



Amazon Web Services: Overview of Security Processes

May 2011

(Please consult <http://aws.amazon.com/security> for the latest version of this paper)



Amazon Web Services (AWS) delivers a scalable cloud computing platform with high availability and dependability, offering the flexibility to enable customers to build a wide range of applications. Helping to protect the confidentiality, integrity, and availability of our customers' systems and data is of the utmost importance to AWS, as is maintaining customer trust and confidence. This document is intended to answer questions such as, "How does AWS help me protect my data?" Specifically, AWS physical and operational security processes are described for network and server infrastructure under AWS' management, as well as service-specific security implementations. This document provides an overview of security as it pertains to the following areas relevant to AWS:

- Shared Responsibility Environment

- Control Environment Summary

- Secure Design Principles

- Backup

- Monitoring

- Information and Communication

- Employee Lifecycle

- Physical Security

- Environmental Safeguards

- Configuration Management

- Business Continuity Management

- Backups

- Fault Separation

- Amazon Account Security Features

- Network Security

- AWS Service Specific Security

- Amazon Elastic Compute Cloud (Amazon EC2) Security

- Amazon Virtual Private Cloud (Amazon VPC)

- Amazon Simple Storage Service (Amazon S3) Security

- Amazon SimpleDB Security

- Amazon Relational Database Service (Amazon RDS) Security

- Amazon Simple Queue Service (Amazon SQS) Security

- Amazon Simple Notification Service (SNS) Security

- Amazon CloudWatch Security

- Auto Scaling Security

- Amazon CloudFront Security

- Amazon Elastic MapReduce Security

Shared Responsibility Environment

Moving IT infrastructure to AWS creates a shared responsibility model between the customer and AWS. This shared model can relieve customer operational burden as AWS operates, manages and controls the components from the host operating system and virtualization layer down to the physical security of the facilities in which the services operate. The customer assumes responsibility and management of, but not limited to, the guest operating system (including updates and security patches), other associated application software, as well as the configuration of the AWS provided security group firewall. Customers should carefully consider the services they choose as their responsibilities vary depending on the services used, the integration of those services into their IT environment and applicable laws and regulations. It is possible for customers to enhance security and/or meet more stringent compliance requirements by leveraging technology such as host based firewalls, host based intrusion detection/prevention, encryption and key management. The nature of this shared responsibility also provides the flexibility and customer control that permits the deployment of solutions that meet industry-specific certification requirements.

Control Environment Summary

AWS manages a comprehensive control environment that includes the necessary policies, processes and control activities for the delivery of each of the web service offerings. The collective control environment encompasses the people, processes, and technology necessary to maintain an environment that supports the effectiveness of specific controls and the control frameworks for which AWS is certified and/or compliant.

AWS is compliant with various certifications and third-party attestations. These include:

- **SAS70 Type II.** This report includes detailed controls AWS operates along with an independent auditor opinion about the effective operation of those controls.
- **PCI DSS Level 1.** AWS has been independently validated to comply with the PCI Data Security Standard as a shared host service provider.
- **ISO 27001.** AWS has achieved ISO 27001 certification of the Information Security Management System (ISMS) covering infrastructure, data centers, and services.
- **FISMA.** AWS enables government agency customers to achieve and sustain compliance with the Federal Information Security Management Act (FISMA). AWS has been awarded an approval to operate at the FISMA-Low level. It has also completed the control implementation and successfully passed the independent security testing and evaluation required to operate at the FISMA-Moderate level. AWS is currently pursuing an approval to operate at the FISMA-Moderate level from government agencies.

Additionally, customers have built healthcare applications compliant with **HIPPA's** Security and Privacy Rules on AWS.

Further information about these certifications and third-party attestations is available in the Risk and Compliance whitepaper available on the website: <http://aws.amazon.com/security>.

Secure Design Principles

AWS' development process follows secure software development best practices, which include formal design reviews by the AWS Security Team, threat modeling, and completion of a risk assessment. Static code analysis tools are run as a part of the standard build process, and all deployed software undergoes recurring penetration testing performed by carefully selected industry experts. Our security risk assessment reviews begin during the design phase and the engagement lasts through launch to ongoing operations.

Monitoring

AWS utilizes automated monitoring systems to provide a high level of service performance and availability. Proactive monitoring is available through a variety of online tools both for internal and external use. Systems within AWS are



extensively instrumented to monitor key operational metrics. Alarms are configured to notify operations and management personnel when early warning thresholds are crossed on key operational metrics. An on-call schedule is used such that 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.

Information and 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 also available to provide customers with a single location to obtain security and compliance details about AWS.

Customers may subscribe to Premium Support offerings that include direct communication with the customer support team and proactive alerts to any customer impacting issues.

Employee Lifecycle

AWS has established formal policies and procedures to delineate the minimum standards for logical access to AWS platform and infrastructure hosts. AWS requires that staff with potential access to customer data undergo an extensive background check (as permitted by law) commensurate with their position and level of data access. The policies also identify functional responsibilities for the administration of logical access and security.

Account Provisioning

The responsibility for provisioning employee and contractor access is shared across Human Resources (HR), Corporate Operations and Service Owners.

A standard employee or contractor account with minimum privileges is provisioned in a disabled state when a hiring manager submits his or her approval. The account is automatically enabled when the employee's record is activated in Amazon's HR system.

Access to other resources including Services, Hosts, Network devices, Windows and UNIX groups must be explicitly approved in Amazon's proprietary permission management system by the appropriate owner or manager. All changes affected in the permissions management tool are captured in an audit. When changes in an employee's job function occur, continued access must be explicitly approved to the resource or it will be automatically revoked.

Account Review

Every access grant is reviewed every 90 days; explicit re-approval is required or access to the resource is automatically revoked.

Access Removal

Access is automatically revoked when an employee's record is terminated in Amazon's Human Resources system. Windows and UNIX accounts are disabled and Amazon's permission management system removes the user from all systems.

Password Policy

Access and administration of logical security for Amazon relies on user IDs, passwords and Kerberos to authenticate users to services, resources and devices as well as to authorize the appropriate level of access for the user. AWS Security has established a password policy with required configurations and expiration intervals.

Physical Security

Amazon has many years of experience in designing, constructing, and operating large-scale datacenters. This experience has been applied to the AWS platform and infrastructure. AWS datacenters 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 datacenter floors. All visitors and contractors are required to present identification and are signed in and continually escorted by authorized staff.

AWS only provides datacenter 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 datacenters by AWS employees is logged and audited routinely.

Environmental Safeguards

Amazon's data centers are state of the art, utilizing innovative architectural and engineering approaches.

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.

Configuration Management

Emergency, non-routine, and other configuration changes to existing AWS infrastructure are authorized, logged, tested, approved, and documented in accordance with industry norms for similar systems. Updates to AWS' infrastructure are done to minimize any impact on the customer and their use of the services. AWS will communicate with customers, either via email, or through the AWS Service Health Dashboard (<http://status.aws.amazon.com/>) when service use is likely to be adversely affected.

Software

AWS applies a systematic approach to managing change so that changes to customer impacting services are thoroughly reviewed, tested, approved and well communicated.

AWS' change management process is designed avoid unintended service disruptions and to maintain the integrity of service to the customer. Changes deployed into production environments are:

- Reviewed: Peer reviews of the technical aspects of a change
- Tested: being applied will behave as expected and not adversely impact performance
- Approved: to provide appropriate oversight and understanding of business impact

Changes are typically pushed into production in a phased deployment starting with lowest impact areas. Deployments are tested on a single system and closely monitored so impact can be evaluated. Service owners have a number of configurable metrics that measure the health of the service's upstream dependencies. These metrics are closely monitored with thresholds and alarming in place. Rollback procedures are documented in the Change Management (CM) ticket.

