ALIEN VAULT

WHITE PAPER

Beginner's Guide to Hybrid Cloud Security

From the Data Center to the Cloud

Today, organizations are rapidly shifting their IT workloads to public cloud infrastructure to recognize operational and cost-savings benefits. Cloud infrastructure is typically easier to deploy, easier to manage, and more cost-effective than deploying everything in an on-premises network or data center. In fact, a recent IDC report predicted that cloud adoption in organizations will grow 45% by 2018.

But, even as applications and workloads migrate to the cloud, on-premises monitoring remains a critical part of an organization's security and compliance management program. In addition, as cloud migration reaches a critical mass, IT security teams are realizing the challenges of monitoring their "hybrid cloud" environment—a mix of public cloud, private cloud, and physical network infrastructure—in a complete, yet efficient way. This makes hybrid cloud security a big concern—and often a big question mark—for IT security teams.

Public cloud environments face many of the same security challenges as on-premises deployments, including familiar attack strategies. However, attack strategies manifest in the cloud somewhat differently than in on-premises environments, thanks to the separation of security concerns in the cloud as well as the unique architecture and scalability of cloud environments.

To deal with security challenges in the cloud, as you do (or should do) in your on-premises environments, you must consider how the threat landscape changes as you move from the data center to the cloud. You must also consider the security resources provided by cloud service providers and how they can augment your own security tools and measures to deliver complete **hybrid cloud security**.

In this AlienVault® Beginner's Guide, we will explore the security challenges that impact cloud and hybrid cloud infrastructure environments, and discuss the best methods of detecting them. The four main areas of focus will be:

- **Part One:** The Cloud Security Shared Responsibility Model
- **Part Two:** Familiar Network Security Challenges Persist in the Cloud
- **Part Three:** New Security Challenges Emerge in the Cloud
- **Part Four:** Hybrid Cloud Security: A Smarter Approach to Threat Detection

Part One: The Cloud Security Shared Responsibility Model

Any discussion on hybrid cloud security requires a fundamental understanding of the **Shared Responsibility Model** and how it applies to cloud infrastructure-as-a-service (IaaS) security concerns.

When you run your systems and applications and store your data on-premises—whether in a physical or virtual network environment—the full weight of IT security sits squarely on your shoulders.

By contrast, when cloud service providers like Amazon and Microsoft own the network infrastructure and make it available to you as a service—whether as Infrastructure as a Service (IaaS) or Platform as a Service (PaaS)—you no longer have visibility of the underlying hardware, computing, and networking resources that support your cloud workloads. Instead, the cloud service provider takes responsibility for maintaining and securing the cloud infrastructure, from building access to the securing of network and server hardware, and including oversight of the hypervisor hosting virtual machines. You are responsible for securing the operating systems, applications, and data running on cloud accounts. This separation of security concerns is known as the Shared Responsibility Model—a main pillar of cloud and hybrid cloud operations and security.

While you are responsible for securing anything that you deploy on the cloud, cloud service providers have a shared interest in your security success and provide services to help you more easily implement security best practices for controlling access and limiting network exposures. In fact, many cloud services provide a level of visibility into the cloud environment that IT managers can only dream of from their on-premises infrastructure.

Cloud service providers supply tools to help you better defend your virtual environments. For example, leveraging cloud environment logging and monitoring capabilities like AWS CloudTrail provides you with the ability to see the actions being taken by both legitimate users and bad actors operating in your cloud environment.

These services are designed to work in conjunction with your cloud-based security management tools. While many traditional security tools, such as firewalls, file integrity monitoring, and centralized logging, remain effective as you expand your perimeter and move data into the cloud, adding layers of security measures that are purpose-built for the cloud can help you to better secure and monitor the full environment. We'll look at this more closely in part three of this beginner's guide.

