

9 CONSIDERATIONS FOR SECURING WEB APPS IN THE CLOUD



**WEB APPS ARE THE NO. 1 SOURCE OF DATA BREACHES;
PROPERLY SECURING THEM IS A PRIORITY.**

Effective security for cloud-hosted web applications requires full visibility into the environment in which the apps live and the potential exposure to vulnerabilities — and to do so consistently, while proactively monitoring for attacks without causing delays in application development and delivery.

Cloud adoption means that a focus on perimeter security is not sufficient and may even be obstructive, because it can impact application performance and availability.

Similarly, relying on your cloud service provider's security services is insufficient; any provider will tell you that securing the cloud environment is a shared responsibility between cloud service provider and customer, and the responsibility for specifically securing web applications is the sole responsibility of the customer. And although the industry as a whole has become better at protecting lower-level network and server resources, as attackers look for targets, they are moving up the application stack.

Indeed, web application attacks are the No. 1 source of data breaches according to Verizon's [2017 Data Breach Investigations Report](#), up by over 300% since 2014. They accounted for 30% of all breaches, far outpacing any other form of attack.



Given the rapid rate at which companies now roll out new applications and updates to existing ones, they need to think about security more comprehensively — starting in the application development process — and consider how changes after deployment may affect an application's security profile.

Securing cloud-based web applications is a challenge that requires a rethinking of security best practices, adoption of automated security tools, and a renewed focus on security expertise.

To help guide the way, following are nine key considerations for providing sound web application security in the cloud:

1

UNDERSTAND THE LIMITS OF CLOUD PROVIDER SECURITY SERVICES

Most cloud service providers will offer some level of security services. In general, they are responsible for securing the infrastructure services upon which their offering is built, including the host operating system and hypervisor.

The provider should likewise secure its foundational services, including compute, storage, network, and database infrastructure. That includes perimeter security services and protecting against spoofing, scanning, or any attack that could take down an entire platform. And of course, the provider should offer physical security for the data centers it operates.

Customers, however, are responsible for securing the content, data, and applications that the cloud platform hosts. The exact parameters may vary depending on the type of cloud service in question — for example, infrastructure-as-a-service versus platform-as-a-service. But at a base level, customer



responsibilities may include network security services such as threat detection and monitoring. At the server level, the customer may be responsible for access management, patch management, configuration hardening, and log analysis.

At the application layer, responsibility for security likely falls squarely on you, the customer. That means you must provide sound identity and access management (IAM), and follow best practices for secure coding and configuration management. You also need to monitor for attacks that specifically target the application layer, such as SQL injections, cross-site scripting, or those that target known weaknesses in various applications.

You should ask your cloud provider to clearly detail which security services are provided and which ones are your responsibility. Don't assume the provider is doing all the heavy lifting when it comes to security.



2

CONSIDER APPLICATION VALUE AND RISK

To provide proper security, customers must understand the business value of each of their applications. Companies generally can't afford to provide top-tier security for every single application — and they don't need to — because some are simply more important than others.

The important point is to identify those applications that are most critical to the organization. That will likely include apps that account for revenue generation, as well as those that support intellectual property or other sensitive information. Talk with application business owners to get a handle on which ones have the most value.

In addition, consider how applications are used and accessed, with the goal being to identify those that are at greater risk of a breach. Apps that can be accessed by any customer, for example, present a greater risk than those that are used only by a small group of employees. Similarly, complex web applications that include lots of third-party components also present a higher degree of risk because any one of those components may include unknown vulnerabilities.

With this information in hand, you can go the extra mile toward applying stringent security to high-value, high-risk applications and more lenient measures to others.

BE COGNIZANT OF CONFIGURATION VULNERABILITIES

3

Virtual machines in cloud environments come standard with what is essentially a firewall. But it's up to the customer to configure that firewall correctly, by deciding which ports should accept traffic and from which sources.

But if you open every port and allow traffic from any source, you can't blame the provider if an attacker takes advantage.

Another key component is properly protecting who has access to your cloud-based applications and infrastructure, and ensuring individuals have only the level of permission they need.

At the base level in Amazon Web Services (AWS), for example, is the root key, which allows full access to all resources in the account. The root key should be locked down so that it is not accessible by anyone, and AWS recommends deleting root access keys and creating IAM user credentials instead.

Then you have application-level APIs and access keys to your cloud services. These come with varying levels of privilege, which must be carefully considered as you decide who can use them. Again, the principle of least privilege applies: give people access only to what they absolutely need to do their jobs.

It's also important to watch how authorized credentials are used to look for patterns that indicate anomalous behavior. If there's a configuration change in an area of the system where such changes are not expected, that should raise a red flag. For example, an application that must comply with the Payment Card Industry (PCI) Data Security Standard (DSS) standard requires that a change control process be followed. Any change to those systems should be documented with a change ticket attached showing what the change was, why it was made, and who made it.

