

AWS Whitepaper

Architecting for HIPAA Security and Compliance on Amazon Web Services



Architecting for HIPAA Security and Compliance on Amazon Web Services: AWS Whitepaper

Copyright © 2024 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

Amazon's trademarks and trade dress may not be used in connection with any product or service that is not Amazon's, in any manner that is likely to cause confusion among customers, or in any manner that disparages or discredits Amazon. All other trademarks not owned by Amazon are the property of their respective owners, who may or may not be affiliated with, connected to, or sponsored by Amazon.

Table of Contents

Abstract	i
Introduction	2
Encryption and protection of PHI in AWS	4
Amazon API Gateway	8
Amazon AppFlow	9
Amazon AppStream 2.0	9
Amazon Athena	10
Amazon Aurora	10
Amazon Aurora PostgreSQL	11
Amazon CloudFront	11
Lambda@Edge	12
Amazon CloudWatch	12
Amazon CloudWatch Events	12
Amazon CloudWatch Logs	12
Amazon Comprehend	13
AWS Identity and Access Management	13
Data protection and secrets management	14
Network segmentation and hardening	16
Host and image hardening	17
Multi-tenancy	17
Cross-service confused deputy prevention	18
Amazon Comprehend Medical	18
Amazon Connect	18
Amazon DocumentDB (with MongoDB compatibility)	18
Amazon DynamoDB	19
Amazon Elastic Block Store	19
Amazon EC2	20
Amazon Elastic Container Registry	20
Amazon ECS	21
Amazon EFS	21
Amazon EKS	22
Amazon ElastiCache for Redis	22
Encryption at Rest	23
Transport Encryption	24

Authentication	24
Applying ElastiCache Service Updates	24
Amazon OpenSearch Service	25
Amazon EMR	25
Amazon EventBridge	25
Amazon Forecast	26
Amazon FSx	26
Amazon GuardDuty	27
Amazon HealthLake	27
Amazon Inspector	28
Amazon Managed Service for Apache Flink	28
Amazon Data Firehose	29
Amazon Kinesis Streams	29
Amazon Kinesis Video Streams	29
Amazon Lex	30
Amazon Managed Streaming for Apache Kafka (Amazon MSK)	30
Amazon MQ	31
Amazon Neptune	32
AWS Network Firewall	32
Amazon Pinpoint	32
Amazon Polly	33
Amazon Quantum Ledger Database (Amazon QLDB)	34
Amazon QuickSight	35
Amazon RDS for MariaDB	35
Amazon RDS for MySQL	35
Amazon RDS for Oracle	36
Amazon RDS for PostgreSQL	36
Amazon RDS for SQL Server	37
Encryption at Rest	37
Transport Encryption	38
Auditing	38
Amazon Redshift	38
Amazon Rekognition	39
Amazon Route 53	39
Amazon S3 Glacier	39
Amazon S3 Transfer Acceleration	39

Amazon SageMaker	40
Amazon SNS	40
Amazon Simple Email Service (Amazon SES)	41
Amazon SQS	41
Amazon S3	42
Amazon Simple Workflow Service	43
Amazon Textract	43
Amazon Transcribe	43
Amazon Translate	44
Amazon Virtual Private Cloud	44
Amazon WorkDocs	44
Amazon WorkSpaces	45
AWS App Mesh	45
AWS Application Migration Service	46
AWS Auto Scaling	46
AWS Backup	47
AWS Batch	47
AWS Certificate Manager	48
AWS Cloud Map	49
AWS CloudFormation	50
AWS CloudHSM	50
AWS CloudTrail	50
AWS CodeBuild	51
AWS CodeDeploy	51
AWS CodeCommit	52
AWS CodePipeline	52
AWS Config	52
AWS Data Exchange	53
AWS Database Migration Service	53
AWS DataSync	54
AWS Directory Service	54
AWS Directory Service for Microsoft AD	54
Amazon Cloud Directory	54
AWS Elastic Beanstalk	55
AWS Elastic Disaster Recovery	55
AWS Fargate	56