Learn more about securing your hybrid environment with AlienVault USM™ Anywhere >

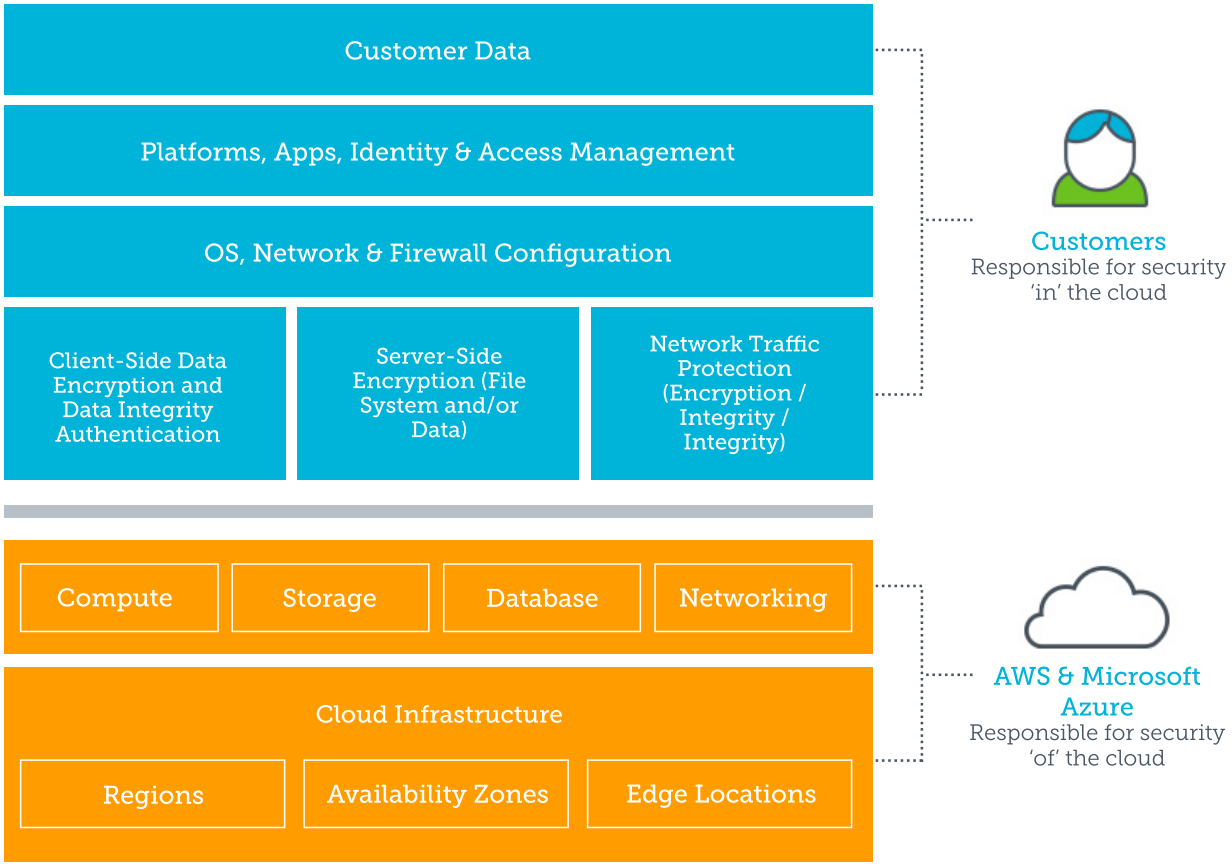

Part Two: Familiar Network Security Challenges Persist in the Cloud

Cloud environments face many of the same security challenges as on-premises deployments, including familiar attack strategies. Many of the attack strategies that target on-premises infrastructure, such as code injection and cross-site scripting (XSS), persist in the cloud and can be dealt with using traditional tools like firewalls and proxy servers.

However, attack strategies manifest in the cloud somewhat differently than in on-premises environments, thanks to the separation of security concerns in the cloud as well as the unique architecture and scalability of cloud environments.

Let's look at four well-known types of attacks and consider how they manifest in cloud environments.

Distributed Denial of Service (DDoS)

DDoS attacks work on a simple premise: flood a service or website with so much network traffic that it effectively crashes the service or site. DDoS attackers orchestrate a horde of botnet hosts to send requests repeatedly to a target at the same time. Because the hosts are distributed across many locations—or IoT devices, as witnessed in the recent Mirai botnet attacks—traditional defense tactics like blocking a particular domain or IP range are not effective. This attack strategy remains the same whether the service is hosted on-premises or in the cloud.

