Network Security Policy Across Heterogeneous Hosts

Context

Big Switch Networks' clients typically operate high scale, high security data center environments. With workloads spanning bare metal, virtual machines and containers, many face the challenge of creating homogenous security policy designs across heterogeneous host environments.

A decade ago, the physical network would have been considered the bastion of security policy. VLANs, assigned at the network edge, were the core primitive for segmenting bare metal hosts in to logical groups. ACLs or stateful firewalls enforced security policy in between. Today, however, the physical network is no longer the data center network edge. For virtualized and containerized hosts, the vSwitch edge presents both opportunities and constraints for enforcing policy in between network segments.

Unfortunately, the vSwitch edge in ESX environments has a different approach to security policy from the vSwitch edge in KVM/OpenStack environments. Container environments are again different. Advanced technologies like micro-segmentation (sub-subnet segments) are available on some vSwitch environments but there is no practical way to apply the same design to include bare metal server workloads, and even new 'guest firewall' technologies typically have some pockets of host operating system or guest software versions that are not on their compatibility lists.

Across Big Switch Networks' clients, the bad news is that we have yet to see a unified technology for policy enforcement that spans the full diversity of common host environments. Sophisticated clients use a portfolio approach to match technologies to organizational needs, turning this into an advantage.

Technologies in the Portfolio

We typically see three different categories of security policy technologies in use by advanced, security conscious organizations: physical firewalls, microsegments and SDN-style ACLs. While an exhaustive discussion of each is out of scope, a few highlights are below.

Micro-segmentation	SDN-Style ACLs	Physical/Virtual Firewalls
Configured in central controllers; enforced in vSwitches or host agents; configured by workload rather than switch port; examples include NSX, Illumio, Contrail	Configured in central controllers; enforced on leaf switches; configuration is tied to workload rather than switch port; examples include Big Cloud Fabric Logical Router ACLs or Cisco ACI Group Policy (contracts between EPG)	Configured and enforced on physical or virtual appliances; configuration is a combination of workload and port; examples include Palo Alto Networks, Fortinet, Checkpoint
(+) Policy follows workload as it moves, expands/contracts or is removed (+) Typically capable of sub-subnet segmentation (-) Typically restricted to a specific hypervisor technology	(+) Policy follows workload as it moves, expands/contracts or is removed (+) Policy is agnostic to host technology (hypervisor/bare metal/containers) (+) Easy to migrate old ACLs to SDN-style ACLs (-) Segmentation is done at the subnet level	(+) Rich advanced features (L7/deep packet) for stateful firewalling (+) Rich features for audit (-) Introduces network configuration / workload constraints (sub-optimal paths)

Firewalls vs Micro-segments: A Tale of Two Extremes

A small percentage of Big Switch Networks clients fall in one of two camps: an everything-in-the-firewall camp and an everything-in-a-micro-segment camp. Both have pros and cons.
The first approach eschews all policy that can't be implemented in a firewall. Some traffic flows may hairpin, and we hear stories about flows between VMs on the same host that traverse the entire data center network to reach a firewall for processing. While these designs make preparation for security audits a straightforward exercise of collecting the firewall configuration, these configurations also tend to get complex rapidly with critical organization-wide rules intertwined with minor, application-specific details. We have seen firewalls with over 10,000 firewall rules in complex data center environments that follow this approach. With this complexity comes longer lead times to add necessary protection or introduce the inevitable pinholes needed by application owners. This can also create bottlenecks if the firewall is not correctly 'right-sized' to the environment.

The second approach is a hero's journey of micro-segmenting an entire data center. The subnets in which both hosts and VMs are instantiated in these environments are orthogonal to security policy, a powerful argument in favor of this approach. However, in these environments, we see security policy driving host technology decisions, e.g. no software is allowed in the data center unless it is running on ESX (and segmented by NSX). For many data centers that have not only new applications but also legacy applications with no time or business case to perform a migration, and for whom subnet-level segmentation is a well worn operational path, this approach is quite heavy handed and often has only limited appeal.
The vast majority of Big Switch Networks' clients who operationalize highly secure environments at scale fall in to a third, more pragmatic bucket. Blending the two approaches above with SDN-style ACLs (details below), the technology advantages and limitations across all three serve as a blessing.

Three Tiers of Network Security: A Pragmatic Portfolio Approach

For clients in verticals such as financial services, government and security-conscious SaaS providers with complex policy and significant auditing requirements, a pragmatic three tier approach has emerged as the common case.

- **TIER 1: Security Appliances**
  Inter-business unit traffic is routed through stateful firewalls. This can be achieved through routing or L4-L7 service insertion in products like BCF. In this context, "inter-business unit traffic" is a short-hand to imply traffic that is subject not only to security best practices but also to material audit, charge-back or monitoring requirements. The policies often have a large number of internal and external stakeholders, many of them non-technical. The inefficiencies of packet pathways are often a reasonable trade-off for the conceptual simplicity these designs represent, specifically the ability to point a large group of stakeholders to a small number of physical boxes and say, "there is the policy you care about."

