3 Critical Capabilities for Cloud Security and Compliance

A Practical Guide to Getting Started Securing Cloud Workloads

ATOMICORP

A move to the cloud isn't simply shifting and outsourcing where your enterprise applications run.

Cloud architectures are highly complex and dynamic. As the state of each component changes, it impacts other components. The environment changes as services spin up and down. There is no "normal" baseline in the cloud; everything blurs together in completely software-defined layers. This obviously presents new security and compliance challenges that are not easily met with legacy on-premise technologies.

But the opportunity to invest in cloud security technologies can seem endless. There's always a new threat, a new technology, a new need – but what do you really need to spend on to ensure regulatory compliance and workload security. In Gartner's Market Guide to Cloud Workload Protection Platforms, they list 8 separate categories of protection technologies to ensure cloud security.

CWPP Controls Hierarchy

Gartner, Market Guide to Cloud Workload Protection Platforms, April 2019

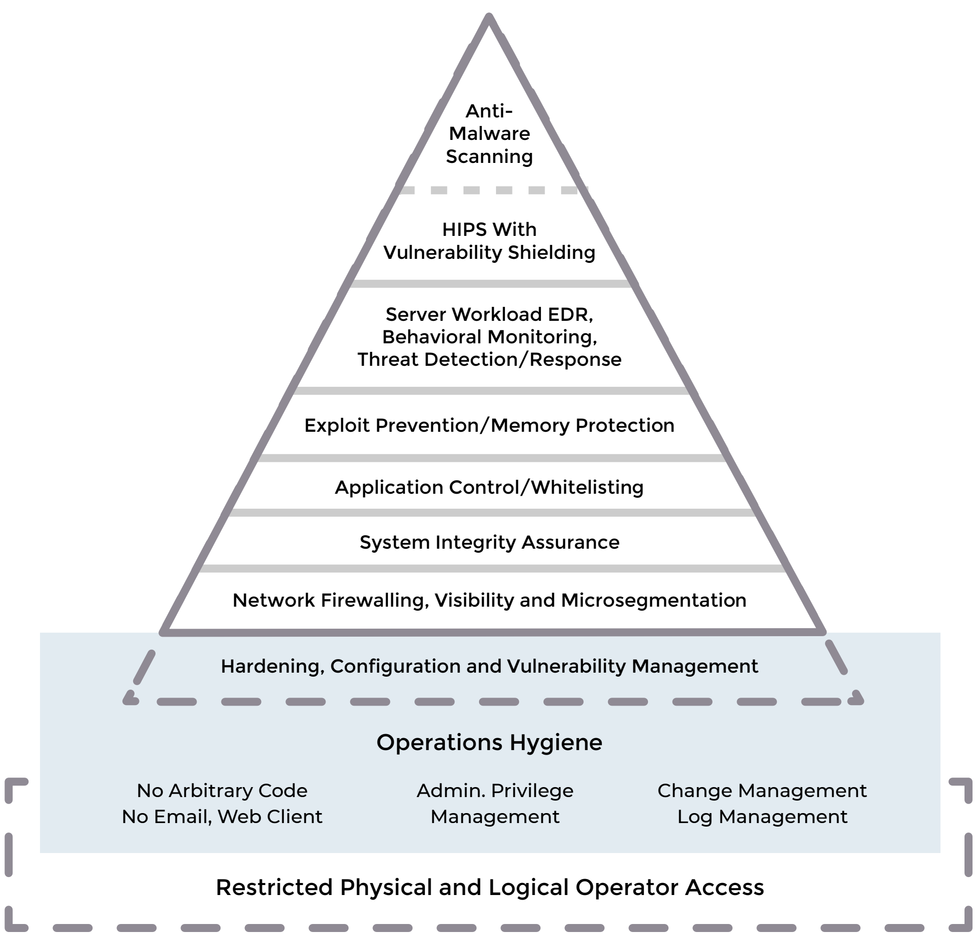

While all steps are valid protections, prioritization is necessary for most organizations. As with any security investment, your objective should be to maximize return on investment. That is, start first with measures that protect against the widest range of attacks across the broadest possible surface area.

