The American Association of Professional Landmen

CONTRACT CENTER

System Architecture and Security Model for ContractRoom

Updated Q4-2016

System Architecture

Development

- Built on Django 1.6.x
- Written in Python 2.7.x

Databases

- Amazon RDS
- Toku MX Enterprise 2.x

Hosting

- Hosted on Amazon Web Service (AWS)
- Ubuntu LTS
- Highly available: redundant servers in different data centers
- Automatic failover for all components
- Scalability based on Amazon RDS cluster and multiple TokuMX instances
- All vital components are being monitored

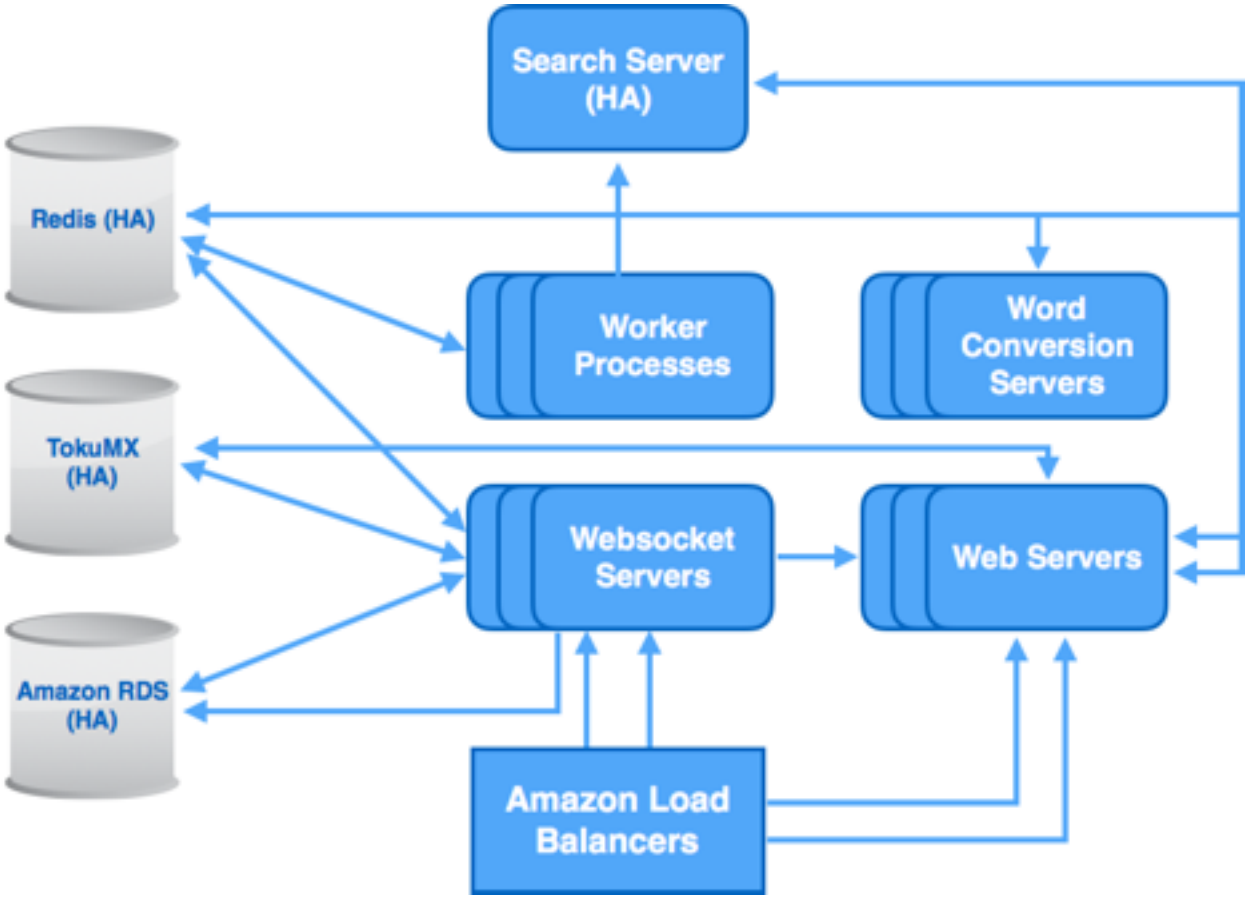

Privacy & Security Features

Data Protection

- Connection to the ContractRoom is achieved via secure socket layer/transport layer security (EV SSL/TLS), guaranteeing that our customers and their counterparties have a secure connection to their data.
- Individual user sessions are uniquely identified and re-verified with each transaction.
- Customer passwords are not accessible by ContractRoom personnel.
- Application logs record user actions (creators, updaters, timestamps), as well as originating IP address for every record and transaction completed.

Extended Validation SSL Certificate

An Extended Validation Certificate (EV) is an X.509 public key certificate issued according to a specific set of identity verification criteria. These criteria require extensive verification of the requesting entity's identity by the certificate authority (CA) before a certificate is issued. Certificates issued by a CA under the EV guidelines are designated with a CA-specific policy identifier so that EV-aware software, such as a web browser, can recognize them. EV certificates are mainly presented by web servers to web browsers for use with Transport Layer Security (TLS) connections.

The criteria for issuing EV certificates are defined by the Guidelines for Extended Validation, currently (as of March 24, 2014) at version 1.4.6. The guidelines are produced by the CA/Browser Forum, a voluntary organization whose members include leading CAs and vendors of Internet software, as well as representatives from the legal and audit professions.

Encryption of data at rest

ContractRoom protects data from all transactions generated and stored in the system using Advanced Encryption Standard (AES) method. Cryptography is implemented on ContractRoom's hybrid database system, housing the data and on the physical storage the databases are stored. Data encryption keys are updated on a regular basis and stored separately from the data. Federation will also be implemented as it becomes needed for international transactions.

