World Wide Technology, Inc.

Issues in Cloud Security

Private, Public, Hybrid

Abstract

This white paper discusses the major computer security issues confronting an organization when moving to the cloud. Even for small companies, migrating to the cloud can be a long and at times difficult process. Regardless of the planned final architecture, there is inevitably an extended period where the network is really a hybrid cloud, a mix of private and public infrastructure that slowly shifts in relative proportions as the migration proceeds. Hence, from a security viewpoint it is absolutely critical to consider and understand issues related to securing the private cloud; securing the public cloud; and securing the fabric connecting the two.

Introduction: Security in the Cloud

Cloud computing is a computing paradigm, an abstraction where data and services are accessible over the network to authorized users and processes. Although cloud computing may be fundamentally driven by a business shift from capital expenditure to operational expense, the implications for cybersecurity are profound. The security concerns and solutions are much the same, but the abstraction of computing away from the physical host entails a loss of control of corporate data and loss of visibility into where the data lives and who has access to it. The need for strong encryption and authentication protocols is as great or greater for cloud computing than for traditional data center and enterprise services.

We normally view cloud computing as providing some mixture of three core types of service: *Infrastructure as a Service (IaaS),* typically servers, storage and network devices provided through virtualization; *Platform as a Service (PaaS),* used to deploy applications that are provided by the cloud provider, provider's partners or the cloud customer; and *Software as a Service (SaaS)*, where software applications are accessed over the network (usually the Internet) as hosted services. Increasingly, a wide range of differentiated cloud offerings are appearing — *Lab as a Service* and *Desktop as a Service* are two examples — leading to growing use of the acronym "XaaS" to refer to any cloud-based service.

IaaS, PaaS and SaaS are very different in terms of cloud access and present distinct security problems. IaaS is what we usually first think of when we hear "cloud," and at the very least there is a clear requirement for strong encryption, key management and access controls. SaaS is actually the more common problem, one raising serious security concerns for *every* enterprise regardless of their intended cloud usage. The increasing prevalence of smart phones, tablets and ubiquitous WiFi, coupled with the borderless corporate network, has led to a corresponding increase in *stealth* or *shadow* IT*,* the use of off-premises services and applications without approval, or even awareness, by the enterprise. This immediately provides an attack vector for bad actors coming into the enterprise as well as an easy channel for data to leave the enterprise.

Another dimension to cloud computing with very serious implications for security is the deployment model — who owns the infrastructure, and how is it accessed? The fundamental divide is *private* vs. *public*. *Private cloud* refers to a collection of resources, usually compute and storage, used by a single organization. This is typically owned and managed by the organization itself, and hence in practical terms is little different than any other data center owned and managed by that enterprise. *Public cloud* refers to resources accessible by anyone (or at least anyone willing to pay), usually over the public Internet, managed and owned by a third party. A third category, *hybrid cloud,* refers to a combination of private and public clouds along with the connecting fabric between the two.

With regard to the different categories of service and deployment, we need to keep several points in mind:

- As a practical matter, most enterprises will use a hybrid cloud as they slowly migrate their data and computation from private to public cloud; moreover, many if not most will always maintain some level of private cloud.

- Securing hybrid cloud means securing the private and public components as well as the flow between the components and the access to the public cloud.

- SaaS presents clear security challenges that every enterprise must address and solve, but these are quite different from IaaS and PaaS. In this paper, we will focus on the challenges of securing IaaS in the context of hybrid cloud, leaving SaaS for a separate treatment.

- Securing PaaS starts with securing the IaaS aspects, followed by traditional web server security with special attention paid to Identity and Access Management. Secure application development is a critical part of PaaS security, and there have been recent developments by companies like Waratek in the area of application runtime security.

In this note we will focus on the question of securing a hybrid cloud. Much of this is no different than security of any enterprise network or data center, but there are a few issues unique to cloud that must be handled correctly.

For clarity, we follow what Amazon calls the *shared security* model: we view the security model for IaaS as divided between security **of** the cloud and security **in** the cloud. The former refers loosely to the security of the infrastructure hosting the virtual machines making up the cloud, while the latter refers to the security of everything that happens within the virtual machines — the platform. So when we discuss the *hosting infrastructure* we mean the physical hardware as well as all of the software involved — this includes at least the operating system, the hypervisor and the necessary storage and network management software.