In this paper we will focus on the three measures that give you the biggest bang for your cloud security and compliance dollars.

- Real-time Monitoring
- Vulnerability Shielding with an Agent-based Web Application Firewall
- Redcing Attack Surfaces by Hardening Cloud and Workload Configurations

Once these pillars are in place, you can gain a true understanding of any remaining gaps and can budget appropriately to fill them. At that point, expect to look at:

- Exploit Protection and Memory Protection
- Application Control & Whitelisting
- Server-Based Endpoint Detection and Response (EDR)
- Antimalware Scanning (though this should always be done on cloud storage like S3 buckets)

Before You Get Started...

Understanding Shared Responsibility

Shared responsibility means that the cloud provider protects the infrastructure of its cloud. But what you put in the cloud - your sensitive data and critical assets - are still on your plate.

When moving to the cloud, engineers recognize that they've been freed from the burden of protecting the infrastructure. They think, "I used to do these ten things. Now my cloud provider does six of them. I can continue to do the other four just like I've always done them."

But this isn't an accurate way to think about the cloud, because the cloud is not just a version of the on-premise environment that is located elsewhere and accessed through a firewall. It's completely different and must be treated differently if security is to be maintained.

Traditional security models are built around the concept of compartmentalizing. This layer is protected differently from that layer. This vendor protects this and that vendor protects something else. But you can't create cloud computing around compartmentalization. The whole point of cloud computing is that the hardware is abstracted away from the business environment. The heart of the business is its data and apps, and those remain solidly the responsibility of the cloud customer.

