



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Implementing and Auditing CIS Controls (Security 566)"
at <http://www.giac.org/registration/gccc>

Implementing the Critical Security Controls in the Cloud

GIAC GCCC Gold Certification

Author: Jon Mark Allen, jm@allensonthe.net

Advisor: Barbara L. Filkins

Accepted: 31 January 2016

Abstract

This paper will explore the advantages and challenges of implementing the latest version of the Critical Security Controls (v6) in a cloud environment. Each control will be evaluated for differences between a cloud versus. on-premises implementation. Where possible, cloud-native products, tools, or services will be added to the existing list of suggested tools included in the SANS course. The example cloud environment will focus on an Amazon Web Services (AWS) installation.

1. Introduction

Amazon refers to cloud computing as “the on-demand delivery of IT resources and applications via the Internet with pay-as-you-go pricing” (Amazon Web Services, 2015). The ability to add or remove service instances based on current demand is certainly attractive and eliminates the need to maintain data center capacity even when utilization is low. A quick glance at Amazon’s AWS Case Studies page¹ testifies to the popularity of cloud computing.

Two factors quickly come to mind which have the potential to cause problems for security: dynamic resources and loss of control. At first, it may seem a daunting task to maintain a valid inventory at all times of which resources need defending since computing resources can come and go rapidly. Add to that the “shared security responsibility model” Amazon touts (Amazon Web Services, 2015, p. 3) and security practitioners may start sweating their ability to secure their new world.

The Critical Security Controls are “a prioritized, highly focused set of actions that have a community support network to make them implementable, usable, scalable, and compliant with all industry or government security requirements” (Center for Internet Security, 2015, p. 6). These controls assume that an organization has management control over the infrastructure of the environment – a condition that is no longer true once cloud resources enter the picture.

But those resources can still be protected even with a move to the cloud. The gaps created by this new model can be accounted for by implementing the proper security controls. The security architect must understand how cloud networks are abstracted from traditional hardware and, therefore, how those networks differ in the way they work versus an on-premise data center (Mogull, 2015, p. 4).

So how *does* a move to the cloud change the way the Critical Security Controls are implemented? Which controls are easier to implement and how? Which controls are more difficult and why?

¹ <https://aws.amazon.com/solutions/case-studies/all/>

This paper groups the Critical Security Controls (v6) into three different categories: cloud positive, cloud negative, and cloud neutral. The cloud positive category includes controls that are simpler, easier, or better (overall) after a move to the cloud. Cloud negative controls are more difficult, complicated, or require a re-evaluation of methods in a cloud environment. Lastly, cloud neutral controls are more or less unchanged in the difficulty of implementation in the cloud versus an on-premise network.

In the interest of time, Amazon Web Services (AWS) will serve as the baseline implementation for this paper, though similar issues and solutions will present themselves on cloud platforms such as Microsoft Azure. The controls discussed are generally operating system (OS) neutral, but details specific to a given OS are called out where needed. The goal here is to evaluate the controls in light of the changes to the network (both good and bad) which a cloud solution brings to the table.

2. Cloud Positive Controls

Let us start with the “good news” first. The first group we will review are those controls that are simpler, easier, or better after a move to the cloud.

2.1. Control 1 - Inventory of Authorized and Unauthorized Devices

In this author’s experience, maintaining an accurate list of computing devices can be a nightmare in a traditional network environment. Laptops and desktops are swapped out regularly, and hardware technicians are too busy, or simply forget, to update device documentation. Sadly, a similar lack of proper documentation can run rampant even in server environments. Application owners and server owners sometimes seem impossible to keep track of.

This control is made significantly simpler in a cloud environment. Since customers are charged based on cloud resource utilization, the full list of devices will be shown at any given time (as well as device history) in the AWS console. Even if the OS build didn’t include the other management frameworks (e.g. Microsoft’s SCCM) that are normally necessary to manage and keep track of assets. Amazon CloudTrail can also be

utilized to log those changes to a dedicated S3 storage bucket for review and analysis as needed.

2.2. Control 5 - Controlled use of Administrative Privileges

With the move to the cloud, a new type of administrative access is added – that of the cloud management console itself. The initial AWS setup involves the creation of an AWS root account, which has access to all functions within the console, including billing and instance management, among other things. Critical Control 5 requires that employees have access to do their job, but no more – or less – access than is actually needed.

Amazon's Identity and Access Manager (IAM) allows for the creation of admin accounts that can be granted granular permissions across the entire cloud infrastructure. (Amazon IAM, 2015). Amazon's "Best Practices" document recommends the root account credentials be stored away safely, and general use accounts created for each system administrator or service that requires access. The IAM administrator can then delegate appropriate permissions for each account as needed (Amazon IAM, 2015). IAM Policies can also enforce multi-factor authentication (MFA) for administrative functions.

As with other Amazon AWS services, IAM supports federated accounts, allowing on-premise Active Directory accounts to be used to login to the AWS console. Using federated accounts helps ensure that access is removed as part of the normal employee termination process, as the employee's AD account is disabled or removed.

Methods of managing elevated privileges within the guest operating system are unchanged from an on-premise installation.

2.3. Control 10 - Data Recovery Capability

If properly designed, advantages to a cloud installation are most apparent when it comes to data recovery capability. For file storage, a combination of S3 storage buckets and Amazon Glacier storage can be a cost-effective way to ensure data is backed up and easily recoverable. For EC2 instances, regularly scheduled system snapshots make for simple backup and recovery procedures.

One thing to keep in mind regarding this control in the cloud is sub-control 10.4: "Ensure that key systems have at least one backup destination that is not continuously

addressable through OS calls.” (Center for Internet Security, 2015, p. 35) Fortunately, EC2 snapshots meet this requirement. As with all backup processes, it is extremely important to regularly test the recovery process – which, of course, is also mandated in sub-control 10.2.

2.4. Control 14 - Controlled Access Based on the Need to Know

There are three areas we need to examine for this control: network design, Data Loss Prevention (DLP), and user permissions.

A solid cloud-aware network design easily allows for the proper data and resource segmentation required by this control. Each resource in the Amazon Virtual Private Cloud (VPC) can be placed in a Security Group, which the EC2 documentation (Amazon EC2, 2015) describes like this:

When you launch an instance in a VPC, you can assign the instance to up to five security groups. Security groups act at the instance level, not the subnet level. Therefore, each instance in a subnet in your VPC could be assigned to a different set of security groups.

These security groups can function as a virtual Private VLAN – specifically, restricting resources in the same security group from communicating with each other unless explicitly granted that permission.

Host-based DLP solutions are still a valid option in the cloud. However, it is critical that the selected solution handles consistently and gracefully the rapid and frequent commissioning and decommissioning of clients.

Lastly, Amazon’s IAM service, previously mentioned in Control 5, can ensure proper permissions are granted and managed to any given set of data.

3. Cloud Negative Controls

Now that the “good news” is out of the way, let us examine some of the “less-good news.” Note that these controls are not made impossible by a move to the cloud, simply more difficult or cost-prohibitive.

3.1. Control 6 - Maintenance, Monitoring, and Analysis of Audit Logs

Amazon charges customers for EC2 instances over three major areas: CPU time, storage, and bandwidth. Control 6 can have a negative impact in any of these three areas.