Cross site scripting (XSS) protection

ContractRoom's technology protects you against XSS attacks. XSS attacks allow a user to inject client side scripts into the browsers of other users. This is usually achieved by storing the malicious scripts in the database where it will be retrieved and displayed to other users, or by getting users to click a link which will cause the attacker's JavaScript to be executed by the user's

browser. However, XSS attacks can originate from any untrusted source of data, such as cookies or Web services, whenever the data is not sufficiently sanitized before including in a page.

Cross site request forgery (CSRF) protection

CSRF attacks allow a malicious user to execute actions using the credentials of another user without that user's knowledge or consent. ContractRoom has built-in protection against CSRF.

CSRF protection works by checking for a nonce in each POST request. This ensures that a malicious user cannot simply "replay" a form POST to ContractRoom and have another logged in user unwittingly submit that form. The malicious user would have to know the nonce, which is user specific.

SQL injection protection

SQL injection is a type of attack where a malicious user is able to execute arbitrary SQL code on a database. This can result in records being deleted or data leakage. In ContractRoom's query sets, the resulting SQL will be properly escaped by the underlying database driver.

Clickjacking protection

Clickjacking is a type of attack where a malicious site wraps another site in a frame. This attack can result in an unsuspecting user being tricked into performing unintended actions on the target site.

ContractRoom contains clickjacking protection in the form of the X-Frame-Options middleware which in a supporting browser can prevent a site from being rendered inside a frame.

Hashed password storage

By applying a hashing algorithm to your user's passwords before storing them in our database, we make it implausible for any attacker to determine the original password, while still being able to compare the resulting hash to the original password in the future.

Without hashing, any passwords that are stored in our application's database can be stolen if the database is compromised, and then immediately used to compromise your users' accounts, if they do not use unique passwords.

Salt password storage

A cryptographic salt is data which is applied during the hashing process in order to eliminate the possibility of the output being looked up in a list of pre-calculated pairs of hashes and their input, known as a rainbow table.

