How Cisco IT Secures Its Data Centers

Cisco IT Methods

Introduction

A cornerstone in the effort to secure enterprise networks is implementing robust data center security capabilities to safeguard sensitive mission-critical applications and data. The complexity of protecting not only physical data centers, but also the intersections of where the data center meets the virtual environment, creates a unique situation. Replicating the same segmentation in the virtual environment that exists in the physical data center is where multitenancy becomes an important innovation in security. Multitenant architecture allows applications to virtually partition data and configurations, enabling each client customized virtual applications. Critical assets in the data center include systems to support identity, authentication and authorization, configuration management, and security, and all require special treatment. To protect the network, data center security employs a number of capabilities and integrates them into the technologies in data centers.

Deployment: security layers in the data center space

Data center security is based on a foundation, a series of processes that are designed as layers, to protect critical assets such as customer data and intellectual property. “Cisco is entrusted with a lot of customer data, and so we take a data-centric approach to security. We focus on the systems containing customer data because they are key to maintaining customer trust,” says Scott Stanton, InfoSec architect at Cisco. The foundation of data center security entails five security capabilities: prevent and mitigate, measure, remediate, detect, and contain. We implement controls to prevent or mitigate known risks; measure security vulnerabilities in the environment; remediate and fix all known issues; detect when an event occurs on the networks, servers, or applications; and lastly, contain compromised systems to prevent further movement from the affected systems (see Figure 1).

Figure 1. Data center security foundation overview

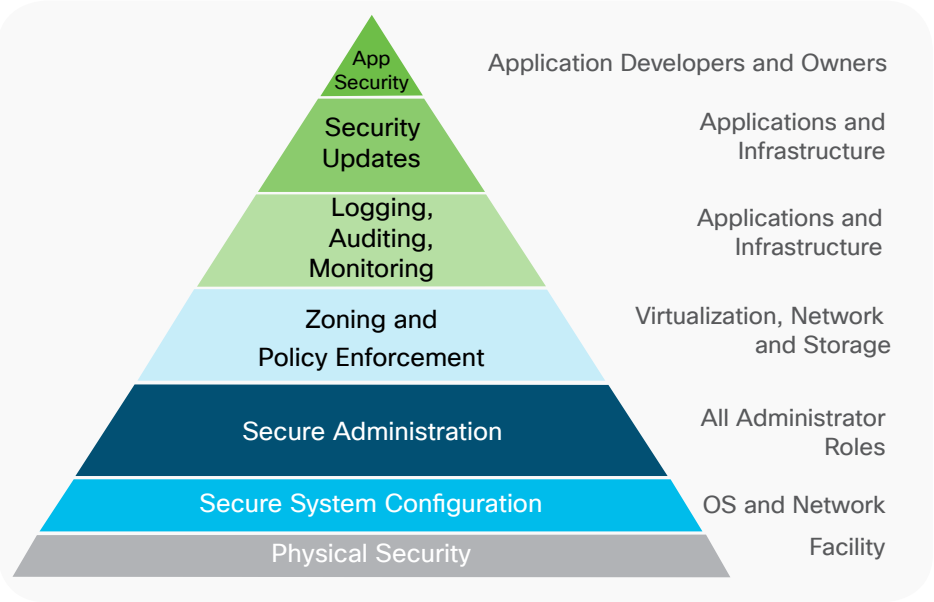

“We have a secure system configuration, secure administration processes, layered zoning and network security policy enforcement, firewalling, intrusion detection system [IDS] or intrusion prevention system [IPS] -type monitoring and blocking, the ability to update the environment, and application security controls and secure software development to protect the data itself,” says Stanton. “Failure or compromise at lower level controls will generally undermine the higher level controls.” Each of these five capabilities is applied to the technology in the data center, including applications on the systems, operating systems running on the hardware, virtualization technologies such as VMware or OpenStack, storage systems, and the network. Application security controls, containment, prevention, and monitoring capabilities are integral to securing that data.

Capturing and examining vital network traffic allows for swift action. To look closely at the network and run anomaly-based detection, Cisco® InfoSec focuses on two areas: the packets on the network itself and captured NetFlow data from the networkplatforms. Information drawn from these sources is fed through remote switch port analyzer (RSPAN) or NetFlow feeds to security devices next to the network equipment receiving a host of information about the quality of that data. Cisco IPS technology from the Cisco acquisition of SourceFire searches for botnets, viruses, and evidence of compromised hosts in the data center, while monitoring is physically performed at the core data center gateway layer. Although a number of technologies are used at that layer, the main component is the Cisco Catalyst® 6500 Series Switch for loadbalancing functionality.