We will use the terms *cloud infrastructure* and *cloud platform,* or often just *infrastructure* and *platform*, to refer to the two cases. This division is especially useful when assessing the security of public cloud offerings, aligning with the division of responsibility between provider and customer. For public cloud instances, the provider is responsible for securing the cloud infrastructure and the customer handles platform security. In this discussion, we will view the "provider" as owning the cloud infrastructure, with the understanding that, in the case of private cloud, this "provider" will just be the customer.

Fundamental Questions

Regardless of who is responsible for the cloud infrastructure, the set of security issues that must be addressed is the same. These questions are fundamental to cybersecurity in general. When evaluating a cloud provider, it will generally not be possible to get full transparency on most of these issues, but we believe that it is important to assess whether or not the provider has at least thought through the main issues and is willing to work with the customer to satisfy security requirements.

Physical Security: Access to the servers must be strictly controlled, which typically means careful control of the data center facilities: badged access and continual CCTV monitoring of entry points form a good starting point.

Isolation of infrastructure: The provider must be able to separate the logical infrastructure, data and applications of its customers. This is especially important because of the prevalence of multi-tenancy, multiple customers on the same physical machine.

Staffing: Continuous security awareness and remediation for cloud infrastructure requires a trained, dedicated security staff. The major cloud providers have all invested significantly in qualified staff. Any enterprise hosting its own private cloud must similarly invest in their own security staff or contract the work out to a third party.

Identity and Access Management: Absolutely critical for controlling risk to data and computing resources. The real challenge is implementing strong IAM controls without unduly hindering employees.

Encryption: Provider must use best practices to encrypt data at rest and in transit within the cloud infrastructure. This has the immediate effect of preventing unauthorized access to enterprise data, and a side benefit of providing some level of implicit authentication.

Key Management: Strong encryption is impossible without strong key management. This usually requires either a certificate-based solution, and hence a trusted certificate authority, or a dedicated appliance subject to strict access controls.

Patching and Software Updates: Fast, automatic patch management is a critical security process. For the cloud infrastructure, this certainly includes normal updates to the operating system and hypervisor software as well as any critical application and management software.

Monitoring: Continuous monitoring of the output of sensors and analytic engines is vital. This includes both manual and automated real-time assessments of security incidents and controls.

Downstream Visibility: The provider must make available to the customer all log data pertaining to real-time detection and analysis of events normally outside the customer's vision (typically events in the outer infrastructure), to enable the customer to correlate such events from activity within the inner infrastructure.

Reporting: The provider should be able to provide the customer with standard reports on security posture and events as well as non-security operational reporting.

Remediation: The provider needs to be prepared to take immediate action to remediate all vulnerabilities identified by either its own staff or a customer. This is in addition to routine patching of errors and vulnerabilities, which will typically be driven by those who maintain third-party software or the operating system.

Retention of Evidence: In the case of a breach, it is critical that the provider can give security related data (logs, packet capture) to the customer so that the customer can determine the nature and extent of the compromise of customer data. A big data strategy may be employed to support this requirement.

Encryption and Key Management

Encryption, the bedrock of infrastructure security, is in many ways the clearest piece of the security puzzle. The principle is simple: encrypt everything, all data both "at rest" and "in motion." Encryption of data at rest can mean encrypting files or databases that contain sensitive data (e.g. PII and HIPAA), although when thinking about protecting infrastructure it can also make sense to think in terms of encrypting file systems, logical volumes or physical disks. Many organizations choose to use full disk encryption because of ease and transparency — many storage solutions include native encryption that require only a password for correct decryption. Unfortunately, while this is an effective tool for limiting exposure of data in the case of physical loss of the asset, the data is still vulnerable to unauthorized access by malware or network intruders whenever an authorized user is logged on.

A much more effective solution is for the encryption to be granular down to the file and user level. When combined with effective key management, this provides strong protection of sensitive data against the greatest risks:

- Insider threats posed by system administrators with legitimate access to the infrastructure; hence the key management must be designed to be strong enough to stop the root user.

- The threat from other customers in a multi-tenant environment, a typical public cloud problem. Even though in a virtualized environment the provider will configure the hypervisor and physical hosts to attempt to guarantee isolation of customers — that is, no customer should be able to view resources like disk and memory that "belong" to another customer — vulnerabilities can, and do, occur which enables an attacker to break out of his environment and gain access to unauthorized resources.

