



**Creating Healthcare Data Applications  
to Promote HIPAA and HITECH Compliance**

*August 2012*

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

## Abstract

This paper briefly outlines how companies can use Amazon Web Services (AWS) to power information processing systems that facilitate HIPAA and HITECH compliance. We will focus on HIPAA's Standards for Privacy of Individually Identifiable Health Information (the Privacy Rule) and HIPAA's Security Standards for the Protection of Electronic Protected Health Information (the Security Rule), and how to encrypt and protect data in the AWS cloud.

## Introduction

In the U.S., certain organizations, called covered entities, that create, maintain, transmit, use, and disclose an individual's protected health information (PHI) are required to meet Health Insurance Portability and Accountability Act of 1996 (HIPAA) requirements. HIPAA was expanded by the Health Information Technology for Economic and Clinical Health Act, which is Title XIII of the American Recovery and Reinvestment Act (HITECH), in 2009. HIPAA and HITECH establish a set of federal standards intended to protect the security and privacy of PHI. These standards affect the use and disclosure of PHI by covered entities (such as health care providers engaged in certain electronic transactions, health plans, and health care clearinghouses) and their business associates. HIPAA and HITECH impose requirements related to the use and disclosure of PHI, appropriate safeguards to protect PHI, individual rights, and administrative responsibilities. For additional information on HIPAA and HITECH, visit <http://www.hhs.gov/ocr/hipaa>.

Covered entities and their business associates subject to HIPAA and HITECH can utilize the secure, scalable, low-cost, IT infrastructure provided by Amazon Web Services (AWS) as part of building applications designed to promote compliance with HIPAA and HITECH. AWS offers a complete set of infrastructure and application services that enable businesses to deploy applications and services cost-effectively and with flexibility, scalability, and reliability. AWS offers a secure, durable technology platform with industry-recognized certifications and audits: PCI DSS Level 1, ISO 27001, FISMA Moderate, and SOC 1/SSAE 16/ISAE 3402. AWS services and data centers have multiple layers of operational and physical security to ensure the integrity and safety of your data. With no minimum fees, no term-based contracts required, and pay-as-you-use pricing, AWS is a reliable and effective solution for growing health care industry applications.

## What Are HIPAA and HITECH and Why Are They Important?

HIPAA and HITECH provide national minimum standards to protect an individual's PHI. HIPAA originally was created to streamline healthcare processes and reduce costs by standardizing certain common health care transactions, while protecting the security and privacy of individuals' PHI. HITECH expanded on the privacy and security requirements of HIPAA. The U.S. Department of Health and Human Services (HHS) manages and enforces these standards.

HIPAA and HITECH focus on PHI, which generally includes any personally identifiable information regarding an individual's physical or mental health, the provision of health care to him or her, or payment for related services. PHI also includes any personally identifiable demographic information, including, for example, name, address, phone numbers, and Social Security numbers.

To promote compliance, covered entities must design their information systems and applications to meet HIPAA's and HITECH's privacy and security standards and related administrative, technical, and physical safeguards.

## Privacy and Security Rules

---

HIPAA's Privacy Rule restricts uses and disclosures of PHI, creates individual rights with respect to their PHI, and mandates administrative requirements. Among other requirements, the privacy rule requires a covered entity to reasonably safeguard PHI from any intentional or unintentional use or disclosure that is in violation of the requirements of HIPAA.

HIPAA's Security Rule requires covered entities to ensure the confidentiality, integrity, and availability of its electronic PHI, to protect against reasonably anticipated threats or hazards to the security or integrity of its electronic PHI, to protect against reasonably anticipated impermissible uses and disclosure of its electronic PHI, and to ensure compliance by their workforce. Additionally, the Security Rule requires covered entities to put in place detailed administrative, physical, and technical safeguards to protect electronic PHI. To do this, covered entities are required to implement access controls and set up back-up and audit controls for electronic PHI in a manner commensurate with the associated risk.