When possible, changes are scheduled during regular change windows. Emergency changes to production systems that require deviations from standard change management procedures are associated with an incident and are logged and approved as appropriate.

Periodically, AWS performs self-audits of changes to key services to monitor quality, maintain high standards and to facilitate continuous improvement of the change management process. Any exceptions are analyzed to determine the root cause and appropriate actions are taken to bring the change into compliance or roll back the change if necessary. Actions are then taken to address and remediate the process or people issue.

Infrastructure

Amazon's Corporate Applications team develops and manages software to automate IT processes for UNIX/Linux hosts in the areas of third-party software delivery, internally developed software and configuration management. The Infrastructure team maintains and operates a UNIX/Linux configuration management framework to address hardware scalability, availability, auditing, and security management. By centrally managing hosts through the use of automated processes that manage change, the Company is able to achieve its goals of high availability, repeatability, scalability, robust security and disaster recovery. Systems and Network Engineers monitor the status of these automated tools on a daily basis, reviewing reports to respond to hosts that fail to obtain or update their configuration and software.

Internally developed configuration management software is installed when new hardware is provisioned. These tools are run on all UNIX hosts to validate that they are configured and that software is installed in compliance with standards determined by the role assigned to the host. This configuration management software also helps to regularly update



packages that are already installed on the host. Only approved personnel enabled through the permissions service may log in to the central configuration management servers.

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. Data center 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 customers 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 uninterruptable 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.

Customers should architect their 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 24 x 7 x 365 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.

Note that on April 21, 2011, EC2 suffered a customer-impacting service disruption in the US East Region. Details about the service disruption are described in "Summary of the Amazon EC2 and Amazon RDS Service Disruption in the US East Region" (<http://aws.amazon.com/message/65648/>).

Backups

Data stored in Amazon S3, Amazon SimpleDB, or Amazon Elastic Block Store (EBS) is redundantly stored in multiple physical locations as part of normal operation of those services and at no additional charge. Amazon S3 and Amazon SimpleDB provide object durability by storing objects multiple times across multiple Availability Zones on the initial write and then actively doing further replication in the event of device unavailability or detected bit-rot. Amazon EBS replication is stored within the same Availability Zone, not across multiple zones and therefore it is highly recommended that customers conduct regular snapshots to Amazon S3 for long-term data durability. For customers that have architected complex transactional databases using EBS, it is recommended that backups to Amazon S3 be performed through the database management system so that distributed transactions and logs can be checkpointed. AWS does not



perform backups of data that are maintained on virtual disks attached to running instances on Amazon EC2.

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. If a hardware device is unable to be decommissioned using these procedures, the device will be degaussed or physically destroyed in accordance with industry-standard practices.

Fault Separation

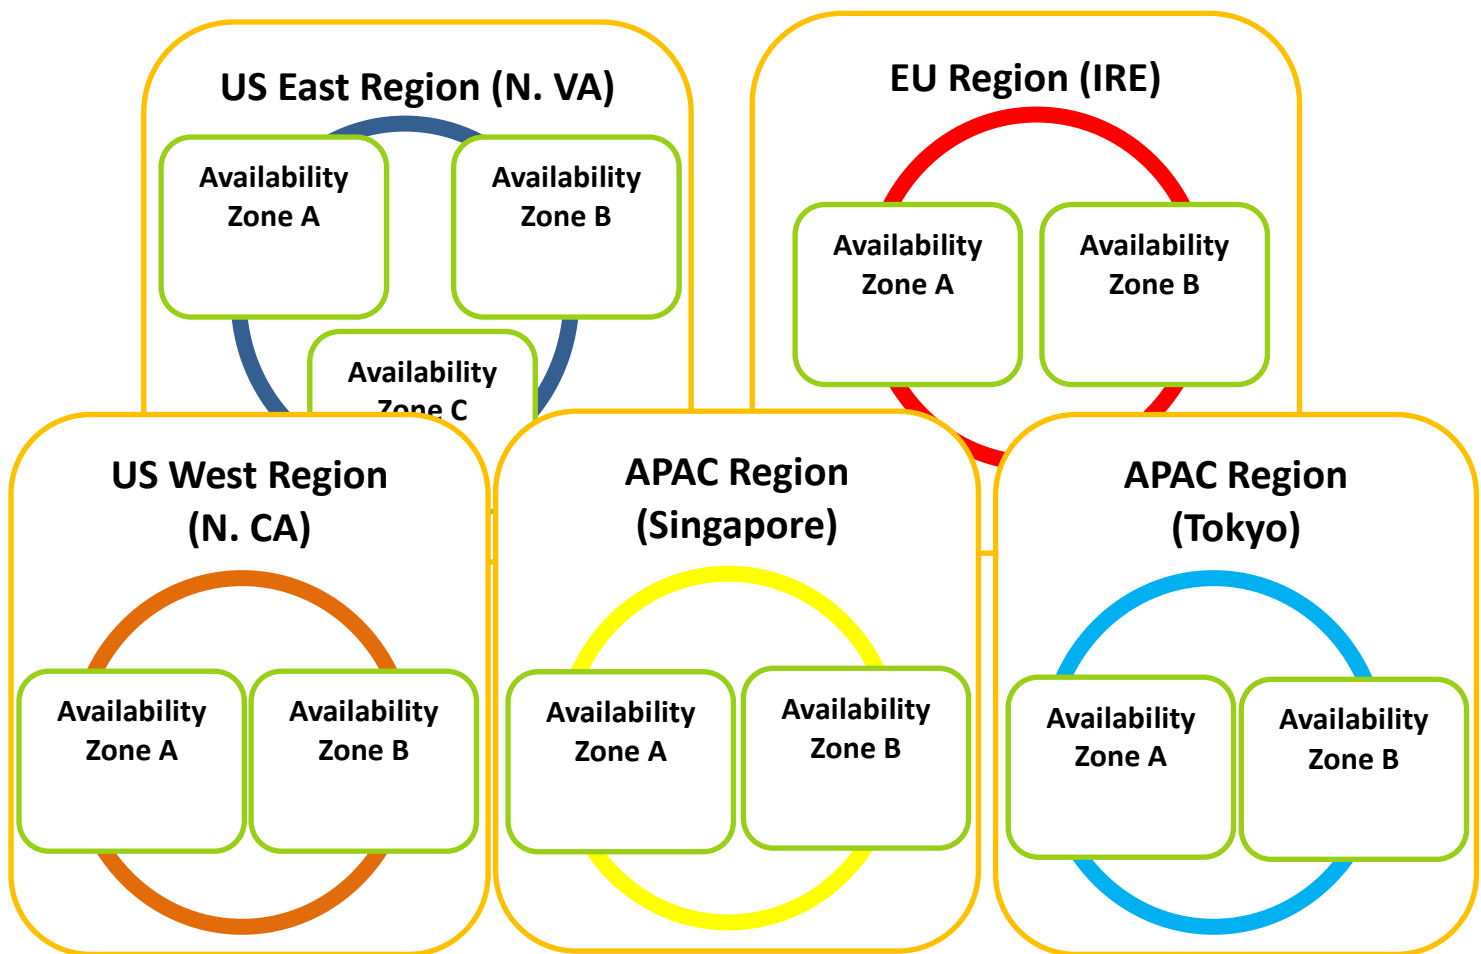
AWS provides customers the flexibility to place instances and store data within multiple geographic Regions. Each Region is an independent collection of AWS resources in a defined geography. AWS currently supports five Regions: US East (Northern Virginia), US West (Northern California), EU (Ireland), Asia Pacific (Singapore) and Asia Pacific (Tokyo). The Amazon S3 US Standard Region includes the US East facilities in Northern Virginia and facilities in Western Washington State.