4

GET A HANDLE ON SOFTWARE VULNERABILITIES

While configuration vulnerabilities are essentially created when someone makes a mistake, software vulnerabilities arise during the application development process. Combatting them means building security controls into that process.

You must be cognizant of known common vulnerabilities and exposures (CVEs) and common weakness enumeration (CWE), which is a list of software weakness types. CVEs and CWEs can apply to major operating systems, OS configurations, server applications, and standard web applications. That gets difficult when an application uses third-party components, which — as mentioned on the previous page — can include vulnerabilities of which the developer isn't aware.

Combatting these threats requires a vulnerability assessment to be built into the application development process. And given the increasing pace at which companies are developing new applications and updating existing ones, the assessment has to be conducted not just once in a while, but continuously.

You must be cognizant of known common vulnerabilities and exposures and common weakness enumeration.



5

ASSESS THE APPLICATION'S THREAT PROFILE

Understanding the application's threat profile can also help you determine how best to secure it because it shows the threats to which an application is most likely to be vulnerable. For example, if an app relies on a SQL database, then you need to protect against SQL injection attacks; but if it doesn't, that's one less risk to be concerned about.

Here again, vulnerability assessments can help you understand the threats to which your applications are most susceptible. Some cloud providers offer vulnerability assessment tools that may help, but they are designed to scan for a limited set of vulnerabilities.

Vulnerability scanners typically work in one of two ways:

- They either check from outside and look at what vulnerabilities an application is exposing, or
- They inspect inside at the package level to determine what components make up the application and whether they contain any known vulnerabilities.

Most cloud providers use the latter approach, relying on an open-source list of CVEs and other rule sets to conduct the assessment. Such scanning tools help you find any known vulnerabilities, such as in any third-party components or applications that use insecure network protocols to move data among application components.

But these sorts of vulnerability scanners may be limited in effectiveness because they must be run manually, which means they do not run continuously. Unless users set up automation to enable scanning at defined intervals, users manually set them up to look at specific applications at specific times. So, while the tool

may catch a vulnerability when an application is first launched, you'd have to run it each time the application is updated.

What's more, many of these tools, do not provide help remediating vulnerabilities that they find. And they report vulnerabilities on the instance that was scanned rather than the image that created it.

A better approach is to use a vulnerability scanner that performs scans both from the outside looking in and from the inside at the package level. The scanner should also support automated features that enable it to run continuously.

Understanding the application's threat profile can also help you determine how best to secure it.



6

ENSURE YOU MEET COMPLIANCE REQUIREMENTS

Organizations that are subject to industry or government regulations such as PCI-DSS, HIPAA, and Sarbanes-Oxley should thoroughly vet how their applications will work in a cloud environment. Some traditional approaches to compliance may not work in the cloud, but you also may have opportunities to improve your compliance posture.

Here again, expect to share responsibility for compliance with your provider. AWS, for example, offers services that are PCI-DSS-compliant. Although it certainly helps, the certification extends to the AWS infrastructure and not customer-deployed workloads. AWS customers can leverage AWS' compliance, but must manage their PCI-DSS obligations and certification for the environment deployed on AWS.

Use the same sort of approach with any cloud provider, for any compliance requirement. Ask for a matrix that shows what the provider is responsible for, what you the customer must handle, and any shared responsibilities.

7

MONITOR MULTIPLE SOURCES FOR SECURITY DATA

When it comes to providing application security there's no single source of data that will allow you to detect all potential attacks if you watch it closely enough. Rather, you need to simultaneously monitor multiple sources.

At the network layer, watch out for attempted remote attacks, such as an intruder using tools designed to exploit a vulnerability in application frameworks, a SQL injection, or the like. That means using firewalls and intrusion detection devices that have visibility to all traffic within the network.

It's also important to monitor server log data to detect any changes that perhaps shouldn't have occurred and patterns that suggest suspicious behavior. And you need to monitor HTTP traffic at the application layer, again to determine whether someone is trying to access applications in a way that suggests nefarious intent.

In security circles, this kind of approach is known as "defense in depth." The idea is to provide multiple layers of defense such that, even if an intruder successfully gets past one defense, another will be there to stop him.

The key, however, is not so much gathering data as making sense of it. Cloud providers often offer access to all sorts of log and other data. AWS, for example, offers AWS CloudTrail, which allows you to log, continuously monitor, and retain events related to API calls across the AWS infrastructure. There is certainly value in such offerings, so long as you have the tools and, most importantly, expertise to help you make sense of the data.

