vmware ON VMWARE

How VMware IT Uses Zero Trust in the Data Center

INDUSTRY

Zero Trust, Data Center Security

LOCATION

Palo Alto, California

KEY CHALLENGES

- Navigating a large enterprise with a complex IT group made up of several teams that must interact to achieve results
- Implementing Zero Trust for the data center using micro-segmentation
- Working on complex projects with third-party applications that are also complex

SOLUTION

To strengthen security, VMware IT imposed a Zero Trust architecture in the data center to restrict communication between workloads to the minimum required for data center applications using the security capabilities in VMware NSX® Distributed Firewall.

In 2015, VMware IT initiated a Zero Trust project for the data center[i]. By 2017, the IT organization had substantially progressed toward achieving Zero Trust by micro-segmenting several already deployed applications. Simultaneously, various business and functional organizations at VMware had continued deploying new applications without micro-segmentation, reducing VMware's progress rate toward Zero Trust. Thus, when VMware decided to replace its enterprise planning and supply chain application, the IT organization sensed an opportunity to accelerate its Zero Trust efforts by micro-segmenting the replacement application before launch.

However, the replacement application had a long testing and deployment cycle of more than 12 months, and only three months remained until the launch date. Also, the replacement application was one of the most complex applications in the data center.

VMware IT raced to micro-segment this new but complex application, aiming for Zero Trust via micro-segmentation to become the new normal in the data center.

Organizational background

VMware is a large enterprise software company with more than $10B in annual revenues and more than 30K employees. VMware IT is the group responsible for creating and maintaining the IT infrastructure to support internal functional organizations, such as finance and semi-autonomous business units.

VMware IT itself is a complex group with several teams. One of these, the information security (IS) team, is responsible for ensuring information security across the entire IT infrastructure. IS team efforts are supported by a separate IT solutions engineering and design team (IT-SE) that translates business and security requirements into product purchase recommendations and reference architectures. The IT-SE team also maintains the reference architecture on an ongoing basis.

VMware IT also includes a network services team responsible for day-to-day (Day 2) operations after a new reference architecture has been adequately vetted by the IT-SE team.

Finally, VMware IT includes an application operations team. This team is responsible for deploying and maintaining enterprise applications, such as the resource planning and supply chain application.

BUSINESS BENEFITS

- With new confidence in the VMware NSX Distributed Firewall and micro-segmentation, progress toward Zero Trust accelerated.
- VMware internal applications now have better security—a clear second line of defense behind the edge firewall based on the architecturally sound principles of Zero Trust and implemented using the mature micro-segmentation capabilities of the NSX Distributed Firewall.
- VMware IT gained operational agility; deploying new instances of existing workload types requires no new security policies—they are automatically secure on creation.
- Expressing policies using dynamic security groups (or security tags) is simpler and more intuitive than doing the same with IP addresses, port numbers, and protocol identifiers.

VMWARE FOOTPRINT

- VMware NSX® Distributed Firewall™
- VMware NSX® Distributed IDS/IPS™
- VMware NSX® Intelligence™
- VMware NSX-T™
- VMware NSX-V™
- VMware vRealize Log Insight™

Enterprise resource planning and supply chain application

To conduct its business, VMware relies on a third-party enterprise resource planning and supply chain application through which all VMware financial transactions flow. This mission-critical application features highly complicated business processes and is extremely complex in nature.

Furthermore, the application is actually an aggregation of modules that provide adjacent functionality. For example, some vendors include modules for business insights and compliance management alongside the core modules for enterprise resource planning and supply chain management. These modules are complex in their own right, often consisting of a business logic portion (a sub-application) and a database portion (to store the sub-application data).

New data center security approach

Before 2015, a set of edge firewalls provided security for the data center. The firewalls processed mostly north-south (to or from the Internet) traffic, with a small amount of east-west (internal) traffic transiting through them [1]. As a result, most applications in the data center did not have a second line of defense (see Figure 1). An attacker who managed to compromise one workload could easily scan for—and attempt to compromise—additional workloads.