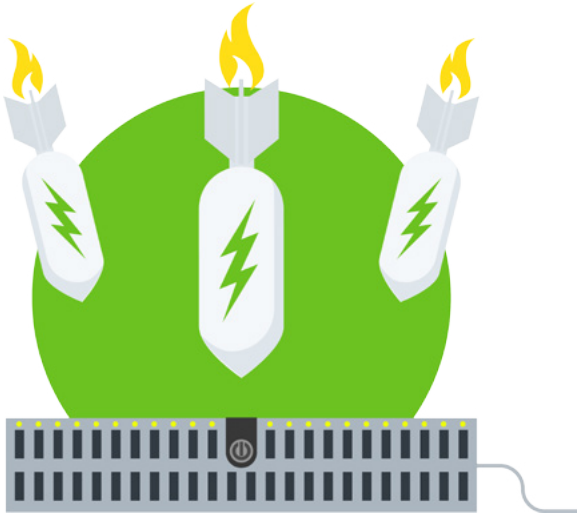

DDoS is a numbers game between an attacker's resources and a victim's computing and networking capabilities. In the cloud, your resources are elastic, so you can dynamically add more resources to meet a sudden spike in demand. This provides some built-in DDoS resilience, but it comes at a cost. As you spin up additional cloud computing resources, you can quickly drive up your monthly bill to your cloud service provider.

Another consideration in cloud environments is that some resources are shared, so a DDoS attack against another user's system has the potential to drain resources from your workloads and cause your services to become slow or unavailable. However, cloud service providers take responsibility for mitigating and protecting against DDoS attacks on shared infrastructure. In addition, cloud service providers protect against low-level network attacks to the cloud infrastructure (e.g. SYN Flood, malformed packets, etc.) as part of the shared responsibility model.

Exploiting Vulnerabilities

The beginning of a malware infection typically starts by an attacker finding a vulnerability in an OS or application and exploiting it to download malware and gain control of the system. This could be on something as incidental as a corporate printer. Once the attacker has a foothold inside your environment, he can move around laterally to find targeted data.

A strong vulnerability management program is essential to minimizing the attack surface of your network environment. By proactively finding and fixing your vulnerabilities, you reduce the likelihood of attackers exploiting them for harm. The same is true in cloud environments.

Cloud providers do provide some vulnerability management support. For example, they typically supply users with libraries of up-to-date patched OS instances that users can deploy into their environments. This is a good starting point, but in the shared responsibility model, automated patching stops at the point of deployment.

Ultimately, cloud users are responsible for identifying and managing vulnerabilities and patching above the hypervisor layer. For example, Amazon Web Services recommends that you map all of your assets to threats and then conduct vulnerability assessment and impact analysis on those assets to get a complete picture of your threat posture. This requires a comprehensive hybrid security management solution that can bring together asset inventory, vulnerability assessment, and threat intelligence into a unified view.

Another point to consider is that many of the services that cloud providers offer to IaaS customers are managed and protected by the cloud providers, for example, AWS S3 (storage) and RDS (database). When you use these services, you are only responsible for protecting your data, not the service itself. So, any patch management work falls on the cloud service provider, saving you time and effort.

Brute Force Attacks (Password Cracking)

Why pick the lock if you can kick in the door? That's the logic behind the brute force attack, one of the most common security exploits. The idea behind the brute force attack is to try all possible combinations of passwords until an attacker finds the one that works. These attacks persist in part because there are many automated tools are available (e.g., John the Ripper, Brutus, Wfuzz) and pre-built digests help them crack accounts. In addition, users continue to be a weak link, often choosing simple, easy-to-crack passwords.

So, does the cloud inherently improve defenses against password compromise? Arguably, readily available services like AWS Identity and Access Management (IAM) and Azure Active Directory (free tier) provide better password security and enable extra security measures like multi-factor authentication (MFA). However, the only real defense against password compromise is to apply good password hygiene, and hygiene applies equally in the cloud as it does on-premises.