In simpler terms, a salt is a bit of additional data which makes your hashes significantly more difficult to crack. There are a number of services online which provide extensive lists of pre-computed hashes, as well as the original input for those hashes. The use of a

salt makes it implausible or impossible to find the resulting hash in one of these lists.

Brute Force attack protection

A brute force attack is a trial-and-error method used to obtain information such as a user password or personal identification number (PIN). In a brute force attack, automated software is used to generate a large number of consecutive guesses as to the value of the desired data. Brute force attacks may be used by criminals to crack encrypted data, or by security analysts to test an organization's network security.

Unlike hacks that focus on vulnerabilities in software, a Brute Force Attack aims at being the simplest kind of method to gain access to a site: it tries usernames and passwords, over and over again, until it gets in. Often deemed 'inelegant', they can be very successful when people use passwords like '123456' and usernames like 'admin.' They are, in short, an attack on the weakest link in any website's security: the user.

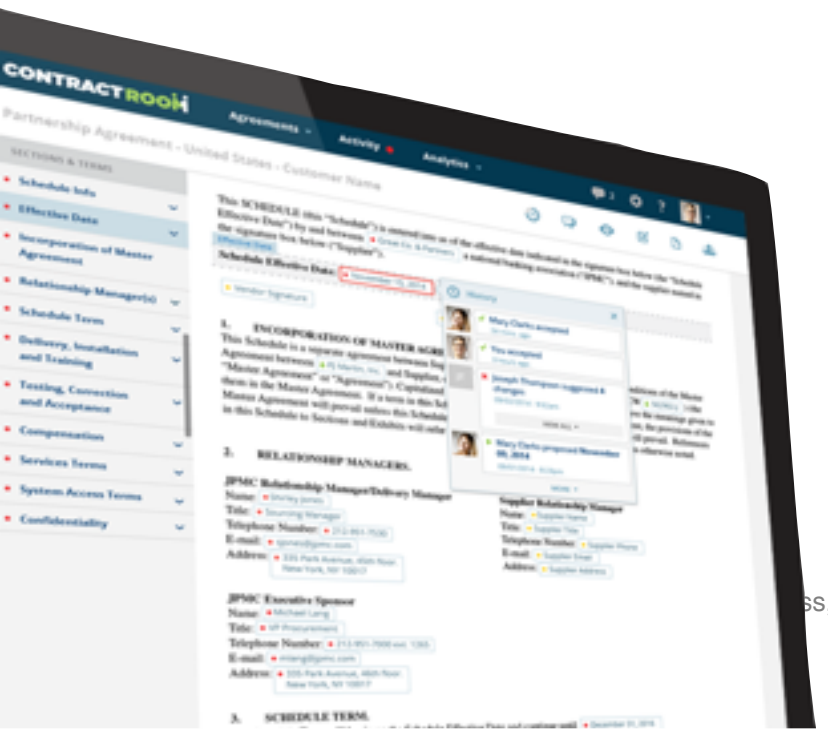

Minimum password strength enforcement

ContractRoom can be configured to use a password policy, which means that new passwords entered by users will be validated according to a certain set of requirements. Passwords that do not meet the specified conditions will be rejected.

Optional Multi-factor authentication

Multi-factor authentication combines two or more independent credentials: what the user knows (password), what the user has (security token on mobile) and what the user is (fingerprint - available in future for iOS devices). It creates a layered defense and make it more difficult for an unauthorized person to access a target such as a physical location, computing device, network or database. If one factor is compromised or broken, the attacker still has at least one more barrier to breach before successfully breaking into the target.

Optional Multi-factor signature authorization

Uses multi-factor authentication elements to complete signature and approval processes.

Hash storage for agreements signed in ContractRoom

The system will detect and verify if someone attempts to modify a document/agreement in the system after it was signed/executed.

Data Separation

- Hardware and software configurations to provide secure logical separations of customer data so each customer can only view its related information.
- Multi-tenant security controls include unique, non-predictable session tokens, configurable session timeout values, password policies, sharing rules, and user profiles.
- The ContractRoom service supports delegated authentication.

