



# **Amazon Web Services: Overview of Security Processes**

**September 2008**

(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, and the flexibility to enable customers to build a wide range of applications. The issues of end-to-end security and end-to-end privacy within the cloud computing world are more sophisticated than within a single data center not facing the Internet. Ensuring the confidentiality, integrity, and availability of customer's systems and data is of the utmost importance to AWS, as is maintaining trust and confidence. This document is intended to answer customer questions such as "How does AWS help me ensure my data is 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
- Physical Security
- Backups
- Amazon Elastic Compute Cloud (EC2) Security
- Amazon Simple Storage Service (S3) Security
- Amazon SimpleDB Security

## **Certifications and Accreditations**

AWS is working with a public accounting firm to ensure continued Sarbanes Oxley (SOX) compliance and attain certifications such as recurring Statement on Auditing Standards No. 70: Service Organizations, Type II (SAS70 Type II) certification. These certifications provide outside affirmation that AWS has established adequate internal controls and that those controls are operating efficiently. 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 AWS platform also permits the deployment of solutions which meet industry-specific certification requirements. For instance, AWS customers have built HIPAA-compliant healthcare applications using S3 and other components.

## **Physical Security**

Amazon has many years of experience in designing, constructing, and operating large-scale data centers. This experience has been applied to the AWS platform and infrastructure. AWS data centers are housed in nondescript facilities, and critical facilities have extensive setback and military grade perimeter control berms as well as other natural boundary protection. 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 no fewer than three times to access

data center floors. All visitors and contractors are required to present identification and are signed in and continually escorted by authorized staff.

Amazon only provides data center access and information to employees 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 and electronic access to data centers by Amazon employees is logged and audited routinely.

## **Backups**

Data stored in Amazon S3, Amazon SimpleDB, or Amazon Elastic Block Store is redundantly stored in multiple physical locations as a normal part of those services and at no additional charge. Data that is maintained within running instances on Amazon EC2, or within Amazon S3 and Amazon SimpleDB, is all customer data and therefore AWS does not perform backups.

## **Amazon Elastic Compute Cloud (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 stateful 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 non-authorized systems or users and that Amazon EC2 instances themselves are as secure as possible without sacrificing the flexibility in configuration that customers demand.

Further details are provided below:

- **Host Operating System:** AWS administrators with a business need are required to use their individual cryptographically strong SSH keys to gain access to a bastion host. These bastion hosts are specifically built systems that are designed and configured to protect the management plane of the cloud. Once connected to the bastion, authorized administrators are able to use a privilege escalation command to gain access to an individual host. All such accesses are logged and routinely audited. When an AWS employee no longer has a business need to

administer EC2 hosts, their privileges on and access to the bastion hosts are revoked.

- **Guest Operating System:** Virtual instances are completely controlled by the customer. They have full root access and all administrative control over additional accounts, services, and applications. AWS administrators do not have access to customer instances, and cannot log into the guest OS. Customers should disable password-based access to their hosts and utilize token or key-based authentication to gain access to unprivileged accounts. Further, customers should employ a privilege escalation mechanism with logging on a per-user basis. For example, if the guest OS is Linux, utilize SSH with keys to access the virtual instance, enable shell command-line logging, and use the 'sudo' utility 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 the Amazon EC2 customer must explicitly open any ports to allow inbound traffic. The traffic may be restricted by protocol, by service port, as well as by source IP address (individual IP or CIDR block).

The firewall can be configured in groups permitting different classes of instances to have different rules, for example the case of a traditional three-tiered web application. The group for the web servers would have port 80 (HTTP) and port 443 (HTTPS) open to the world. 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.

The firewall is controlled not by the host/instance itself, but requires the customer's X.509 certificate and key to authorize changes, thus adding an extra

layer of security. Within EC2, the host administrator and cloud administrator can be separate people, permitting two man rule security policies to be enforced. In addition, AWS encourages customers to apply additional per-instance filters with host-based firewalls such as IPtables. This can restrict both inbound and outbound traffic on each instance.

