

Overview of AWS Security - Application Services

June 2016

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

THIS PAPER HAS BEEN ARCHIVED

For the latest technical content, see <https://docs.aws.amazon.com/security/>



Notices

This document is provided for informational purposes only. It represents AWS' current product offerings and practices 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' 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.

Archived

Application Services

Amazon Web Services offers a variety of managed services to use with your applications, including services that provide application streaming, queueing, push notification, email delivery, search, and transcoding.

Amazon CloudSearch Security

Amazon CloudSearch is a managed service in the cloud that makes it easy to set up, manage, and scale a search solution for your website. Amazon CloudSearch enables you to search large collections of data such as web pages, document files, forum posts, or product information. It enables you to quickly add search capabilities to your website without having to become a search expert or worry about hardware provisioning, setup, and maintenance. As your volume of data and traffic fluctuates, Amazon CloudSearch automatically scales to meet your needs.

An Amazon CloudSearch domain encapsulates a collection of data you want to search, the search instances that process your search requests, and a configuration that controls how your data is indexed and searched. You create a separate search domain for each collection of data you want to make searchable. For each domain, you configure indexing options that describe the fields you want to include in your index and how you want to use them, text options that define domain-specific stopwords, stems, and synonyms, rank expressions that you can use to customize how search results are ranked, and access policies that control access to the domain's document and search endpoints.

Access to your search domain's endpoints is restricted by IP address so that only authorized hosts can submit documents and send search requests. IP address authorization is used only to control access to the document and search endpoints. All Amazon CloudSearch configuration requests must be authenticated using standard AWS authentication.

Amazon CloudSearch provides separate endpoints for accessing the configuration, search, and document services:

- You use the configuration service to create and manage your search domains. The region-specific configuration service endpoints are of the form: `cloudsearch.region.amazonaws.com`. For example, `cloudsearch.us-east-1.amazonaws.com`. For a current list of supported regions, see [Regions and Endpoints](#) in the AWS General Reference. The document service endpoint is used to submit documents to the domain for indexing and is accessed through a domain-specific endpoint: `http://doc-domainname-domainid.us-east-1.cloudsearch.amazonaws.com`
- The search endpoint is used to submit search requests to the domain and is accessed through a domain-specific endpoint: <http://search-domainname-domainid.us-east-1.cloudsearch.amazonaws.com>

Note that if you do not have a static IP address, you must re-authorize your computer whenever your IP address changes. If your IP address is assigned dynamically, it is also likely that you're sharing that address with other computers on your network. This means that when you authorize the IP address, all computers that share it will be able to access your search domain's document service endpoint.

Like all AWS Services, Amazon CloudSearch requires that every request made to its control API be authenticated so only authenticated users can access and manage your CloudSearch domain. API requests are signed with an HMAC-SHA1 or HMAC-SHA256 signature calculated from the request and the user's AWS Secret Access key. Additionally, the Amazon CloudSearch control API is accessible via SSL-encrypted endpoints. You can control access to Amazon CloudSearch management functions by creating users under your AWS Account using AWS IAM, and controlling which CloudSearch operations these users have permission to perform.

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

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 is 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/HTTPS, 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. 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 transmission 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 for which they have been granted access via policy. By default, access to each individual topic is restricted to the AWS Account that created it. However, you can allow other access to SNS, using either an SNS-generated policy or a policy you write.

Amazon Simple Workflow Service (Amazon SWF) Security

The Amazon Simple Workflow Service (SWF) makes it easy to build applications that coordinate work across distributed components. Using Amazon SWF, you can structure the various processing steps in an application as “tasks” that drive work in distributed applications, and Amazon SWF coordinates these tasks in a reliable and scalable manner. Amazon SWF manages task execution dependencies, scheduling, and concurrency based on a developer's application logic. The service stores tasks, dispatches them to application components, tracks their progress, and keeps their latest state.

Amazon SWF provides simple API calls that can be executed from code written in any language and run on your EC2 instances, or any of your machines located anywhere in the world that can access the Internet. Amazon SWF acts as a coordination hub with which your application hosts interact. You create desired workflows with their associated tasks and any conditional logic you wish to apply and store them with Amazon SWF.

Amazon SWF access is granted based on an AWS Account or a user created with AWS IAM. All actors that participate in the execution of a workflow—deciders, activity workers, workflow administrators—must be IAM users under the AWS Account that owns the Amazon SWF resources. You cannot grant users associated with other AWS Accounts access to your Amazon SWF workflows. An AWS IAM user, however, only has access to the workflows and resources for which they have been granted access via policy.

