Akamai

WHITE PAPER

Software-Defined Segmentation for Data Center Operators

For operators of multi-tenant data centers, segmentation of computing environments is not just important — it is fundamental to their operating model. First, they need to separate their own infrastructure from their clients' environments, and share certain resources while preventing access to others. Second, they need to prevent "cross-contamination" among their clients' respective environments, whether accidental or nefarious. That includes preventing successful breaches or malware infections from spreading from one client's environment to others. Finally, within the owned operational applications, a good level of separation is required to limit the impact of a potential breach. Looking deeper into data center providers' operational networks, there are three scenarios in which segmentation, if achieved efficiently, can significantly improve security posture and reduce costs.

1. **Separating operational networks** (DCIM, BMS, etc.) from the enterprise network (the provider's internal systems, which include billing) and customer networks

2. **Reducing the risk of lateral movement inside the operational network,** which has many hard-to-patch systems and introduces risks if not properly segmented

3. **Creating efficient and secure connectivity between customer-facing networks** — such as the DMZ, where the custom portal is located, which needs secure access to data from operational networks (reading the power status, for instance) and from enterprise networks (for reading the billing information)

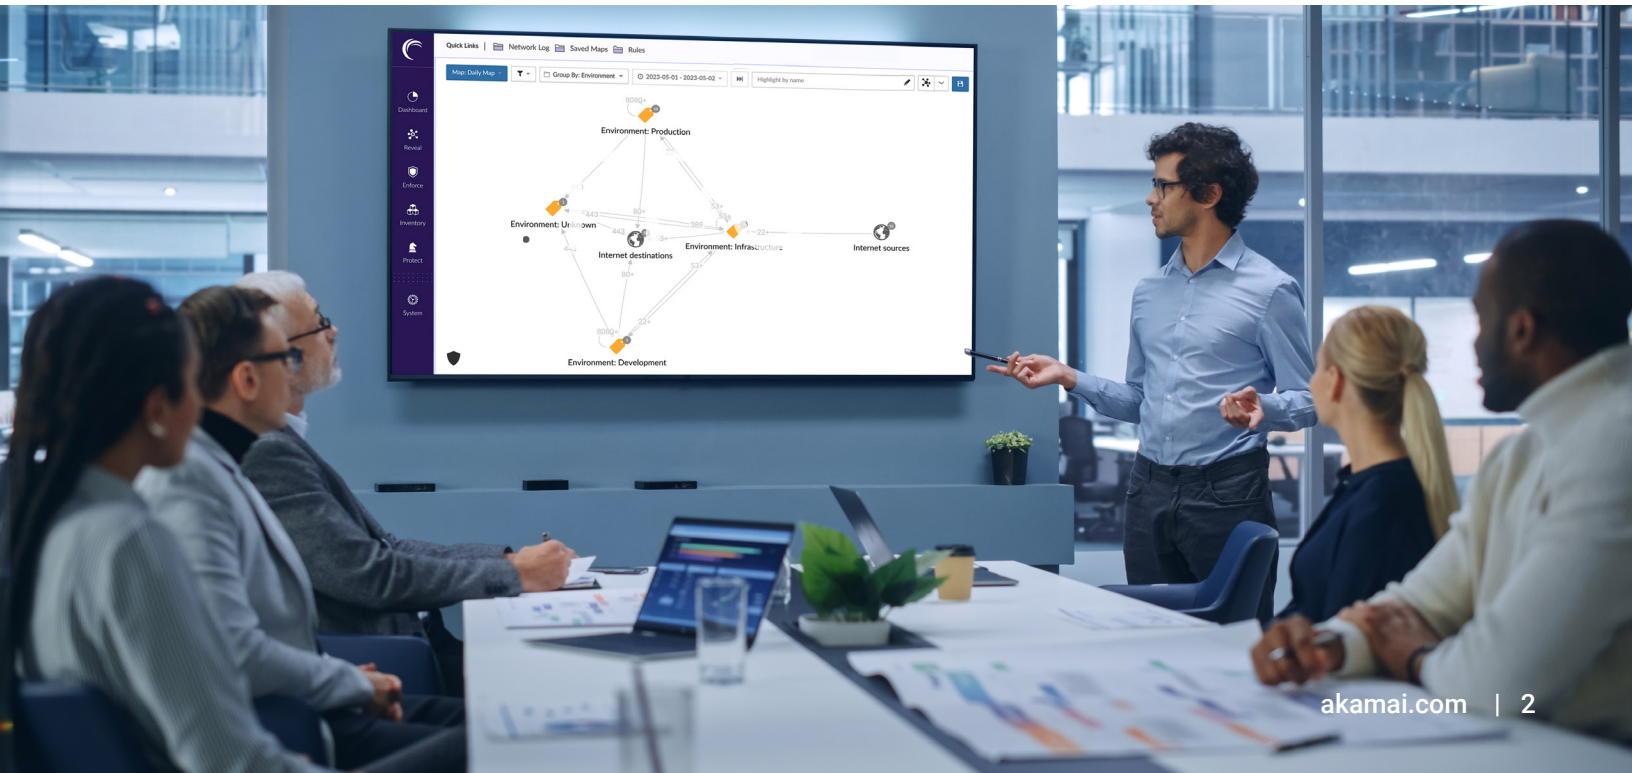

Those are handled today through very complex, slow-to-implement, and inefficient networking constructs, VLANs, interim networks, etc. Implementing a software-defined solution without relying on any complex network configurations will have significant cost reductions and also introduce tighter and more robust control on connectivity.

In addition, customers struggle to implement and maintain a strong level of segmentation within their applications (hosted or on-premises). This introduces an important opportunity for data center operators to leverage their internal segmentation expertise, tools, and operational models to offer managed services to their customers and create a very attractive revenue stream around a segmentation practice. Furthermore — with the ability to extend security policies to customer premises with the right methodology, tooling, and processes — the operator will be able to gain access and visibility to the non-hosted applications, which can help accelerate their secure migration to the hosted data center, thus contributing to the core business.