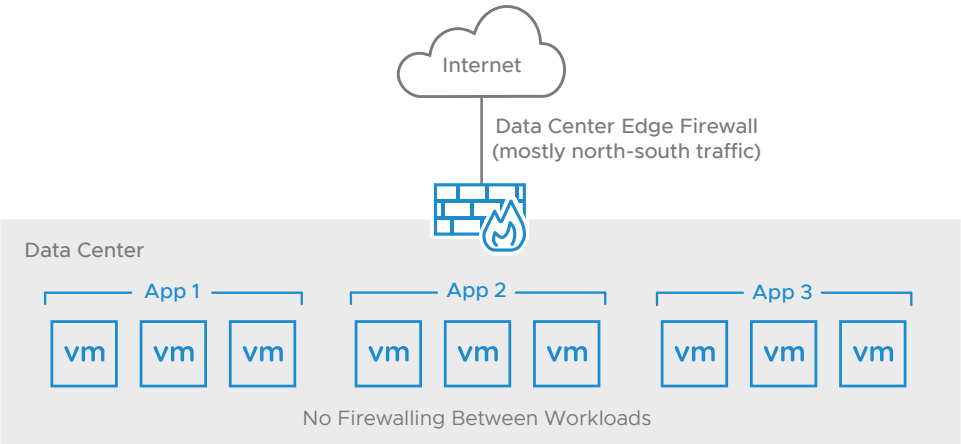

FIGURE 1: Starting state

The IS team wanted to strengthen security by imposing a Zero Trust architecture in the data center [2]. With Zero Trust, communication between workloads would be restricted to the minimum required for the data center applications. The IS team requested the IT-SE team create an architecture for Zero Trust in the data center.

The IT-SE team investigated traditional security technologies, including hardware firewall appliances. The team was concerned about several operational challenges with appliances. First, network changes would be required to insert firewall appliances in the data center. Second, network traffic would have to be hair-pinned to the appliances and back, adding communication latency and creating a traffic choke point. Third, dynamic security policy management (to keep up with the ever-changing data center applications) would be difficult. The IT-SE team also worried about the budget needed to procure adequate firewall capacity to inspect all east-west traffic using such appliances.

The team concluded that the emerging micro-segmentation technology gave it the best chance to cost-effectively achieve Zero Trust in the data center [3].

Zero Trust expectations

With the Zero Trust effort, the IS team expected every application in the data center to be micro-segmented from every other application. Workloads belonging to an application would be held down to the minimum communication needed for the workload to function (see Figure 2).

In addition, the application operations team expected that micro-segmentation would not interfere with any application's availability or performance—and they expected the enterprise resource planning and supply chain application would launch on time.