One element that is unique to cloud is that root account credentials, if not handled properly, can be publicly accessible from the internet. A compromise of this credential gives attackers "the key to the kingdom," giving them control over your cloud environment and the ability to spin up cloud resources (perhaps while you're asleep), leaving you with the bill. There's no direct parallel of this type of compromise in your on-premises environment, since the resources in your data center are likely owned, static, and finite.

Web Application Attacks

Securing applications from attacks is clearly the responsibility of cloud users in the cloud security shared responsibility model. Web application attacks can usually be mitigated with better coding practices or supplemented with security technologies like web application firewalls (WAF) and proxy servers. Today, most security vendors offer licensed products for the cloud similar to the products they provide for on-premises environments. Some cloud vendors have also added free tools to their offerings (e.g. AWS WAF for CloudFront) that defend against common attacks like cross-site scripting and code injection.

Part Three: New Security Challenges Emerge in the Cloud

When you migrate workloads and services to a public cloud like AWS or Microsoft Azure, new security challenges begin to take shape that didn't previously exist in on-premises network environments. Also, because cloud computing infrastructure is dynamic and scalable, it changes more frequently, adding a layer of complexity in monitoring your cloud security.

In this section, we'll look at some of the unique security challenges that appear in public cloud computing and why legacy security monitoring tools that were built for the data center may not be sufficient in addressing these challenges.

Learn more about overcoming AWS security & compliance challenges with AlienVault >

Learn more about overcoming Microsoft Azure security challenges with AlienVault >

Safeguarding the Keys to Your Kingdom

Access keys and root account credentials are a major security concern in public cloud environments. They are the proverbial "keys to your kingdom," and if an attacker compromises them, they can gain access and control over your cloud account. Once inside your account and with full permissions, an attacker can spin up cloud resources on your dime, steal your data, or run malicious software on your resources and with your reputation.

Compared to physical network environments wherein the infrastructure is ultimately finite and static, cloud environments are super elastic and can be scaled rapidly from a central management console. The only real limitation is the size of your wallet. A malicious actor with your root account credentials could easily spin up an enormous amount of resources (to mine bitcoins, for example), leaving you with an enormous bill.

While it seems like a no-brainer to not publically share your root account credentials, there have been many cautionary tales in recent years of web developers and even security industry analysts who have accidently published their AWS access keys to GitHub or other public locations, resulting in thousands of dollars of fraudulent charges racked up overnight. And, although the cloud service providers in these tales often come to the rescue to notify victims of fraudulent activity and to remediate charges, it's important to remember that it is ultimately your responsibility to keep your credentials and access keys secure.

You can read more on how to secure your AWS root access keys here.

While it's an important first step, it's not enough to hide your keys and hope for the best. In many of these cautionary tales, cloud users had initially tried to scrub their encryption keys from publically shared data, only to miss an instance and later discover the mistake after their bills had skyrocketed. To avoid "alert by bill shock," you should constantly monitor your cloud environment for suspicious root account logins, changes in security policies and privileges, and other anomalous activities. A cloud-native SIEM solution enables granular security monitoring and analysis of cloud activities by integrating directly into your cloud environment.

Managing Cloud User Activities

You probably aren't the only person in your organization accessing your cloud resources. But, do you know who is and what they're doing in your cloud environment? The nature of the cloud lends itself to a greater number of unsanctioned or "shadow IT" projects and IT decentralization—whether intentional or unintentional. In fact, as a recent Cisco report found, "companies are using up to 15 times more cloud services to store critical company data than CIOs were aware of or had authorized."

It's essential to your cloud security management to know who (users and services) are using your cloud resources so that you can identify the account activities that constitute "normal user behavior" and investigate the activities that do not. However, trying to implement organization-wide cloud security controls can leave security folks feeling like professional cat herders.

Fortunately, within your cloud accounts (the sanctioned ones that you're aware of anyway), you have multiple methods and tools to optimize your identity and access management (IAM) as well as to protect your user accounts from threats like phishing attacks. These include creating role-based permission groups and enforcing multi-factor authentication policies for your users and APIs.