Shared Responsibility in the Cloud

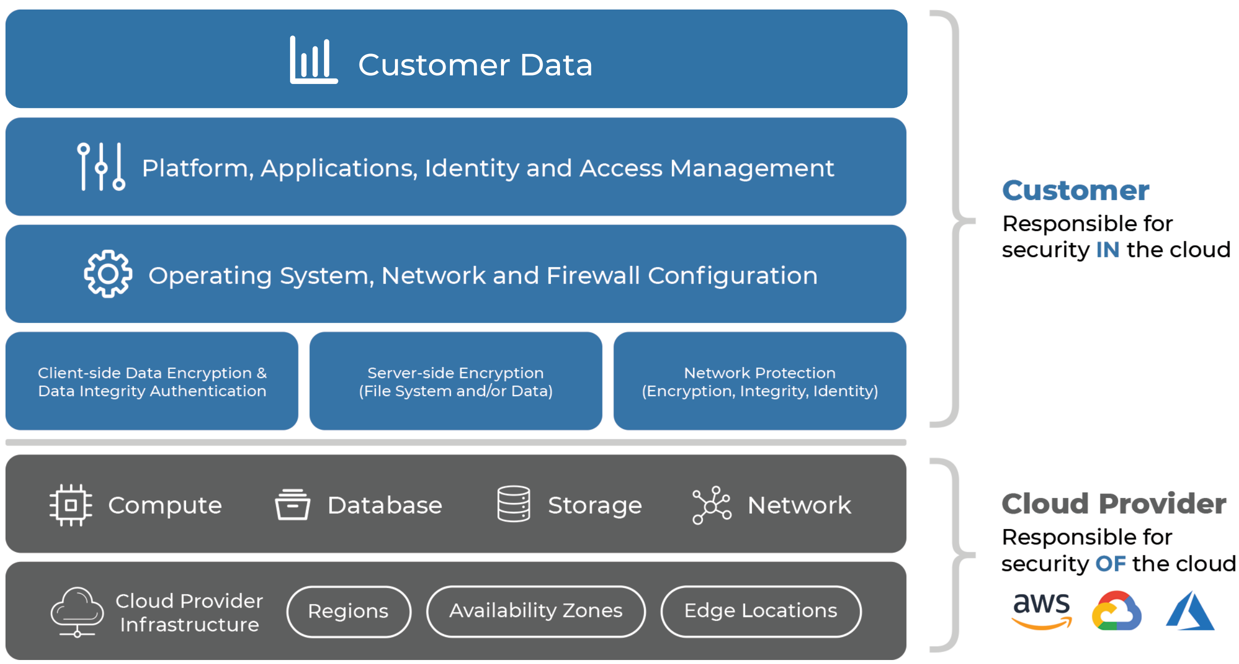

System and File Integrity Monitoring in Real-Time

A FIM monitors both files containing sensitive data and files native to the operating system, as well as files controlling the configuration of network devices, web servers, and firewalls. Today's FIMs also serve as forensic tools that can be used for security operations. A FIM is also a key compliance requirement for standards like PCI-DSS. If your applications are subject to this regulation, you must deploy at FIM anywhere where PCI data is stored.

A common tactic employed by malicious actors is the fileless attack, in which changes are made to the system without making changes to the files, including the operating system or critical applications. The use of this type of attack is exploding: in just the first half of last year, fileless attacks increased by [94 percent](#).

Because a FIM only detects changes to files, a system also needs a SIM to protect against fileless attacks.

But this constant monitoring is more than a security tool. It also aids in compliance, enabling organizations to conduct "continuous compliance" on all systems touching regulated data, instead of relying on periodic audits.

Choosing a FIM

An enterprise-grade FIM should be able to:

- Indicate not just that a file has changed, but precisely what has changed in that file
- Keep a copy of those changes in order to support forensics and configuration management capabilities
- Keep copies of as many changes as desired to enable roll backs and support deep levels of analysis
- Include logging, malware detection, and active response features
- Be capable of functioning alongside orchestration tools to roll back configuration changes and act as an automatic backup system.

Choosing a SIM

A cloud aware SIM should be able to:

- Use as little memory and CPU as possible, legacy SIM products can significantly raise cloud computing costs and should be avoided in cloud environments
- Block unauthorized changes to the system in real time
- Detect malicious activity by unauthorized and authorized users

Vulnerability Shielding with an Agent-Based Web Application Firewall

Today's organizations rely on hundreds or even thousands of applications and APIs to conduct their daily business. Each of those applications contains vulnerabilities that present an opportunity for hackers – and a nightmare for security teams who struggle to download, patch, and test each application before it can be exploited.

To fix each application would be impossible; there are simply too many, and no organization can be sure its apps' developers are knowledgeable about security. A solution is needed that protects applications without the need to modify any source code.

That's what a web application firewall (WAF) does. It sits in front of an application, filtering, monitoring, and blocking traffic that does not comply with the application's security rules.

The types of attacks a WAF can prevent include SQL injections, Distributed Denial-of-Service (DDOS), cross-site scripting, malware, and many others.

How is a WAF different than a traditional firewall?

Traditional firewalls are a network security technology that acts as a gate between servers, allowing or disallowing connections. While it shares the word "firewall," a WAF provides a more nuanced function, monitoring and filtering web traffic to protect against a variety of attacks. While traditional firewalls can only accept or reject traffic, WAFs apply sophisticated policy rules based on user roles and specific applications. And unlike traditional firewalls that are costly and time-consuming to keep up to date on the latest malware and AV protections, WAFs are updated automatically whenever new threats are discovered. A WAF can be deployed on physical or virtual machines, discern between legitimate and fraudulent traffic, and determine which applications are being targeted.

Why a WAF should be located on the workload

In the on-premise environments of the past, workloads were part of a private network. They could be protected by protecting the network perimeter with a firewall.

While it's tempting to try to port this architecture to the cloud by putting a "firewall" in front of your cloud infrastructure, to do so is both inefficient and ineffective. Cloud environments lack traditional network boundaries, leverage APIs and shared compute resources, and scale elastically. For these reasons, a WAF should be deployed as part of the workload — not on the infrastructure.

Organizations using a public cloud provider like AWS, Google, or Azure may expect their provider to secure their applications and data. However, those providers base their service level agreements on the shared responsibility model that requires them to only secure the cloud infrastructure. Their customers are responsible for securing their own workloads. A WAF is a necessity for these customers, especially those covered by regulations such as PCI. Considering the likelihood of succumbing to one of the many types of attacks that a WAF can discourage, any organization that doesn't have a WAF and is breached will be in a bad position when the assessors, lawyers, and insurance reps visit.

