FIREEYE™

Top 5 Cloud Security

Myths Debunked

The Great Cloud Migration

Five years ago, almost every organization on earth appeared to be just weeks away from moving all workloads into the cloud—or at least it seemed that way. But reality has moved at a far more measured pace. Instead of a ubiquitous, all-in phenomenon, most organizations turned to a cloud-ish approach in which many new applications are developed on the cloud, but most older apps and infrastructure either stay put or have long-term migration plans.

So, the promise of moving to the cloud is alive and well; it's just happening in slow motion. This means:

- Almost everyone is new at this.
- Developers are expecting to be unencumbered by organizational process.
- Technology changes from month to month.

The combination of these factors creates a situation that is simultaneously exciting for organizations as they reap the benefits of the agility cloud solutions provide, and disconcerting because they are inexperienced with the cloud.

These positive and negative tensions have generated several myths about and around shifting business operations to the cloud.

- The cloud is unsafe
- My organization doesn't use the cloud
- My cloud provider will keep me secure
- The cloud is just someone else's computer
- Advanced adversaries aren't attacking the cloud

This is the true origin of cloud fears: The internet at large has access to the cloud.

Myth 1: The cloud is unsafe

The cloud itself is not inherently unsafe. When used properly, it is no less safe than a typical datacenter. In all the FireEye Mandiant incident responses conducted on public clouds, we have yet to see a case where the cloud infrastructure itself was exploited. We've discovered improper cloud configuration or vulnerable customer code, but not flaws in the cloud provider's code or infrastructure.

It's through customization of a cloud environment that vulnerabilities are introduced. For instance, if you create a storage bucket using AWS S3 or Azure Blob Storage, they are locked down by default and only accessible to administrators and the bucket's creator. However, to use the data in the bucket, access must be provided to servers or directly to users, and granting and administrating these permissions is where many organizations are challenged.

This challenge isn't inherently different than traditional datacenter security, where misconfiguration, stolen credentials and vulnerable code are also the primary ways to gain unauthorized access. However, by default, a typical data center's firewalls prevent inbound access to assets, whereas cloud services (excluding virtual machines), allow authenticated traffic from anywhere. Some services can be configured by policy to only allow certain IP ranges, but this is not the default setting.

When we examine this fear, we find a rebuttal in the form of the "Zero Trust Model" made famous by Google and Microsoft in which firewalls are removed completely, and all access is authenticated and authorized.

Zero Trust networks eliminate the concept of trust based on network location within a perimeter. Instead, Zero Trust architectures leverage device and user trust claims to gate access to organizational data and resources.[1]

Zero Trust, and the cloud in general, claims that over the years, firewalls have allowed organizations to become complacent in their security by serving as a crutch for vulnerable systems. Organizations incorrectly assume that any intruder is outside the firewall perimeter, and therefore, the vulnerabilities are mitigated. Zero Trust argues that one must assume an intruder is already inside the firewall perimeter and therefore, the firewall does not actually protect you.

While the Zero Trust model may be extreme for servers and virtual machines, it's an excellent example of what it takes to secure cloud services: Without a perimeter, services must assume all clients are untrusted, and while the underlying code for those services is not necessarily vulnerable, the configuration or credentials may be.

1 https://cloudblogs.microsoft.com/microsoftsecure/2018/06/14/building-zero-trust-networks-with-microsoft-365/

Myth 2: My organization doesn't use the cloud

The term "cloud" includes the category of software as a service, and virtually every organization uses some form of web service, whether it is for human resources, banking, shipping, content management, web hosting or any of the other activities that take place in a modern business. Even if organizational policy does not explicitly permit cloud services, or no overt evidence of cloud service usage exists, your organization may still rely on the cloud.

One of the great challenges for security practitioners is not only securing these services against accidental misuse but detecting their use in the first place. Visibility into "shadow IT" or any unauthorized use of cloud services is a common pain point for organizations. Even cloud access service brokers (CASBs) cannot spot the occasional (but sometimes critical) use of cloud services by employees in their homes.

When security empowers developers to get work done, there is a far higher likelihood that new cloud services being used will be visible, which makes them defensible.