THINK THROUGH YOUR THREAT DETECTION TEAM AND PROCESSES

8

Security tools and applications alone will not keep you safe from threats, because tools alone can't discern whether a given pattern merely looks like a threat or indeed is a threat. Making those kinds of decisions requires human beings with specialized expertise.

Security tools generate enormous amounts of information, maybe tens of thousands of alerts per day for a large enterprise. Most of those alerts are benign, but the key to defending against intruders is being able to identify the handful that really are out of the

ordinary and potentially represent a real attack. Security tools can help bubble up those threats, but only an experienced professional can determine for sure whether any given alert (or group of alerts) warrants investigation.

Similarly, it takes experienced professionals to properly tune the various tools that protect your networks, servers, and applications, including cloud-based web applications. Whether it's ensuring proper configuration for devices, servers, and applications, examining log data for evidence of intrusions, or ensuring compliance, you need people who understand how to provide security.

Assess the level of your staff's expertise and determine whether it includes security experts in areas including network, applications, and cloud. Will they be able to properly configure, tune, research new threats, and monitor the various components of your security infrastructure on a 24x7 basis?

If there are any holes in your security know-how, you've essentially got three options:

- Hire new staffers with the knowledge that's currently lacking
- Train existing personnel in required disciplines
- Hire a security-as-a-service provider to do the job for you

Here the decision is much the same as the one that led you to cloud-based services in the first place. You can go to the expense of hiring or training someone to provide security, but it won't likely differentiate your business. Hiring a security-as-a-service provider may well prove to be both less costly and more effective.

9

CONSIDER AUTOMATION REQUIREMENTS

In addition to the sheer volume of data that must be examined to ensure security, cloud-based applications bring changes to certain processes that raise security

issues. In addition to experienced security professionals, dealing with them requires a healthy dose of automation.

Consider the way application development has changed. Previously, before an application was released into production, a change control board would meet to examine the application and ensure it will work as intended, without any ill effects on any other applications. Security was always a primary concern and applications had to go through a thorough risk analysis process.

Today, companies are investing heavily in rapid application development processes and DevOps. As a result, application releases happen in minutes, perhaps several times per day. That doesn't leave time for the sort of work those change control boards used to do. In fact, a [recent study](#) found that more than 70% of application developers admitted business pressures to quickly release application updates often drove them to override security concerns.

To avoid such actions, security must be automated and baked in to DevOps processes.

Successfully automating the application development process, including the security risk analysis, presents a significant opportunity for companies. If you can successfully automate the delivery of security in applications, you're removing what was historically a gating factor with respect to rapid application deployment.

The cloud environment presents just such an opportunity through development and deployment tools such as [Jenkins](#), [Chef](#), and [Puppet](#). These tools automate much of the application development process, including risk assessments. And many of them are configured to work seamlessly with popular cloud services such as AWS and other popular cloud services.

Another characteristic of the cloud environment is the idea of immutable infrastructure. This means that infrastructure is regenerating on a regular basis and, therefore, it can't be compromised for very long. Companies can build security into the app development process, and in so doing, they can remove the possibility that an application can be changed either inadvertently or willfully by anyone without proper authorization — thus removing one of the biggest sources of security risks.



MEETING YOUR RESPONSIBILITY

For all sorts of organizations, cloud-based web applications are becoming a crucial component in their IT strategy. They help enable an agile, scalable, cost-effective IT environment that can meet the demands of today's always-on businesses.

But they also present challenges in terms of providing security, and it's important to recognize that cloud providers do not shoulder all the security burden; customers must share it.

Part of your responsibility is protecting against application-layer attacks, which are becoming increasingly more common, while remaining vigilant against lower-level threats. The responsibility also extends to following secure application development practices.

Highly automated tools and processes are required to meet these responsibilities, but even they will take you only so far in identifying actual threats to your environment. It requires humans with deep security expertise to determine which of the threats that your tools identify need to be acted on.

Be honest about whether you have the experience in-house to take on these security tasks, or if you're better off turning to a security-as-a-service provider and your cloud provider for help.

Alert Logic Security-as-a-Service solutions integrate cloud-based software, analytics and 24x7 expert services to assess, detect, and block workload threats and help you comply with mandates like PCI, HIPAA and SOX COBIT. With more than 4,000 customers worldwide, we focus on threats most relevant to AWS hosted applications by defending your full web application and infrastructure stack.

At BSI Cybersecurity and Information Resilience, we provide expertise to our clients on the identification, protection, compliance and management of their information assets.

We partner with Alert Logic to provide full stack security-as-a-service to our clients looking to secure their application infrastructure in a cloud, hybrid or on-premises environment. If you are struggling with the security of your web applications, contact us.

To learn more, visit: www.alertlogic.com/aws