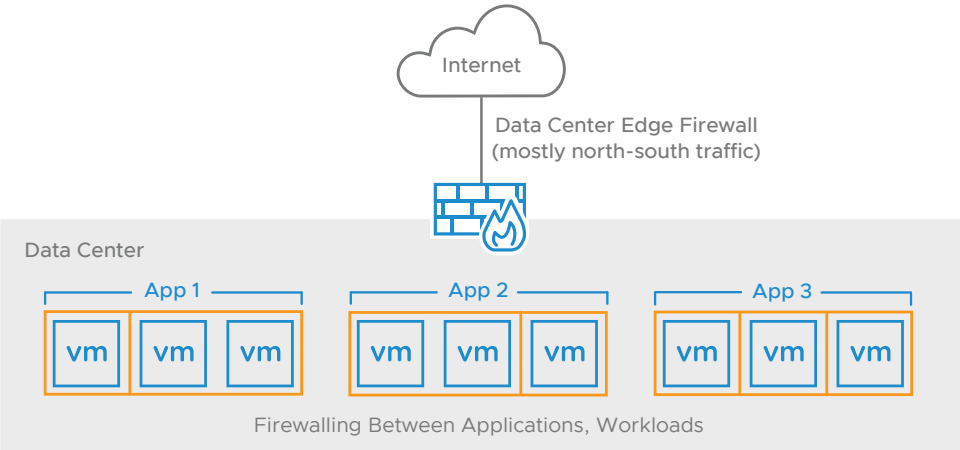

FIGURE 2: Desired end state

Fortunately, there was a credible implementation of micro-segmentation available, and it was homegrown—VMware NSX® Distributed Firewall[ii]. However, the presence of just one credible solution increased the pressure on the IT-SE team to thoroughly validate the security architecture and test the solution on its own. If the security architecture came up short, the IT-SE team would have to undertake an expensive re-architecture cycle. Similarly, if the NSX Distributed Firewall ended up deficient in features, performance, or stability, the IT-SE team would have to take a subsequent risk by moving to a less mature vendor.

Organizational buy-in for greenfield micro-segmentation

Until 2017, the IT-SE team had focused mostly on protecting already deployed applications with micro-segmentation (i.e., on brownfield deployments). While this approach was sound, various functional organizations and business groups in VMware were deploying new applications without micro-segmentation. This reduced the rate of progress against the goal of achieving Zero Trust in the data center.

Then, a new opportunity arose.

In 2017, VMware was switching from one provider of the enterprise planning and supply chain software platform to another. The conversion from one complex application to another is itself complex. The process was already underway with a cutover date three months into the future when the IT-SE team approached the application operations team to micro-segment the new application. The application operations team was concerned

that adding micro-segmentation requirements to the project would put the application's launch date at risk.

In response, the IT-SE and IS teams maintained that while there were risks, three months was adequate to secure the new application. This confidence derived from their previous experience micro-segmenting other applications in the data center. The two teams were also motivated to make more significant strides in achieving Zero Trust. Eventually, they convinced the application operations team to micro-segment at launch instead of postponing the work and risk security incidents.

Journey to a successful launch

The IT-SE team swung into action to add security to the new application prior to launch. They had to determine the application's topology (communication pattern), design robust and extensible micro-segmentation policies based on the communication pattern, and thoroughly test micro-segmentation with the application.

Understanding the application's communication pattern

From previous experience, the IT-SE team knew that the most challenging part of micro-segmenting an application lay in determining the application's communication pattern. The more complex the application, the greater the number and variety of valid traffic flows. They also knew that determining the communication pattern is an iterative process.

The team typically proceeded by using the log analysis tool VMware vRealize® Log Insight™ [4] to aggregate traffic logs into traffic flows from a test installation of the application. Then, they worked with the application operations team to determine valid flows, create temporary policies to permit those flows, and block everything else. If the application stopped working, then the IT-SE and application operations teams studied the blocked flows further to determine the nature of the flows that had been erroneously blocked.

For the enterprise resource planning and supply chain application, the application operations team was already using a testbed for functional and load testing. The IT-SE team turned on traffic logging over this testbed and systematically determined the valid traffic flows for each workload type.

Creating micro-segmentation policies for the application

The new application consisted of 16 modules. Most modules had a business logic component and a database component. Between the business logic and the databases, the application had approximately 300 workloads. There was also some duplication of workload type to support performance and redundancy requirements.

The IT-SE team named the workloads algorithmically for easy identification. All workloads started with the prefix of the module that they belonged to. Thus, the enterprise resource planning module workloads started with VM_ERP_, the business insights workloads started with VM_BI_, and so on.

The business logic and database workloads within a module were also named algorithmically. Therefore, the business logic for the enterprise resource planning module had the prefix VM_ERP_APP_, the module's database had the prefix VM_ERP_DB_, etc. (see Figure 3).