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 developers should plan carefully what they will open when building and securing their applications. Well-informed traffic management and security design is still required on a per-instance basis.

- **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 their behalf. In addition, API calls can be encrypted in transit 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. Because paravirtualized guests rely on the hypervisor to provide support for operations that normally require privileged access, it is possible to run the guest OS with no elevated access to the CPU. This explicit virtualization of the physical resources leads to a clear separation between guest and hypervisor, resulting in strong security separation between the two.

### ***Instance Isolation***

Different instances running on the same physical machine are isolated from each other utilizing the Xen hypervisor. Amazon is an active participant and contributor within the Xen community, which ensures awareness of potential pending issues. In addition, the aforementioned firewall resides within the hypervisor layer, between the physical

interface and the instance's virtual interface. All packets must pass through this layer, thus an instance's neighbors have no additional access to that instance, 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 wipes every block of storage used by the customer, and guarantees that one customer's data is never exposed to another. Note that unintentionally leaving data on disk devices is only one possible breach of confidentiality; many others exist, and for this reason AWS recommends that customers further protect their data using appropriate means. One common solution is to run an encrypted filesystem on top of the virtualized disk device.

### ***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 API endpoints are hosted on the same Internet-scale, world class infrastructure that supports the Amazon.com retail site. Standard DDoS mitigation techniques such as syn cookies and connection limiting are used. To further mitigate the effect of potential DDoS attacks, Amazon maintains internal bandwidth which exceeds its provider-supplied Internet bandwidth.
- **Man In the Middle (MITM) Attacks:** All of the AWS APIs are available via SSL-protected endpoints which provides server authentication. Amazon EC2 AMIs automatically generate new SSH host keys on first boot and log them to the console. Customers can then use the secure APIs to call the console and access the host keys before logging into the instance for the first time. Customers are encouraged to use the SSL endpoints for all of their interactions with AWS.

- **IP Spoofing:** Amazon EC2 instances cannot send spoofed 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. 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.

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. This includes two virtual instances that are owned by the same customer, even if they are located on the same physical host. Attacks such as ARP cache poisoning do not work within 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.

## **Amazon Simple Storage Service (Amazon S3) Security**

With any shared storage system, the biggest question is whether unauthorized users can access information either intentionally or by mistake. To ensure that customers have the utmost in 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. Write and Delete permission is controlled by an Access Control List (ACL) associated with the bucket. Permission to modify the bucket ACLs is itself controlled by an ACL, and it defaults to creator-only access. Therefore, the customer maintains full control over who has access to their data. Amazon S3 access can be granted based on AWS Account ID, DevPay Product ID, or open to everyone.

### ***Data Management***

Another potential concern is whether or not data can be intercepted while “in transit” from one node on the Internet to AWS. For maximum security, Amazon S3 is accessible via SSL encrypted endpoints. The encrypted endpoints are accessible from both the Internet and from within Amazon EC2, ensuring that data is transferred securely both within AWS and to and from sources outside of AWS.

Customers may wish to secure data even when it is being stored within Amazon S3. Data stored within Amazon S3 is not encrypted at rest by AWS. However, users can encrypt their data before it is uploaded to Amazon S3 so that the data cannot be accessed or tampered with by unauthorized parties.

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 external access to the deleted object. That storage area is then made available only for write operations and the data is overwritten by newly stored data.



## **Amazon SimpleDB Security**

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

SimpleDB access can be granted based on an AWS Account ID. Once authenticated, a subscriber has full access to all user operations in the system. Access to each individual domain is controlled by an independent Access Control List (ACL) that maps authenticated users to the domains they own.

SimpleDB is accessible via SSL-encrypted endpoints. The encrypted endpoints are accessible from both the Internet and from within EC2. Data stored within SimpleDB is not encrypted by AWS; however the user can encrypt data before it is uploaded to 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 SimpleDB guarantees that no party, including AWS, has access to sensitive customer data.

### ***SimpleDB Data Management***

When a domain is deleted from 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 external access to the deleted domain.

When item and attribute data is 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 external access to the deleted data. That storage area is then made available only for write operations and the data is overwritten by newly stored data.