Equifax: A worst-case scenario

If you are wondering "what's the worst that could happen" with weak, ineffective, or nonexistent environment segmentation, the highly publicized Equifax breach of 2017 stands as a prime historical example. The breach resulted in the compromise of 143 million Americans' highly sensitive personal information. According to the U.S. Government Accountability Office (GAO) investigation, the attackers initially broke into the giant credit bureau's customer dispute resolution portal by exploiting a vulnerability, known as CVE 2017-5638, in the Apache Struts web framework. Once inside, they essentially had free run of the company's systems for 76 days. The GAO report attributed this freedom of lateral movement to a lack of segmentation, which allowed easy access to databases at will — a virtually unlimited attack surface.

The question is how to achieve this kind of segmentation most effectively, efficiently, and economically. Operators have historically relied on traditional firewalling or VLANs to separate environments within a multi-tenant or multi-user architecture. Implementing and maintaining such measures, however, is typically arduous, highly manual, time-consuming, and costly. Moreover, these techniques are by no means airtight and can leave a substantial amount of attack surface exposed. The efficacy of solutions designed for perimeter defense is particularly problematic within the data center, especially since most of these environments include a variety of virtual machines, hypervisors, containers, and even cloud components, and workloads dynamically spin up and down automatically. Another important note is that segmentation with VLANs requires downtime of an application, which for critical operational controls can be a showstopper.

For all these reasons, operators of shared environments are taking a closer look at modern, software-defined segmentation techniques, including microsegmentation. Advances in microsegmentation technologies have made it a viable option for all types of companies, and, arguably, the optimal choice for achieving a Zero Trust security model. Of equal importance, with the right tools and a little thoughtful planning, microsegmentation can be implemented more quickly and easily than the aforementioned methods, and is easier to manage and maintain as well. In fact, recent testing has demonstrated that microsegmentation can reduce time to deployment by as much as 30 times compared with traditional firewall implementation. An additional crucial benefit: With software-defined segmentation, no networking changes or application downtime are required. Those time savings and efficiencies translate to significantly lower costs over the deployment lifecycle.