From a network perspective, Cisco operates mainly in the prevent and mitigate, detect, and contain areas, with firewalls being one of the primary controls. Ben Kelly, a member of the technical staff, says, “In the data center network security space, we build containers, and we put applications or tiers of applications into those containers. Typically, these are subnetbased boundaries, with access control lists , or ACLs, and firewalls to secure traffic. We host this on the Cisco Nexus platform.” Cisco has many firewalls, including a corporate firewalls and dedicated firewalls that are data center–to–data center specific. Where the internal and demilitarized zone (DMZ) data center networks cross over, traffic enters a security control chokepoint, a dedicated firewall.

The data center is based on the Cisco Nexus® platform running a threelayer access, distribution, and core network model. Cisco Nexus Switches are used at the access layer. At the distribution layer, Cisco Nexus functions as a boundary between Layers 2 and 3. From a network perspective, the distribution layer is also referred to as a pod. Cisco Nexus s distribution switches and anything below the distribution, including the access layer, servers, and storage connected to that layer, make up the pods. Above the pods sits a high-speed, highly resilient core based on Cisco Nexus switches0. Compute is based on layered Cisco Unified Computing System™ (Cisco UCS®) servers. The Cisco UCS FabricExtender connects service from the Cisco UCS domain uplink directly to the Cisco Nexus distribution layer (see Figure 2).

“That’s the traditional network,” says Kelly. “However, with technologies like FabricPath and Cisco Application Centric Infrastructure , or Cisco ACI, solution, we’re seeing a trend to what is broadly referred to as ‘data center fabrics.’ ” From a network architecture topology, networks are flattening out, and rather than having access, distribution, and core layers, networks have only two layers: a leaf that connects to endpoints and a spine that interconnects all leafs. This particular topology scales out horizontally rather than vertically, matching the trend in data center traffic patterns.

“Historically, we saw more traffic north to south, which means from a server to access, distribution, and core and out of the data center. What we see now is a lot more east-to-west traffic, or server to server in the data center.” Network fabrics are highly resilient, have high throughput and low latency, and have consistent latency between any two endpoints in the data center.

Figure 2. Architecture overview of multitenant data center security

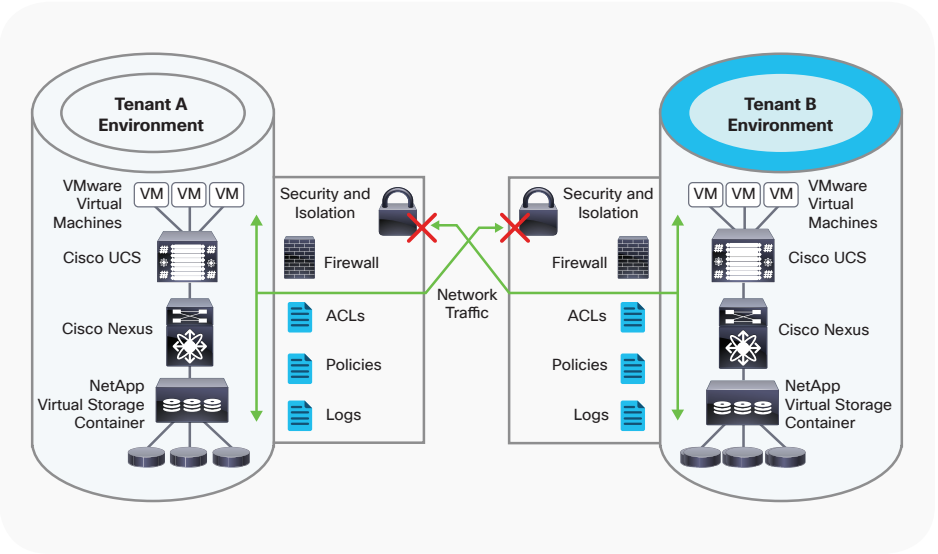

Segmenting the data center network

At a high level, data center network and compute pods are segmented into production, non production, and DMZ environments, which enable the subdivision of internal zones to isolate a sensitive system without rezoning. The Cisco network is divided into two main security zones: systems that are Internet facing, which are labeled DMZ, and systems that are internal and not directly reachable from the Internet. Internal networks have an additional type of secure network known as a protected network, which has ACLs to enforce security policy. These protected networks are used for sensitive hosts and also serve as a buffer between DMZ and internal networks. The Cisco data center has thousands of ACL rules on the Cisco Nexus platform to protect interfaces between different application components.