Optional Configurable Customer Privacy and Security

- Access designation to different categories of data.
- Customizable password rules.
- Log-off times for inactivity.
- Re-verification of unrecognized IP addresses or devices.

Network Security

- Multiple layers of external firewalls
- Intrusion-detection sensors
- Security event management system

- Continuous external vulnerability scanning

AWS Security Policy

AWS Compliance Program

The AWS Compliance Program enables customers to understand the robust security in place and then helps them streamline their compliance with industry and government requirements for security and data protection. The IT infrastructure that AWS provides to its customers is designed and managed in alignment with best security practices and a variety of IT security standards, including:

- SOC 1/SSAE 16/ISAE 3402 (formerly SAS 70 Type II)
- SOC 2
- SOC 3
- FISMA, DIACAP, and FedRAMP
- PCI DSS Level 1
- ISO 27001
- ITAR
- FIPS 140-2

In addition, the flexibility and control that the AWS platform provides allows customers to deploy solutions that meet several industry-specific standards, including:

- HIPAA
- Cloud Security Alliance (CSA)
- Motion Picture Association of America (MPAA)

AWS provides a wide range of information regarding its IT control environment to customers through white papers, reports, certifications, accreditations, and other third-party attestations. More information is available in the Risk and Compliance white paper available on the website: http://aws.amazon.com/security.

Physical and Environmental Security

AWS's data centers are state of the art, utilizing innovative architectural and engineering approaches. Amazon has many years of experience in designing, constructing, and operating large-scale data centers. This experience has been applied to the AWS platform and infrastructure. AWS data centers are housed in nondescript facilities. Physical access is strictly controlled both at the perimeter and at building ingress points by professional security staff utilizing video surveillance, intrusion detection systems, and other electronic means. Authorized staff must pass two-factor authentication a minimum of two times to access data center floors. All visitors and contractors are required to present identification and are signed in and continually escorted by authorized staff.

AWS only provides data center access and information to employees and contractors who have a legitimate business need for such privileges. When an employee no longer has a business need for these

privileges, his or her access is immediately revoked, even if they continue to be an employee of Amazon or Amazon Web Services. All physical access to data centers by AWS employees is logged and audited routinely.

Fire Detection and Suppression

Automatic fire detection and suppression equipment has been installed to reduce risk. The fire detection system utilizes smoke detection sensors in all data center environments, mechanical and electrical infrastructure spaces, chiller rooms and generator equipment rooms. These areas are protected by either wet-pipe, double-interlocked pre-action, or gaseous sprinkler systems.

Power

The data center electrical power systems are designed to be fully redundant and maintainable without impact to operations, 24 hours a day, and seven days a week. Uninterruptible Power Supply (UPS) units provide back-up power in the event of an electrical failure for critical and essential loads in the facility. Data centers use generators to provide back-up power for the entire facility.

Climate and Temperature

Climate control is required to maintain a constant operating temperature for servers and other hardware, which prevents overheating and reduces the possibility of service outages. Data centers are conditioned to maintain atmospheric conditions at optimal levels. Personnel and systems monitor and control temperature and humidity at appropriate levels.

Management

AWS monitors electrical, mechanical, and life support systems and equipment so that any issues are immediately identified. Preventative maintenance is performed to maintain the continued operability of equipment.

Storage Device Decommissioning

When a storage device has reached the end of its useful life, AWS procedures include a decommissioning process that is designed to prevent customer data from being exposed to unauthorized individuals. AWS uses the techniques detailed in DoD 5220.22-M (“National Industrial Security Program Operating Manual “) or NIST 800-88 (“Guidelines for Media Sanitization”) to destroy data as part of the decommissioning process. All decommissioned magnetic storage devices are degaussed and physically destroyed in accordance with industry-standard practices.

Business Continuity Management

Amazon’s infrastructure has a high level of availability and provides customers the features to deploy a resilient IT architecture.