System level logs can be shipped back down to an on-premise centralized log collection system, but bandwidth utilization could very quickly grow to be cost-prohibitive depending on the systems running in AWS and the volume of logs generated.

Alternatively, logs could be stored locally in S3 storage buckets and migrated to lower-cost storage buckets based on a reasonable data retention policy. However, this may still be a problem, depending on what solution is in use to read or correlate logs from cloud-hosted systems. How will these logs be imported into or accessed from the log query engine or SIEM? Does that log query engine or SIEM support direct access to S3 buckets, or do the logs need to be imported into the tool's native storage format for proper indexing? Can the query engine run in the cloud? Possibly, but that would incur significant CPU cycle costs and potentially negate the value of moving to the cloud in the first place.²

An entirely new set of logs must also be managed and monitored when running in the cloud: cloud provider account activity logs. Logging must include administrator actions such as spinning up or shutting down an EC2 instance, access to storage buckets, and account administration. Amazon provides this ability through their CloudTrail service and advertises several partners with the ability to read these logs natively (Amazon CloudTrail, 2015).

One other method to reduce costs for transferring logs (or other high-bandwidth items) is an AWS Direct Connect agreement. Direct Connect allows network communications to traverse a dedicated link, resulting in reduced Internet bandwidth

² Rich Mogull has an excellent blog post, "Why I design for one cloud at a time", which nicely explains how an architecture like this could negate the benefits of moving to the cloud. <https://securosis.com/blog/why-i-design-for-one-cloud-at-a-time>

requirements as well as a lower bandwidth fee from Amazon³. This option is not feasible for all customers but is currently one of the better solutions to help limit the impact of implementing this control.

To avoid unnecessary costs, log levels will need to be carefully tuned and a strong data retention policy/process developed to ensure that the right logs are generated, kept, and deleted appropriately. The data retention policy will also need to balance the need to maintain the right logs for incident response, troubleshooting, and compliance auditing. Admittedly, traditional on-premise data retention policies must have this same balance, but it is important to keep in mind these new factors may contribute additional costs and/or risks to the organization.

3.2. Control 7 – Email and Web Browser Protections

Brand new in version 6 of the Critical Security Controls, Control 7 focuses on client email and browser protections. If browsing the web directly from servers is permitted by organizational policies, then client Internet traffic will need to be logged according to sub-control 7.4. The most likely method for logging these requests would be the use of a web filter or proxy, required in sub-control 7.6, which also limits Internet access to approved sites or categories. If browsing the web from servers is not permitted by policy, implement egress firewall filtering rules to prevent such activity.

3.3. Control 8 - Malware Defenses

When it comes to malware defenses, the network-based sub-controls present the most difficulty. Inspecting network traffic for cloud resources will require a network routing table that either funnels traffic through a virtual security appliance (running as another cloud instance) or routes all Internet traffic back down through existing on-premise security filters. Either of these options could prove to be problematic or cost-prohibitive, depending on the function of the cloud resources. If Direct Connect is an option for the organization, this is another control that would significantly benefit from

³ Microsoft offers a similar network connectivity option for Azure customers, though the specifics of those features were out of scope for this paper.

that arrangement. Unfortunately, this is not the only area where network-based controls may prove difficult.

Sub-control 8.6 calls for “query logging to detect hostname lookup[s] for known malicious C2 domains” (Center for Internet Security, 2015, p. 31). This should only present a new issue if the DNS server is running in the cloud. Since DNS query logs can grow very large in a very short amount of time, all of the problems already discussed around logging will apply.

3.4. Control 12 - Boundary Defense

The placement of intrusion detection/prevention sensors may be the biggest challenge for this control. Not all of the options available in an on-premise installation are available in the cloud (e.g. NetFlow or network taps to feed an IDS sensor), so the organization may need to make adjustments to account for any gaps.

Firewalls and blacklists can still be utilized with native AWS controls, but implementing certain sub-controls (such as those seen in Table 1 below) requires maintaining a security appliance instance as a network gateway. If not done with great care, this configuration has the potential to impact the ability to scale instances properly.

Sub-control	Description
12.2	“[R]ecord at least packet header information, and preferably full packet header and payloads of the traffic destined for or passing through the network border.”
12.3	“Deploy network-based IDS sensors on Internet and extranet DMZ systems and networks...”
12.4	“Network-based IPS devices should be deployed to complement IDS...”
12.5	“[A]ll outgoing network traffic to the Internet must pass through at least one application layer filtering proxy server.”
12.9	“Deploy NetFlow collection and analysis to DMZ network flows...”

12.10	“[C]onfigure the built-in firewall session tracking mechanisms included in many commercial firewalls to identify TCP sessions that last an unusually long time...”
--------------	--

Table 1 - Cloud-negative sub-controls for Boundary Defense (Center for Internet Security, 2015, p. 41-42)

3.5. Control 13 - Data Protection

Network-based defenses are again problematic when it comes to Data Protection, especially if the implemented cloud model requires direct Internet access as opposed to routing all traffic, not just management traffic, through an Amazon Direct Connect link or something similar. Unfortunately, direct Internet access is a very common use case for cloud services.

The problem is that the cloud provider controls the network perimeter, and does not typically allow for network taps or SPANs for visibility to security appliances (Mogull, 2015, p. 11). The solution to this problem may begin to sound like a broken record: either implement a cloud security appliance as a customer-managed perimeter device or route all Internet traffic back down to an existing on-premise solution.

Again, host-based DLP is still a viable solution and may be required here to make up for gaps in network visibility.

There is good news when looking at sub-control 13.2, “Deploy approved hard drive encryption...[on] systems that hold sensitive data.” (Critical Security Controls, 2015, p. 46). Many of the commonly used Amazon storage platforms (S3, Elastic Block Store, Storage Gateway, Relational Database Service, and Elastic MapReduce) support several traditional encryption libraries, such as OpenSSL, Bitlocker, dm-lock, or ecryptfs, to name a few. The array of options for key management in AWS is vast and more complicated than can be accommodated here, but is clearly laid out in Amazon’s whitepaper on the subject (Beer & Holland, 2014).

3.6. Control 19 - Incident Response and Management

Incident response for cloud-based systems is interesting in that, on paper, it is no different from an on-premise installation. Ensure that Incident Response plans have clearly defined roles and responsibilities, and list the necessary actions to take when an

incident occurs. But in the author's personal experience, gathering relevant system-level logs from cloud resources involved in an incident is more difficult. Proper attention must be paid to logging, as discussed under Control 6, but it is mentioned here as a reminder that incident response is much harder without appropriate logs to analyze.

4. Cloud Neutral Controls

A move to the cloud does not significantly change this next group of controls, but there are just a few points that should be briefly addressed.

4.1. Control 4 - Continuous Vulnerability Assessment and Remediation

Vulnerability scans and penetration tests can still be conducted against corporate resources in the AWS cloud, but explicit permission must be obtained prior to executing the scan. Amazon's permission request form "requires you to submit information about the instances you wish to test, identify the expected start and end dates/times of your test, and requires you to read and agree to Terms and Conditions specific to penetration testing and to the use of appropriate tools for testing." (Amazon AWS, 2015). However, scanning the "small" or "micro" EC2 instances is explicitly prohibited by Amazon's policy, due to performance concerns on shared resources (ibid).

