



Amazon Web Services: Overview of Security Processes

June 2009

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

Amazon Web Services (AWS) delivers a highly scalable cloud computing platform with high availability and dependability, offering the flexibility to enable customers to build a wide range of applications. Ensuring the confidentiality, integrity, and availability of our customers' systems and data is of the utmost importance to AWS, as is maintaining trust and confidence. This document is intended to answer questions such as "How does AWS help me ensure my data are secure?" Specifically, AWS physical and operational security processes are described for network and 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:

- Certifications and Accreditations
- Secure Design Principles
- Physical Security
- Backups
- Network Security
- AWS Security
 - Amazon Elastic Compute Cloud (Amazon EC2) Security
 - Amazon Simple Storage Service (Amazon S3) Security
 - Amazon SimpleDB Security
 - Amazon Simple Queue Service (Amazon SQS) Security
 - Amazon CloudFront Security
 - Amazon Elastic MapReduce Security

Certifications and Accreditations

To provide customers with assurance of the security measures implemented, AWS is working with a public accounting firm to ensure continued Sarbanes Oxley (SOX) compliance, and attain certifications and unbiased Audit Statements such as recurring Statement on Auditing Standards No. 70: Service Organizations, Type II (SAS70 Type II). AWS will continue efforts to obtain the strictest of industry certifications in order to verify its commitment to provide a secure, world-class cloud computing environment. The flexibility and customer control that the AWS platform provides permits the deployment of solutions that meet industry-specific certification requirements. For instance, customers have built HIPAA-compliant healthcare applications on AWS.

Secure Design Principles

Amazon's development process follows secure software development best practices, which include formal design reviews by our internal Amazon Security team, threat modeling, completion of a risk assessment, and static code analysis as well as recurring

penetration testing by carefully selected industry experts. Our security risk assessment reviews begin during the design phase and the engagement lasts through post-launch.

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, state of the art 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.

Amazon 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 Amazon employees is logged and audited routinely.

Amazon 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 access to data. Amazon understands that security and privacy of your confidential data is of paramount concern to you.

Backups

Data stored in Amazon S3, Amazon SimpleDB, or Amazon Elastic Block Store 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 ensure object durability by storing objects multiple times across multiple datacenters 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 in order to ensure long-term data durability. Snapshots of Amazon EBS can be taken without quiescing the file system, which is an important feature for those utilizing Amazon EBS for databases. AWS does not perform backups of data that are maintained on virtual disks attached to running instances on Amazon EC2.

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, Amazon'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 Amazon -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:** Port scans by Amazon EC2 customers are a violation of the Amazon EC2 Acceptable Use Policy (AUP). Violations of the AUP are taken seriously, and every reported violation is investigated. Customers can report suspected abuse via the contacts available here: <http://aws.amazon.com/contact->

[us/report-abuse/](#) When 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 ensuring the security of the HTTP server software, such as Apache.

- **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. 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.

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 in such a manner that in the vast majority of cases they will not impact the customer and their Service use. AWS will communicate with customers, either via email, or through the AWS Service Health Dashboard (<http://status.aws.amazon.com/>) when there is a chance that their Service use may be affected.

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 ensure that data contained within Amazon EC2 cannot be intercepted by unauthorized systems or users and that Amazon EC2 instances themselves 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 including: customers should disable password-based access to their hosts, and utilize 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 mode and Amazon EC2 customers must explicitly open all 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 IPsec. 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 an X.509 certificate or the customer's Amazon Secret Access Key. Without access to the customer's Secret Access Key or X.509 certificate, 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.

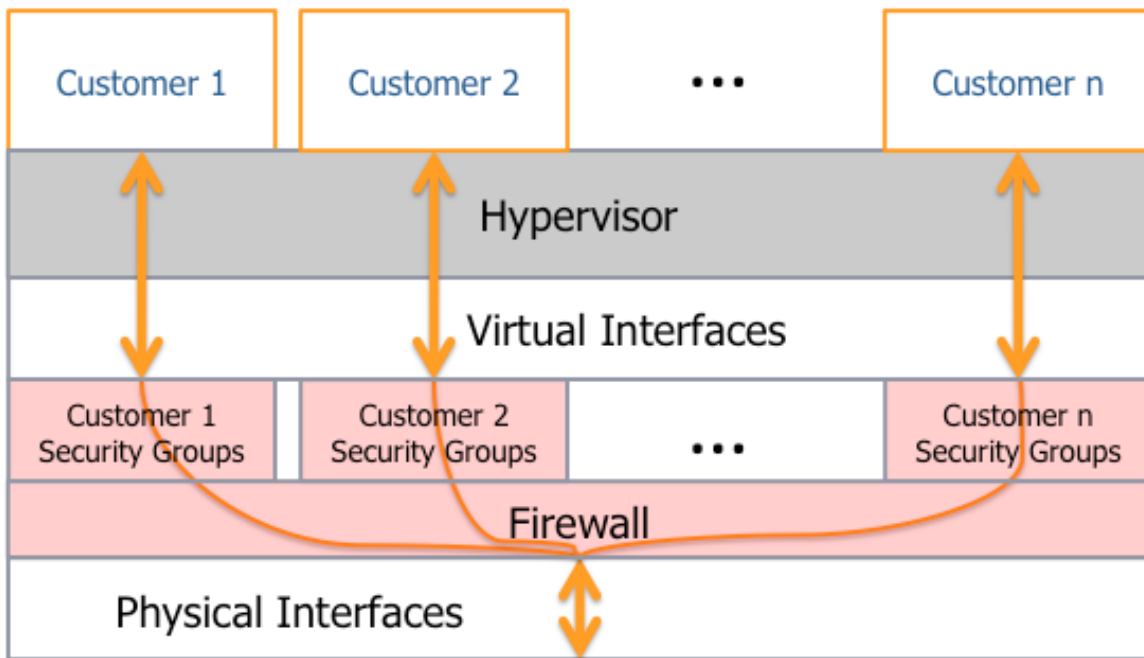
The Hypervisor

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 ensures 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.