WAFs Protect Against These Types of Attacks

- Basic web attacks
- Brute force attacks
- Advanced attacks (SSRF, XXE)
- Geoblocking
- Content scraping protection
- Proxy abuse

Agent-Based WAF architecture for cloud environments

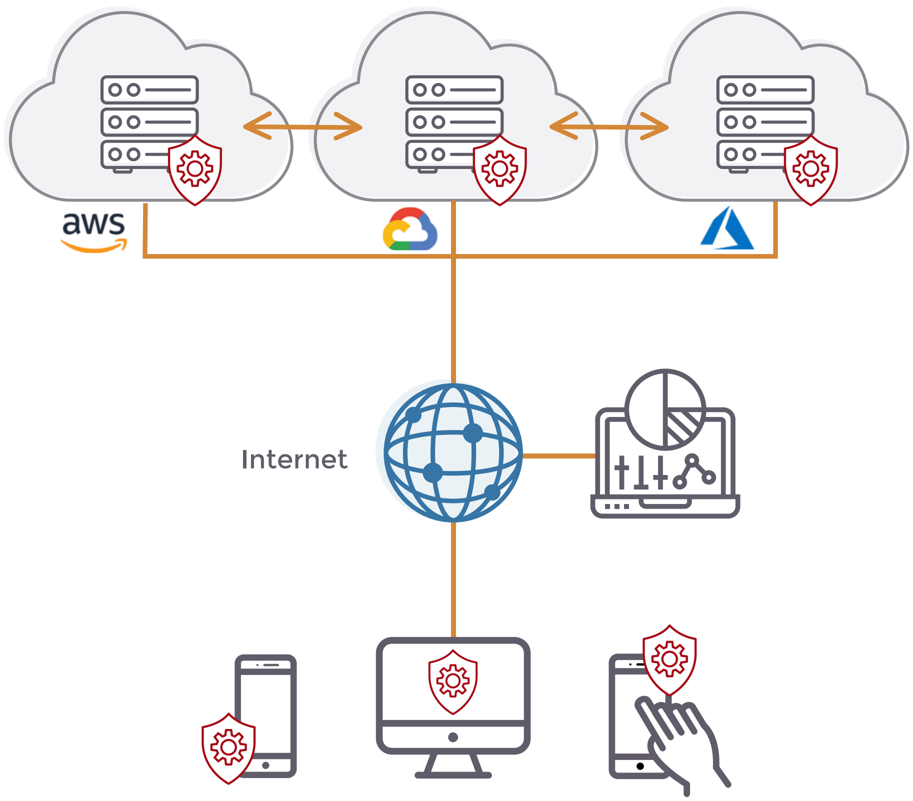

WAF Rules Determine WAF Effectiveness

A WAF by itself will not protect your applications or APIs. The comprehensiveness and currency of your WAF ruleset will determine its ability to defend your workloads. This means you must be extremely thoughtful in choosing which WAF technology to deploy. For maximum protection, choose a WAF ruleset that protects both applications and APIs, contains rules to protect a wide variety of advanced attacks, and is updated very frequently to address new attacks. Atomicorp offers a comprehensive comparison of common WAF rulesets on its website.

Correct Configuration is Critical in the Cloud

The more complex a system, the more points of vulnerability it presents to a potential attacker. Keeping these attack surfaces as small as possible is a basic security principle, but one that has become increasingly impractical in today's cloud-based business environments. Instead of minimizing the size of an attack surface, organizations must strengthen them by eliminating misconfigurations that could expose private data to the public cloud. That process is called hardening.

Where does your attack surface begin the cloud?

Configuration management is another of those tasks that organizations may assume is the responsibility of their public cloud provider. However, cloud providers only offer tools to help customers secure their workloads – the actual work of securing those workloads remains on the customers' plates. That can lead to a lot of confusion on the customers' side, and confusion leads to leaks. In fact, Gartner notes that 99% of cloud security failures will be the customer's fault. Many of these failures are the direct result of misconfigurations.