AWS Firewall Manager	56
AWS Global Accelerator	57
AWS Glue	57
AWS Glue DataBrew	57
AWS IoT Core and AWS IoT Device Management	58
AWS IoT Greengrass	58
AWS Lambda	58
AWS Managed Services	59
AWS OpsWorks for Chef Automate	59
AWS OpsWorks for Puppet Enterprise	59
AWS OpsWorks Stack	60
AWS Organizations	60
AWS RoboMaker	60
AWS SDK Metrics	61
AWS Secrets Manager	61
AWS Security Hub	62
AWS Server Migration Service	62
AWS Serverless Application Repository	63
Service Catalog	63
AWS Shield	63
AWS Snowball	64
AWS Snowball Edge	64
AWS Step Functions	65
AWS Storage Gateway	65
File Gateway	65
Volume Gateway	65
Tape Gateway	65
AWS Systems Manager	66
AWS Transfer for SFTP	66
AWS WAF – Web Application Firewall	66
AWS X-Ray	67
Elastic Load Balancing	67
FreeRTOS	67
Using AWS KMS for Encryption of PHI	68
VM Import/Export	68
Auditing, backups, and disaster recovery	70

Document revisions

72

Notices

77

Architecting for HIPAA Security and Compliance on Amazon Web Services

Publication date: **September 28, 2022** ([Document revisions](#))

This paper briefly outlines how customers can use Amazon Web Services (AWS) to run sensitive workloads regulated under the U.S. Health Insurance Portability and Accountability Act (HIPAA). We will focus on the HIPAA Privacy and Security Rules for protecting Protected Health Information (PHI), how to use AWS to encrypt data in transit and at-rest, and how AWS features can be used to run workloads containing PHI.

Introduction

The Health Insurance Portability and Accountability Act of 1996 (HIPAA) applies to “covered entities” and “business associates.” HIPAA was expanded in 2009 by the Health Information Technology for Economic and Clinical Health (HITECH) Act.

HIPAA and HITECH establish a set of federal standards intended to protect the security and privacy of PHI. HIPAA and HITECH impose requirements related to the use and disclosure of protected health information (PHI), appropriate safeguards to protect PHI, individual rights, and administrative responsibilities. For more information on HIPAA and HITECH, go to the [Health Information Privacy Home](#).

Covered entities and their business associates can use the secure, scalable, low-cost IT components provided by Amazon Web Services (AWS) to architect applications in alignment with HIPAA and HITECH compliance requirements. AWS offers a commercial-off-the-shelf infrastructure platform with industry-recognized certifications and audits such as [ISO 27001](#), [FedRAMP](#), and the Service Organization Control Reports ([SOC1](#), [SOC2](#), and [SOC3](#)). AWS services and data centers have multiple layers of operational and physical security to help ensure the integrity and safety of customer 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 healthcare industry applications.

AWS enables covered entities and their business associates subject to HIPAA to securely process, store, and transmit PHI. Additionally, as of July 2013, AWS offers a standardized Business Associate Addendum (BAA) for such customers. Customers who execute an AWS BAA may use any AWS service in an account designated as a HIPAA Account, but they may only process, store and transmit PHI using the HIPAA-eligible services defined in the AWS BAA. For a complete list of these services, see the [HIPAA Eligible Services Reference](#) page.

AWS maintains a standards-based risk management program to ensure that the HIPAA-eligible services specifically support HIPAA administrative, technical, and physical safeguards. Using these services to store, process, and transmit PHI helps our customers and AWS to address the HIPAA requirements applicable to the AWS utility-based operating model.

AWS’s BAA requires customers to encrypt PHI stored in or transmitted using HIPAA-eligible services in accordance with guidance from the Secretary of Health and Human Services (HHS): [Guidance to Render Unsecured Protected Health Information Unusable, Unreadable, or Indecipherable to Unauthorized Individuals](#) (“Guidance”). Please refer to this site because it may be updated, and may be made available on a successor (or related) site designated by HHS.