Most employees have needed to use a service such as Google Drive at some point to share a file with an external partner. There may be an organization-approved method, but many times security authorization processes create hurdles for getting business concluded quickly. Human nature can cause well-meaning employees to operate outside of the purview of their security staff.

Two things must happen to solve this issue:

Employees must be able to accomplish their goals using the organization-approved method.

Proper visibility and controls must be in place around the sanctioned cloud services.

Cloud security starts with visibility. Whenever a developer spins up a new cloud service, they should be able to easily centralize telemetry. Analysts can then go to this single location to review the security status of all cloud services.

The best way to ensure compliance with this security measure is to incentivize developers to approach the security team, instead of the other way around. This can be accomplished by making the telemetry useful for operational monitoring, so it serves the dual purpose of security and operations.

It may not always be possible to get visibility into every cloud service. Some of the more esoteric services may not offer telemetry, and some require additional purchases to enable it. For example, Microsoft Azure Active Directory can only view sign-in logs with its "Premium" edition, and Salesforce requires the purchase of its "Shield" product to see audit data. No matter how easy you make it for developers in your organization to centralize their telemetry, they still may have reasons (whether good or bad) not to do so. Simply knowing that there may be cloud services in use you aren't aware of is helpful, especially for threat modeling (thinking about how an attack might happen).

Myth 3: My cloud provider will keep me secure

Under the shared responsibility model, the cloud tenant is the ultimate custodian of their data and is responsible for safeguarding it. The cloud provider ensures that the facilities are secure, the hardware is not compromised, and the underlying software and operating systems of any services offered are secure. But it's up to the customer to make sure that virtual machines are patched, applications are not vulnerable and permissions are appropriate.

Safeguarding the cloud consists of three high-level activities:

Protect credentials used to access resources and monitor for compromise.

Be vigilant for and guard against misconfiguration.

Centralize telemetry data for visibility to support security monitoring to audit trails

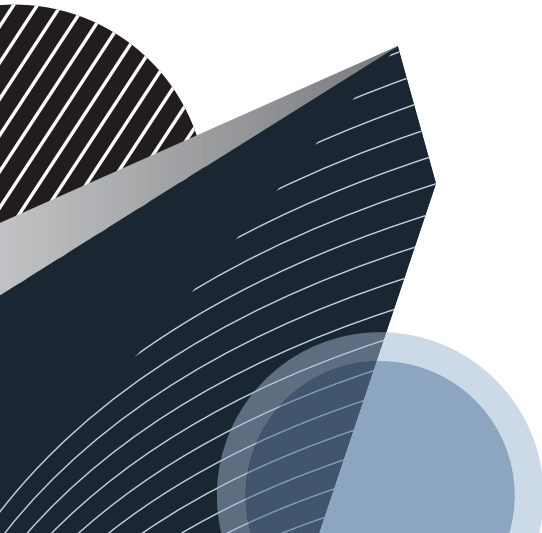

Cloud providers have many helpful native tools, but cloud consumers still have many more responsibilities. For instance, a given cloud provider may have tools for detecting or enforcing rollbacks on risky configurations, but consumers must integrate those tools into their organization's policies and procedures. The tools also need to be merged with security controls and visibility for any other clouds the organization may be using.

Since cloud providers probably aren't intimately familiar with every consumer's line of business, it is ultimately the consumer organization's responsibility to verify that all is as it should be.

Any points at which data is shared across organizational boundaries are potentially vulnerable seams in defenses, because access to data must be granted and then audited rather than uniformly denied. Each point of access becomes an audit task for the security organization, and complete visibility is required to ensure that all data being accessed is truly authorized.

Complete visibility means that the organization's security analysts have direct, programmatic access to all relevant telemetry data and can operationalize it into their standard operating procedures. This goes beyond point-in-time audit capabilities to continuous monitoring for anomalies.

Myth 4: The cloud is just someone else's computer

Securing the cloud is not like securing a computer in someone else's data center. There are storage services, containers and other non-traditional services to consider in addition to more familiar virtual machines. These services may be comprised of hundreds or thousands of real servers spread across many data centers, all to fulfill a single service request.