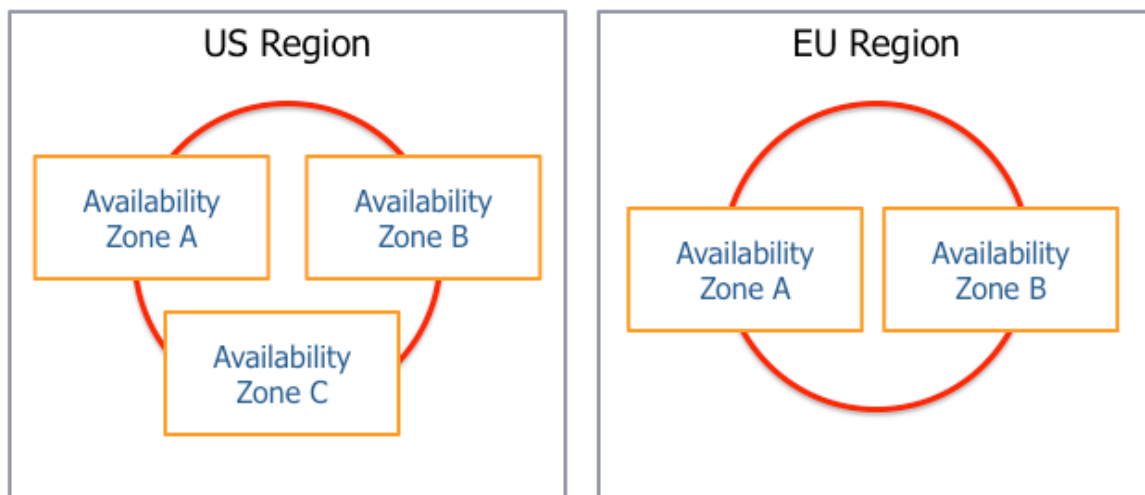


Customer instances have no access to raw disk devices, but instead are presented with virtualized disks. The AWS proprietary disk virtualization layer automatically resets every block of storage used by the customer, ensuring that one customer's data are never unintentionally exposed to another. AWS recommends that 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.

Fault Separation

Amazon EC2 provides customers the flexibility to place instances within multiple geographic regions as well as across multiple Availability Zones. Each Availability Zone is designed with fault separation. This means that Availability Zones are physically separated within a typical metropolitan region, on different flood plains, in seismically stable areas. In addition to discrete uninterruptable power source (UPS) and onsite backup generation facilities, they are each fed via different grids from independent utilities to further reduce single points of failure. They are all redundantly connected to multiple tier-1 transit providers.



It should be noted that although traffic flowing across the private networks between Availability Zones in a single region is on AWS-controlled infrastructure, all communications between regions is across public Internet infrastructure, so appropriate encryption methods should be used to protect sensitive data. Data are not replicated between regions unless proactively done so by the customer.

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. To ensure

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 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. 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 users by AWS Account ID or email, or DevPay Product ID. Customers can also grant access to their Amazon S3 data to all AWS users 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, ensuring 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 S3 datacenters is limited to an audited list of Amazon personnel. Encryption of sensitive data is generally a good security practice, and Amazon 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.

Storage Device Decommissioning

When a storage device has reached the end of its useful life, AWS procedures include a decommissioning process that ensures customer data are not 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 Sanitation”) to destroy data, as part of the decommissioning process.

Amazon 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, a subscriber has full access to all user operations. Access to each individual domain is controlled by an independent Access Control List (ACL) that maps authenticated users to the domains they own.

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 to Amazon SimpleDB helps ensure that no party, including AWS, has access to sensitive customer data.

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 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 ID. Once authenticated, a user has full access to all user operations. By default, access to each individual queue is restricted to the AWS account ID 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 ensure that no party, including AWS, has access to sensitive customer data.

Amazon CloudFront Security

Amazon CloudFront requires every request made to its control API be authenticated. This ensures that 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.

Amazon CloudFront currently is designed to deliver publicly available files. Objects delivered through Amazon CloudFront must be stored in Amazon S3 with ACL policies allowing public read access. Likewise, there are no access controls or other authentication features available through Amazon CloudFront's edge locations.

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.

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 Security

Amazon Elastic MapReduce requires every request made to its API be authenticated. This ensures that 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.

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

© 2008-2009 Amazon.com, Inc., or its affiliates. This whitepaper is provided for informational purposes only. Amazon Web Services LLC is not responsible for any damages related to the information in this whitepaper, which is provided “as is” without warranty of any kind, whether express, implied, or statutory. Nothing in this whitepaper creates any warranties or representations from Amazon Web Services LLC, its affiliates, suppliers, or licensors. This whitepaper does not modify the applicable terms and conditions governing your use of Amazon Web Services technologies, including the Amazon Web Services website. This whitepaper represents Amazon Web Services' current product offerings as of the date of issue of this document, which are subject to change without notice.