In one infamous example, Verizon lost the personal data of 6 million of its customers when a contractor misconfigured the security settings on an S3 storage bucket. Misconfigured S3 buckets are a common problem, and one that Amazon has addressed with new options that can be applied retroactively to secure existing buckets, but other misconfigurations around things like default passwords, unpatched software, unencrypted network traffic, lax access privileges, and misconfigured BIOS, firewalls, ports, etc., can leave sensitive data exposed as well.

Controlling the workload in the Cloud

A comprehensive Cloud Workload Protection Platform (CWPP) makes breaches based on misconfigurations harder for hackers to execute and easier for security teams to detect. An enterprise-grade CWPP should:

- Look for stored credentials in order to monitor and block the use of cloud provider credentials to access resources like S3 storage
- Look for suspicious sources of access and data egresses like TOR nodes and anonymizing VPNs
- Prevent exfiltration of sensitive data like AWS stored credentials and tokens, credit card numbers, and social security numbers

> **Cloud Workload Protection Platforms**
>
> CWPPs replace endpoint protection by focusing on protecting workloads from attack, no matter where they're located or how large or small they might be. A CWPP provides consistent visibility and control of all server workloads in all locations, including virtual machines and container-based application architectures.

- Detect and block system access from unusual sources and anomalous behavior between systems
- Provide defense-in-depth, allowing a check on cloud infrastructure configuration and an additional check on cloud provider security facilities.

And Then Fill in Your Gaps...

Once you've implemented the foundational strategies for cloud security and compliance, it's time to examine the gaps and implement additional layers of protection.

1. Exploit Prevention & Memory Protection

Exploits use weaknesses in legitimate software like Adobe Flash to gain access to systems. This is a common problem: up to 90 percent of breaches leverage an exploit at some point in the attack chain. Exploits are easy for unsophisticated attackers to use because they can buy exploit kits on the black market that do the hard work for them. Despite the prevalence of this hacking technique, most leading security products do not include exploit prevention as a function. And in the case of those that do, coverage can be weak.

The Importance of a Single Agent

Your security needs to be integrated and holistic. Cloud-based security should use a single agent that has visibility of all the interdependencies in the fluid environment.

This approach saves resources and also eases the burden on the security ops team, who can perform all administration and deployment operations through the single agent.

Virtual patching is a way to implement a security policy to mitigate or eliminate a vulnerability taking without actually taking the time to download, install, and test a patch. Virtual patches make a fast response possible without the need to make changes to the software itself. Also, be sure to implement memory protection to prevent the exploitation of buffer overruns and code reuse.

2. Application Control ("Whitelisting")

Application control, also known as whitelisting, enables organizations to control which executables are run on their servers. This means they can default to a zero trust security model, in which no one inside or outside the network is trusted until they are verified — so fileless malware is always blocked. Application control is a foundational requirement that should be built into the CWPP.

3. Server-Based Endpoint Detection and Response (EDR), Behavioral Monitoring, Threat Detection and Response

Server-based EDR is an advanced form of EDR that is more powerful than traditional system integrity monitoring or host-based intrusion detection systems. A server-based approach to EDR looks at behaviors, such as the communications that are crossing the network, the processes that are launched , and the files that are opened, to expose malicious activity, even in containers. The most robust EDR capabilities available today are based on OSSEC, the world's most popular open source intrusion detection system.

4. Antimalware Scanning

For most applications, viruses and malware do not present a major threat and signature based technologies do not offer much protection. However, on storage services such as Amazon S3 buckets, they are a serious risk. Anti-malware products should always be run against storage services. Over time, S3 or other storage buckets become contaminated with malware and other threats. These threats don't stay contained in the bucket - they spread to other users, apps, and databases, raising risk levels across the entire organization and eating up security resources. To prevent S3 storage buckets from becoming threat vectors, a cloud security solutions needs to automatically scan the buckets so any threats they contain can be remediated before they can propagate.