This obviously requires some advance notice (and time allotted to answer any questions Amazon presents back to the business), but this requirement should not present an issue so long as a proper level of project planning is included in the assessment and remediation activities.

4.2. Control 16 - Account Monitoring and Control

A move to the cloud also introduces some new account types that must be managed, namely AWS console accounts. Fortunately, if Active Directory is already in use, those AD accounts can be federated into AWS using either SAML 2.0 or OpenID Connect (OIDC) (Amazon AWS, 2015).

Even if native AWS accounts are required, IAM can still help manage permissions. Inactive AWS accounts can be reported via the `password_last_used`

field in the AWS Credential report. Amazon IAM started tracking last password usage as of October 20, 2014 (Amazon Web Services, 2012, loc. 2562). Unfortunately, AWS does not appear to support account expiration dates natively, so implementing sub-control 16.2 cannot be achieved without the benefit of federated accounts.

4.3. Control 17 - Security Skills Assessment and Training to Fill Gaps

Additional training should be provided to cover cloud implementations for system and security administrators. Otherwise, there is no real change or challenge specific to cloud implementations in this control.

4.4. Control 18 – Application Software Security

Software development best practices do not change, even if the resulting application is running in the cloud. Therefore, sub-control 18.2 is the only potentially problematic issue: “Protect web applications by deploying web application firewalls (WAFs) that inspect all traffic flowing to the web application...” (Center for Internet Security, 2015, p. 67). Once again, a security appliance must be inserted into the traffic flow for this control. Fortunately, implementing a WAF as a reverse proxy is already a standard security configuration (Beechey, 2009, p. 2) and there is no shortage of cloud-based WAF providers, such as Incapsula or Cloudflare to choose from. Amazon has even introduced their own native WAF as an option in this space (Amazon AWS Blog, 2015).

4.5. All the Rest

Controls 2, 3, and 11 (“Inventory of Authorized and Unauthorized Software”, “Secure Configurations for Hardware and Software”, and “Secure Configurations for Network Devices”) revolve around software and configuration management and are controlled via the same frameworks as an on-premise installation (with the exception of whatever applications are included in the AMI, which is analogous to an install image anyway).

Controls 9, and 15 are not significantly different for cloud environments versus traditional network environments. Control 9 focuses on firewall configurations, which are

unchanged in the cloud. (One might even make the argument that Control 9 would fall into the “cloud positive” category with the use of cloud zones.) And there’s no such thing as wireless access points in the cloud (Control 15)⁴.

Steps to implement Control 20 (“Penetration Tests and Red Team Exercises”) in the cloud would be the same as for Control 4 (“Continuous Vulnerability Assessment and Remediation”): request Amazon’s approval in advance via their request form.

5. Concluding Thoughts

A move to “the Cloud” does not signal doomsday for security teams. In fact, moving to a system that is so natively and tightly integrated with automation in mind makes some controls significantly simpler. Security and network architects need to update their knowledge base to take into account the special oddities and advantages presented by a cloud architecture. Organizations can take advantage of security improvements in the “cloud positive” controls and actually strengthen their overall security posture, so long as they pay careful attention to the problem areas introduced by the “cloud negative” controls. Understanding the *purposes* of the Critical Security Controls allows an organization to effectively adopt those controls (perhaps with some tweaks) even in a cloud environment.

⁴ Access *Controllers* certainly run in the cloud (i.e. Meraki), just not Access *Points*.

Appendix A: List of Controls and Relevant Minimum Control Sensors for the Cloud

Control	Critical Security Control #1: Inventory of Authorized and Unauthorized Devices	AWS Tool
1.1	Deploy an automated asset inventory discovery tool and use it to build a preliminary inventory of systems connected to an organization's public and private network(s). Both active tools that scan through IPv4 or IPv6 network address ranges and passive tools that identify hosts based on analyzing their traffic should be employed.	AWS Console/APIs
1.2	If the organization is dynamically assigning addresses using DHCP, then deploy dynamic host configuration protocol (DHCP) server logging, and use this information to improve the asset inventory and help detect unknown systems.	AWS Console/APIs
1.3	Ensure that all equipment acquisitions automatically update the inventory system as new, approved devices are connected to the network.	AWS Console/APIs
1.4	Maintain an asset inventory of all systems connected to the network and the network devices themselves, recording at least the network addresses, machine name(s), purpose of each system, an asset owner responsible for each device, and the department associated with each device. The inventory should include every system that has an Internet protocol (IP) address on the network, including but not limited to desktops, laptops, servers, network equipment (routers, switches, firewalls, etc.), printers, storage area networks, Voice Over-IP telephones, multi-homed addresses, virtual addresses, etc. The asset inventory created must also include data on whether the device is a portable and/or personal device. Devices such as mobile phones, tablets, laptops, and other portable electronic devices that store or process data must be identified, regardless of whether they are attached to the organization's network.	AWS Console/APIs
1.5	Deploy network level authentication via 802.1x to limit and control which devices can be connected to the network. The 802.1x must be tied into the inventory data to determine authorized versus unauthorized systems.	n/a
1.6	Use client certificates to validate and authenticate systems prior to connecting to the private network.	n/a

Control	Critical Security Control #2: Inventory of Authorized and Unauthorized Software	AWS Tool
2.1	Devise a list of authorized software and version that is required in the enterprise for each type of system, including servers, workstations, and laptops of various kinds and uses. This list should be monitored by file integrity checking tools to validate that the authorized software has not been modified.	n/a
2.2	Deploy application whitelisting technology that allows systems to run software only if it is included on the whitelist and prevents execution of all other software on the system. The whitelist may be very extensive (as is available from commercial whitelist vendors), so that users are not inconvenienced when using common software. Or, for some special-purpose systems (which require only a small number of programs to achieve their needed business functionality), the whitelist may be quite narrow.	n/a
2.3	Deploy software inventory tools throughout the organization covering each of the operating system types in use, including servers, workstations, and laptops. The software inventory system should track the version of the underlying operating system as well as the applications installed on it. The software inventory systems must be tied into the hardware asset inventory so all devices and associated software are tracked from a single location.	n/a
2.4	Virtual machines and/or air-gapped systems should be used to isolate and run applications that are required for business operations but based on higher risk should not be installed within a networked environment.	n/a