AWS has designed its systems to tolerate system or hardware failures with minimal customer impact. Datacenter Business Continuity Management at AWS is under the direction of the Amazon Infrastructure Group.

Availability

Data centers are built in clusters in various global regions. All data centers are online and serving customers; no data center is "cold." In case of failure, automated processes move customer data traffic away from the affected area. Core applications are deployed in an N+1 configuration, so that in the event of a data center failure, there is sufficient capacity to enable traffic to be load-balanced to the remaining sites.

AWS provides you with the flexibility to place instances and store data within multiple geographic regions as well as across multiple availability zones within each region. Each availability zone is designed as an independent failure zone. This means that availability zones are physically separated within a typical metropolitan region and are located in lower risk flood plains (specific flood zone categorization varies by Region). In addition to discrete uninterruptible power supply (UPS) and onsite backup generation facilities, they are each fed via different grids from independent utilities to further reduce single points of failure. Availability zones are all redundantly connected to multiple tier-1 transit providers.

You should architect your AWS usage to take advantage of multiple regions and availability zones. Distributing applications across multiple availability zones provides the ability to remain resilient in the face of most failure modes, including natural disasters or system failures.

Incident Response

The Amazon Incident Management team employs industry-standard diagnostic procedures to drive resolution during business-impacting events. Staff operators provide 24x7x365 coverage to detect incidents and to manage the impact and resolution.

Company-Wide Executive Review

Amazon's Internal Audit group has recently reviewed the AWS services resiliency plans, which are also periodically reviewed by members of the Senior Executive management team and the Audit Committee of the Board of Directors.

Communication

AWS has implemented various methods of internal communication at a global level to help employees understand their individual roles and responsibilities and to communicate significant events in a timely manner. These methods include orientation and training programs for newly hired employees; regular management meetings for updates on business performance and other matters; and electronics means such as

video conferencing, electronic mail messages, and the posting of information via the Amazon intranet.

AWS has also implemented various methods of external communication to support its customer base and the community. Mechanisms are in place to allow the customer support team to be notified of operational issues that impact the customer experience. A "Service Health Dashboard" is available and maintained by the customer support team to alert customers to any issues that may be of broad impact. A “Security and Compliance Center” is available to provide you with a single location to obtain security and compliance details about AWS. You can also subscribe to AWS Support offerings that include direct communication with the customer support team and proactive alerts to any customer impacting issues.

Network Security

The AWS network has been architected to permit you to select the level of security and resiliency appropriate for your workload. To enable you to build geographically dispersed, fault-tolerant web architectures with cloud resources, AWS has implemented a world-class network infrastructure that is carefully monitored and managed.

Secure Network Architecture

Network devices, including firewall and other boundary devices, are in place to monitor and control communications at the external boundary of the network and at key internal boundaries within the network. These boundary devices employ rule sets, access control lists (ACL), and configurations to enforce the flow of information to specific information system services.

ACLs, or traffic flow policies, are established on each managed interface, which manage and enforce the flow of traffic. ACL policies are approved by Amazon Information Security. These policies are automatically pushed using AWS’s ACL-Manage tool, to help ensure these managed interfaces enforce the most up-to-date ACLs.

Secure Access Points

AWS has strategically placed a limited number of access points to the cloud to allow for a more comprehensive monitoring of inbound and outbound communications and network traffic. These customer access points are called API endpoints, and they allow secure HTTP access (HTTPS), which allows you to establish a secure communication session with your storage or compute instances within AWS. To support customers with FIPS 140-2 requirements, the Amazon Virtual Private Cloud VPN endpoints and SSL-terminating load balancers in AWS GovCloud (US) operate using FIPS 140-2 level 2-validated hardware.

In addition, AWS has implemented network devices that are dedicated to managing interfacing communications with Internet service providers (ISPs). AWS employs a

redundant connection to more than one communication service at each Internet-facing edge of the AWS network. These connections each have dedicated network devices.

Transmission Protection

You can connect to an AWS access point via HTTP or HTTPS using Secure Sockets Layer (SSL), a cryptographic protocol that is designed to protect against eavesdropping, tampering, and message forgery.