In recent years, the trend toward network optimization, flexibility, and agility has inspired the use of virtual route forwarding (VRF). VRF is the segmentation of the network at Layer 3, but allows a shared network infrastructure. Essentially, pods can be DMZ and internal but still maintain the logical separation needed between the two networks.

"We've seen that trend continue into the future, where we have the fabric-based solutions such as Cisco ACI™ as well as Dynamic Fabric Automation, or DFA," says Stanton. This scalable, high-performance, automated infrastructure is an important feature of cloud computing.

Data centers supporting multitenancy need to have the ability to separate servers, network, and storage services between different business units or customers. The predominant way to do this is through virtualization rather than physical separation. In addition to providing a transparent, virtual environment, having logical controls over each infrastructure component enforces segmentation between tenants and provides access control, provisioning, monitoring, and resource consumption measurement.

Cisco Application Centric Infrastructure policy and security in the data center

Cisco ACI transforms the way that network management and security policy are done in the data center. In the current policy model of VLAN or subnet-based containers, using VRFs to perform logical separation between raw types of compute workload is cumbersome. ACLs are distributed throughout the data center and require a great deal of manual labor to locate and modify. This model affects Operating Expenses (OpEx) and agility. An application-centric view of the network means that the underlying configuration of the network will always be based on the application's network profile, and detailed knowledge about network technology, topology, and configuration is not required to understand the network dependencies. Cisco ACI provides a common operational model shared by all teams in defining application requirements.Drawing on predefined application requirements and policy profiles, Cisco ACI automates the provisioning of the network, network application services, security policies, and workload placement. Cisco ACI enables the Cisco security organization to more easily manage security policy and controls in the data center network.

A significant benefit of Cisco ACI is that it uses a policy-based model that defines applications in terms of relationships. For example, if an application has a web server, an application server, and a database server, these are defined as policy objects that can be shared with other applications. The applications engineer can create mappings between those systems in the Cisco Application Policy Infrastructure Controller (APIC), which then creates the underlying network to enable those relationships. Although it defines the network, subnet, and VLANs, these topology concepts are not used to enforce policy. Policy enforcement is separate, allowing the Cisco ACI fabric to forward traffic efficiently while controlling which systems can communicate at a very granular level. Cisco ACI follows a "whitelist" model and has security included by default, tying directly into InfoSec's prevent and mitigate capabilities and preventing unnecessary access. The fabric blocks traffic between two systems unless specifically configured to permit that traffic. In the Cisco ACI environment, security policy is based on fabric endpoints (that is ports or VXLANs) and is decoupled from IP addressing. As a result, endpoint mobility has a minimal effect on security policy enforcement. Security management is simplified, and risks are mitigated through automated, centralized compliance and auditing.

Cisco ACI and multitenancy in the data center

Tenant segregation, resource isolation, and identity-based resource entitlements are three pillars of cloud multitenancy. Multitenant environments must keep tenants isolated and segregated from each other to make sure of security between them. Shared resources in a multitenant environment must also be provisioned based on entitlements or subscription, and resource guarantees must be enforced such that high consumption or utilization does not adversely affect other tenants.

From a compute virtualization perspective, multitenant constructs must be enforced in the compute hypervisor. Provisioning of virtual CPUs and RAM must guarantee proper entitlements and isolation as well as administrative segregation. At the network and storage level, logical segmentation is the primary means of performing tenant segregation and resource isolation. In that environment, processes, automation, and orchestration make sure of resource entitlements and align the network and storage with the compute

and virtualization platforms. When performing network-based tenant segregation, Cisco IT previously used the Virtual Security Gateway (VSG) platform, which integrates with the Cisco Nexus 1000V Virtual Switch. Because VSG performs security policy enforcement at the virtual switch port (Layer 2), it allows for the allocation of a large IP subnet across multiple tenants, resulting in a very efficient and highly elastic use of IP address space. In this single subnet, the different tenants cannot access each other unless VSG security policy allows them to do so.

Cisco ACI natively translates tenancy constructs to underlying network constructs, such as VRFs, to perform logical separation at the network level. The Cisco ACI policy model supports each tenant, and tenant configuration occurs in the APIC interface. Cisco ACI natively supports multitenancy with its own constructs in the policy model. Cisco ACI policy objects can have different levels of scope (application, tenant, or global) and multiple levels of tenant granularity (endpoint group, application network profile, bridge domain, and VRF) at the network level. IT teams and business units can spin up a multitenant compute environment and be online with their application in less than an hour, which is a huge time improvement over traditional IT provisioning processes.