The pitfalls of conventional approaches

To understand the advantages of software-defined segmentation or microsegmentation, it is useful for comparative purposes to look at some of the drawbacks and limitations of standard techniques employed both on-premises and in the cloud. These might include some combination of physical or virtualized firewalls and network configurations such as VLANs. In general, these methods are resource- and labor-intensive. Creating security policies is a cumbersome process. Additions and modifications need to be performed manually, creating a drag on ongoing operational efficiency and raising the risk of vulnerability.

Internal firewalls, in particular, are expensive to acquire and complex to set up. They also interfere with the normal flow of traffic, altering patterns and creating circuitous "hairpins" that ultimately impede system performance. As the industry is learning, firewalls are not intended for segmentation within the data center — some providers will readily admit that firewalls simply don't belong there.

One of the most painful challenges when trying to introduce segmentation to an existing, running production environment is that traditional methods require downtime for an application. Downtime is costly. It can only happen within specific time windows, and oftentimes it is not possible at all.

An additional challenge worth noting is that creating any internal segmentation requires good knowledge of east-west application dependencies. This insight is usually nonexistent. Without a simple way to map application dependencies, it is extremely hard and risky to separate a brownfield environment.

Why software-defined segmentation is more effective

Operational efficiency, better security posture: Software-defined segmentation overcomes the inherent inefficiencies of traditional techniques, and perhaps more important, results in stronger security for multi-user environments. As the name implies, software-defined segmentation takes the concept of network segmentation and implements it without any need for infrastructure change. It entails the creation of security policies around individual or logically grouped applications, regardless of where they reside in the hybrid data center. These policies dictate which applications can and cannot communicate with one another — true Zero Trust.

No manual changes or downtime: Software-defined segmentation does not require any networking changes or for any VLANs to be created, which results in significant operational savings. It also does not require any application downtime or changes due to a migration to a new VLAN. This is important. In many applications for which downtime is very expensive or impossible, this is the only way to provide this crucial security measure.

Extensive visibility: Additionally, advanced software-defined segmentation solutions — designed to address the east-west traffic segmentation challenges — provide an integrated visibility tool that helps identify the segment boundaries and application dependencies. This results in an efficient process and eliminates operational errors when creating the policies.

Policies and controls automation: Software-defined segmentation also makes it possible to apply policies in a dynamic fashion, so that as workloads are spun up or down, they are attributed to the correct policy automatically. This saves considerable resources by eliminating the need for manual moves, additions, or changes.

Infrastructure agnostic: A key advantage of software-defined segmentation is that it is infrastructure independent. The same tool provides visibility and segmentation across any infrastructure: bare metal, virtualized, PaaS, cloud, containers, etc. All under one pane of glass and with a singular workflow. This results in significant operational freedom in which security standards can be achieved without any constraints on the underlying infrastructure choice.

More revenue, stickier relationships: Most important, this introduces a significant opportunity for data center operators. While they manage and provide the internal segmentation, they can leverage the training, tooling, and processes to offer a very-much-needed managed service to its customers — managing segmentation not only for the hosted applications, but also for applications that are on customer premises or in the cloud — within the same tool, within the same pane of glass. This not only results in additional revenue potential, but also creates a stronger dependency on the operator, resulting in longer relationships and higher profits.

Why Akamai

To deliver on these benefits, a software-defined segmentation solution must meet a number of essential criteria. It must allow deep, process-level visibility into all applications running in the computing environment and the ability to map all the data flows among them. The flexibility to properly label assets for policy creation and automatically modify labels as workloads autoscale are also key to efficient deployment and management. And the solution needs to be platform- and infrastructure-agnostic. Policies must be able to follow their respective applications and perform consistently among multiple environments. Finally, the solution should allow for an automated and simplified operational model for policy creation, management, and enforcement.

Only Akamai Guardicore Segmentation meets all these criteria. Software-defined segmentation is our core capability. The solution provides an unprecedented graphic visualization of all assets in the environment and the dependencies among them — whether bare metal, virtual machines, public cloud, containers, or IoT devices. This deep visibility dramatically accelerates the process of identifying, grouping, and creating security policies around microsegments of applications.

For more information, visit akamai.com/guardicore.