The selection of a Region within an acceptable geographic jurisdiction to the customer provides a solid foundation to meeting 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. Appropriate encryption methods should be used to protect sensitive data.

Within a given Region, Amazon EC2, Amazon EBS and Amazon Relational Database Service (RDS) allow customers to place instances and store data across multiple Availability Zones. See the “Business Continuity Management” section for more information on availability.

Amazon S3, Amazon SimpleDB, Amazon Simple Notification Service (SNS), and Amazon Simple Queue Service (SQS) do not expose the concept of Availability Zones to customers. With these services, data is automatically stored on multiple devices across multiple facilities within a Region.

The diagram below demonstrates the Regions and Availability Zones within each Region for Amazon EC2, Amazon EBS and Amazon RDS.



Amazon Account Security Features

AWS provides a number of ways for customers to identify themselves and securely access their AWS Account. A complete list of credentials supported by AWS can be found on the Security Credentials page under Your Account. AWS also provides additional security options that enable customers to further protect their AWS Account and control access: AWS Identity and Access Management (AWS IAM), Multi-Factor Authentication (MFA) and Key Rotation.

AWS Identity and Access Management (AWS IAM)

AWS Identity and Access Management (AWS IAM) enables a customer to create multiple users and manage the permissions for each of these users within their AWS Account. A user is an identity (within a customer AWS Account) with unique security credentials that can be used to access AWS Services. AWS IAM eliminates the need to share passwords or access keys, and makes it easy to enable or disable a user's access as appropriate.

AWS IAM enables customers to implement security best practices, such as least privilege, by granting unique credentials to every user within their AWS Account and only granting permission to access the AWS Services and resources required for the users to perform their job. AWS IAM is secure by default; new users have no access to AWS until permissions are explicitly granted.

AWS IAM enables customers to minimize the use of their AWS Account credentials. Instead all interactions with AWS Services and resources should be with AWS IAM user security credentials. More information about AWS Identity and Access Management (AWS IAM) is available on the AWS website: <http://aws.amazon.com/iam/>

AWS Multi-Factor Authentication (AWS MFA)

AWS Multi-Factor Authentication (AWS MFA) is an additional layer of security that offers enhanced control over AWS Account settings and the management of the AWS Services and resources for which the account is subscribed. When customers enable this opt-in feature, they will need to provide a six-digit single-use code in addition to their standard username and password credentials before access is granted to their AWS Account settings or AWS Services and resources. Customers get this single use code from an authentication device that they keep in their physical possession. This is called Multi-Factor Authentication because two factors are checked before access is granted: customers need to provide both their username (Amazon e-mail in the case of the AWS Account) and password (the first “factor”: something you know) and the precise code from their authentication device (the second “factor”: something you have). Customers can enable MFA devices for their AWS Account as well as for the users they have created under their AWS Account with AWS IAM.

It is easy to obtain an authentication device from a participating third party provider and to set it up for use via the AWS website. More information about Multi-Factor Authentication is available on the AWS website:
<http://aws.amazon.com/mfa/>

Key Rotation

For the same reasons as it is important to change passwords frequently, AWS recommends that customers rotate their access keys and certificates on a regular basis. To let customers do this without potential impact to their applications’ availability, AWS supports multiple concurrent access keys and certificates. With this feature, customers can rotate keys and certificates into and out of operation on a regular basis without any downtime to their application. This can help to mitigate risk from lost or compromised access keys or certificates. The AWS IAM APIs enables a customer to rotate the access keys of their AWS Account as well as for users created under their AWS Account using AWS IAM.

Network Security

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

Distributed Denial Of Service (DDoS) Attacks

AWS Application Programming Interface (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. Customers can then use the secure APIs to call the console and access the host certificates before logging into the instance for the first time. Customers are encouraged to use SSL for all of their 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 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 the customer. The customer's strict management of security groups can further mitigate the threat of port scans. If the customer configures 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, the customer 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. Customers may request permission to conduct vulnerability scans as required to meet their specific compliance requirements. These scans must be limited to the customer's own instances and must not violate the AWS Acceptable Use Policy. Advanced approval for these types of scans can be initiated by submitting a request via the website at: <https://aws-portal.amazon.com/gp/aws/html-forms-controller/contactus/AWSecurityPenTestRequest>

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 customers can place their 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 customers should encrypt sensitive traffic.

Amazon Elastic Compute Cloud (Amazon EC2) Security

Security within Amazon EC2 is provided on multiple levels: the operating system (OS) of the host system, the virtual instance operating system or guest OS, a firewall, and signed API calls. Each of these items builds on the capabilities of the others. The goal is to protect against data contained within Amazon EC2 from being intercepted by unauthorized systems or users and to provide Amazon EC2 instances themselves that are as secure as possible without sacrificing the flexibility in configuration that customers demand.

Multiple Levels of Security

Host Operating System: Administrators with a business need to access the management plane are required to use multi-factor authentication to gain access to purpose-built administration hosts. These administrative hosts are systems that are specifically designed, built, configured, and hardened to protect the management plane of the cloud. All such access is logged and audited. When an employee no longer has a business need to access the management plane, the privileges and access to these hosts and relevant systems are revoked.

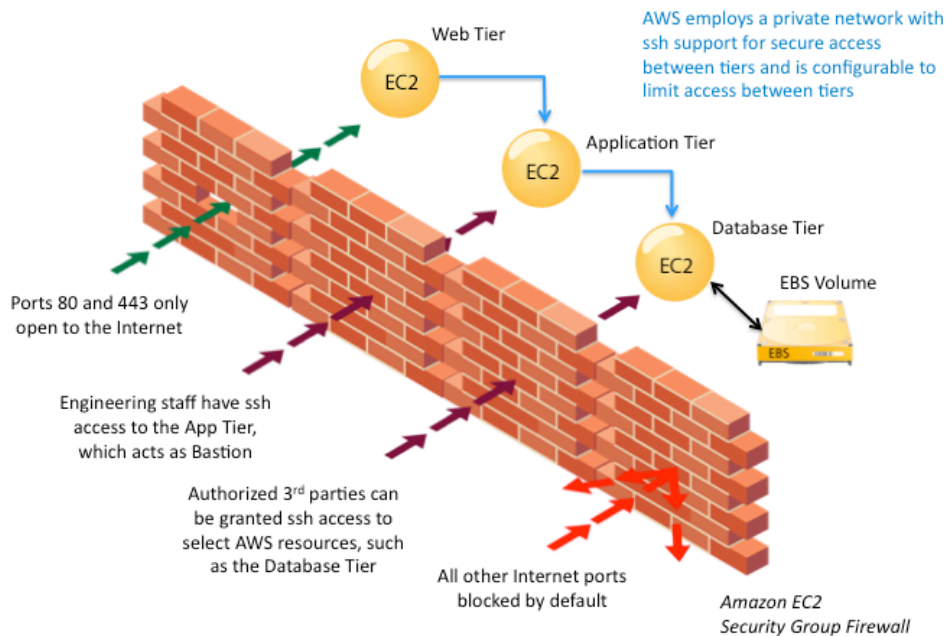
Guest Operating System: Virtual instances are completely controlled by the customer. Customers have full root access or administrative control over accounts, services, and applications. AWS does not have any access rights to customer instances and cannot log into the guest OS. AWS recommends a base set of security best practices to include disabling password-only access to their hosts, and utilizing some form of multi-factor authentication to gain access to their instances (or at a minimum certificate-based SSH Version 2 access). Additionally, customers should employ a privilege escalation mechanism with logging on a per-user basis. For example, if the guest OS is Linux, after hardening their instance, they should utilize certificate-based SSHv2 to access the virtual instance, disable remote root login, use command-line logging, and use 'sudo' for privilege escalation. Customers should generate their own key pairs in order to guarantee that they are unique, and not shared with other customers or with AWS.