Control	Critical Security Control #3: Secure Configurations for Hardware and Software	AWS Tool
3.1	Establish standard secure configurations of your operating systems and software applications. Standardized images should represent hardened versions of the underlying operating system and the applications installed on the system. These images should be validated and refreshed on a regular basis to update their security configuration in light of recent vulnerabilities and attack vectors.	Amazon AMIs
3.2	Follow strict configuration management, building a secure image that is used to build all new systems that are deployed in the enterprise. Any existing system that becomes compromised should be re-imaged with the secure build. Regular updates or exceptions to this image should be integrated into the organization's change management processes. Images should be created for workstations, servers, and other system types used by the organization.	n/a
3.3	Store the master images on securely configured servers, validated with integrity checking tools capable of continuous inspection, and change management to ensure that only authorized changes to the images are possible. Alternatively, these master images can be stored in offline machines, air-gapped from the production network, with images copied via secure media to move them between the image storage servers and the production network.	n/a
3.4	Perform all remote administration of servers, workstation, network devices, and similar equipment over secure channels. Protocols such as telnet, VNC, RDP, or others that do not actively support strong encryption should only be used if they are performed over a secondary encryption channel, such as SSL, TLS or IPSEC.	n/a
3.5	Use file integrity checking tools to ensure that critical system files (including sensitive system and application executables, libraries, and configurations) have not been altered. The reporting system should: have the ability to account for routine and expected changes; highlight and alert on unusual or unexpected alterations; show the history of configuration changes over time and identify who made the change (including the original logged-in account in the event of a user ID switch, such as with the su or sudo command). These integrity checks should identify suspicious system alterations such as: owner and permissions changes to files or directories; the use of alternate data streams which could be used to hide malicious activities; and the introduction of extra files into key system areas (which could indicate malicious payloads left by attackers or additional files inappropriately added during batch distribution processes).	n/a
3.6	Implement and test an automated configuration monitoring system that verifies all remotely testable secure configuration elements, and alerts when unauthorized changes occur. This includes detecting new listening ports, new administrative users, changes to group and local policy objects (where applicable), and new services running on a system. Whenever possible use tools compliant with the Security Content Automation Protocol (SCAP) in order to streamline reporting and integration.	n/a
3.7	Deploy system configuration management tools, such as Active Directory Group Policy Objects for Microsoft Windows systems or Puppet for UNIX systems that will automatically enforce and redeploy configuration settings to systems at regularly scheduled intervals. They should be capable of triggering redeployment of configuration settings on a scheduled, manual, or event-driven basis.	Amazon APIs

Control	Critical Security Control #4: Continuous Vulnerability Assessment and Remediation	AWS Tool
4.1	Run automated vulnerability scanning tools against all systems on the network on a weekly or more frequent basis and deliver prioritized lists of the most critical vulnerabilities to each responsible system administrator along with risk scores that compare the effectiveness of system administrators and departments in reducing risk. Use a SCAP-validated vulnerability scanner that looks for both code-based vulnerabilities (such as those described by Common Vulnerabilities and Exposures entries) and configuration-based vulnerabilities (as enumerated by the Common Configuration Enumeration Project).	No new tool, but permission must be granted prior to running the test
4.2	Correlate event logs with information from vulnerability scans to fulfill two goals. First, personnel should verify that the activity of the regular vulnerability scanning tools is itself logged. Second, personnel should be able to correlate attack detection events with prior vulnerability scanning results to determine whether the given exploit was used against a target known to be vulnerable.	n/a

Control	Critical Security Control #4: Continuous Vulnerability Assessment and Remediation	AWS Tool
4.3	Perform vulnerability scanning in authenticated mode either with agents running locally on each end system to analyze the security configuration or with remote scanners that are given administrative rights on the system being tested. Use a dedicated account for authenticated vulnerability scans, which should not be used for any other administrative activities and should be tied to specific machines at specific IP addresses. Ensure that only authorized employees have access to the vulnerability management user interface and that roles are applied to each user.	No new tool, but permission must be granted prior to running the test
4.4	Subscribe to vulnerability intelligence services in order to stay aware of emerging exposures, and use the information gained from this subscription to update the organization's vulnerability scanning activities on at least a monthly basis. Alternatively, ensure that the vulnerability scanning tools you use are regularly updated with all relevant important security vulnerabilities.	n/a
4.5	Deploy automated patch management tools and software update tools for operating system and software/applications on all systems for which such tools are available and safe. Patches should be applied to all systems, even systems that are properly air gapped.	n/a
4.6	Monitor logs associated with any scanning activity and associated administrator accounts to ensure that this activity is limited to the timeframes of legitimate scans.	n/a
4.7	Compare the results from back-to-back vulnerability scans to verify that vulnerabilities were addressed either by patching, implementing a compensating control, or documenting and accepting a reasonable business risk. Such acceptance of business risks for existing vulnerabilities should be periodically reviewed to determine if newer compensating controls or subsequent patches can address vulnerabilities that were previously accepted, or if conditions have changed, increasing the risk.	n/a
4.8	Establish a process to risk-rate vulnerabilities based on the exploitability and potential impact of the vulnerability, and segmented by appropriate groups of assets (example, DMZ servers, internal network servers, desktops, laptops). Apply patches for the riskiest vulnerabilities first. A phased rollout can be used to minimize the impact to the organization. Establish expected patching timelines based on the risk rating level.	n/a

Control	Critical Security Control #5: Controlled Use of Administrative Privileges	AWS Tool
5.1	Minimize administrative privileges and only use administrative accounts when they are required. Implement focused auditing on the use of administrative privileged functions and monitor for anomalous behavior.	AWS Console/APIs; IAM
5.2	Use automated tools to inventory all administrative accounts and validate that each person with administrative privileges on desktops, laptops, and servers is authorized by a senior executive.	AWS Console/APIs; IAM
5.3	Before deploying any new devices in a networked environment, change all default passwords for applications, operating systems, routers, firewalls, wireless access points, and other systems to have values consistent with administration-level accounts.	n/a
5.4	Configure systems to issue a log entry and alert when an account is added to or removed from a domain administrators' group, or when a new local administrator account is added on a system.	n/a
5.5	Configure systems to issue a log entry and alert on any unsuccessful login to an administrative account.	n/a
5.6	Use multifactor authentication for all administrative access, including domain administrative access. Multi-factor authentication can include a variety of techniques, to include the use of smart cards, certificates, One Time Password (OTP) tokens, biometrics, or other similar authentication methods.	AWS IAM
5.7	Where multi-factor authentication is not supported, user accounts shall be required to use long passwords on the system (longer than 14 characters).	AWS Console/APIs; IAM
5.8	Administrators should be required to access a system using a fully logged and non-administrative account. Then, once logged on to the machine without administrative privileges, the administrator should transition to administrative privileges using tools such as Sudo on Linux/UNIX, RunAs on Windows, and other similar facilities for other types of systems.	n/a

Control	Critical Security Control #5: Controlled Use of Administrative Privileges	AWS Tool
5.9	Administrators shall use a dedicated machine for all administrative tasks or tasks requiring elevated access. This machine shall be isolated from the organization's primary network and not be allowed Internet access. This machine shall not be used for reading e-mail, composing documents, or surfing the Internet.	n/a