Multitenancy also requires integrated identity. IT needs to know who has access to a particular environment. Entitlement and privilege management allow specific users to manage their resources in the multitenant environment and to be billed only for the resources that they consume.

Management, service, and support

Currently, data center network security has two major management categories. The first centers on the firewalls and the Cisco Nexus ACLs. The second focuses on the management of the IPS/IDS monitoring and incident response capabilities. Configurations and devices for firewalls and ACLs are managed by the Cisco network team under Global Infrastructure Services (GIS). GIS governs and manages changes to the firewalls and ACLs. After a ticket is opened with GIS, several iterations between InfoSec and GIS occur before the network team implements requested changes. For example, if an employee requires a specific server to communicate with another server on a certain port, InfoSec investigates whether it is a valid request before routing the request to GIS for implementation.

Management and monitoring of the IPS/IDS are the responsibility of the Cisco Computer Security Incident Response Team (CSIRT). CSIRT operates the sensors that generate data around network traffic and devices. When an anomaly arises, it is flagged as an event related to that packet or session and sent to a Tier 1 analyst for review. If the analyst finds something about the event worth further investigation, Tier 1 escalates the alert to an investigator for deeper analysis and response. If an investigator is unable to track down the offending system's owner, a tool in the network can "black hole" the offending system until long-term actions can be taken.

"For example, if an infected system on our network tries to spread to other systems in the network, we have the ability in our core network through a process called Real-Time Blacklist, or RBL, to black hole at the network level," says Stanton. "The offending packets will never reach their destination."

CSIRT may take a host off the network without knowing where it is plugged into the network and essentially remove the host's Internet connectivity. Depending on how close the host is to the network core, one of these black hole routers can take them off the Cisco internal network as well. After a host is added to the RBL or the black hole, packets from that host go into the black hole and do not come out.

Processes and management

Described earlier are the two main types of network-centric management of the controls. Much of our application and data security relies on the top of the pyramid (see Figure 1), and these focus on process. From a software perspective, InfoSec has security requirements for applications in our environment. In addition, Web Application Firewalls (WAFs) protect web applications from malicious requests, typically on a web server. Much like a network intrusion detection system watches for unusual network activity, a WAF analyzes web traffic in front of a specific web server or set of web servers for unusual web requests such as cross-site scripting, SQL injection, parameter tampering, brute forcing, logic evasion, and other classes of web application attacks. WAF provides a monitoring and prevention capability at the application security layer.Process and governance around testing web applications for security vulnerabilities are managed by InfoSec web application security programs. These programs include Basic Application Vulnerability Assessment (BAVA) and Deep Application Vulnerability

For more information

To read additional Cisco IT case studies about a variety of business solutions, visit Cisco on Cisco: Inside Cisco IT at www.cisco.com/go/ciscoit.

Assessment (DAVA). The first is a self-service program that enables every developer to use BAVA, an assessment based on IBM Appscan Enterprise. Developers use this tool to scan their web application as they build. If they follow the process, they are actively testing, developing, and scanning code for vulnerabilities as they build. After all their code components are stitched together and they have a working application, if it is a business-critical application, it goes through DAVA. DAVA is the Cisco red team/tiger team. They can hack through the application to detect and locate vulnerabilities for the developer to remediate. Because this is largely a manual process, it requires developers and InfoSec to work together. It is a key component of the Cisco application security program in the data center.

InfoSec has realized the value of partnership across IT and other organizations and has set up a governance program to expand security responsibility into those organizations. InfoSec deputizes a number of security primes, a director from each IT service area, to be responsible for security program execution and to provide security visibility and awareness in IT services.

"It goes beyond corporate IT and includes engineering and services IT," says Stanton, "expanding our governance beyond what InfoSec could do alone. In addition, the Unified Security Metrics program is able to measure those primes and IT service owners on security compliance or if applications have undergone BAVA or DAVA, for example, to make sure that our people are doing the right thing."

Security is always evolving, and IT, likewise, has to evolve and adapt its security in the data center. Cisco ACI can rapidly deliver the network infrastructure onto which applications are deployed, at large scale with high security and full visibility into the applications. Cisco ACI will integrate into secure cloud environments, enabling consistently secure policies for both physical and virtual workloads, and allow InfoSec to protect the data center more efficiently.

Note

This publication describes how Cisco has benefited from the deployment of its own products. Many factors may have contributed to the results and benefits described. Cisco does not guarantee comparable results elsewhere.

CISCO PROVIDES THIS PUBLICATION AS IS WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING THE IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Some jurisdictions do not allow disclaimer of express or implied warranties; therefore, this disclaimer may not apply to you.

C36-740863-00 06/18