For other applications, consider anti-malware as part of overall good security hygiene and to meet specific regulatory burdens.

Factors to Look For in All Cloud Security Solutions

Workloads in modern data centers are hybridized across on-premises and multiple public cloud environments. Security has to be as distributed as the workload. A cloud security solution has to deliver consistent visibility and control of all workloads, regardless of size or location.

An enterprise-grade cloud security solution should include these capabilities:

1. Small Footprint, Low Resource Utilization

An upside of the cloud computing model is that you only pay for what you consume. But if your footprint is large, that model becomes a downside. Organizations coming from on-premises computing environments are likely to have some unnecessary or bloated legacy software in their technology stacks that should be pruned or removed before a move to the cloud. Otherwise, consumption costs will rise, even doubling in some cases.

Legacy technologies present other issues in the cloud as well. A major benefit of the cloud is the ability it provides for organizations to scale up and down elastically. Legacy technologies don't do this well. Attempting to carry them over to the cloud will hobble agility and negate the very reason most businesses chose the cloud in the first place.

2. Elastic Scaling

Traditional scalability adapts to increases in workloads by provisioning resources incrementally. Elastic scaling adapts to both increases and decreases in workloads by provisioning and de-provisioning resources automatically. This not only enables the appropriate allocation of resources and budget, but also strengthens security by reducing the attack surface when workload demands decrease.

3. Hybrid and Cross Cloud Compatibility

Many organizations are using many types of software in more than one public cloud and often in combination with hybrid and private clouds and on-premise environments. Every one of these environments has to be secured, and the security as a whole must work in concert with no gaps or conflicts.

Security solutions for hybrid and cross-cloud environments should include automated detection, prevention, and analytics for all environments, automated workload protections for all environments, and system security for legacy systems such as Solaris, HP-UX, and IBM AIX.

4. Devops Integration

Most organizations moving to the public cloud have also adopted a DevOps process with hundreds or even thousands of releases a week. Yet few DevOps professionals possess an the necessary level of security expertise to ensure these releases are safe. Organizations are left reliant on pre-launch vulnerability scans that only measure the security of a release at a single point in time, not on an ongoing basis. DevOps needs to be integrate with security processes to ensure that security is treated as a continual process, not as a momentary event.

To effectively secure releases, security tools must be deployed that easily use DevOps orchestration tools, such as Ansible, Puppet, and Chef. Open-source OSSEC integrates with these tools and with custom scripts, so it can be deployed with every release.

5. Secured Containers

A lot of organizations try to secure their containers in the same way they secure their virtual machines (VMs) – in a reactive manner from the outside in. But containers require a different approach. Security needs to focus on visibility into the app through detection of scanning attempts, anomalous probes of a service across more than one machine, whitelisting, automatic WAF-to-app configuration, and blocks and alerts for suspicious actions.

About Atomicorp

Atomicorp leverages the power of the world's most widely-used open source security solutions, including OSSEC, ModSecurity, OpenSCAP, and dozens of others to protect critical workloads, no matter where they reside.

Atomicorp's Unified Workload Security Platform delivers detection, protection, and compliance through a single lightweight agent. The platform protects workloads across physical and virtual machines, containers, and multiple public clouds with a single management interface. It enables organizations to scale cost effectively in dynamic public cloud environments, where costs are driven by resource utilization. And it addresses key security and compliance requirements without the need for a battery of point solutions.

Unified Workload Security Platform Capabilities

- Advanced file integrity monitoring (FIM)
- Compliance reporting
- System hardening
- Vulnerability shielding
- WAF protection
- Log-based intrusion detection (LIDS)
- Behavioral monitoring
- Deception
- Memory protection
- Application control and visibility
- Antivirus
- Workload-based and cloud-based machine learning

Contact Atomicorp to see how we can help you protect what matters today.

ATOMICORP

Headquarters:

10549 Conference Center Drive, Suite 180, Chantilly VA 20151

(703) 299-6667
www.atomicorp.com
info@atomicorp.com