Control	Critical Security Control #6: Maintenance, Monitoring, and Analysis of Audit Logs	AWS Tool
6.1	Include at least two synchronized time sources from which all servers and network equipment retrieve time information on a regular basis so that timestamps in logs are consistent.	n/a
6.2	Validate audit log settings for each hardware device and the software installed on it, ensuring that logs include a date, timestamp, source addresses, destination addresses, and various other useful elements of each packet and/or transaction. Systems should record logs in a standardized format such as syslog entries or those outlined by the Common Event Expression initiative. If systems cannot generate logs in a standardized format, log normalization tools can be deployed to convert logs into such a format.	n/a
6.3	Ensure that all systems that store logs have adequate storage space for the logs generated on a regular basis, so that log files will not fill up between log rotation intervals. The logs must be archived and digitally signed on a periodic basis.	Amazon S3
6.4	Have security personnel and/or system administrators run biweekly reports that identify anomalies in logs. They should then actively review the anomalies, documenting their findings.	AWS APIs; CloudTrail
6.5	Configure network boundary devices, including firewalls, network-based IPS, and inbound and outbound proxies, to verbosely log all traffic (both allowed and blocked) arriving at the device.	Amazon S3
6.6	Deploy a SIEM (Security Information and Event Management) or log analytic tools for log aggregation and consolidation from multiple machines and for log correlation and analysis. Using the SIEM tool, system administrators and security personnel should devise profiles of common events from given systems so that they can tune detection to focus on unusual activity, avoid false positives, more rapidly identify anomalies, and prevent overwhelming analysts with insignificant alerts.	Amazon S3

Control	Critical Security Control #7: Email and Web Browser Protections	AWS Tool
7.1	Ensure that only fully supported web browsers and email clients are allowed to execute in the organization, ideally only using the latest version of the browsers provided by the vendor in order to take advantage of the latest security functions and fixes.	n/a
7.2	Uninstall or disable any unnecessary or unauthorized browser or email client plugins or add-on applications. Each plugin shall utilize application / URL whitelisting and only allow the use of the application for pre-approved domains.	n/a
7.3	Limit the use of unnecessary scripting languages in all web browsers and email clients. This includes the use of languages such as ActiveX and JavaScript on systems where it is unnecessary to support such capabilities.	n/a
7.4	Log all URL requests from each of the organization's systems, whether onsite or a mobile device, in order to identify potentially malicious activity and assist incident handlers with identifying potentially compromised systems.	n/a
7.5	Deploy two separate browser configurations to each system. One configuration should disable the use of all plugins, unnecessary scripting languages, and generally be configured with limited functionality and be used for general web browsing. The other configuration shall allow for more browser functionality but should only be used to access specific websites that require the use of such functionality.	n/a
7.6	The organization shall maintain and enforce network based URL filters that limit a system's ability to connect to websites not approved by the organization. The organization shall subscribe to URL categorization services to ensure that they are up-to-date with the most recent website category definitions available. Uncategorized sites shall be blocked by default. This filtering shall be enforced for each of the organization's systems, whether they are physically at an organization's facilities or not.	n/a

Control	Critical Security Control #7: Email and Web Browser Protections	AWS Tool
7.7	To lower the chance of spoofed e-mail messages, implement the Sender Policy Framework (SPF) by deploying SPF records in DNS and enabling receiver-side verification in mail servers.	n/a
7.8	Scan and block all e-mail attachments entering the organization's e-mail gateway if they contain malicious code or file types that are unnecessary for the organization's business. This scanning should be done before the e-mail is placed in the user's inbox. This includes e-mail content filtering and web content filtering.	n/a

Control	Critical Security Control #8: Malware Defenses	AWS Tool
8.1	Employ automated tools to continuously monitor workstations, servers, and mobile devices with anti-virus, anti-spyware, personal firewalls, and host-based IPS functionality. All malware detection events should be sent to enterprise anti-malware administration tools and event log servers.	n/a
8.2	Employ anti-malware software that offers a centralized infrastructure that compiles information on file reputations or have administrators manually push updates to all machines. After applying an update, automated systems should verify that each system has received its signature update.	n/a
8.3	Limit use of external devices to those with an approved, documented business need. Monitor for use and attempted use of external devices. Configure laptops, workstations, and servers so that they will not auto-run content from removable media, like USB tokens (i.e., "thumb drives"), USB hard drives, CDs/DVDs, FireWire devices, external serial advanced technology attachment devices, and mounted network shares. Configure systems so that they automatically conduct an anti-malware scan of removable media when inserted.	No new tool, but risk is greatly reduced, since all machines are virtual. :-)
8.4	Enable anti-exploitation features such as Data Execution Prevention (DEP), Address Space Layout Randomization (ASLR), virtualization/containerization, etc. For increased protection, deploy capabilities such as Enhanced Mitigation Experience Toolkit (EMET) that can be configured to apply these protections to a broader set of applications and executables.	n/a
8.5	Use network-based anti-malware tools to identify executables in all network traffic and use techniques other than signature-based detection to identify and filter out malicious content before it arrives at the endpoint.	Limited to Security Appliance installation
8.6	Enable domain name system (DNS) query logging to detect hostname lookup for known malicious C2 domains.	n/a

Control	Critical Security Control #9: Limitation and Control of Network Ports	AWS Tool
9.1	Ensure that only ports, protocols, and services with validated business needs are running on each system.	n/a
9.2	Apply host-based firewalls or port filtering tools on end systems, with a default-deny rule that drops all traffic except those services and ports that are explicitly allowed.	n/a
9.3	Perform automated port scans on a regular basis against all key servers and compared to a known effective baseline. If a change that is not listed on the organization's approved baseline is discovered, an alert should be generated and reviewed.	n/a
9.4	Verify any server that is visible from the Internet or an untrusted network, and if it is not required for business purposes, move it to an internal VLAN and give it a private address.	n/a
9.5	Operate critical services on separate physical or logical host machines, such as DNS, file, mail, web, and database servers.	n/a
9.6	Place application firewalls in front of any critical servers to verify and validate the traffic going to the server. Any unauthorized services or traffic should be blocked and an alert generated.	n/a

Control	Critical Security Control #10: Data Recovery Capability	AWS Tool
10.1	Ensure that each system is automatically backed up on at least a weekly basis, and more often for systems storing sensitive information. To help ensure the ability to rapidly restore a system from backup, the operating system, application software, and data on a machine should each be included in the overall backup procedure. These three components of a system do not have to be included in the same backup file or use the same backup software. There should be multiple backups over time, so that in the event of malware infection, restoration can be from a version that is believed to predate the original infection. All backup policies should be compliant with any regulatory or official requirements.	Amazon S3; Amazon Glacier; EC2 snapshots
10.2	Test data on backup media on a regular basis by performing a data restoration process to ensure that the backup is properly working.	n/a
10.3	Ensure that backups are properly protected via physical security or encryption when they are stored, as well as when they are moved across the network. This includes remote backups and cloud services.	Amazon S3; Amazon Glacier; EC2 snapshots
10.4	Ensure that key systems have at least one backup destination that is not continuously addressable through operating system calls. This will mitigate the risk of attacks like CryptoLocker which seek to encrypt or damage data on all addressable data shares, including backup destinations.	EC2 snapshots