Amazon Simple Email Service (Amazon SES) Security

Amazon Simple Email Service (SES) is an outbound-only email-sending service built on Amazon's reliable and scalable infrastructure. Amazon SES helps you maximize email deliverability and stay informed of the delivery status of your emails. Amazon SES integrates with other AWS services, making it easy to send emails from applications being hosted on services such as Amazon EC2.

Unfortunately, with other email systems, it's possible for a spammer to falsify an email header and spoof the originating email address so that it appears as though the email originated from a different source. To mitigate these problems, Amazon SES requires users to verify their email address or domain in order to confirm that they own it and to prevent others from using it. To verify a domain, Amazon SES requires the sender to publish a DNS record that Amazon SES supplies as proof of control over the domain. Amazon SES periodically reviews domain verification status, and revokes verification in cases where it is no longer valid.

Amazon SES takes proactive steps to prevent questionable content from being sent, so that ISPs receive consistently high-quality email from our domains and therefore view Amazon SES as a trusted email origin. Below are some of the features that maximize deliverability and dependability for all of our senders:

- Amazon SES uses content-filtering technologies to help detect and block messages containing viruses or malware before they can be sent.
- Amazon SES maintains complaint feedback loops with major ISPs. Complaint feedback loops indicate which emails a recipient marked as spam. Amazon SES provides you access to these delivery metrics to help guide your sending strategy.
- Amazon SES uses a variety of techniques to measure the quality of each user's sending. These mechanisms help identify and disable attempts to use Amazon SES for unsolicited mail, and detect other sending patterns that would harm Amazon SES's reputation with ISPs, mailbox providers, and anti-spam services.
- Amazon SES supports authentication mechanisms such as Sender Policy Framework (SPF) and DomainKeys Identified Mail (DKIM). When you authenticate an email, you provide evidence to ISPs that you own the domain. Amazon SES makes it easy for you to authenticate your emails. If you configure your account to use Easy DKIM, Amazon SES will DKIM-sign your emails on your behalf, so you can focus on other aspects of your email-sending strategy. To ensure optimal deliverability, we recommend that you authenticate your emails.

As with other AWS services, you use security credentials to verify who you are and whether you have permission to interact with Amazon SES. For information about which credentials to use, see [Using Credentials with Amazon SES](#). Amazon SES also integrates with AWS IAM so that you can specify which Amazon SES API actions a user can perform.

If you choose to communicate with Amazon SES through its SMTP interface, you are required to encrypt your connection using TLS. Amazon SES supports two mechanisms for establishing a TLS-encrypted connection: STARTTLS and TLS Wrapper. If you choose to communicate with Amazon SES over HTTP, then all communication will be protected by TLS through Amazon SES's HTTPS endpoint. When delivering email to its final

destination, Amazon SES encrypts the email content with opportunistic TLS, if supported by the receiver.

Amazon Elastic Transcoder Service Security

The Amazon Elastic Transcoder service simplifies and automates what is usually a complex process of converting media files from one format, size, or quality to another. The Elastic Transcoder service converts standard-definition (SD) or high-definition (HD) video files as well as audio files. It reads input from an Amazon S3 bucket, transcodes it, and writes the resulting file to another Amazon S3 bucket. You can use the same bucket for input and output, and the buckets can be in any AWS region. The Elastic Transcoder accepts input files in a wide variety of web, consumer, and professional formats. Output file types include the MP3, MP4, OGG, TS, WebM, HLS using MPEG-2 TS, and Smooth Streaming using fmp4 container types, storing H.264 or VP8 video and AAC, MP3, or Vorbis audio.

You'll start with one or more input files, and create transcoding jobs in a type of workflow called a transcoding pipeline for each file. When you create the pipeline you'll specify input and output buckets as well as an IAM role. Each job must reference a media conversion template called a transcoding preset, and will result in the generation of one or more output files. A preset tells the Elastic Transcoder what settings to use when processing a particular input file. You can specify many settings when you create a preset, including the sample rate, bit rate, resolution (output height and width), the number of reference and keyframes, a video bit rate, some thumbnail creation options, etc.

A best effort is made to start jobs in the order in which they're submitted, but this is not a hard guarantee and jobs typically finish out of order since they are worked on in parallel and vary in complexity. You can pause and resume any of your pipelines if necessary.