AWS offers a comprehensive set of features and services to make key management and encryption of PHI easy to manage and simpler to audit, including the AWS Key Management Service (AWS KMS). Customers with HIPAA compliance requirements have a great deal of flexibility in how they meet encryption requirements for PHI.

When determining how to implement encryption, customers can evaluate and take advantage of the encryption features native to the HIPAA-eligible services. Or customers can satisfy the encryption requirements through other means consistent with the guidance from HHS.

Encryption and protection of PHI in AWS

The HIPAA Security Rule includes addressable implementation specifications for the encryption of PHI in transmission (“in transit”) and in storage (“at rest”). Although this is an addressable implementation specification in HIPAA, AWS requires customers to encrypt PHI stored in or transmitted using HIPAA-eligible services in accordance with guidance from the Secretary of Health and Human Services (HHS): [Guidance to Render Unsecured Protected Health Information Unusable, Unreadable, or Indecipherable to Unauthorized Individuals \(“Guidance”\)](#). Please refer to this site because it may be updated, and may be made available on a successor (or related site) designated by HHS.

AWS offers a comprehensive set of features and services to make key management and encryption of PHI easy to manage and simpler to audit, including the AWS Key Management Service (AWS KMS). Customers with HIPAA compliance requirements have a great deal of flexibility in how they meet encryption requirements for PHI.

When determining how to implement encryption, customers may evaluate and take advantage of the encryption features native to the HIPAA-eligible services, or they can satisfy the encryption requirements through other means consistent with the guidance from HHS. The following sections provide high-level details about using available encryption features in each of the HIPAA-eligible services and other patterns for encrypting PHI, and how AWS KMS can be used to encrypt the keys used for encryption of PHI on AWS.

Topics

- [Amazon API Gateway](#)
- [Amazon AppFlow](#)
- [Amazon AppStream 2.0](#)
- [Amazon Athena](#)
- [Amazon Aurora](#)
- [Amazon Aurora PostgreSQL](#)
- [Amazon CloudFront](#)
- [Amazon CloudWatch](#)
- [Amazon CloudWatch Events](#)
- [Amazon CloudWatch Logs](#)

- [Amazon Comprehend](#)
- [Amazon Comprehend Medical](#)
- [Amazon Connect](#)
- [Amazon DocumentDB \(with MongoDB compatibility\)](#)
- [Amazon DynamoDB](#)
- [Amazon Elastic Block Store](#)
- [Amazon Elastic Compute Cloud](#)
- [Amazon Elastic Container Registry](#)
- [Amazon Elastic Container Service](#)
- [Amazon Elastic File System \(Amazon EFS\)](#)
- [Amazon Elastic Kubernetes Service \(Amazon EKS\)](#)
- [Amazon ElastiCache for Redis](#)
- [Amazon OpenSearch Service](#)
- [Amazon EMR](#)
- [Amazon EventBridge](#)
- [Amazon Forecast](#)
- [Amazon FSx](#)
- [Amazon GuardDuty](#)
- [Amazon HealthLake](#)
- [Amazon Inspector](#)
- [Amazon Managed Service for Apache Flink](#)
- [Amazon Data Firehose](#)
- [Amazon Kinesis Streams](#)
- [Amazon Kinesis Video Streams](#)
- [Amazon Lex](#)
- [Amazon Managed Streaming for Apache Kafka \(Amazon MSK\)](#)
- [Amazon MQ](#)
- [Amazon Neptune](#)
- [AWS Network Firewall](#)