Control	Critical Security Control #11: Secure Configurations for Network Devices	AWS Tool
11.1	Compare firewall, router, and switch configuration against standard secure configurations defined for each type of network device in use in the organization. The security configuration of such devices should be documented, reviewed, and approved by an organization change control board. Any deviations from the standard configuration or updates to the standard configuration should be documented and approved in a change control system.	n/a
11.2	All new configuration rules beyond a baseline-hardened configuration that allow traffic to flow through network security devices, such as firewalls and network-based IPS, should be documented and recorded in a configuration management system, with a specific business reason for each change, a specific individual's name responsible for that business need, and an expected duration of the need.	n/a
11.3	Use automated tools to verify standard device configurations and detect changes. All alterations to such files should be logged and automatically reported to security personnel.	n/a
11.4	Manage network devices using two-factor authentication and encrypted sessions.	n/a
11.5	Install the latest stable version of any security-related updates on all network devices.	n/a
11.6	Network engineers shall use a dedicated machine for all administrative tasks or tasks requiring elevated access. This machine shall be isolated from the organization's primary network and not be allowed Internet access. This machine shall not be used for reading e-mail, composing documents, or surfing the Internet.	n/a
11.7	Manage the network infrastructure across network connections that are separated from the business use of that network, relying on separate VLANs or, preferably, on entirely different physical connectivity for management sessions for network devices.	n/a

Control	Critical Security Control #12: Boundary Defense	AWS Tool
12.1	Deny communications with (or limit data flow to) known malicious IP addresses (black lists), or limit access only to trusted sites (whitelists). Tests can be periodically carried out by sending packets from bogon source IP addresses (non-routable or otherwise unused IP addresses) into the network to verify that they are not transmitted through network perimeters. Lists of bogon addresses are publicly available on the Internet from various sources, and indicate a series of IP addresses that should not be used for legitimate traffic traversing the Internet.	Limited to Security Appliance installation
12.2	On DMZ networks, configure monitoring systems (which may be built in to the IDS sensors or deployed as a separate technology) to record at least packet header information, and preferably full packet header and payloads of the traffic destined for or passing through the network border. This traffic should be sent to a properly configured Security Information Event Management (SIEM) or log analytics system so that events can be correlated from all devices on the network.	Limited to Security Appliance installation

Control	Critical Security Control #12: Boundary Defense	AWS Tool
12.3	Deploy network-based IDS sensors on Internet and extranet DMZ systems and networks that look for unusual attack mechanisms and detect compromise of these systems. These network-based IDS sensors may detect attacks through the use of signatures, network behavior analysis, or other mechanisms to analyze traffic.	Limited to Security Appliance installation
12.4	Network-based IPS devices should be deployed to complement IDS by blocking known bad signatures or the behavior of potential attacks. As attacks become automated, methods such as IDS typically delay the amount of time it takes for someone to react to an attack. A properly configured network-based IPS can provide automation to block bad traffic. When evaluating network-based IPS products, include those using techniques other than signature-based detection (such as virtual machine or sandbox-based approaches) for consideration.	Limited to Security Appliance installation
12.5	Design and implement network perimeters so that all outgoing network traffic to the Internet must pass through at least one application layer filtering proxy server. The proxy should support decrypting network traffic, logging individual TCP sessions, blocking specific URLs, domain names, and IP addresses to implement a black list, and applying whitelists of allowed sites that can be accessed through the proxy while blocking all other sites. Organizations should force outbound traffic to the Internet through an authenticated proxy server on the enterprise perimeter.	Limited to Security Appliance installation
12.6	Require all remote login access (including VPN, dial-up, and other forms of access that allow login to internal systems) to use two-factor authentication.	n/a
12.7	All enterprise devices remotely logging into the internal network should be managed by the enterprise, with remote control of their configuration, installed software, and patch levels. For third-party devices (e.g., subcontractors/vendors), publish minimum security standards for access to the enterprise network and perform a security scan before allowing access.	n/a
12.8	Periodically scan for back-channel connections to the Internet that bypass the DMZ, including unauthorized VPN connections and dual-homed hosts connected to the enterprise network and to other networks via wireless, dial-up modems, or other mechanisms.	Limited to Security Appliance installation
12.9	Deploy NetFlow collection and analysis to DMZ network flows to detect anomalous activity.	Limited to Security Appliance installation
12.10	To help identify covert channels exfiltrating data through a firewall, configure the built-in firewall session tracking mechanisms included in many commercial firewalls to identify TCP sessions that last an unusually long time for the given organization and firewall device, alerting personnel about the source and destination addresses associated with these long sessions.	Limited to Security Appliance installation

Control	Critical Security Control #13: Data Protection	AWS Tool
13.1	Perform an assessment of data to identify sensitive information that requires the application of encryption and integrity controls	n/a
13.2	Deploy approved hard drive encryption software to mobile devices and systems that hold sensitive data.	n/a
13.3	Deploy an automated tool on network perimeters that monitors for sensitive information (e.g., personally identifiable information), keywords, and other document characteristics to discover unauthorized attempts to exfiltrate data across network boundaries and block such transfers while alerting information security personnel.	Limited to Security Appliance installation
13.4	Conduct periodic scans of server machines using automated tools to determine whether sensitive data (e.g., personally identifiable information, health, credit card, or classified information) is present on the system in clear text. These tools, which search for patterns that indicate the presence of sensitive information, can help identify if a business or technical process is leaving behind or otherwise leaking sensitive information.	n/a
13.5	If there is no business need for supporting such devices, configure systems so that they will not write data to USB tokens or USB hard drives. If such devices are required, enterprise software should be used that can configure systems to allow only specific USB devices (based on serial number or other unique property) to be accessed, and that can automatically encrypt all data placed on such devices. An inventory of all authorized devices must be maintained.	No new tool, but risk is greatly reduced, since all machines are virtual. :-)

Control	Critical Security Control #13: Data Protection	AWS Tool
13.6	Use network-based DLP solutions to monitor and control the flow of data within the network. Any anomalies that exceed the normal traffic patterns should be noted and appropriate action taken to address them.	Limited to Security Appliance installation
13.7	Monitor all traffic leaving the organization and detect any unauthorized use of encryption. Attackers often use an encrypted channel to bypass network security devices. Therefore it is essential that organizations be able to detect rogue connections, terminate the connection, and remediate the infected system.	Limited to Security Appliance installation
13.8	Block access to known file transfer and e-mail exfiltration websites.	Limited to Security Appliance installation
13.9	Use host-based data loss prevention (DLP) to enforce ACLs even when data is copied off a server. In most organizations, access to the data is controlled by ACLs that are implemented on the server. Once the data have been copied to a desktop system, the ACLs are no longer enforced and the users can send the data to whomever they want.	n/a

Control	Critical Security Control #14: Controlled Access Based on the Need to Know	AWS Tool
14.1	Segment the network based on the label or classification level of the information stored on the servers. Locate all sensitive information on separated VLANs with firewall filtering to ensure that only authorized individuals are only able to communicate with systems necessary to fulfill their specific responsibilities.	Amazon EC2 security groups
14.2	All communication of sensitive information over less-trusted networks should be encrypted. Whenever information flows over a network with a lower trust level, the information should be encrypted.	n/a
14.3	All network switches will enable Private Virtual Local Area Networks (VLANs) for segmented workstation networks to limit the ability of devices on a network to directly communicate with other devices on the subnet and limit an attackers ability to laterally move to compromise neighboring systems.	Amazon EC2 security groups
14.4	All information stored on systems shall be protected with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principal that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.	Amazon IAM
14.5	Sensitive information stored on systems shall be encrypted at rest and require a secondary authentication mechanism, not integrated into the operating system, in order to access the information.	n/a
14.6	Enforce detailed audit logging for access to nonpublic data and special authentication for sensitive data.	Amazon CloudTrail
14.7	Archived data sets or systems not regularly accessed by the organization shall be removed from the organization's network. These systems shall only be used as stand alone systems (disconnected from the network) by the business unit needing to occasionally use the system or completely virtualized and powered off until needed.	n/a