## Encrypting Data in the Cloud

Amazon Elastic Compute Cloud (Amazon EC2) provides resizable compute capacity in the cloud, and Amazon Simple Storage Service (Amazon S3) provides a virtually unlimited cloud-based data object store. The flexibility of AWS allows businesses to choose the programming models, languages, and operating systems they are already using or that are best suited for their project.

HIPAA's Security Rule includes addressable implementation specifications regarding the encryption of PHI in transmission ("in-flight") and in storage ("at-rest"). The same data encryption mechanisms used in a traditional computing environment, such as a local server or a managed hosting server, also can be used in a virtual computing environment, such as Amazon EC2 and Amazon S3. Amazon EC2 provides the customer with full root access and administrative control over virtual servers. To protect data security during electronic transmission, files containing PHI should be encrypted using technologies such as 256-bit AES algorithms. Furthermore, to reduce the risk of exposing PHI and to reduce bandwidth usage, any data, including PHI, not required by applications running in the cloud should be removed prior to transmission.

Using AWS, customers' system administrators can utilize token or key-based authentication to access their virtual servers. Amazon EC2 creates a 2048-bit RSA key pair, with private and public keys and a unique identifier for each key pair to help facilitate secure access. Administrators also can utilize a command-line shell interface, Secure Shell (SSH) keys, or sudo to enable additional security and privilege escalation.

A complete firewall solution can be created in the cloud by utilizing Amazon EC2's default deny-all mode, which automatically denies all inbound traffic unless the customer explicitly opens an EC2 port. Administrators can create multiple security groups to enforce different ingress policies as needed. They can control each security group with a PEM- encoded X.509 certificate and restrict traffic to each EC2 instance by protocol, service port, or source IP address. For more information on encryption and firewalls, see the [AWS Security Whitepaper](#).

When sending data to Amazon S3 for either short term or long term storage, we highly recommend encrypting data before transmission. We also recommend against putting any PHI or other sensitive data, including keys, in Amazon S3 metadata. Amazon S3 can be accessed via Secure Socket Layer (SSL)-encrypted endpoints over the Internet and from within Amazon EC2. Following these practices helps keep PHI and other sensitive data secure.

## High-Level Data Protection

While data flowing to and from the AWS cloud should be safeguarded with encryption, data that comes in contact with administrators or third-party partners may require different control mechanisms. To help customers comply with HIPAA's Security Rule, this section discusses AWS security policies and processes regarding data and how customers can implement authentication, access consent processes, and audit controls to reduce the risk of outside compromise. These controls, among others, allow customers to restrict access to their systems, carefully and constantly monitor them, and quickly lock them down in case of threat or attack.

### AWS Security Policies

---

For Amazon EC2, AWS employees do not look at customer data, do not have access to customer EC2 instances, and cannot log into the guest operating system. AWS internal security controls limit data access.

For Amazon S3, AWS employees' access to customer data is highly restricted and not necessary for customer support or maintenance. Despite these internal AWS controls, we strongly suggest that customers encrypt all sensitive data.

### Access Control Processes

---

AWS provides a number of mechanisms to control access to data while in-flight and at-rest in the AWS cloud. The customer's system administrator should set user and computer access controls to restrict data access and secure data. Using Amazon EC2, SSH network protocols can be used to authenticate remote users or computers through public-key cryptography. Public-key cryptography or key pairs are used to protect confidentiality by issuing a private key for decryption and a public key for encryption. The administrator also can allow or block access at the account or instance level and can set security groups, which restrict network access from instances not residing in that same group.

Using Amazon S3, access can be easily controlled down to the object level. The customer's system administrator maintains full control over who has access to the data at all times and the default setting permits authenticated access only to the creator. Read, write, and delete permissions are controlled by an Access Control List (ACL) associated with each object.