Cloud service providers also deliver services that enable you to monitor environmental activities and changes. For example, you can leverage AWS CloudTrail to see all of the user activities and API calls made to your account. With a cloud-native SIEM solution that has direct hooks into these services' APIs, you can readily monitor this data and perform security analysis in correlation to your other data sources and threat intelligence.

Finally, as with on-premises security management, it's important to employ the principle of least privilege in your cloud environment. While it's a seemingly obvious practice—to give your users only the absolute minimum level of access needed to do their jobs—in practice, it can be slowly chipped away at as admins and developers ask for small exceptions here and there.

Navigating Your Blind Spots

The big benefit of cloud computing with Infrastructure as a Service (IaaS) is that you no longer have to deal with the capital and operational expenses of managing your infrastructure. The network infrastructure is abstracted away from you, maintained and secured by cloud service providers, and delivered as a beautiful economy-of-scale price.

The trade-off of cloud IaaS is that, by relinquishing your responsibility for the underlying network infrastructure, you also relinquish some of the deep network traffic visibility that security professionals are accustomed to having in on-premises networks. In the cloud, it's no longer feasible to drop a passive tap or SPAN port on the wire to monitor traffic to detect threats and intrusions as you would in your own network. This means that legacy network intrusion detection systems (NIDS) are no longer effective tools for cloud security monitoring. This calls for...wait for it...yes, a paradigm shift.

While legacy security monitoring tools like NIDS cannot be readily shoehorned to fit your cloud security monitoring needs, you can still navigate the security blind spots in the cloud to get a complete picture of your security posture. You just need a new paradigm around cloud security monitoring.

Part Four: Hybrid Cloud Security, A Smarter Approach to Threat Detection Across All Environments

Regardless of the environment—whether cloud or on-premises—the goal of threat detection is the same: to prevent data loss, financial loss, and business disruption. Yet, as the environment and the infrastructure changes, you need a new approach—and perhaps new tools—to tackle this familiar challenge.

In this section, we'll examine how and why you should approach hybrid cloud security with different methods, tools, and best practices from those in the data center.

Best Practice: Develop Good Identity and Access Management Practices

In the previous section, we looked at how mismanaging your cloud credentials can be an expensive mistake. It's also a common pitfall in cloud security. According to Gartner analysts Neil MacDonald and Greg Young, "Through 2020, 80% of cloud breaches will be due to customer misconfiguration, mismanaged credentials or insider theft, not cloud provider vulnerabilities." They recommend building your cloud security on a solid foundation of identity and access management (IAM) practices, and I agree.

Cloud identity and access management best practices include both the use of cloud provider IAM services as well as establishing organizational policies around those services. For example, in AWS, you can use IAM groups to more easily manage cloud users who need the same permissions to AWS resources to do their jobs, and, you should define IAM groups based on the principle of least privilege.

Whether you're just starting out with public cloud computing or you already have production systems in play, make the effort now to establish your IAM guidelines and policies and establish a routine to ensure that your IAM services are continually configured and working accordingly.

You can find more in-depth best practices on AWS and Azure identity and access management below, respectively.

AWS Security Best Practices White Paper >

Azure Identity Management and Access Control Security Best Practices >

Furthermore, for your hybrid cloud environment, you can streamline identity and access management with cloud provider services that either synchronize, consolidate, or federate your cloud identity management with your on-premises directory. These include services like Azure AD Connect, Azure AD Federated Services, AWS Directory Service, and AWS AD Connector. As one Microsoft Azure article points out, integrating your on-premises identity and cloud identity not only reduces administrative overhead, but also decreases the likelihood of mistakes and security breaches.

IAM is not a "set it and forget it" configuration. Rather, it's important to constantly monitor your hybrid cloud environment for suspicious root account logins, changes in security policies and privileges, and other anomalous account activities. By enabling a cloud-native SIEM solution to collect and analyze your cloud access logs and API calls, you can identify compromised account credentials sooner to prevent or mitigate the damage of a cloud breach.

Best Practice: Know What Security Data to Look for in the Cloud and Where to Find It

In section one, we looked at common threats, like DDoS and brute force attacks, and how they persist across on-premises and public cloud environments. In fact, many of the threats facing the cloud today are not unique to public cloud environments. Rather, attackers use many of the same attack methods and infrastructure in cloud-based attacks as they do in on-premises attacks. You just need a new approach to recognize indicators of threat in the cloud and to know where to look for them.