Control	Critical Security Control #15: Wireless Access Control	AWS Tool
15.1	Ensure that each wireless device connected to the network matches an authorized configuration and security profile, with a documented owner of the connection and a defined business need. Organizations should deny access to those wireless devices that do not have such a configuration and profile.	n/a
15.2	Configure network vulnerability scanning tools to detect wireless access points connected to the wired network. Identified devices should be reconciled against a list of authorized wireless access points. Unauthorized (i.e., rogue) access points should be deactivated.	n/a
15.3	Use wireless intrusion detection systems (WIDS) to identify rogue wireless devices and detect attack attempts and successful compromises. In addition to WIDS, all wireless traffic should be monitored by WIDS as traffic passes into the wired network.	n/a

Control	Critical Security Control #15: Wireless Access Control	AWS Tool
15.4	Where a specific business need for wireless access has been identified, configure wireless access on client machines to allow access only to authorized wireless networks. For devices that do not have an essential wireless business purpose, disable wireless access in the hardware configuration (basic input/output system or extensible firmware interface).	n/a
15.5	Ensure that all wireless traffic leverages at least Advanced Encryption Standard (AES) encryption used with at least Wi-Fi Protected Access 2 (WPA2) protection.	n/a
15.6	Ensure that wireless networks use authentication protocols such as Extensible Authentication Protocol-Transport Layer Security (EAP/TLS), which provide credential protection and mutual authentication.	n/a
15.7	Disable peer-to-peer wireless network capabilities on wireless clients.	n/a
15.8	Disable wireless peripheral access of devices (such as Bluetooth), unless such access is required for a documented business need.	n/a
15.9	Create separate virtual local area networks (VLANs) for BYOD systems or other untrusted devices. Internet access from this VLAN should go through at least the same border as corporate traffic. Enterprise access from this VLAN should be treated as untrusted and filtered and audited accordingly.	n/a

Control	Critical Security Control #16: Account Monitoring and Control	AWS Tool
16.1	Review all system accounts and disable any account that cannot be associated with a business process and owner.	Amazon IAM
16.2	Ensure that all accounts have an expiration date that is monitored and enforced.	Amazon IAM
16.3	Establish and follow a process for revoking system access by disabling accounts immediately upon termination of an employee or contractor. Disabling instead of deleting accounts allows preservation of audit trails.	Amazon IAM
16.4	Regularly monitor the use of all accounts, automatically logging off users after a standard period of inactivity.	n/a
16.5	Configure screen locks on systems to limit access to unattended workstations.	n/a
16.6	Monitor account usage to determine dormant accounts, notifying the user or user's manager. Disable such accounts if not needed, or document and monitor exceptions (e.g., vendor maintenance accounts needed for system recovery or continuity operations). Require that managers match active employees and contractors with each account belonging to their managed staff. Security or system administrators should then disable accounts that are not assigned to valid workforce members.	Amazon IAM
16.7	Use and configure account lockouts such that after a set number of failed login attempts the account is locked for a standard period of time.	n/a
16.8	Monitor attempts to access deactivated accounts through audit logging.	Amazon IAM, CloudTrail
16.9	Configure access for all accounts through a centralized point of authentication, for example Active Directory or LDAP. Configure network and security devices for centralized authentication as well.	Amazon IAM, CloudTrail
16.10	Profile each user's typical account usage by determining normal time-of-day access and access duration. Reports should be generated that indicate users who have logged in during unusual hours or have exceeded their normal login duration. This includes flagging the use of the user's credentials from a computer other than computers on which the user generally works.	Amazon IAM, CloudTrail
16.11	Require multi-factor authentication for all user accounts that have access to sensitive data or systems. Multi-factor authentication can be achieved using smart cards, certificates, One Time Password (OTP) tokens, or biometrics.	Amazon IAM, CloudTrail
16.12	Where multi-factor authentication is not supported, user accounts shall be required to use long passwords on the system (longer than 14 characters).	Amazon IAM, CloudTrail
16.13	Ensure that all account usernames and authentication credentials are transmitted across networks using encrypted channels.	n/a
16.14	Verify that all authentication files are encrypted or hashed and that these files cannot be accessed without root or administrator privileges. Audit all access to password files in the system.	n/a

Control	Critical Security Control #17: Security Skills Assessment and Appropriate Training to Fill Gaps	AWS Tool
17.1	Perform gap analysis to see which skills employees need and which behaviors employees are not adhering to, using this information to build a baseline training and awareness roadmap for all employees.	n/a
17.2	Deliver training to fill the skills gap. If possible, use more senior staff to deliver the training. A second option is to have outside teachers provide training onsite so the examples used will be directly relevant. If you have small numbers of people to train, use training conferences or online training to fill the gaps.	AWS Training & Certification ⁵
17.3	Implement an security awareness program that (1) focuses only on the methods commonly used in intrusions that can be blocked through individual action, (2) is delivered in short online modules convenient for employees (3) is updated frequently (at least annually) to represent the latest attack techniques, (4) is mandated for completion by all employees at least annually, and (5) is reliably monitored for employee completion.	n/a
17.4	Validate and improve awareness levels through periodic tests to see whether employees will click on a link from suspicious e-mail or provide sensitive information on the telephone without following appropriate procedures for authenticating a caller; targeted training should be provided to those who fall victim to the exercise.	n/a
17.5	Use security skills assessments for each of the mission-critical roles to identify skills gaps. Use hands-on, real-world examples to measure mastery. If you do not have such assessments, use one of the available online competitions that simulate real-world scenarios for each of the identified jobs in order to measure skills mastery.	n/a

Control	Critical Security Control #18: Application Software Security	AWS Tool
18.1	For all acquired application software, check that the version you are using is still supported by the vendor. If not, update to the most current version and install all relevant patches and vendor security recommendations.	n/a
18.2	Protect web applications by deploying web application firewalls (WAFs) that inspect all traffic flowing to the web application for common web application attacks, including but not limited to cross-site scripting, SQL injection, command injection, and directory traversal attacks. For applications that are not web-based, specific application firewalls should be deployed if such tools are available for the given application type. If the traffic is encrypted, the device should either sit behind the encryption or be capable of decrypting the traffic prior to analysis. If neither option is appropriate, a host-based web application firewall should be deployed.	Limited to Security Appliance installation; or 3rd party cloud-based WAF (e.g. Incapsula, CloudFlare, Securi)
18.3	For in-house developed software, ensure that explicit error checking is performed and documented for all input, including for size, data type, and acceptable ranges or formats.	n/a
18.4	Test in-house-developed and third-party-procured web applications for common security weaknesses using automated remote web application scanners prior to deployment, whenever updates are made to the application, and on a regular recurring basis. In particular, input validation and output encoding routines of application software should be reviewed and tested.	n/a
18.5	Do not display system error messages to end-users (output sanitization).	n/a
18.6	Maintain separate environments for production and nonproduction systems. Developers should not typically have unmonitored access to production environments.	n/a
18.7	For applications that rely on a database, use standard hardening configuration templates. All systems that are part of critical business processes should also be tested.	n/a
18.8	Ensure that all software development personnel receive training in writing secure code for their specific development environment.	n/a