- A remote threat to the cloud infrastructure is the classic hacker threat. This includes anything from gaining access using stolen credentials to more esoteric 0-day attacks, possibly breaking through the front door of the cloud provider or gaining elevated privileges within the (less secure) customer network and extending access to the cloud infrastructure.

In each of these cases, an authorized user or process will have access to data at rest, strong encryption and properly configured key management to prevent data from being compromised by the attacker.

Encryption of data in motion typically means SSL (ubiquitous even for the home user), IPSec or layer 2/layer 3 VPNs. This is required in many channels, from host to host within the cloud infrastructure, between cloud customer and cloud provider, from client to server in the case of private cloud, and between all the virtual appliances hosted on the cloud platform. By the very nature of cloud computing, data in motion can just as easily pass between data centers on opposite sides of the world as from one VM to another on the same physical host, which makes channel encryption between all VMs absolutely necessary.[1]

There are a number of robust commercial and open-source solutions for encryption. Companies like Vormetrics and Intel Security Group (formerly McAfee) have specialized key management and encryption solutions. Most storage vendors offer native disk encryption, a very attractive option for infrastructure providers despite its limitations. File system and volume encryption is also incorporated within the kernels of Microsoft Windows, all Linux distributions and Mac OS X. Similarly, for encryption of data in motion, nearly every computing platform now supports management of cryptographic certificates and SSL/TLS. Most also support IPSec, especially in the burgeoning world of IPv6, and there are several strong VPN solutions (for example, Cisco AnyConnect).

Identity and Access Management

The principles underlying strong Identity and Access Management (IAM) used to manage access to resources in the cloud are essentially the same as in any network.

One difference in the case of cloud infrastructure, however, is that access is usually at the system level, typically a system administrator logging on to the hypervisor or underlying host, which is a high-risk case. Here the provider is responsible for limiting the number of employees with root access, while the customer is concerned with users who have access to the cloud account management interface. In the case of the cloud platform, access might be to the virtual network via a VPN connection or directly to any number of VMs within that virtual network. In the case of PaaS applications, access maybe to a web or application server.

[1] As Google discovered to its chagrin with the Snowden leaks.

The best practices for IAM in all of these cases are the same as for any system:

- Rigorous password policy: strong passwords, frequent change, auditing and automatic policy enforcement.

- Multi-factor authentication: Ideally this includes a hardware piece, such as the DoD CAC or RSA SecurID token, but using SMS or a smartphone-based authenticator app significantly improves access control.

- X.509 certificate management: Critical for sound key management, this must include a Certificate Authority, whose certificate needs to be installed on all endpoints.

- Role-based access control (RBAC) with the concept of *least privilege*: Each user has only the read and write permissions required to fulfill his role. This also extends to carefully determining the fewest number of users with administrative privileges.

- Separation of duties: For example, the administrative users who manage the public cloud account should be tightly restricted, hence the importance of associating encryption and key management with users; system administrators should not be able to decrypt any data or volumes.

Host Defense: Toward a Trusted Cloud Platform

The most sophisticated malware is sometimes installed at a very low level of the operating system by modifying key system files, including the master boot record (MBR) and BIOS firmware. The motivation of the malware author is to gain control of the system before the host defense is active, which makes it quite simple to render such defense useless. Further, in the case of BIOS infection, the system files on the host are untouched when at rest and so pass even advanced monitoring (e.g. Tripwire). This led to an industry-wide collaboration beginning in the early 2000s to define and develop a *Trusted Computing* platform that would guarantee a trusted environment from the ground up – or, more aptly, from the boot-up on. This is usually described in terms of *secure boot* and the *chain of trust.*

As part of the Trusted Computing consortium, Intel has led the way, addressing this problem at the chipset level with their *Trusted Execution Technology,* or TXT. The TXT coordinates with an integrated third-party component called the Trusted Platform Module (TPM) to ensure that the operating environment cannot be tampered with by using a variety of strong cryptographic primitives: public key infrastructure, strong hash functions, digital signatures and symmetric encryption. For encryption, Intel has provided AES hardware instructions, and in 2015, the company released the instructions to access their new on-chipset physical random noise source, a critical piece in secure key generation and key exchange protocols.