Log Collection in the Cloud

To detect threats in your public cloud environments, first you need to know what log data sources are available to you and may be "interesting" or useful from a security standpoint. Then, you must be able to collect and send the log data to your SIEM for correlation and security analysis.

As with on-premises infrastructure, cloud log data sources include system logs and the log files of the assets and applications that you launch in the cloud. As an additional layer of information, you can collect access logs from your cloud infrastructure (e.g., AWS EC2, ELB, S3) to know:

- Which users are accessing my cloud resources and workloads?
- Where and when are they signing in?
- What resources or instances are being spun up or down?
- Has anyone altered my security groups or IAM roles?

These logs can be collected through cloud services like AWS CloudWatch and CloudTrail, and in Azure, Diagnostics and Monitor. AlienVault's cloud-native security monitoring solution, AlienVault USM Anywhere directly hooks into these services to collect log data from your AWS and Azure environments, ensuring the most comprehensive set of data for security analysis while significantly reducing the complexity of cloud log collection.

Learn more about AWS log management with AlienVault >

Learn more about Azure log management with AlienVault >

Intrusion Detection in the Cloud

Once you've established a centralized cloud log management solution, your cloud is secure, right? Not quite. Cloud log management is only a prerequisite to complete hybrid cloud security. After you've gathered your log data, you still need a way to perform cloud intrusion detection—to correlate and analyze your cloud log data in context of the latest threat intelligence to identify intrusions in your cloud environment. This requires a SIEM that's built to natively perform cloud intrusion detection.

AlienVault USM Anywhere is continuously updated with **cloud-specific correlation rules** based on the latest threat intelligence, so even without having to write your own correlation rules, you have the assurance of up-to-date cloud security analytics and alarms that give you actionable insight about threats and intrusions in your cloud infrastructure.

For example, the creation of an AWS EBS snapshot could mean that someone within your organization is taking an incremental data backup. However, attackers can use data replication services like EBS snapshot to get access to production data. In another example, if your cloud access credentials are used from an IP address that's external to your cloud environment, it may indicate that your credentials have been compromised and are being used by a malicious user. In both examples, USM Anywhere would generate alarms to alert you to these activities.

Learn more about intrusion detection in Azure with AlienVault >

Learn more about intrusion detection in the AWS with AlienVault >

Best Practice: Dismantle Siloes Between Your Cloud and On-Premises Security Monitoring

For today's resource-constrained IT teams, the explosion of public cloud services has only increased the complexity of securing critical infrastructure. IT Professionals who are tasked with deploying and managing security monitoring and threat detection tools across dynamic cloud and on-premises may take one of many approaches. You might try to extend the use of your legacy on-premises security tools to monitor your public cloud environments. However, as discussed in section two, many legacy security tools are not readily adaptable or optimized for cloud environments. Another approach—to maintain separate, siloed security monitoring solutions for public cloud and on-premises environments—is fraught with extra costs, complexity, and potential security blind spots.

A better approach is to centralize your public cloud, private cloud, virtual and physical on-premises security monitoring on a single cloud-based, SaaS-delivered security solution. Cloud-native security monitoring tools like AlienVault USM Anywhere take full advantage of cloud architectures, services, and APIs in ways that legacy solutions were not built for, while also providing complete threat detection for on-premises physical and virtual infrastructure.

By centralizing your security monitoring, you can effectively reduce the cost, time, effort, and complexity of managing your security posture across your multiple IT environments. In addition, this can help you to eliminate your security blind spots and ensure continuous monitoring as you migrate services from the data center to the cloud.

Best Practice: Take a Unified Approach to Hybrid Cloud Security

Finally, consider a **unified approach to your hybrid cloud security**. Security operations centers traditionally worked to weave together multiple point security solutions for asset management, vulnerability scanning, intrusion detection, SIEM and event correlation, behavioral monitoring, and log management in their on-premises networks. This typically required an extensive amount of integration, fine-tuning, and management to create a single source of threat detection and incident response. Recreating this process in hybrid cloud environments is often too cumbersome and error-prone for most IT teams.