- **TIER 2: SDN-style ACLs**
  Intra-business unit traffic policies are enforced by SDN-style ACLs (sometimes referred to as "macro-segmentation"). Within a business unit, the chance of having a heterogeneous set of host technologies is high. We typically see a mix of bare metal databases, ESX-based VMs, KVM/OpenStack-based VMs and container technologies within the same business unit, with applications of many different vintages. A policy-follows-workload model is critical as these workloads must plan for movement due to maintenance and data center changes, but typically either applications or general business unit security zones can be easily segmented by subnet as a practical concern. The stakeholders responsible for policy here are typically a far smaller number than the example above, often the infrastructure team rather than a corporate compliance or risk team, and often are closer to the infrastructure technology evolution. We see teams at this level digesting more abstract policy enforcement than the "it is in that box over there" physical metaphor needed above.

- **TIER 3: Micro-segmentation**
  Intra-application traffic polices (for new applications) are enforced by micro-segmentation. Rather than try to back-port older applications in to micro-segments, we typically see this used for new applications. As these newer applications typically have homogenous host technologies (e.g. all on ESX, or all on Kubernetes), the downside of micro-segmentation technologies is limited. Unlike the two technologies above, in many Big Switch clients, it is the application owners themselves who are responsible for defining, configuring and maintaining the micro-segments rather than infrastructure providers or some security team.

SDN-style ACLs versus Traditional ACLs

SDN-style ACLs may not be a crisply defined term in the industry, but most of the modern switching/routing "underlay" SDN providers have these in some form. Examples include Cisco ACI's Contracts and Big Switch's Big Cloud Fabric ACLs.

Like traditional ACLs, SDN-style ACLs are typically configured between subnets. (Other options are available, but this is the common configuration.) Unlike traditional ACLs that are configured on a specific physical port of a specific switch or router, SDN-style ACLs are typically configured on central controllers in reference to a "logical" router. While the configuration at the CLI may feel identical to traditional ACLs, these SDN-style ACLs are typically enforced when a packet first ingresses into the switching fabric at the top-of-rack switch. This allows VMs within that subnet to be placed anywhere in the switching fabric with no changes to the running configuration, i.e. a policy-follows-workload model that makes these far more practical than traditional ACLs configured in reference to a physical port.

These SDN-style ACLs have a series of operational advantages over their predecessor, the port-based ACL. Using Big Cloud Fabric's ACLs as an example, a few of these are shown below.

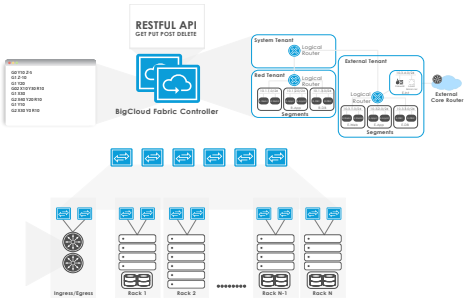

CENTRALIZED LOGICAL CONFIGURATION

As opposed to having ACL configuration spread across numerous switch/routers in the traditional port-centric ACL model, SDN-style ACLs are configured in one central location (Big Cloud Fabric Controllers), making it far easier to express the intent of the logical design, and maintain these complex configurations over the long term.

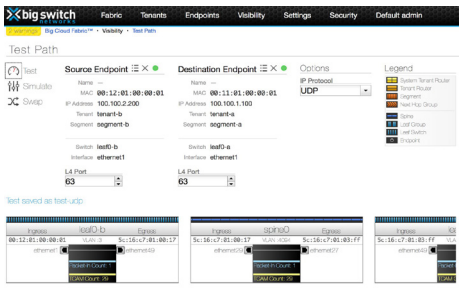

ADVANCED TROUBLESHOOTING TOOLS ("TEST PATH")

In the SDN-style model, since all ACLs along with the network topology are known in software in a single location, troubleshooting tools like Big Cloud Fabric's Test Path are available. In this screenshot, the SDN-style ACL configuration, logical router configuration and physical path relevant to any source/destination pair are available with a single command. Simulation commands are also available. Compared to legacy ACLs, troubleshooting these at scale is vastly simpler.

POLICY LOGS

While traditionally log-on-drop policy auditing was a feature only found in stateful firewalls, this dividing line is blurry in the SDN-style model. Big Cloud Fabric features a relatively advanced policy logging feature, as any ACL can send counters of matching packets to the central controller.

Conclusion

As briefly noted above, Big Switch's sophisticated clients actively use a combination of two or three of the approaches listed above to match technologies to organizational needs. Firewalls are a good match with non-technical compliance and risk management teams, SDN-style ACLs are a good match with infrastructure teams and micro-segments are a good match with the application owners themselves. Rather than try to shoe-horn all policy needs into one technology and leave many critical stakeholders out of the loop, the portfolio approach seems to result in organizational efficiencies that far outweigh the learning curve of managing these three disparate technologies.