Indiana University School of Infromatics and Computing

HATS Researach Group

SDN: Data Center Case Study Writeup

Data Center Types

Modern datacenters exhibit vast architectural diversity, however, the differences can be abstracted by two sets of models that capture relevant security aspects. The first one focuses on defining administrative domains, while the latter is more concerned with identifying and defining stakeholders and their incentives.

Boundaries between different administrative domains dictate how a specific security mechanism is implemented, while, in turn, the assumptions on trust (as well as other factors) dictate where that boundary is.

Structural Models

- *End-host virtualization* allows an operator to retain control over the end. This is very important since it allows complexity to be moved out of the network. For instance, an isolation or a denial-ofservice prevention mechanisms can be both implemented inside of a hypervisor (push-back filters) and the network (VLANs). Although functionally equivalent, they have varying effects on performance, scalability, cost and security.
- *Network virtualization* allows a tenant to gain control over its network. The meaning of control is somewhat ambiguous, since it can refer to multitude of things including resource control, routing control, and other things.

Organizational Models

- *Multi-Tenant* data centers may have hundreds of thousands of customers all utilizing the same resources. Competitors may be using the same infrastructure posing difficult issues with separation. Denial of Service attacks are exacerbated as a significant 'insider' threat. As the infrastructure is available for any customer, an attacker may have privileged access if it can be collocated with their victim. Stakeholders not only include the tenants of data centers but also their customers. A security incident effecting a multi-tenant data center may not only affect the direct tenants but all of the data of the customers of applications running on the infrastructure.
- *Private/Enterprise* data centers may seem to have fewer stakeholders than a multi-tenant data center but there are still issues for their customers and the threat of data loss and security. Individual enterprises may have specific requirements dependent on their business (financial, medical, etc.). Many businesses, especially small or mid-sized enterprises, do not have experienced personnel capable of understanding the security implications of a network change. In short, enterprise data centers are very similar to the multi-tenant ones, however, more emphasis is put towards cost and/or security.
- *Content Providers* may have millions of customers using their services. What differentiates a content provider is that is has no notion of tenant, and consequently, security is shifted towards availability and resource allocation. Any disruption to the service can have widely reaching effects. Content providers such as social networks and cloud storage deliver increasingly large volumes of data. The loss of a day's service by (say) a video rental firm can have a major business impact.

Incentives

Initially, when the first data-center operators started building their infrastructure, they ran into an awkward problem of adopting the existing enterprise tools (such as VLANs and firewalls) and practices to an ill-suited environment. Consequently to this day we are still lacking proper primitives to describe desired security and isolation policies between tenants.

Drivers for SDN deployment include:

- *Network virtualization and isolation* are major drivers of possible SDN deployments. It provides mechanisms to segregate traffic for both security and traffic engineering purposes. It also covers attempts to rectify the limitations of existing protocols.
- *Fine-grain control over network resource allocation* would provide the ability to enforce variety of policies between tenants, therefore, mitigating certain denial of service attacks. The choice of policy is left to the tenants or the application.
- *Consolidation of middlebox and network* in networks that contain large numbers of 'middleboxes' such as firewalls, load balancers, wan accelerators. These boxes not only complicate the network and increase costs; they also can provide performance bottlenecks. The behavior of traffic becomes much more complicated because of multiple different devices affecting data transmission in ways that are hard to analyse and can interact in subtle ways, sometimes causing serious failures or creating security vulnerabilities.
- *Merging of L2 and L3* in a unified network fabric. An example is the use of SDN technologies to reduce the problems inherent with large broadcast domains. Large Layer 2 domains allow for transparent live migration of services from one physical host to another without requiring renumbering. The large number of broadcast messages and risks of broadcast storms however makes this impractical beyond a certain point due to issues of reliability and performance. SDN technologies could provide the means to maintain core connectivity for established virtual networks during such a storm; however, this is a function of configuration and not innate to the technology.
- *Reduced management complexity and cost* are a core value proposition for SDN. Established routes can be established as flows, with less vulnerability to route churn or other influences from external parties. High value or highly sensitive routes can be isolated to the point of invisibility from other network participants.
- *Open integration* allows operators to rapidly develop and deploy new applications services without having to wait for the lengthy process of standardization and adoption.

Weaknesses

Increased control granularity and flexibility comes at a cost. Unless the system is designed in a scalable manner, two new vectors of attacks (resource exhaustion) became possible.

- **Data-plane resource exhaustion**: The total number of actions installed can potentially render the system unusable.
- **Control-plane resource exhaustion**: In a reactive model, a controller can become unresponsive in case of a high churn.

Although not related to data-centers, it is worth mentioning the following two aspects that are not present in the traditional networks.

- **Failure and recovery**: Inconsistent views due to out-of-band control.

- **Software bugs**: Now that the network control plane has been moved to software, it becomes theoretically possible for an attacker to take control of a controller.
- **Control Plane Injection**
- **Control Plane DoS**
- **Cross Application interactions**

Technical Operations

Multi-Tenant Data Centers have high levels of redundancy and multiple availability zones where failures are expected not to be cross domains. However extremely large numbers of users use these services so when a service outage effects one data center it may effect thousands of applications tens of and millions of users, even for one availability zone. So multi-tenant data centers could be high value targets for a strategic attacker because of the number of services that use them.

Compliance issues for enterprises may also be salient. A financial institution may not only have operational security goals but may also require mechanisms to audit and verify any network security properties that a bank relies on for compliance reasons, for example to enforce a Chinese Wall between retail and investment operations.

Other systems with significant compliance requirements may include medical records systems, and other systems holding sensitive personal information (in terms of EU data protection law and sector-specific US laws such as HIPAA). This may require that systems be separated so that only clinical personnel and approved support staff have access to them. Similar though generally more stringent provisions hold in respect of classified information held not just by government departments but defence contractors. Health and defense information may have geographical limitations, in that it may not be stored or transmitted outside a givcen country of alliance without special protection measures. Such protection properties must to be capable of being audited.

As well as separation, transparency is becoming steadily more important. At present, customers of cloud enterprises have little or no visibility of into their data storage and transfer. Just as a cloud service provider might face a business demand for a 'Switzerland only' virtual network for a bank or an 'EU only' network for personal health information, so there may be further demands from firms who do not want any corporate information stored in jurisdictions where it might be more vulnerable to subpoena or to governmental coercion.