Firewall: Amazon EC2 provides a complete firewall solution; this mandatory inbound firewall is configured in a default



deny-all mode and Amazon EC2 customers must explicitly open the ports needed to allow inbound traffic. The traffic may be restricted by protocol, by service port, as well as by source IP address (individual IP or Classless Inter-Domain Routing (CIDR) block).

The firewall can be configured in groups permitting different classes of instances to have different rules. Consider, for example, the case of a traditional three-tiered web application. The group for the web servers would have port 80 (HTTP) and/or port 443 (HTTPS) open to the Internet. The group for the application servers would have port 8000 (application specific) accessible only to the web server group. The group for the database servers would have port 3306 (MySQL) open only to the application server group. All three groups would permit administrative access on port 22 (SSH), but only from the customer's corporate network. Highly secure applications can be deployed using this expressive mechanism. See diagram below:



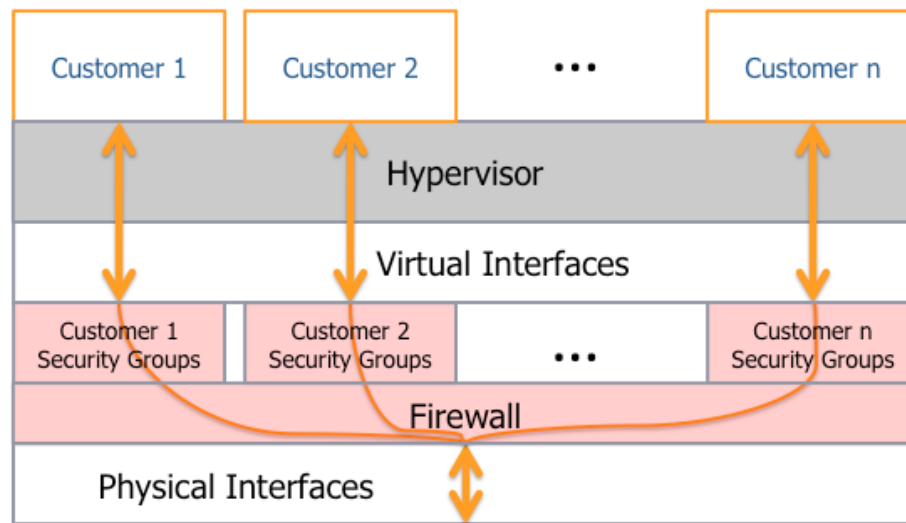
The firewall isn't controlled through the Guest OS; rather it requires the customer's X.509 certificate and key to authorize changes, thus adding an extra layer of security. AWS supports the ability to grant granular access to different administrative functions on the instances and the firewall, therefore enabling the customer to implement additional security through separation of duties. The level of security afforded by the firewall is a function of which ports are opened by the customer, and for what duration and purpose. The default state is to deny all incoming traffic, and customers should plan carefully what they will open when building and securing their applications. Well-informed traffic management and security design are still required on a per-instance basis. AWS further encourages customers to apply additional per-instance filters with host-based firewalls such as IPtables or the Windows Firewall and VPNs. This can restrict both inbound and outbound traffic on each instance. API calls to launch and terminate instances, change firewall parameters, and perform other functions are all signed by the customer's Amazon Secret Access Key, which could be either the AWS Accounts Secret Access Key or the Secret Access key of a user created with AWS IAM. Without access to the customer's Secret Access Key, Amazon EC2 API calls cannot be made on his/her behalf. In addition, API calls can be encrypted with SSL to maintain confidentiality. Amazon recommends always using SSL-protected API endpoints. AWS IAM also enables a customer to further control what APIs a user created with AWS IAM has permissions to call.

The Hypervisor

Amazon EC2 currently utilizes a highly customized version of the Xen hypervisor, taking advantage of paravirtualization (in the case of Linux guests). Because paravirtualized guests rely on the hypervisor to provide support for operations that normally require privileged access, the guest OS has no elevated access to the CPU. The CPU provides four separate privilege modes: 0-3, called *rings*. Ring 0 is the most privileged and 3 the least. The host OS executes in Ring 0. However, rather than executing in Ring 0 as most operating systems do, the guest OS runs in a lesser-privileged Ring 1 and applications in the least privileged Ring 3. This explicit virtualization of the physical resources leads to a clear separation between guest and hypervisor, resulting in additional security separation between the two.

Instance Isolation

Different instances running on the same physical machine are isolated from each other via the Xen hypervisor. Amazon is active in the Xen community, which provides awareness of the latest developments. In addition, the AWS firewall resides within the hypervisor layer, between the physical network interface and the instance's virtual interface. All packets must pass through this layer, thus an instance's neighbors have no more access to that instance than any other host on the Internet and can be treated as if they are on separate physical hosts. The physical RAM is separated using similar mechanisms.



Customer instances have no access to raw disk devices, but instead are presented with virtualized disks. The AWS proprietary disk virtualization layer automatically resets every block of storage used by the customer, so that one customer's data are never unintentionally exposed to another. AWS recommends customers further protect their data using appropriate means. One common solution is to run an encrypted file system on top of the virtualized disk device.

Elastic Block Storage (Amazon EBS) Security

Amazon EBS volume access is restricted to the AWS Account that created the volume, and to the users under the AWS Account created with AWS IAM if the user has been granted access to the EBS operations, thus denying all other AWS Accounts and users the permission to view or access the volume. However, a customer can create Amazon S3 snapshots of their Amazon EBS volume and enable other AWS Accounts the ability to use the shared snapshot as the basis for creating their own volumes. Customers also have the ability to make Amazon EBS volume snapshots publicly available to all AWS Accounts. Sharing Amazon EBS volume snapshots does not provide other AWS Accounts with the permission to alter or delete the original snapshot as that right is explicitly reserved for the AWS Account that created the volume. An EBS snapshot is a block level view of an entire EBS volume. Data which is not visible through the file system on the volume, such as files which have been deleted, may be present in the EBS snapshot. Customers that want to create

shared snapshots should do so carefully. If a volume has held sensitive data or has had files deleted from it, a new EBS volume should be created. The data to be contained in the shared snapshot should be copied to the new volume, and the snapshot created from the new volume.

Amazon EBS volumes are presented to the customer as raw unformatted block devices, which have been wiped prior to being made available for use. Customers that have procedures requiring that all data be wiped via a specific method, such as those detailed in DoD 5220.22-M (“National Industrial Security Program Operating Manual”) or NIST 800-88 (“Guidelines for Media Sanitization”), have the ability to do so on Amazon EBS. Customers should conduct a specialized wipe procedure prior to deleting the volume for compliance with their established requirements. Encryption of sensitive data is generally a good security practice, and AWS encourages users to encrypt their sensitive data via an algorithm consistent with their stated security policy.

Amazon Virtual Private Cloud (Amazon VPC) Security

Security within Amazon Virtual Private Cloud begins with the very concept of a VPC and extends to include the security groups, network access control lists (ACLs), routing, and external gateways. Each of these items is complementary to providing a secure, isolated network that can be extended through selective enabling of direct Internet access or private connectivity to another network. Below we describe the multiple levels of security in Amazon VPC. This is followed by a diagram depicting how the Amazon VPC components relate.