⁵ The Amazon AWS training site is available at <https://aws.amazon.com/training/>

Control	Critical Security Control #18: Application Software Security	AWS Tool
18.9	For in-house developed applications, ensure that development artifacts (sample data and scripts; unused libraries, components, debug code; or tools) are not included in the deployed software, or accessible in the production environment.	n/a

Control	Critical Security Control #19: Incident Response and Management	AWS Tool
19.1	Ensure that there are written incident response procedures that include a definition of personnel roles for handling incidents. The procedures should define the phases of incident handling.	n/a
19.2	Assign job titles and duties for handling computer and network incidents to specific individuals.	n/a
19.3	Define management personnel who will support the incident handling process by acting in key decision-making roles.	n/a
19.4	Devise organization-wide standards for the time required for system administrators and other personnel to report anomalous events to the incident handling team, the mechanisms for such reporting, and the kind of information that should be included in the incident notification. This reporting should also include notifying the appropriate Community Emergency Response Team in accordance with all legal or regulatory requirements for involving that organization in computer incidents.	n/a
19.5	Assemble and maintain information on third-party contact information to be used to report a security incident (e.g., maintain an e-mail address of security@organization.com or have a web page http://organization.com/security).	n/a
19.6	Publish information for all personnel, including employees and contractors, regarding reporting computer anomalies and incidents to the incident handling team. Such information should be included in routine employee awareness activities.	n/a
19.7	Conduct periodic incident scenario sessions for personnel associated with the incident handling team to ensure that they understand current threats and risks, as well as their responsibilities in supporting the incident handling team.	n/a

Control	Critical Security Control #20: Penetration Tests and Red Team Exercises	AWS Tool
20.1	Conduct regular external and internal penetration tests to identify vulnerabilities and attack vectors that can be used to exploit enterprise systems successfully. Penetration testing should occur from outside the network perimeter (i.e., the Internet or wireless frequencies around an organization) as well as from within its boundaries (i.e., on the internal network) to simulate both outsider and insider attacks.	No new tool, but permission must be granted prior to running the test
20.2	Any user or system accounts used to perform penetration testing should be controlled and monitored to make sure they are only being used for legitimate purposes, and are removed or restored to normal function after testing is over.	n/a
20.3	Perform periodic Red Team exercises to test organizational readiness to identify and stop attacks or to respond quickly and effectively.	No new tool, but permission must be granted prior to running the test
20.4	Include tests for the presence of unprotected system information and artifacts that would be useful to attackers, including network diagrams, configuration files, older penetration test reports, e-mails or documents containing passwords or other information critical to system operation.	n/a
20.5	Plan clear goals of the penetration test itself with blended attacks in mind, identifying the goal machine or target asset. Many APT-style attacks deploy multiple vectors—often social engineering combined with web or network exploitation. Red Team manual or automated testing that captures pivoted and multi-vector attacks offers a more realistic assessment of security posture and risk to critical assets.	n/a
20.6	Use vulnerability scanning and penetration testing tools in concert. The results of vulnerability scanning assessments should be used as a starting point to guide and focus penetration testing efforts.	n/a

Control	Critical Security Control #20: Penetration Tests and Red Team Exercises	AWS Tool
20.7	Wherever possible, ensure that Red Teams results are documented using open, machine-readable standards (e.g., SCAP). Devise a scoring method for determining the results of Red Team exercises so that results can be compared over time.	n/a
20.8	Create a test bed that mimics a production environment for specific penetration tests and Red Team attacks against elements that are not typically tested in production, such as attacks against supervisory control and data acquisition and other control systems.	n/a

References

- Amazon AWS Blog. (2015, October 6). Introducing AWS WAF. Retrieved January 27, 2016, from <https://aws.amazon.com/about-aws/whats-new/2015/10/introducing-aws-waf/>
- Amazon AWS. (2015). Simulated event testing - Amazon Web Services (AWS). Retrieved December 13, 2015, from <http://aws.amazon.com/security/penetration-testing/>
- Amazon AWS. (2015). Identity providers and federation - AWS Identity and Access Management. Retrieved December 20, 2015, from http://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_providers.html
- Amazon CloudTrail. (2015). AWS CloudTrail | Partners. Retrieved January 25, 2016, from <https://aws.amazon.com/cloudtrail/partners/>
- Amazon Direct Connect. (2015). AWS Direct Connect - Private network connection to the AWS cloud. Retrieved December 15, 2015, from <https://aws.amazon.com/directconnect/>
- Amazon EC2. (2015). Security groups for your VPC - Amazon Virtual Private Cloud. Retrieved December 19, 2015, from http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_SecurityGroups.html
- Amazon IAM. (2015). IAM best practices - AWS Identity and Access Management. Retrieved December 13, 2015, from <http://docs.aws.amazon.com/IAM/latest/UserGuide/best-practices.html>
- Amazon IAM. (2015). What Is IAM? - AWS Identity and Access Management. Retrieved December 13, 2015, from <http://docs.aws.amazon.com/IAM/latest/UserGuide/introduction.html>
- Amazon Web Services. (2012). AWS Identity and Access Management (IAM) user guide [Kindle Edition]. Retrieved from <http://www.amazon.com/gp/product/B007S3WJYW>
Last updated 20 November 2015

- Amazon Web Services. (2013). AWS security best practices. Retrieved from <https://d0.awsstatic.com/whitepapers/aws-security-best-practices.pdf>
- Amazon Web Services. (2015). Introduction to AWS security. Retrieved from https://d0.awsstatic.com/whitepapers/Security/Intro_to_AWS_Security.pdf
- Amazon Web Services. (2015). What is Cloud Computing? - Amazon Web Services. Retrieved December 19, 2015, from https://aws.amazon.com/what-is-cloud-computing/?nc1=f_cc
- Beechey, J. (2009). Web Application Firewalls: Defense in depth for your web infrastructure. Retrieved from SANS Technology Institute website: http://www.sans.edu/student-files/projects/200904_01.doc
- Beer, K., & Holland, R. (2014). Encrypting data at rest. Retrieved from Amazon website: http://d0.awsstatic.com/whitepapers/AWS_Securing_Data_at_Rest_with_Encryption.pdf
- Center for Internet Security. (2015). The CIS Critical Security Controls for effective cyber defense (v6). Retrieved from <https://www.cisecurity.org/critical-controls.cfm>
- Mogull, R. (2015). Pragmatic security for cloud and hybrid networks. Retrieved from Securosis website: <https://securosis.com/assets/library/reports/PragmaticNetSec.v.1.1.1.final.pdf>
- Mogull, R. (2015, November 4). Securosis Blog | Why I design for one cloud at a time [Blog post]. Retrieved from <https://securosis.com/blog/why-i-design-for-one-cloud-at-a-time>