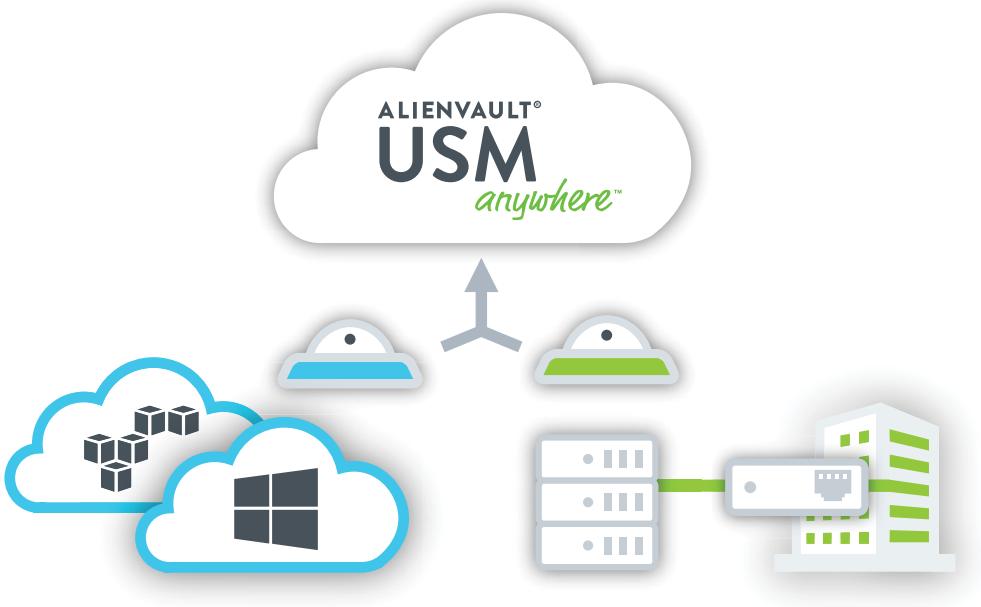

AlienVault disrupted this piecemeal approach with the introduction of Unified Security Management™ (USM™), first for on-premises networks via USM Appliance, and today for hybrid cloud and on-premises environments via USM Anywhere. USM brings together multiple essential security capabilities onto a unified platform, so it can be launched quickly, cost effectively, and without complex integration requirements. With a built-in library of correlation rules that are continuously updated by the AlienVault Labs Security Research Team, USM starts detecting threats within minutes of installation and continues to detect emerging threats as they appear "in the wild."

In short, a unified approach to hybrid cloud security can significantly reduce the amount of resources needed (time, budget, staffing) to monitor your security posture across your cloud and on-premises critical infrastructure.

Check out **USM Anywhere in an online demo** or use it for **free for 14 days**.

*Gartner, Best Practices for Securing Workloads in Amazon Web Services, Neil MacDonald and Greg Young, 15 April 2015, Foundational

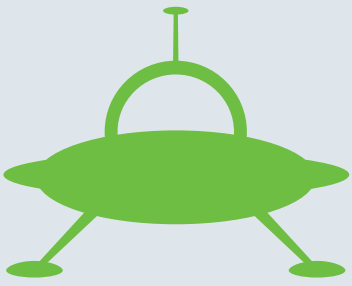

About AlienVault

AlienVault has simplified the way organizations detect and respond to today's ever evolving threat landscape. Our unique and **award-winning approach**, trusted by **thousands of customers**, combines the essential security controls of our all-in-one platform, AlienVault **Unified Security Management**, with the power of AlienVault's **Open Threat Exchange**, the world's largest crowd-sourced threat intelligence community, making effective and affordable threat detection attainable for resource-constrained IT teams. AlienVault is a privately held company headquartered in Silicon Valley and backed by Trident Capital, Kleiner Perkins Caufield & Byers, Institutional Venture Partners, GGV Capital, Intel Capital, Jackson Square Ventures, Adara Venture Partners, Top Tier Capital and Correlation Ventures.