- [Amazon Pinpoint](#)
- [Amazon Polly](#)
- [Amazon Quantum Ledger Database \(Amazon QLDB\)](#)
- [Amazon QuickSight](#)
- [Amazon RDS for MariaDB](#)
- [Amazon RDS for MySQL](#)
- [Amazon RDS for Oracle](#)
- [Amazon RDS for PostgreSQL](#)
- [Amazon RDS for SQL Server](#)
- [Amazon Redshift](#)
- [Amazon Rekognition](#)
- [Amazon Route 53](#)
- [Amazon S3 Glacier](#)
- [Amazon S3 Transfer Acceleration](#)
- [Amazon SageMaker](#)
- [Amazon Simple Notification Service \(Amazon SNS\)](#)
- [Amazon Simple Email Service \(Amazon SES\)](#)
- [Amazon Simple Queue Service \(Amazon SQS\)](#)
- [Amazon Simple Storage Service \(Amazon S3\)](#)
- [Amazon Simple Workflow Service](#)
- [Amazon Textract](#)
- [Amazon Transcribe](#)
- [Amazon Translate](#)
- [Amazon Virtual Private Cloud](#)
- [Amazon WorkDocs](#)
- [Amazon WorkSpaces](#)
- [AWS App Mesh](#)
- [AWS Application Migration Service](#)
- [AWS Auto Scaling](#)

- [AWS Backup](#)
- [AWS Batch](#)
- [AWS Certificate Manager](#)
- [AWS Cloud Map](#)
- [AWS CloudFormation](#)
- [AWS CloudHSM](#)
- [AWS CloudTrail](#)
- [AWS CodeBuild](#)
- [AWS CodeDeploy](#)
- [AWS CodeCommit](#)
- [AWS CodePipeline](#)
- [AWS Config](#)
- [AWS Data Exchange](#)
- [AWS Database Migration Service](#)
- [AWS DataSync](#)
- [AWS Directory Service](#)
- [AWS Elastic Beanstalk](#)
- [AWS Elastic Disaster Recovery](#)
- [AWS Fargate](#)
- [AWS Firewall Manager](#)
- [AWS Global Accelerator](#)
- [AWS Glue](#)
- [AWS Glue DataBrew](#)
- [AWS IoT Core and AWS IoT Device Management](#)
- [AWS IoT Greengrass](#)
- [AWS Lambda](#)
- [AWS Managed Services](#)
- [AWS OpsWorks for Chef Automate](#)
- [AWS OpsWorks for Puppet Enterprise](#)

- [AWS OpsWorks Stack](#)
- [AWS Organizations](#)
- [AWS RoboMaker](#)
- [AWS SDK Metrics](#)
- [AWS Secrets Manager](#)
- [AWS Security Hub](#)
- [AWS Server Migration Service](#)
- [AWS Serverless Application Repository](#)
- [Service Catalog](#)
- [AWS Shield](#)
- [AWS Snowball](#)
- [AWS Snowball Edge](#)
- [AWS Step Functions](#)
- [AWS Storage Gateway](#)
- [AWS Systems Manager](#)
- [AWS Transfer for SFTP](#)
- [AWS WAF – Web Application Firewall](#)
- [AWS X-Ray](#)
- [Elastic Load Balancing](#)
- [FreeRTOS](#)
- [Using AWS KMS for Encryption of PHI](#)
- [VM Import/Export](#)

Amazon API Gateway

Customers can use Amazon API Gateway to process and transmit protected health information (PHI). While Amazon API Gateway automatically uses HTTPS endpoints for encryption in-flight, customers can also choose to encrypt payloads client-side. API Gateway passes all non-cached data through memory and does not write it to disk. Customers can use AWS Signature Version 4 for authorization with API Gateway. For more information, see the following:

- [Amazon API Gateway FAQs: Security and Authorization](#)

- [Controlling and managing access to a REST API in API Gateway](#)

Customers can integrate with any service that is connected to API Gateway, provided that when PHI is involved, the service is configured consistent with the Guidance and BAA. For information on integrating API Gateway with backend services, see [Set up REST API methods in API Gateway](#).