For customers who require additional layers of network security, AWS offers the Amazon Virtual Private Cloud (VPC), which provides a private subnet within the AWS cloud, and the ability to use an IPsec Virtual Private Network (VPN) device to provide an encrypted tunnel between the Amazon VPC and your data center. For more information about VPC configuration options, refer to the Amazon Virtual Private Cloud (Amazon VPC) Security section below.

Amazon Corporate Segregation

Logically, the AWS Production network is segregated from the Amazon Corporate network by means of a complex set of network security / segregation devices. AWS developers and administrators on the corporate network who need to access AWS cloud components in order to maintain them must explicitly request access through the AWS ticketing system. All requests are reviewed and approved by the applicable service owner.

Approved AWS personnel then connect to the AWS network through a bastion host that restricts access to network devices and other cloud components, logging all activity for security review. Access to bastion hosts require SSH public-key authentication for all user accounts on the host.

Fault-Tolerant Design

Amazon's infrastructure has a high level of availability and provides you with the capability to deploy a resilient IT architecture. AWS has designed its systems to tolerate system or hardware failures with minimal customer impact. Data centers are built in clusters in various global regions. All data centers are online and serving customers; no data center is "cold." In case of failure, automated processes move customer data traffic away from the affected area. Core applications are deployed in an N+1 configuration, so that in the event of a data center failure, there is sufficient capacity to enable traffic to be load-balanced to the remaining sites. AWS provides you with the flexibility to place instances and store data within multiple geographic regions as well as across multiple availability zones within each region. Each availability zone is designed as an independent failure zone. This means that availability zones are physically separated within a typical metropolitan region and are located in lower risk flood plains (specific flood zone categorization varies by region). In addition to utilizing discrete uninterruptible power supply (UPS) and onsite backup

generators, they are each fed via different grids from independent utilities to further reduce single points of failure.

Availability zones are all redundantly connected to multiple tier-1 transit providers. You should architect your AWS usage to take advantage of multiple regions and availability zones. Distributing applications across multiple availability zones provides the ability to remain resilient in the face of most failure scenarios, including natural disasters or system failures.

However, you should be aware of location-dependent privacy and compliance requirements, such as the EU Data Privacy Directive. Data is not replicated between regions unless proactively done so by the customer, thus allowing customers with these types of data placement and privacy requirements the ability to establish compliant environments. It should be noted that all communications between regions is across public Internet infrastructure; therefore, appropriate encryption methods should be used to protect sensitive data.

The AWS Cloud infrastructure is built around Regions and Availability Zones ("AZs"). A Region is a physical location in the world where we have multiple Availability Zones. Availability Zones consist of one or more discrete data centers, each with redundant power, networking and connectivity, housed in separate facilities. These Availability Zones offer you the ability to operate production applications and databases which are more highly available, fault tolerant and scalable than would be possible from a single data center. The AWS Cloud operates 33 Availability Zones within 12 geographic Regions around the world.

Network Monitoring and Protection

AWS utilizes a wide variety of automated monitoring systems to provide a high level of service performance and availability. AWS monitoring tools are designed to detect unusual or unauthorized activities and conditions at ingress and egress communication points. These tools monitor server and network usage, port scanning activities, application usage, and unauthorized intrusion attempts. The tools have the ability to set custom performance metrics thresholds for unusual activity.

Systems within AWS are extensively instrumented to monitor key operational

metrics. Alarms are configured to automatically notify operations and management personnel when early warning thresholds are crossed on key operational metrics. An on-call schedule is used so personnel are always available to respond to operational issues. This includes a pager system so alarms are quickly and reliably communicated to operations personnel.

Documentation is maintained to aid and inform operations personnel in handling incidents or issues. If the resolution of an issue requires collaboration, a conferencing system is used which supports communication and logging capabilities. Trained call leaders facilitate communication and progress during the handling of operational issues that require collaboration. Post-mortems are convened after any significant operational issue, regardless of external impact, and Cause of Error (COE) documents are drafted so the root cause is captured and preventative actions are taken in the future. Implementation of the preventative measures is tracked during weekly operations meetings.