This means there are additional visibility requirements and more planning required to provide security controls and instrumentation around distributed and non-discrete compute offerings. These services may have an API to use, but the concepts of IP addresses and operating systems often don't apply.

Security configuration and controls for these services won't use traditional security implementations like firewalls and antivirus.

Encryption requirements may also dictate that individual services are encrypting data at rest. All major cloud services use SSL/TLS for communication, but understanding where data is written between services is important for regulated data.

The elastic nature of the cloud can also present new challenges when defending against resource misuse. A traditional app that lives in a data center will have a finite amount of resources attached to it. This can act as a safety net in the event of misuse of a publicly exposed service: The app will eventually fail and be reduced to a state of denied service. An app built in the cloud on elastic resources such as autoscaling groups and serverless functions or containers may continue to scale out to a high capacity instead of failing. This is generally a good thing for a legitimate application but can cause a serious issue if the requests are illegitimate.

For example, if an app processes uploaded photos and an attacker submits a deluge of unsolicited photos to process, the app may scale out, leaving the app owner responsible for considerable additional costs. Cloud infrastructure has limits to help with this, but they have to be well understood and set in advance.

Myth 5: Advanced adversaries aren't attacking the cloud

Attackers follow data. As data goes into the cloud, so will the attackers. Approximately one quarter of our Mandiant incident response engagements involves assets housed on a public cloud, and almost every IR we perform involves public cloud in some way. The cloud does not hinder threat actors—they easily adapt their tools, tactics, and procedures to compromise cloud accounts to get access to confidential data, steal computing resources and spy on targets.

The average organization can move more quickly and lower costs by moving to the cloud, but they should understand that anything of value they put there will be a target, and they need to protect resources accordingly. This means they should not only implement basic best practices for cloud security, but also have their security operations ready to actively hunt down advanced attackers that pursue data into the cloud.

Threat mitigation in the cloud requires two things: telemetry data on which to apply security operations and threat models that dictate what adversarial tools, tactics and procedures should be hunted. This can be broken down into four categories of capability:

Intelligence

Application of threat intelligence indicators to telemetry

Rules

Application of known threat patterns to telemetry

Analytics

Organization of telemetry to show anomalies

Hunting

Hypothesis-based searching

Together, these four capabilities can protect an organization from known-threats and suggest how to find new threats by forecasting behavior. Some activities (intelligence, rules and analytics) are automated, but all of them require human interaction at some level to fully scope high-severity events.

Hunting is special: It requires experienced security operators to use what they know about specific threat actors or what is "normal" for an environment to find points of interest in the telemetry data. For example, if a policy states that all cloud assets are to be instantiated using templates, a hunting activity may be to look through all assets created outside of templates. If the hunt becomes sufficiently valuable or simplistic, it can be converted into a rule to automate the action.

Approximately one quarter of our Mandiant incident response engagements involves assets housed on a public cloud...

Conclusion

The cloud brings great promise of allowing organizations to do more with less, democratizing resources so that startups can enjoy the same capabilities as juggernauts of industry, and allowing apps to grow to a scale not seen before. But this means that cloud resources require dedicated, specialized attention to ensure they are not being misused and that their data stays secure. Organizations must keep up with security tools and training they need to ensure that they are ready for the big shift.

To learn more about cloud security, visit: **www.FireEye.com/cloud**

FireEye, Inc.
601 McCarthy Blvd. Milpitas, CA 95035
408.321.6300/877.FIREEYE (347.3393)
info@FireEye.com

About FireEye, Inc.

FireEye is the intelligence-led security company. Working as a seamless, scalable extension of customer security operations, FireEye offers a single platform that blends innovative security technologies, nation-state grade threat intelligence and world-renowned Mandiant® consulting. With this approach, FireEye eliminates the complexity and burden of cyber security for organizations struggling to prepare for, prevent and respond to cyber attacks. FireEye has over 6,600 customers across 67 countries, including more than 45 percent of the Forbes Global 2000.

FIREEYE™