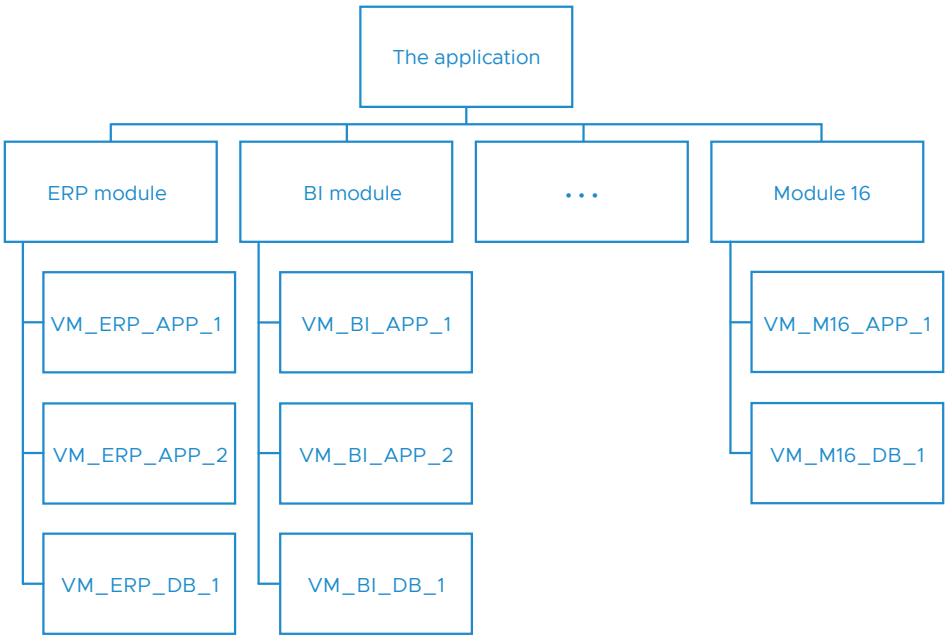

FIGURE 3: Workload naming scheme

The IT-SE team created dynamic security groups using the workload naming scheme. For example, all the VM_ERP_APP_ workloads belonged to the SG_ERP_APP security group.

With the security groups in place, the IT-SE team then mapped valid traffic flows to micro-segmentation policies. The team permitted traffic between security groups only when the traffic mapped to one of the valid traffic flows. For example, intra-module traffic between workloads belonging to SG_ERP_APP and SG_ERP_DB was allowed.

The IT-SE team soon noticed that the business logic components for all the modules needed to communicate with each other. To simplify policy expression, they created nested security groups. An SG_APP group consisting of all the module-specific business logic security groups was defined. That is, SG_APP included SG_ERP_APP, SG_BI_APP, etc.

Communication between business logic workloads (indirectly) belonging to SG_APP was permitted (see Figure 4).

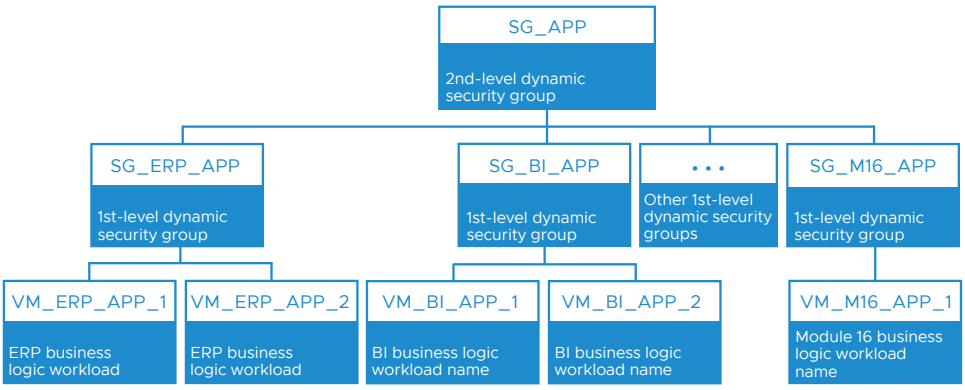

FIGURE 4: Nested security groups

More testing

With the security policies identified, the IT-SE and application operations teams replaced the testbed's temporary rules with the new micro-segmentation policies. They ran additional load tests to convince themselves that the security policies did not interfere with the application's availability or performance.

Deploying in production

With all testing completed within the allocated time, the new application was launched as planned, but with micro-segmentation protection in place.

