud-scale networking.

USING AGILIO SMARTNICS TO ACCELERATE ZERO-TRUST STATEFUL SECURITY WITH OPENSTACK

The Netronome Agilio line of SmartNICs running Agilio Firewall Software is designed to accelerate server-based networking in environments requiring stateful security. Agilio SmartNICs fully offload the overlay tunneling and match/action processing of Open vSwitch (OVS), as well as stateful security processing, while supporting extremely high rule and flow counts. There is a two-fold benefit to the server when offloading the security workload to Agilio SmartNICs: First, the networking I/O bottleneck is eliminated, allowing VMs or containers to receive as much data as they can process. Second, there is significant CPU savings realized by taking the OVS and security workload and executing it on the Agilio SmartNIC instead of the server CPUs.

The Agilio Firewall Software offloads OVS as well as Linux Conntrack to Netronome's family of SmartNICs. The Agilio solution is a drop-in accelerator for OVS with seamless integration with existing network tools and controllers and orchestration software such as OpenStack. Use cases for Agilio Firewall Software include public and private cloud compute nodes, telco and service provider NFV, and enterprise data centers. In these use cases it is common to have a large number of network overlays and/or security policies that are enforced on the server with potentially several thousands of policies per VM. Agilio provides the ability to support very high flow and security policy capacities without degradation in performance.

BENCHMARKS

As we have noted, native OVS and Conntrack running on servers coupled with traditional server adapters that cannot perform security offload, struggle with packet processing which ties up valuable server CPU resources and creates a bottleneck that starves applications. Our benchmarks have shown that Agilio SmartNICs can reclaim up to 50% of the server CPU resources previously dedicated to OVS and security, while at the same time delivering up to 4X or more of the packet data throughput to application workloads.

Netronome has performed benchmarking comparisons of throughput and server CPU utilization for various packets sizes, to compare the performance of a native software-based

implementation of OVS 2.5 with Conntrack security enabled, against an Agilio-based implementation. The server hardware was comprised of a dual socket XEON totaling 24 physical CPU cores. The testing was done with a total of ten thousand stateful security rules populated in the server networking datapath — a realistic number supporting 100 rules per workload, and 100 workloads per server, as an example The results of the benchmarking are shown in the graphs below:

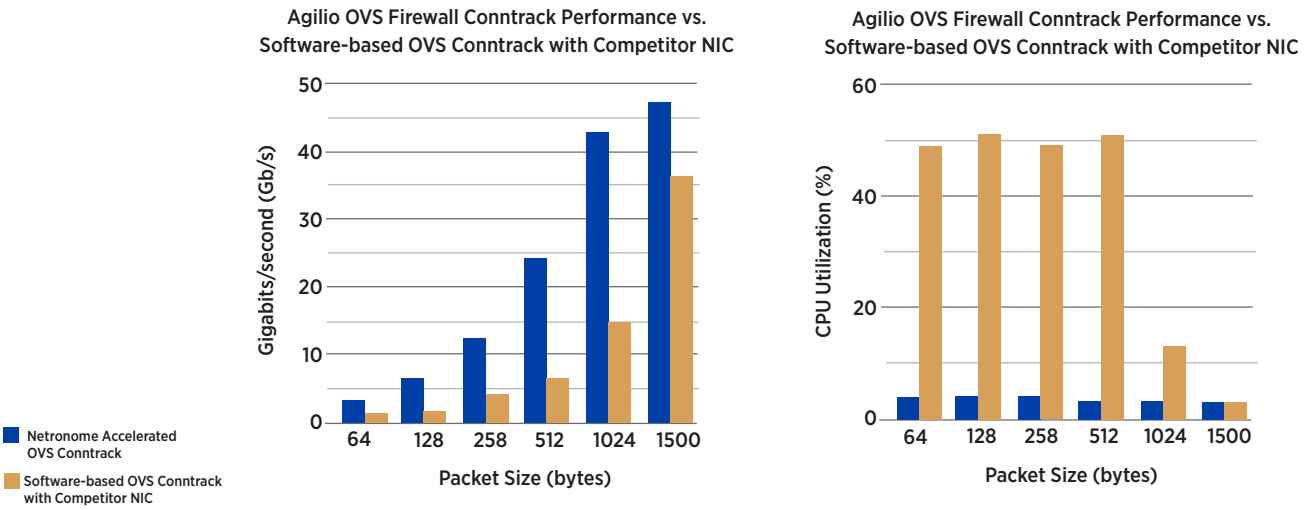

In terms of throughput improvement we can see that, at an average packet size of 512 bytes, a purely software-based OVS with Conntrack solution can only process about 6Gb/s, however with Agilio acceleration, this number reaches up to 24Gb/s, resulting in a 4X improvement in packet delivery to target workloads. At the same time, we can see that about 12 of the 24 server CPU cores were completely freed up from OVS and security processing, and are now available to run more applications.

CONCLUSION

The Agilio SmartNICs have proven to be the optimal solution for offloading stateful security with OpenStack. By offloading isolation, access control and stateful firewalling via Conntrack, Agilio SmartNICs and software prove to eliminate networking bottlenecks offering up to 4X the performance for Conntrack at high rule and flow capacities while simultaneously restoring up to 50% of the CPU cycles which can be repurposed for VM and container deployments and applications. This ability to run more VMs and containers, and deliver more data to the application workloads at the same time, is critically important to realizing the vision and benefits of cloud-scale networking.

NETRONOME

Netronome Systems, Inc.
2903 Bunker Hill Lane, Suite 150 Santa Clara, CA 95054
Tel: 408.496.0022 | Fax: 408.586.0002
www.netronome.com