Multiple Levels of Security

Virtual Private Cloud: Each VPC is a distinct, isolated network within the cloud. At creation time, an IP address range for each VPC is selected by the customer. Network traffic within each VPC is isolated from all other VPCs; therefore, multiple VPCs may use overlapping (even identical) IP address ranges without loss of this isolation. By default, VPCs have no external connectivity. Customers may create and attach an Internet Gateway, VPN Gateway, or both to establish external connectivity, subject to the controls below.

API: Calls to create and delete VPCs, change routing, security group, and network ACL parameters, and perform other functions are all signed by the customer’s Amazon Secret Access Key, which could be either the AWS Accounts Secret Access Key or the Secret Access key of a user created with AWS IAM. Without access to the customer’s Secret Access Key, Amazon VPC API calls cannot be made on the customer’s behalf. In addition, API calls can be encrypted with SSL to maintain confidentiality. Amazon recommends always using SSL-protected API endpoints. AWS IAM also enables a customer to further control what APIs a newly created user has permissions to call.

Subnets: Customers create one or more subnets within each VPC; each instance launched in the VPC is connected to one subnet. Traditional Layer 2 security attacks, including MAC spoofing and ARP spoofing, are blocked.

Route Tables and Routes: Each Subnet in a VPC is associated with a routing table, and all network traffic leaving a subnet is processed by the routing table to determine the destination.

VPN Gateway: A VPN Gateway enables private connectivity between the VPC and another network. Network traffic within each VPN Gateway is isolated from network traffic within all other VPN Gateways. Customers may establish VPN Connections to the VPN Gateway from gateway devices at the customer premise. Each connection is secured by a pre-shared key in conjunction with the IP address of the customer gateway device.

Internet Gateway: An Internet Gateway may be attached to a VPC to enable direct connectivity to Amazon S3, other AWS services, and the Internet. Each instance desiring this access must either have an Elastic IP associated with it or route traffic through a NAT instance. Additionally, network routes are configured (see above) to direct traffic to the Internet Gateway. AWS provides reference NAT AMLs that can be extended by customers to perform network logging,

deep packet inspection, application-layer filtering, or other security controls.

This access can only be modified through the invocation of Amazon VPC APIs. AWS supports the ability to grant granular access to different administrative functions on the instances and the Internet Gateway, therefore enabling the customer to implement additional security through separation of duties.

Amazon EC2 Instances: Amazon EC2 instances running with an Amazon VPC contain all of the benefits described above related to the Host Operating System, Guest Operating System, Hypervisor, Instance Isolation, and protection against packet sniffing.

Tenancy: VPC allows customers to launch Amazon EC2 instances that are physically isolated at the host hardware level; they will run on single tenant hardware. A VPC can be created with 'dedicated' tenancy, in which case all instances launched into the VPC will utilize this feature. Alternatively, a VPC may be created with 'default' tenancy, but customers may specify 'dedicated' tenancy for particular instances launched into the VPC.

Firewall (Security Groups): Like Amazon EC2, Amazon VPC supports a complete firewall solution enabling filtering on both ingress and egress traffic from an instance. The default group enables inbound communication from other members of the same group and outbound communication to any destination. Traffic can be restricted by any IP protocol, by service port, as well as source/destination IP address (individual IP or Classless Inter-Domain Routing (CIDR) block).

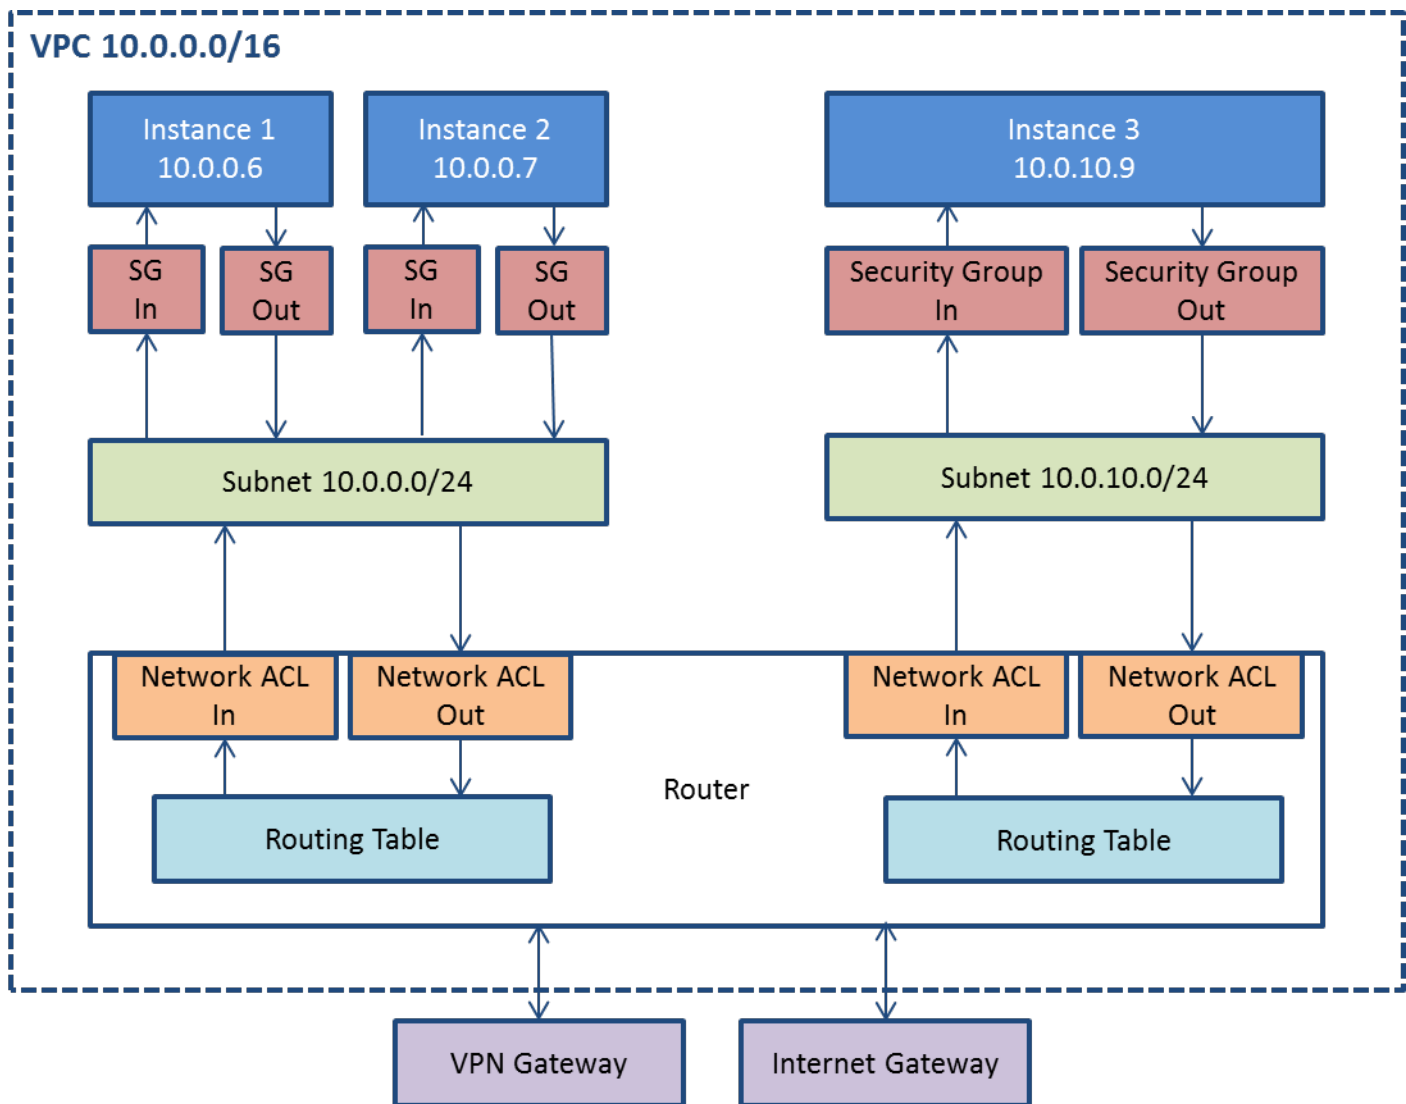
The firewall isn't controlled through the Guest OS; rather it can be modified only through the invocation of Amazon VPC APIs. AWS supports the ability to grant granular access to different administrative functions on the instances and the firewall, therefore enabling the customer to implement additional security through separation of duties. The level of security afforded by the firewall is a function of which ports are opened by the customer, and for what duration and purpose. Well-informed traffic management and security design are still required on a per-instance basis. AWS further encourages customers to apply additional per-instance filters with host-based firewalls such as IPtables or the Windows Firewall.