Customers can use AWS CloudTrail and Amazon CloudWatch to enable logging that is consistent with their logging requirements. Ensure that any PHI sent through API Gateway (such as in headers, URLs, and request/response) is only captured by HIPAA-eligible services that have been configured to be consistent with the Guidance. For more information on logging with API Gateway, see [How do I enable CloudWatch Logs for troubleshooting my API Gateway REST API or WebSocket API?](#)

Amazon AppFlow

Amazon AppFlow is a fully managed integration service that enables customers to securely transfer data between Software-as-a-Service (SaaS) applications such as Salesforce, Marketo, Slack, and ServiceNow, and AWS services such as Amazon S3 and Amazon Redshift. AppFlow can run data flows at a frequency the customer chooses - on a schedule, in response to a business event, or on demand. Customers can also configure data transformation capabilities like filtering and validation to generate rich, ready-to-use data as part of the flow itself, without additional steps.

Amazon AppFlow can be used to process and transfer data containing PHI. Encryption of data while in transit between AppFlow and the configured source/destination is provided by default using TLS 1.2 or later. Data stored at-rest in S3 is automatically encrypted using an AWS KMS key (formerly CMK) that is specified by the customer. For PHI data transferred to non S3 destinations, customers must ensure the at-rest storage for the chosen destination meets their security needs. AppFlow enables application monitoring by integrating with AWS CloudTrail to log API calls and Amazon EventBridge to emit flow execution events.

Amazon AppStream 2.0

Amazon AppStream 2.0 is a fully managed application streaming service. Customers own their data and must configure the necessary Windows applications in a manner that meets their regulatory requirements. Customers are able to configure persistent storage via Home Folders. Files and folders are encrypted in transit using Amazon S3's SSL endpoints. Files and folders are encrypted at-rest using Amazon S3-managed encryption keys. For more information, see [Enable and Administer Persistent Storage for Your AppStream 2.0 Users](#) If customers choose to use a

third-party storage solution they are responsible for ensuring the configuration of that solution is consistent with the guidance. All public API communication with Amazon AppStream 2.0 is encrypted using TLS. For more information, please see [Amazon AppStream 2.0 Documentation](#).

Amazon AppStream 2.0 is integrated with AWS CloudTrail, a service that logs API calls made by or on behalf of Amazon AppStream 2.0 in customer's AWS account and delivers the log files to the specified Amazon S3 bucket. CloudTrail captures API calls made from the Amazon AppStream 2.0 console or from the Amazon AppStream 2.0 API. Customers can also use Amazon CloudWatch to log resource usage metrics. For more information, see [Monitoring Amazon AppStream 2.0 Resources](#) and [Logging AppStream 2.0 API Calls with AWS CloudTrail](#).

Amazon Athena

Amazon Athena is an interactive query service that makes it easy to analyze data directly in Amazon Simple Storage Service (Amazon S3) using standard SQL. Athena helps customers analyze unstructured, semi-structured, and structured data stored in Amazon S3. Examples include CSV, JSON, or columnar data formats such as Apache Parquet and Apache ORC. Customers can use Athena to run ad hoc queries using ANSI SQL, without the need to aggregate or load the data into Athena.

Amazon Athena can now be used to process data containing PHI. Encryption of data while in transit between Amazon Athena and S3 is provided by default using SSL/TLS. Encryption of PHI while at-rest on S3 should be performed according to the guidance provided in the S3 section. Encryption of query results from and within Amazon Athena, including staged results, should be enabled using server-side encryption with Amazon S3 managed keys (SSE-S3), AWS KMS-managed keys (SSE-KMS) or client-side encryption with AWS KMS-managed keys (CSE-KMS). Amazon Athena uses AWS CloudTrail to log all API calls.

Amazon Aurora

Amazon Aurora allows customers to encrypt Aurora database clusters and snapshots at rest using keys that they manage through AWS KMS. On a database instance running with Amazon Aurora encryption, data stored at-rest in the underlying storage is encrypted, as are automated backups, read replicas, and snapshots.

Because the Guidance might be updated, customers should continue to evaluate and determine whether Amazon Aurora encryption satisfies their compliance and regulatory requirements.