The chief working principle espoused by Intel to achieve a trusted computing environment is *measured launch.* At boot, a few authenticated code modules are validated using TXT and values stored securely within the TPM. Once validated, these modules are used to validate the server platform: BIOS code and configuration, system state, master boot record and so on. At this point, kernel code and modules can be measured, loaded and executed. This overall process can be viewed as a chain of trust; at each step in the process, the environment has been verified as trusted by earlier links in the chain.

Since around 2009, researchers have proposed a number of architectures designed to extend the chain of trust to a trusted cloud computing platform. The basic idea is to extend the chain of trust on each host in the cloud infrastructure to the hypervisor, and from there to the VMs. Then a trusted server *attests* to the integrity of each node using strong digital signatures. The Public Key Infrastructure required here is somewhat intricate, but at the level of complexity of modern enterprise certificate designs and within the capabilities of most cybersecurity architects. We believe that this concept of a trusted cloud computing platform should be required for both private and public clouds.

Skyport Systems has recently entered the cloud security marketplace with an appliance that fully implements this concept of chain of trust using Intel TPM and TXT technology, extending the concept to stand up a hardened VM. Their target use case is secure servers within a private cloud, but the appliance can be extended easily to the concept of throwing the VM up to a public cloud platform. This provides an extremely high level of host defense, especially when combined with products like Tripwire, a host-based anti-virus program, and even more sophisticated tools like Cisco AMP for Endpoints and Tanium's solution for increased visibility into the state of the host and efficient real-time reporting.

Automatic Provisioning

A new development arising from increased use of cloud services is the need for rapid, on-demand provisioning of secure servers in the form of virtual machines in the cloud. One motivation is fluctuating customer demand coupled with the pricing structure of public cloud providers: the customer is charged only for the time when his VM is up and running, so it makes sense to tear down a VM as soon as possible. Such servers might have a life span measured not in years, but in hours or even minutes. Coordinating the resources to handle such fast build-up and tear-down is challenging enough, but doing so securely is extraordinarily difficult. The solution developed by Skyport Systems described in the previous section handles the problem of provisioning hardened VMs from a trusted platform quite effectively.

There is also an interesting open-source solution to consider as well. In 2015, two researchers (Valtman and Ferber) presented a methodology for solving this problem[2]. The workflow, which they call the *Cloudefigo lifecycle,* can be broken down as

- Launch
- Configure and harden
- Encrypt disk volumes
- Scan for vulnerabilities
- Move to production

At launch, the virtual machine is automatically updated and all current appropriate patches are applied — this avoids the common security hole that exists between release and application of a critical patch. The keys for the encryption need to be generated strongly and stored securely, not a trivial exercise. Vulnerability scanning is obviously an important step, but it requires tuning to avoid using excessive computing resources as well as to decide on the allowable level of potential vulnerabilities.

When, on termination, the virtual machine is deleted, special care must be taken to delete all of the encryption keys as well ("special care" because the keys are stored off the VM). During the time the server

[2] They also developed a scripting solution based completely on open-source tools; see www.cloudefigo.com.

spends in production, the VM data will invariably be backed up *somewhere*. Since the critical data lives on an encrypted drive, destroying the keys will prevent that data from being exposed.

Network Defense and Monitoring

Once the host defenses are in place, identity and access control is properly implemented and all the channels from private to public cloud, as well as all the connections between all pairs for virtual machines on the cloud platform are encrypted, most of the job of network defense is done. What remains is to stand up firewalls at all access points, establish DMZs between the firewalls and the internal networks, and monitor everything.

In the case of public or hybrid cloud, it's critical to follow best practices outlined by the public cloud provider. For example, Amazon recommends the following for users of AWS:

- Use the Amazon Virtual Private Cloud (VPC)
- Enforce RBAC using security groups with least privilege
- Access the VPC using a secure mechanism: for example, one or more of IPSec, VPN and AWS Direct Connect
- Use firewalls, IDS, IPS and other host-based security tools with appropriate (minimal) authorizations
- Protect the integrity of the operating and file systems
- Use secured configuration management tools

On the private cloud side, network security is essentially typical data center security, with the added needs of providing secure connections to the private cloud (VPN, IPSec) and encrypting all the connections between servers within and between data centers.