Network Access Control Lists: To add a further layer of security within Amazon VPC, customers can configure Network ACLs. These are stateless traffic filters that apply to all traffic inbound or outbound from a subnet within VPC. These ACLs can contain ordered rules to allow or deny traffic based upon IP protocol, by service port, as well as source/destination IP address.

Like security groups, network ACLs are managed through Amazon VPC APIs, adding an additional layer of protection and enabling additional security through separation of duties.

Network Security Summary

The diagram below depicts how the security controls above inter-relate to enable flexible network topologies while providing complete control over network traffic flows.



Amazon Simple Storage Service (Amazon S3) Security

With any shared storage system, the most common security question is whether unauthorized users can access information either intentionally or by mistake. So that customers have flexibility to determine how, when, and to whom they wish to expose the information they store in AWS, Amazon S3 APIs provide both bucket- and object-level access controls, with defaults that only permit authenticated access by the bucket and/or object creator. Unless a customer grants anonymous access to their data, the first step before a user, either an AWS Account, or a user created with AWS IAM, can access data is to be authenticated using an HMAC-SHA1 signature of the request using the user's private key. An authenticated user can read an object only if the user has been granted Read permissions in an Access Control List (ACL) at the object level. An authenticated user can list the keys and create or overwrite objects in a bucket only if the user has been granted Read and Write permissions in an ACL at the bucket level or via permissions granted to them with AWS IAM. Bucket and object level ACLs are independent; an object does not inherit ACLs from its bucket. Permissions to read or modify the bucket or object ACLs are themselves controlled by ACLs that default to creator-only access. Therefore, the customer maintains full control over who has access to their data. Customers can grant access to their Amazon S3 data to other AWS Accounts by AWS Account ID or email, or DevPay Product ID. Customers can also grant access to their Amazon S3 data to all AWS Accounts or to everyone (enabling anonymous access).

Data Management

For maximum security, Amazon S3 is accessible via SSL endpoints. The encrypted endpoints are accessible from both the Internet and from within Amazon EC2, so that data are transferred securely both within AWS and to and from sources outside of AWS.

Securing data at rest involves physical security and data encryption. As mentioned in detail in “Physical Security,” Amazon employs multiple layers of physical security measures to protect customer data at rest. For example, physical access to Amazon datacenters is limited to an audited list of Amazon personnel. Encryption of sensitive data is generally a good security practice, and AWS encourages users to encrypt their sensitive data before it is uploaded to Amazon S3.

When an object is deleted from Amazon S3, removal of the mapping from the public name to the object starts immediately, and is generally processed across the distributed system within several seconds. Once the mapping is removed, there is no remote access to the deleted object. The underlying storage area is then reclaimed for use by the system.

Amazon S3 is designed to provide 99.999999999% durability and 99.99% availability of objects over a given year. Objects are redundantly stored on multiple devices across multiple facilities in an Amazon S3 Region. To help provide durability, Amazon S3 PUT and COPY operations synchronously store your data across multiple facilities before returning SUCCESS. Once stored, Amazon S3 helps maintain the durability of your objects by quickly detecting and repairing any lost redundancy. Amazon S3 also regularly verifies the integrity of data stored using checksums. If corruption is detected, it is repaired using redundant data. In addition, Amazon S3 calculates checksums on all network traffic to detect corruption of data packets when storing or retrieving data.

Amazon S3 provides further protection via Versioning. You can use Versioning to preserve, retrieve, and restore every version of every object stored in your Amazon S3 bucket. With Versioning, you can easily recover from both unintended user actions and application failures. By default, requests will retrieve the most recently written version. Older versions of an object can be retrieved by specifying a version in the request. You can further protect your versions using Amazon S3 Versioning's MFA Delete feature, once enabled for an S3 bucket, each version deletion request must include the six-digit code and serial number from your multi factor authentication device.

Access Logging

An Amazon S3 bucket can be configured to log access to the bucket and objects within it. The access log contains details about each access request including request type, the requested resource, the requestor's IP, and the time and date of the request. When logging is enabled for a bucket, log records are periodically aggregated into log files and delivered to the specified Amazon S3 bucket.

Amazon Simple Data Base (SimpleDB) Security

Amazon SimpleDB APIs provide domain-level controls that only permit authenticated access by the domain creator, therefore the customer maintains full control over who has access to their data.

Amazon SimpleDB access can be granted based on an AWS Account ID. Once authenticated, an AWS Account has full access to all operations. Access to each individual domain is controlled by an independent Access Control List that maps authenticated users to the domains they own. A user created with AWS IAM only has access to the operations and domains for which they have been granted permission via policy.

Amazon SimpleDB is accessible via SSL-encrypted endpoints. The encrypted endpoints are accessible from both the Internet and from within Amazon EC2. Data stored within Amazon SimpleDB is not encrypted by AWS; however the customer can encrypt data before it is uploaded to Amazon SimpleDB. These encrypted attributes would be retrievable



as part of a Get operation only. They could not be used as part of a query filtering condition. Encrypting before sending data to Amazon SimpleDB helps protect against access to sensitive customer data by anyone, including AWS.

Amazon SimpleDB Data Management

When a domain is deleted from Amazon SimpleDB, removal of the domain mapping starts immediately, and is generally processed across the distributed system within seconds. Once the mapping is removed, there is no remote access to the deleted domain.

When item and attribute data are deleted within a domain, removal of the mapping within the domain starts immediately, and is also generally complete within seconds. Once the mapping is removed, there is no remote access to the deleted data. That storage area is then made available only for write operations and the data are overwritten by newly stored data.

Amazon Relational Database Service (Amazon RDS) Security

Amazon RDS allows you to quickly create a relational database instance and flexibly scale the associated compute resources and storage capacity to meet application demand. Amazon RDS manages the database instance on your behalf by performing backups, handling failover, and maintaining the database software.

Amazon RDS DB Instance access is controlled by the customer via Database Security Groups which are akin to Amazon EC2 Security Groups, but are not interchangeable. Database Security Groups default to a “deny all” access mode and customers must specifically authorize network ingress. There are two ways of doing this: authorizing a network IP range, or authorizing an existing Amazon EC2 Security Group. Database Security Groups only allow access to the database server port (all others are blocked) and can be updated without restarting the Amazon RDS DB Instance, which allows a customer seamless control of their database access.

With AWS IAM a customer can further control access to their RDS DB instances. AWS IAM enables a customer to control what RDS operations each individual AWS IAM user has permission to call.

Amazon RDS generates an SSL certificate for each DB Instance, allowing customers to encrypt their DB Instance connections for enhanced security.