For more information on encryption at-rest using Amazon Aurora, see [Protecting data using encryption](#).

Connections to DB clusters running Aurora MySQL must use transport encryption, utilizing Secure Socket Layer (SSL) or Transport Layer Security (TLS). For more information on implementing SSL/TLS, see [Using SSL/TLS with Aurora MySQL DB clusters](#).

Amazon Aurora PostgreSQL

Amazon Aurora allows customers to encrypt Aurora database clusters and snapshots at rest using keys that they manage through AWS KMS. On a database instance running with Amazon Aurora encryption, data stored at-rest in the underlying storage is encrypted, as are automated backups, read replicas, and snapshots.

Because the Guidance might be updated, customers should continue to evaluate and determine whether Amazon Aurora encryption satisfies their compliance and regulatory requirements. For more information on encryption at-rest using Amazon Aurora, see [Protecting data using encryption](#).

Connections to DB clusters running Aurora PostgreSQL must use transport encryption, utilizing Secure Socket Layer (SSL) or Transport Layer Security (TLS). For more information on implementing SSL/TLS, see [Securing Aurora PostgreSQL data with SSL](#).

Amazon CloudFront

Amazon CloudFront is a global content delivery network (CDN) service that accelerates delivery of customer websites, APIs, video content, or other web assets. It integrates with other Amazon Web Services products to give developers and businesses an easy way to accelerate content to end users with no minimum usage commitments. To ensure encryption of PHI while in transit with CloudFront, customers must configure CloudFront to use HTTPS end-to-end from the origin to the viewer.

This includes traffic between CloudFront and the viewer, CloudFront re-distributing from a custom origin, and CloudFront distributing from an Amazon S3 origin. Customers should also ensure that the data is encrypted at the origin to ensure it remains encrypted at-rest while cached in CloudFront. If using Amazon S3 as an origin, customers can make use of S3 server-side encryption features. If customers distribute from a custom origin, they must ensure that the data is encrypted at the origin.

Lambda@Edge

Lambda@Edge is a compute service that allows for the execution of Lambda functions at AWS edge locations. Lambda@Edge can be used to customize content delivered through CloudFront. When using Lambda@Edge with PHI, customers should follow the Guidance for the use of CloudFront. All connections into and out of Lambda@Edge should be encrypted using HTTPS or SSL/TLS.

Amazon CloudWatch

Amazon CloudWatch is a monitoring service for AWS Cloud resources and the applications that customers run on AWS. Customers can use Amazon CloudWatch to collect and track metrics, collect and monitor log files, and set alarms. Amazon CloudWatch itself does not produce, store, or transmit PHI. Customers can monitor CloudWatch API calls with AWS CloudTrail. For more information, see [Logging Amazon CloudWatch API Calls with AWS CloudTrail](#).

For more details on configuration requirements, see the Amazon CloudWatch Logs section.

Amazon CloudWatch Events

Amazon CloudWatch Events delivers a near-real-time stream of system events that describe changes in AWS resources. Customers should ensure that PHI does not flow into CloudWatch Events, and any AWS resource emitting a CloudWatch event that is storing, processing, or transmitting PHI is configured in accordance with the Guidance.

Customers can configure Amazon CloudWatch Events to register as an AWS API call in CloudTrail. For more information, see [Creating a CloudWatch Events Rule That Triggers on an AWS API Call Using AWS CloudTrail](#).

Amazon CloudWatch Logs

Customers can use Amazon CloudWatch Logs to monitor, store, and access their log files from Amazon Elastic Compute Cloud (Amazon EC2) instances, AWS CloudTrail, Amazon Route 53, and other sources. They can then retrieve the associated log data from CloudWatch Logs. Log data is encrypted while in transit and while it is at-rest. As a result, it is not necessary to re-encrypt PHI emitted by any other service and delivered to CloudWatch Logs.