Monitoring the hybrid cloud is a question of obtaining as much security-related information as possible from the public cloud provider and combining it correctly with the logs and other information gathered from the private side. The public provider should make available, either directly or via an API, logging for all the relevant events pertaining to the piece of their infrastructure being used by the customer, while the customer is responsible for the information from the cloud platform (e.g. operating system logs from VMs deployed to the public cloud). Given the volume of information and the relative complexity of correlating events throughout the hybrid cloud, more sophisticated monitoring tools, like Tanium for endpoint logging and Splunk for aggregation and analysis, are even more necessary than in a traditional enterprise network.

Using Public Cloud: Amazon AWS and Microsoft Azure

Beginning in 2006, Amazon has offered a cloud computing platform called Amazon Web Services, usually abbreviated AWS. The two main thrusts of AWS are cloud storage (Simple Storage Service, or S3) and remote computing (Elastic Compute Cloud, or EC2). In the decade or so of cloud business, Amazon has developed AWS to provide a wide range of services. In addition to compute and storage, AWS can be leveraged for networking, content delivery, database requirements, application services, big data and analytics, and several other services that defy categorization.

Microsoft Azure is somewhat newer than AWS (announced in 2008 and released in 2010), but it has developed into a strong contender for leadership in the public cloud arena. Azure offers storage and

computing, as well as a long list of other business services, such as data management, business analytics, content delivery networking and development support.

Both AWS and Azure offer a comprehensive suite of security tools, monitoring assistance and output, and compliance certifications (e.g. HIPAA, ISO-27001). Here is only a partial list of security features in both AWS and Azure that can be leveraged by a customer:

- Network security: availability of built-in firewalls and encryption of data in motion (TLS).
- Private virtual subnets:. Inventory and configuration management tool for deployment of hardened VMs as well as tools to identify, track and manage resources.
- Data encryption available in storage and database services.
- Key management and storage can be done using cloud-based HSM.
- IAM services, including multifactor authentication for privileged accounts.
- IAM integration with corporate directories.
- DDoS prevention and mitigation.
- Monitoring: log aggregation, alert notifications and visibility into API calls.
- Frequent independent auditing to ensure compliance is maintained with all major certifications.

Beyond these mostly traditional security controls, AWS and Azure provide a fundamental level of security derived from the use of software-defined networking: every packet sent by a customer is handled by a custom hypervisor. This enables the provider to guarantee that packets simply cannot travel between VMs in different virtual subnets, regardless of their distribution with regard to physical hosts. Even on the same physical host, there is simply no logical path between VMs unless explicitly authorized. This means that no matter the size of the attacker's footprint within the public cloud or how many packets he throws, he will *not* get even one to reach the intended victim. Bottom line: he can't attack what he can't see.

Conclusion and Recommendations

Although the principles of cloud security are fundamentally the same as in traditional enterprise networks, there are several key areas that need to be understood and addressed distinctly as cloud issues. Multi-tenancy has been, and will continue to be, a major issue, especially in light of recent geo-political shifts in privacy laws. Integration of public cloud security infrastructure with the private cloud and overall enterprise security infrastructures is also a critical area. Once the computing and storage environment is stretched, the danger of losing visibility into key enterprise resources escalates, which means even more focus has to be turned to incorporating information from *all* security sensors into a central SIEM or database for real-time inspection and analytics. Once identity and access management is similarly extended and integrated between enterprise and cloud networks, the overall security picture can be augmented to provide a complete understanding of entity behavior, and, in particular, raise awareness of when a given entity starts engaging in anomalous behavior that warrants investigation.

Fortunately, public cloud security has matured significantly in the past two years, especially the measures and services offered by public cloud providers like Amazon and Microsoft in their AWS and Azure environments. With the use of specially purposed, streamlined hypervisors and software-defined networking, the virtualization inherent in cloud provides a fundamental level of security that is actually not possible in traditional computing environments. For example, every packet transmitted in AWS is intercepted by a network monitor, wrapped with information tying it to a particular virtual private cloud and transmitted *only* within that Virtual Private Cloud; it is simply impossible for an attacker within the AWS world to even send a packet to an intended victim. Moreover, cloud service providers have recognized the

importance of giving customers continuous visibility into cloud security events and low-level access via well-defined, well-supported APIs.

The main lesson, then, is that public cloud security is today advanced enough to no longer be a deterrent for companies looking to reap the numerous business benefits of moving to the cloud. In fact, providers like Microsoft and Amazon claim, with some justification, that with proper planning, assessment and coordination of the security environments throughout the larger hybrid environment, the overall result is actually more secure than the original private network.