Once an Amazon RDS DB Instance deletion API (DeleteDBInstance) is run, the DB Instance is marked for deletion and once the instance no longer indicates ‘deleting’ status, it has been removed. At this point the instance is no longer accessible and unless a final snapshot copy was asked for, it cannot be restored and will not be listed by any of the tools or APIs.

Amazon Simple Queue Service (Amazon SQS) Security

Amazon SQS is a highly reliable, scalable message queuing service that enables asynchronous message-based communication between distributed components of an application. The components can be computers or Amazon EC2 instances or a combination of both. With Amazon SQS you can send any number of messages to an Amazon SQS queue at any time from any component. The messages can be retrieved from the same component or a different one right away or at a later time (within 4 days). Messages are highly durable; each message is persistently stored in highly available, highly reliable queues. Multiple processes can read/write from/to an Amazon SQS queue at the same time without interfering with each other.

Amazon SQS access is granted based on an AWS Account or a user created with AWS IAM. Once authenticated, the AWS Account has full access to all user operations. An AWS IAM user however only has access to the operations and queues



which they have been granted access to via policy. By default, access to each individual queue is restricted to the AWS Account that created it. However, a customer can allow other access to a queue, using either an SQS-generated policy or a policy written by the user.

Amazon SQS is accessible via SSL-encrypted endpoints. The encrypted endpoints are accessible from both the Internet and from within Amazon EC2. Data stored within Amazon SQS are not encrypted by AWS; however the user can encrypt data before it is uploaded to Amazon SQS, provided that the application utilizing the queue has a means to decrypt the message when retrieved. Encrypting messages before sending them to Amazon SQS helps protect against access to sensitive customer data by unauthorized persons, including AWS.

Amazon Simple Notification Service (Amazon SNS) Security

Amazon Simple Notification Service (Amazon SNS) is a web service that makes it easy to set up, operate, and send notifications from the cloud. It provides developers with a highly scalable, flexible, and cost-effective capability to publish messages from an application and immediately deliver them to subscribers or other applications.

Amazon SNS provides a simple web services interface that can be used to create topics that customers want to notify applications (or people) about, subscribe clients to these topics, publish messages, and have these messages delivered over clients' protocol of choice (i.e., HTTP, email, etc.). Amazon SNS delivers notifications to clients using a "push" mechanism that eliminates the need to periodically check or "poll" for new information and updates. Amazon SNS can be leveraged to build highly reliable, event-driven workflows and messaging applications without the need for complex middleware and application management. The potential uses for Amazon SNS include monitoring applications, workflow systems, time-sensitive information updates, mobile applications, and many others. As with all Amazon Web Services, there are no up-front investments required, and you pay only for the resources you use.

Amazon SNS provides access control mechanisms so that topics and messages are secured against unauthorized access. Topic owners can set policies for a topic that restrict who can publish or subscribe to a topic. Additionally, topic owners can encrypt notifications by specifying that the delivery mechanism must be HTTPS.

Amazon SNS access is granted based on an AWS Account or a user created with AWS IAM. Once authenticated, the AWS Account has full access to all user operations. An AWS IAM user however only has access to the operations and topics which they have been granted access to via policy. By default, access to each individual topic is restricted to the AWS Account that created it. However, a customer can allow other access to a queue, using either an SNS-generated policy or a policy written by the user.

Amazon CloudWatch Security

Amazon CloudWatch is a web service that provides monitoring for AWS cloud resources, starting with Amazon EC2. It provides customers with visibility into resource utilization, operational performance, and overall demand patterns—including metrics such as CPU utilization, disk reads and writes, and network traffic.

Amazon CloudWatch requires, like all AWS Services, every request made to its control API be authenticated so only authenticated users can access and manage CloudWatch. Requests are signed with an HMAC-SHA1 signature calculated from the request and the user's private key. Additionally, the Amazon CloudWatch control API is only accessible via SSL-encrypted endpoints.

A customer can further control access to Amazon CloudWatch by creating users under their AWS Account using AWS IAM, and controlling what CloudWatch operations these users have permission to call.

Auto Scaling Security

Auto Scaling allows customers to automatically scale their Amazon EC2 capacity up or down according to conditions they define so that the number of Amazon EC2 instances they are using scales up seamlessly during demand spikes to maintain performance, and scales down automatically during demand lulls to minimize costs.

Auto Scaling requires, like all AWS Services, every request made to its control API be authenticated so only authenticated users can access and manage Auto Scaling. Requests are signed with an HMAC-SHA1 signature calculated from the request and the user's private key.

A customer can further control access to Auto Scaling by creating users under their AWS Account using AWS IAM, and controlling what Auto Scaling APIs these users have permission to call.

Amazon CloudFront Security

Amazon CloudFront requires every request made to its control API be authenticated so only authenticated users can create, modify or delete their own Amazon CloudFront distributions. Requests are signed with an HMAC-SHA1 signature calculated from the request and the user's private key. Additionally, the Amazon CloudFront control API is only accessible via SSL-encrypted endpoints.

There is no guarantee of durability of data held in Amazon CloudFront edge locations. The service may from time to time remove objects from edge locations if those objects are not requested frequently. Durability is provided by Amazon S3, which works as the origin server for Amazon CloudFront holding the original, definitive copies of objects delivered by Amazon CloudFront.

If you want control over who is able to download content from Amazon CloudFront, you can enable the service's private content feature. This feature has two components: the first controls how the Amazon CloudFront edge locations access your objects in Amazon S3. The second controls how content is delivered from the Amazon CloudFront edge location to viewers on the internet.

To control access to the original copies of your objects in Amazon S3, Amazon CloudFront allows you to create one or more "Origin Access Identities" and associate these with your distributions. When an Origin Access Identity is associated with an Amazon CloudFront distribution, the distribution will use that identity to retrieve objects from Amazon S3. You can then use Amazon S3's ACL feature, which limits access to that Origin Access Identity so the original copy of the object is not public readable.

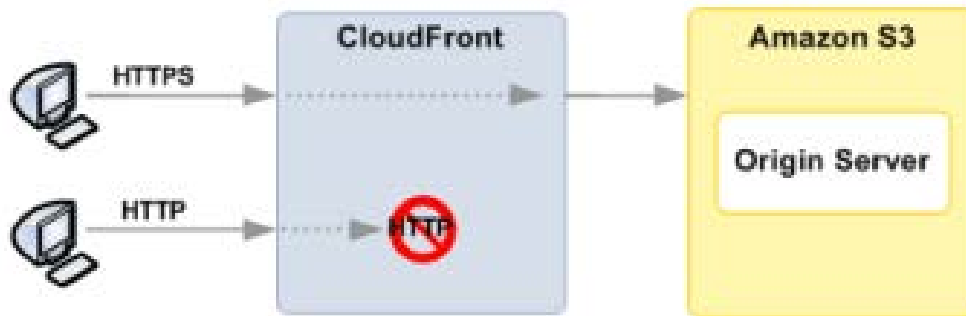
To control who is able to download your objects from Amazon CloudFront edge locations, the service uses a signed-URL verification system. To use this system, you first create a private-key public-key pair, and upload the public key to your account via the Amazon Web Services website. Second, you configure your Amazon CloudFront distribution to indicate which accounts you would authorize to sign requests – you can indicate up to five AWS Accounts you trust to sign requests. Third, as you receive requests you will create policy documents indicating the conditions under which you want Amazon CloudFront to serve your content. These policy documents can specify the name of the object that is requested, the date and time of the request, and the source IP (or CIDR range) of the client making the request. You then calculate the RSA-SHA1 encoding of your policy document and sign this using your private key. Fourth, you include both the encoded policy document and the signature as query string parameters when you reference your objects. When Amazon CloudFront receives a request, it will decode the signature using your public key. Amazon CloudFront will only serve requests that have valid policy document and matching signature.