Amazon Comprehend

Amazon Comprehend uses natural language processing to extract insights about the content of documents. Amazon Comprehend processes any text file in UTF-8 format. It develops insights by recognizing the entities, key phrases, language, sentiments, and other common elements in a document. Amazon Comprehend can be used with data containing PHI. Amazon Comprehend does not retain or store any data and all calls to the API are encrypted with SSL/TLS. Amazon Comprehend uses CloudTrail to log all API calls.

AWS Identity and Access Management

Security access functions such as authentication and authorization are required for accessing Amazon Comprehend and can be controlled with [AWS Identity and Access Management \(IAM\)](#), and credentials can be used to access the IAM. For more information, see [Authentication and Access Control for Amazon Comprehend](#) in the [Amazon Comprehend User Guide](#).

Account management

By default, IAM users don't have permission to create or modify Amazon Comprehend resources, or perform tasks using the Amazon Comprehend API. To allow users to create or modify resources and perform tasks, customers are responsible for leveraging IAM policies that grant users permissions for the specific resources (such as Amazon Comprehend and API actions) users need to use, and then attach policies to the users or groups that require specific permissions.

With Amazon Comprehend you can use AWS Identity and Access Management (IAM) to create a user with an attached policy to enable Amazon Comprehend permissions. Optionally, you can choose to create custom policies to attach to a role. Then, you can add administrators to the role with the ability to invoke API's for Amazon Comprehend administration in accordance with organization-defined role-based access and least privilege principles.

Identity and access

With Amazon Comprehend you can require user to authenticate to AWS using multi-factor authentication in accordance with their organizational requirements for authentication.

Using the AWS Management Console, IAM administrators can create a customer-managed policy that denies all permissions except those required for users to manage their own credentials and MFA devices. A JSON policy template is available on the *My Security Credential page* in the IAM console.

Optionally, you can leverage compatible third-party MFA capabilities with IAM partners. For additional information, see [IAM Partners](#).

Administration

We recommend that you Amazon Comprehend select identity-based policies in which account administrators can attach permissions policies to IAM identities (users, groups, and roles) and thereby grant permissions to perform operations on Amazon Comprehend resources.

A list of [API actions](#) for Amazon Comprehend can be found in the *API Reference guide*. You should also consider authorizing access to predefined IAM policies, customer IAM policies and API actions to users or roles in accordance with their least privilege and role-based organizational requirements. For more information, see [Using the Amazon Comprehend API](#) in the *Developer Guide*.

External authentication

Amazon Comprehend is compatible with identity federation using IAM roles. This enables Amazon Comprehend your users to authenticate to AWS by assuming a role that administrators have provisioned. Users accessing AWS using credentials from their organization or a third-party assume a role indirectly.

AWS support for Kerberos and Active Directory provides the benefits of single sign-on and centralized authentication of database users. AWS users can choose to manage and store user credentials either in AWS Directory Service for Microsoft Active Directory or in the customer on-premises Active Directory.

Data flow enforcement

AWS customers and APN partners, acting either as data controllers or data processors, are responsible for any personal data that they put in the AWS Cloud and Amazon Comprehend. You are responsible for controlling the flow to data inputs and outputs for Amazon Comprehend using IAM policies.

Data protection and secrets management

The AWS [shared responsibility model](#) applies to data protection in Amazon Comprehend. As described in this model, AWS is responsible for protecting the global infrastructure that runs all of the AWS Cloud. You are responsible for maintaining control over your content that is hosted on this

infrastructure. This content includes the security configuration and management tasks for the AWS services that you use. For more information about data privacy, see the [Data Privacy FAQ](#).

The [Data Protection in Amazon Comprehend](#) section in the [Amazon Comprehend Developer Guide](#) provides tips that you should consider in protecting data such as using TLS for transmission and avoiding placement of sensitive information into tags or free-form fields.