For both Amazon S3 and EC2, each account has a secret key that is crucial for maintaining the security of customer accounts. We recommend keeping access keys and account credentials in a secure location. Do not embed secret keys in a web page or other publicly accessible source code and do not transmit them over insecure channels. Customers should use Secure HTTP (HTTPS) connections for web applications running in the cloud to protect any PHI presented in the interface as it travels from AWS to the users' browsers.

### Auditing, Back-Ups, and Disaster Recovery

---

HIPAA's Security Rule also requires in-depth auditing capabilities, data back-up procedures, and disaster recovery mechanisms. The services in AWS contain many features that help customers address these requirements.

In designing an information system that is consistent with HIPAA and HITECH requirements, customers should put auditing capabilities in place to allow security analysts to drill down into detailed activity logs or reports to see who had access, IP address entry, what data was accessed, etc. This data should be tracked, logged, and stored in a central location for extended periods of time, in case of an audit. Using Amazon EC2, customers can run activity log files and audits down to the packet layer on their virtual servers, just as they do on traditional hardware. They also can track any IP traffic that reaches their virtual server instance. A customer's administrators can back up the log files into Amazon S3 for long-term, reliable storage.

Under HIPAA, covered entities must have a contingency plan to protect data in case of an emergency and must create and maintain retrievable exact copies of electronic PHI. To implement a data back-up plan on AWS, Amazon Elastic Block Store (EBS) offers persistent storage for Amazon EC2 virtual server instances. These volumes can be exposed as standard block devices and offer off-instance storage that persists independently from the life of an instance. To adhere to HIPAA guidelines, customers can create point-in-time snapshots of EBS volumes that automatically are stored in Amazon S3 and are replicated across multiple Availability Zones, which are distinct locations engineered to be insulated from failures in other zones (Availability Zones). These snapshots can be accessed at any time and can protect data for long-term durability. Amazon S3 also provides a highly available solution for data storage and automated back-ups. By simply loading a file or image into Amazon S3, multiple redundant copies are automatically created and stored in separate data centers. These files can be accessed at any time, from anywhere (based on permissions) and are stored until intentionally deleted by the customer's system administrator.

Disaster recovery, the process of protecting an organization's data and IT infrastructure in times of disaster, is typically one of the more expensive HIPAA requirements to comply with. It involves maintaining highly available systems, keeping both the data and system replicated off-site, and enabling continuous access to both. AWS inherently offers a variety of disaster recovery mechanisms. With Amazon EC2, administrators can start server instances very quickly and can use an Elastic IP address (a static IP address for the cloud computing environment) for elegant failure from one machine to another. Amazon EC2 also offers Availability Zones. Administrators can launch Amazon EC2 instances in multiple Availability Zones to create geographically diverse, fault tolerant systems that are highly resilient in the event of network failures, natural disasters, and most other probable sources of downtime. Using Amazon S3, a customer's data is replicated and automatically stored in separate data centers to provide reliable data storage with a service level of 99.9% availability and no single points of failure.

For more information on disaster recovery, see the [AWS Disaster Recovery Whitepaper](#).

## The AWS Solution

AWS provides a reliable, scalable, and inexpensive computing platform "in the cloud" that can support health care customers' applications in a manner consistent with HIPAA and HITECH. This platform is built on the same robust technology that Amazon.com uses to run its global web properties. Amazon EC2 offers a flexible computing environment with root access to virtual machines and the ability to scale computing resources up or down depending on demand. Amazon S3 offers a simple, reliable storage infrastructure for data, images, and back-ups. These services change the way organizations deploy, manage, and access computing resources by utilizing simple API calls and pay-as-you-use pricing. To learn more about the AWS solutions, visit <http://aws.amazon.com/solutions>.

## Disclaimer

*This white paper is not intended to constitute legal advice. You are advised to seek the advice of legal counsel regarding compliance with HIPAA, HITECH, and other laws that may be applicable to you and your business. AWS and its affiliated entities make no representations or warranties that your use of AWS services will assure compliance with applicable laws, including but not limited to HIPAA and HITECH.*