AWS security monitoring tools help identify several types of denial of service (DoS) attacks, including distributed, flooding, and software/logic attacks. When DoS attacks are identified, the AWS incident response process is initiated. In addition to the DoS prevention tools, redundant telecommunication providers at each region as well as additional capacity protect against the possibility of DoS attacks.

The AWS network provides significant protection against traditional network security issues, and you can implement further protection. The following are a few examples:

- Distributed Denial Of Service (DDoS) Attacks. AWS API endpoints are hosted on large, Internet-scale, world-class infrastructure that benefits from the same engineering expertise that has built Amazon into the world's largest online retailer. Proprietary DDoS mitigation techniques are used. Additionally, AWS's networks are multi-homed across a number of providers to achieve Internet access diversity.
- Man in the Middle (MITM) Attacks. All of the AWS APIs are available via SSL-protected endpoints which provide server authentication. Amazon EC2 AMIs automatically generate new SSH host certificates on first boot and log them to the instance's console. You can then use the secure APIs to call the console and access the host certificates before logging into the instance for the first time. We encourage you to use SSL for all of your interactions with AWS.
- IP Spoofing. Amazon EC2 instances cannot send spoofed network traffic. The AWS-controlled, host-based firewall infrastructure will not permit an instance to send traffic with a source IP or MAC address other than its own.
- Port Scanning. Unauthorized port scans by Amazon EC2 customers are a violation of the AWS Acceptable Use Policy. Violations of the AWS Acceptable Use Policy are

taken seriously, and every reported violation is investigated. Customers can report suspected abuse via the contacts available on our website at: http://aws.amazon.com/contact-us/report-abuse/. When unauthorized port scanning is detected by AWS, it is stopped and blocked. Port scans of Amazon EC2 instances are generally ineffective because, by default, all inbound ports on Amazon EC2 instances are closed and are only opened by you. Your strict management of security groups can further mitigate the threat of port scans. If you configure the security group to allow traffic from any source to a specific port, then that specific port will be vulnerable to a port scan. In these cases, you must use appropriate security measures to protect listening services that may be essential to their application from being discovered by an unauthorized port scan. For example, a web server must clearly have port 80 (HTTP) open to the world, and the administrator of this server is responsible for the security of the HTTP server software, such as Apache. You may request permission to conduct vulnerability scans as required to meet your specific compliance requirements. These scans must be limited to your own instances and must not violate the AWS Acceptable Use Policy. An explanation of advanced approval for these types of scans can be found here: https://aws.amazon.com/security/penetration-testing/

- Packet sniffing by other tenants. It is not possible for a virtual instance running in promiscuous mode to receive or "sniff" traffic that is intended for a different virtual instance. While you can place your interfaces into promiscuous mode, the hypervisor will not deliver any traffic to them that is not addressed to them. Even two virtual instances that are owned by the same customer located on the same physical host cannot listen to each other's traffic. Attacks such as ARP cache poisoning do not work within Amazon EC2 and Amazon VPC. While Amazon EC2 does provide ample protection against one customer inadvertently or maliciously attempting to view another's data, as a standard practice you should encrypt sensitive traffic.

In addition to monitoring, regular vulnerability scans are performed on the host operating system, web application, and databases in the AWS environment using a variety of tools. Also, AWS Security teams subscribe to news feeds for applicable vendor flaws and proactively monitor vendors' websites and other relevant outlets for new patches. AWS customers also have the ability to report issues to AWS via the AWS Vulnerability Reporting website at: http://aws.amazon.com/security/vulnerability-reporting/

Click on this link for additional details on AWS Security Process

Contact Details:

800 Fournier Street Fort Worth, TX 76102
Tel: 817.847.7700 Fax: 817.847.7704!
aapl@landman.org
landman.org