Note that private content is an optional feature that must be enabled when you set up your CloudFront distribution. Content delivered without this feature enabled will be publicly readable by anyone.



Amazon CloudFront also provides the ability to transfer content over an encrypted connection (HTTPS) to authenticate the content delivered to your users. By default Amazon CloudFront will accept requests over both HTTP and HTTPS protocols.

If you prefer, you can also configure Amazon CloudFront to require HTTPS for all requests and disallow all HTTP requests. For HTTPS requests, Amazon CloudFront will also utilize HTTPS to retrieve your object from Amazon S3, so that your object is encrypted whenever it is transmitted.



Amazon CloudFront Access logs contain a comprehensive set of information about requests for content, including the object requested, the date and time of the request, the edge location serving the request, the client IP address, the referrer, and the user agent. To enable access logs just specify the name of the Amazon S3 bucket to store the logs in when you configure your Amazon CloudFront distribution.

Amazon Elastic MapReduce (Amazon EMR) Security

Amazon Elastic MapReduce requires every request made to its API be authenticated so only authenticated users can create, lookup, or terminate their job flows. Requests are signed with an HMAC-SHA1 signature calculated from the request and the user's private key. Amazon Elastic MapReduce provides SSL endpoints for access to its web service APIs and the console.

When launching job flows on behalf of a customer, Amazon Elastic MapReduce sets up an Amazon EC2 security group of the master node to only allow external access via SSH. The service creates a separate security group of the slaves which does not allow any external access. To protect customer input and output datasets, Amazon Elastic MapReduce transfers data to and from S3 using SSL.

APPENDIX – GLOSSARY OF TERMS

AMI: An Amazon Machine Image (AMI) is an encrypted machine image stored in Amazon S3. It contains all the information necessary to boot instances of a customer's software.

API: Application Programming Interface (API) is an interface in computer science that defines the ways by which an application program may request services from libraries and/or operating systems.

Authentication: Authentication is the process of determining whether someone or something is, in fact, who or what it is declared to be.

Availability Zone: Amazon EC2 locations are composed of regions and Availability Zones. Availability Zones are distinct locations that are engineered to be insulated from failures in other Availability Zones and provide inexpensive, low latency network connectivity to other Availability Zones in the same region.

Bucket: A container for objects stored in Amazon S3. Every object is contained within a bucket. For example, if the object named photos/puppy.jpg is stored in the johnsmith bucket, then it is addressable using the URL <http://johnsmith.s3.amazonaws.com/photos/puppy.jpg>.

CIDR Block: Classless Inter-Domain Routing Block of IP addresses.

EBS: Amazon Elastic Block Store (EBS) provides block level storage volumes for use with Amazon EC2 instances. Amazon EBS volumes are off-instance storage that persists independently from the life of an instance.

HMAC-SHA1: In cryptography, a keyed-Hash Message Authentication Code (HMAC or KMAC), is a type of message authentication code (MAC) calculated using a specific algorithm involving a cryptographic hash function in combination with a secret key. As with any MAC, it may be used to simultaneously verify both the data integrity and the authenticity of a message. Any iterative cryptographic hash function, such as MD5 or SHA-1, may be used in the calculation of an HMAC; the resulting MAC algorithm is termed HMAC-MD5 or HMAC-SHA1 accordingly. The cryptographic strength of the HMAC depends upon the cryptographic strength of the underlying hash function, on the size and quality of the key and the size of the hash output length in bits.

Hypervisor: A hypervisor, also called Virtual Machine Monitor (VMM), is computer software/hardware platform virtualization software that allows multiple operating systems to run on a host computer concurrently.

AWS IAM: AWS Identity and Access Management (AWS IAM) enables a customer to create multiple users and manage the permissions for each of these users within their AWS Account.

IP Address: An Internet Protocol (IP) address is a numerical label that is assigned to devices participating in a computer network utilizing the Internet Protocol for communication between its nodes.

IP Spoofing: Creation of Internet Protocol (IP) packets with a forged source IP address, called spoofing, with the purpose of concealing the identity of the sender or impersonating another computing system.

Object: The fundamental entities stored in Amazon S3. Objects consist of object data and metadata. The data portion is opaque to Amazon S3. The metadata is a set of name-value pairs that describe the object. These include some default metadata such as the date last modified and standard HTTP metadata such as Content-Type. The developer can also

specify custom metadata at the time the Object is stored.

Paravirtualization: In computing, paravirtualization is a virtualization technique that presents a software interface to virtual machines that is similar but not identical to that of the underlying hardware.

Port Scanning: A port scan is a series of messages sent by someone attempting to break into a computer to learn which computer network services, each associated with a "well-known" port number, the computer provides.

Service: Software or computing ability provided across a network (e.g., EC2, S3).

Stateful firewall: In computing, a stateful firewall (any firewall that performs stateful packet inspection (SPI) or stateful inspection) is a firewall that keeps track of the state of network connections (such as TCP streams, UDP communication) traveling across it.

Virtual Instance: Once an AMI has been launched, the resulting running system is referred to as an instance. All instances based on the same AMI start out identical and any information on them is lost when the instances are terminated or fail.

X.509: In cryptography, X.509 is an ITU-T standard for a Public Key Infrastructure (PKI) for Single Sign-On (SSO) and Privilege Management Infrastructure (PMI). X.509 specifies, amongst other things, standard formats for public key certificates, certificate revocation lists, attribute certificates, and a certification path validation algorithm.

Changes since last version (Aug 2010):

- Addition of AWS Identity and Access Management (AWS IAM)
- Addition of Amazon Simple Notification Service (SNS) Security
- Addition of Amazon CloudWatch Security
- Addition of Auto Scaling Security
- Update to Amazon Virtual Private Cloud (Amazon VPC)
- Update to Control Environment
- Removal of Risk Management as it has been expanded in a separate whitepaper

Changes since last version (Nov 2009):

- Major revision

Changes since last version (June 2009):

- Change to Certifications and Accreditations section to reflect SAS70
- Addition of Amazon Virtual Private Cloud (Amazon VPC)
- Addition of Security Credentials section to highlight AWS Multi-Factor Authentication and Key Rotation
- Addition of Amazon Relational Database Service (Amazon RDS) Security

Changes since last version (Sep 2008):

- Addition of Security Design Principles
- Update of Physical Security information and inclusion of background checks
- Backup section updated for clarity with respect to Amazon EBS
- Update of Amazon EC2 Security section to include:
 - Certificate-based SSHv2
 - Multi-tier security group detail and diagram
 - Hypervisor description and Instance Isolation diagram
 - Fault Separation
- Addition of Configuration Management
- Amazon S3 section updated for detail and clarity
- Addition of Storage Device Decommissioning
- Addition of Amazon SQS Security
- Addition of Amazon CloudFront Security
- Addition of Amazon Elastic MapReduce Security

Notices

© 2010-2011 Amazon.com, Inc., or its affiliates. This document is provided for informational purposes only. It represents AWS's current product offerings as of the date of issue of this document, which are subject to change without notice. Customers are responsible for making their own independent assessment of the information in this document and any use of AWS's products or services, each of which is provided "as is" without warranty of any kind, whether express or implied. This document does not create any warranties, representations, contractual commitments, conditions or assurances from AWS, its affiliates, suppliers or licensors. The responsibilities and liabilities of AWS to its customers are controlled by AWS agreements, and this document is not part of, nor does it modify, any agreement between AWS and its customers.