Elastic Transcoder supports the use of SNS notifications when it starts and finishes each job, and when it needs to tell you that it has detected an error or warning condition. The SNS notification parameters are associated with each pipeline. It can also use the List Jobs By Status function to find all of the jobs with a given status (e.g., "Completed") or the Read Job function to retrieve detailed information about a particular job.

Like all other AWS services, Elastic Transcoder integrates with AWS Identity and Access Management (IAM), which allows you to control access to the service and to other AWS resources that Elastic Transcoder requires, including Amazon S3 buckets and Amazon SNS topics. By default, IAM users have no access to Elastic Transcoder or to the resources that it uses. If you want IAM users to be able to work with Elastic Transcoder, you must explicitly grant them permissions.

Amazon Elastic Transcoder requires every request made to its control API be authenticated so only authenticated processes or users can create, modify, or delete their own Amazon Transcoder pipelines and presets. Requests are signed with an HMAC-SHA256 signature calculated from the request and a key derived from the user's secret key. Additionally, the Amazon Elastic Transcoder API is only accessible via SSL-encrypted endpoints.

Durability is provided by Amazon S3, where media files are redundantly stored on multiple devices across multiple facilities in an Amazon S3 region. For added protection against

users accidentally deleting media files, you can use the Versioning feature in Amazon S3 to preserve, retrieve, and restore every version of every object stored in an Amazon S3 bucket. You can further protect versions using Amazon S3 Versioning's MFA Delete feature. Once enabled for an Amazon S3 bucket, each version deletion request must include the six-digit code and serial number from your multi-factor authentication device.

Amazon AppStream Security

The Amazon AppStream service provides a framework for running streaming applications, particularly applications that require lightweight clients running on mobile devices. It enables you to store and run your application on powerful, parallel-processing GPUs in the cloud and then stream input and output to any client device. This can be a pre-existing application that you modify to work with Amazon AppStream or a new application that you design specifically to work with the service.

The Amazon AppStream SDK simplifies the development of interactive streaming applications and client applications. The SDK provides APIs that connect your customers' devices directly to your application, capture and encode audio and video, stream content across the Internet in near real-time, decode content on client devices, and return user input to the application. Because your application's processing occurs in the cloud, it can scale to handle extremely large computational loads.

Amazon AppStream deploys streaming applications on Amazon EC2. When you add a streaming application through the AWS Management Console, the service creates the AMI required to host your application and makes your application available to streaming clients. The service scales your application as needed within the capacity limits you have set to meet demand. Clients using the Amazon AppStream SDK automatically connect to your streamed application.

In most cases, you'll want to ensure that the user running the client is authorized to use your application before letting him obtain a session ID. We recommend that you use some sort of entitlement service, which is a service that authenticates clients and authorizes their connection to your application. In this case, the entitlement service will also call into the Amazon AppStream REST API to create a new streaming session for the client. After the entitlement service creates a new session, it returns the session identifier to the authorized client as a single-use entitlement URL. The client then uses the entitlement URL to connect to the application. Your entitlement service can be hosted on an Amazon EC2 instance or on [AWS Elastic Beanstalk](#).

Amazon AppStream utilizes an AWS CloudFormation template that automates the process of deploying a GPU EC2 instance that has the AppStream Windows Application and Windows Client SDK libraries installed; is configured for SSH, RDC, or VPN access; and has an elastic IP address assigned to it. By using this template to deploy your standalone streaming server, all you need to do is upload your application to the server and run the command to launch it. You can then use the Amazon AppStream Service Simulator tool to test your application in standalone mode before deploying it into production.

Amazon AppStream also utilizes the STX Protocol to manage the streaming of your application from AWS to local devices. The Amazon AppStream STX Protocol is a proprietary protocol used to stream high-quality application video over varying network conditions; it monitors network

conditions and automatically adapts the video stream to provide a low-latency and high-resolution experience to your customers. It minimizes latency while syncing audio and video as well as capturing input from your customers to be sent back to the application running in AWS.

Further Reading

<https://aws.amazon.com/security/security-resources/>

[Introduction to AWS Security Processes](#)

[Overview of AWS Security - Storage Services](#)

[Overview of AWS Security - Database Services](#)

[Overview of AWS Security - Compute Services](#)

[Overview of AWS Security - Application Services](#)

[Overview of AWS Security - Analytics, Mobile and Application Services](#)

[Overview of AWS Security – Network Services](#)

Archived