NSX Manager, the central management console for the NSX Distributed Firewall, was set up to alert the IS team if any policies for the application were changed.

Legitimate policy changes, such as installing a new module, required organizational approval. Such policy changes also engendered a repeat of the iterative communication pattern determination process (albeit on a smaller scale) using vRealize Log Insight, but this time by the network services team who had been trained by the IT-SE team to take over day-to-day micro-segmentation operations.

Once the communication pattern was determined, the network services team created and deployed new policies using the NSX Manager. Thus, new modules were securely deployed, preventing regression in the Zero Trust regime.

However, adding workloads of an existing type (ERP_DB, for example) to expand capacity did not require additional security approvals or policy changes. Given the use of dynamic security groups, adding ERP_DB_2 to reduce the load on an existing ERP_DB_1 would not change any security policies. Similarly, decreasing capacity (by shutting down some workloads) or rebalancing capacity (by moving workloads using VMware vMotion™) required no policy changes.

NSX-V to NSX-T

At the time of the initial deployment of the enterprise resource planning and supply chain application, micro-segmentation was conducted using NSX-V as the current version of NSX (VMware NSX-T™ Data Center was not yet available). Thus, there was no opportunity to use VMware NSX Intelligence™ [5], a distributed analytics engine that automatically determines traffic flows and recommends security policies. Similarly, VMware NSX® Distributed IDS/IPS™ [1], available in NSX-T but not in NSX-V, could not be used with the Zero Trust architecture.

Fortunately, NSX-V and NSX-T can exist side-by-side in the same deployment, aiding migration to NSX-T. In 2021, the enterprise resource planning and supply chain application was migrated from NSX-V to NSX-T. Now, IDS/IPS is turned on, and NSX Intelligence is used to monitor the application's communication pattern and optimize security policies.

Safe and sound data center

The successful launch of the enterprise resource planning and supply chain application under tight time constraints proved that even the most complex and mission-critical applications can be brought under the Zero Trust umbrella with micro-segmentation. After the launch, VMware IT's confidence in the NSX Distributed Firewall and micro-segmentation grew, and progress towards Zero Trust accelerated.

VMware IT is well on its way to achieving Zero Trust in the data center. Although most existing applications have already been micro-segmented, all new application deployments are required to go live with micro-segmentation policies in place.

VMware's internal applications now have better security—a clear second line of defense behind the edge firewall based on the architecturally sound principles of Zero Trust and implemented using the mature micro-segmentation capabilities of the NSX Distributed Firewall.

Lastly, VMware IT has gained operational agility. Deploying new instances of existing workload types requires no new security policies—they are automatically secure on creation. Expressing policies using dynamic security groups (or security tags) is much simpler and more intuitive than doing the same with IP addresses, port numbers, and protocol identifiers.

---

[i] At inception, VMware IT referred to this project as "the micro-segmentation project." Over time, the project came to be known as "the Zero Trust project.

[ii] The security functionality in NSX was rebranded as "NSX Distributed Firewall." See [1] for more details on the rebranded product

LEARN MORE

Contact your sales representative to schedule a 1:1 Customer briefing on this topic with a VMware IT subject matter expert.

FOR MORE INFORMATION OR TO PURCHASE VMWARE PRODUCTS

call 877-4-VMWARE (outside North America, +1-650-427-5000), visit *vmware.com/products*, or search online for an authorized reseller. For detailed product specifications and system requirements, refer to the documentation noted in "References."

References

[1] Internal Firewalls for Dummies, *https://www.vmware.com/content/microsites/learn/en/656351_REG.html*
[2] Zero Trust Architecture: SP 800-207, *https://csrc.nist.gov/publications/detail/sp/800-207/final*
[3] A Practical Path to Zero Trust in the Data Center, *https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/products/vmw-practical-path-zero-trust-data-center.pdf*
[4] vRealize Log Insight, *https://www.vmware.com/products/vrealize-log-insight.html*
[5] VMware NSX Intelligence, *https://www.vmware.com/products/nsx-intelligence-analytics-engine.html*