Encryption of data-at-rest

Amazon Comprehend works with [AWS Key Management Service](#) (AWS KMS) to provide enhanced encryption for your data. [Amazon Simple Storage Service](#) (Amazon S3) already enables you to encrypt your input documents when creating a text analysis, topic modeling, or custom Amazon Comprehend job. Integration with AWS KMS enables you to encrypt the data in the storage volume for **start*** and **create*** jobs, and it encrypts the output results of **start*** jobs using your own AWS KMS key.

It is best practice for Amazon Comprehend users to encrypt Amazon S3 buckets used for input documents using available S3 encryption solutions in accordance with their organizational policies.

The AWS Management Console, encrypts Amazon Comprehend custom models with its own AWS KMS key. For the AWS CLI, Amazon Comprehend can encrypt custom models using either its own AWS KMS key or a provided customer managed key (CMK).

If selecting encryption when using the AWS Management Console, you can choose either or both of the following optional methods:

- **Volume encryption** - ensures that the data on an EBS Volume used by Comprehend is encrypted during training/inference (data is flushed after training/inference, so this key is relevant only while the job is in progress).
- **Output result encryption** - for encrypting the output stored by comprehend in the customer's bucket using a customer provided AWS KMS key.

For more information about encryption types such as volume encryption, see [AWS KMS Encryption in Amazon Comprehend](#).

Personally identifiable information

You can use the Amazon Comprehend console or APIs to detect personally identifiable information (PII) in English text documents. For more information about detecting and labeling PII

entities and operating various PII analysis jobs, see the [Personally identifiable information](#) section in the *Amazon Comprehend Developer Guide*.

Data deletion

If you are an Amazon Comprehend customer using Amazon S3 and choosing to manage your own AWS KMS keys, you should consider revoking AWS KMS keys and defining the procedural justification to do so in accordance with their organizational requirements. Revocation of the AWS KMS key for Amazon S3 renders any data unusable/unreadable.

Network segmentation and hardening

As a managed service, Amazon Comprehend adheres to the [AWS Best Practices for Security, Identity, and Compliance](#).

For recommended network security safeguards, see [Infrastructure Security in Amazon Comprehend](#) in the *Amazon Comprehend Developers Guide*.

Protect jobs using an Amazon Virtual Private Cloud (Amazon VPC)

Amazon Comprehend uses a variety of security measures to ensure the safety of your data with our job containers where it's stored while being used by Amazon Comprehend. However, job containers access AWS resources—such as the Amazon S3 buckets where you store data and model artifacts—over the internet.

To control access to your data, we recommend that you create a virtual private cloud (VPC) and configure it so that the data and containers aren't accessible over the internet. For information about creating and configuring a VPC, see [Getting Started With Amazon VPC](#) in the *Amazon VPC User Guide*. Using a VPC helps to protect your data because you can configure your VPC so that it is not connected to the internet. Using a VPC also allows you to monitor all network traffic in and out of our job containers by using VPC flow logs. For more information, see [VPC Flow Logs](#) in the *Amazon VPC User Guide*.

You specify your VPC configuration when you create a job, by specifying the subnets and security groups. When you specify the subnets and security groups, Amazon Comprehend creates elastic network interfaces (ENIs) that are associated with your security groups in one of the subnets. ENIs allow our job containers to connect to resources in your VPC. For information about ENIs, see [Elastic Network Interfaces](#) in the *Amazon VPC User Guide*.

Note

For jobs, you can only configure subnets with a default tenancy VPC in which your instance runs on shared hardware. For more information on the tenancy attribute for VPCs, see [Dedicated Instances](#) in the *Amazon EC2 User Guide for Linux Instances*.

You can establish a private connection between your VPC and Amazon Comprehend by creating an interface VPC endpoint. For more information, see [Amazon Comprehend and Interface VPC Endpoints \(AWS PrivateLink\)](#).

Host and image hardening

Based on the AWS [shared responsibility model](#), host and image hardening of the AWS environment for Amazon Comprehend is managed by AWS as a provided service.

Multi-tenancy

To help make your recommendation more secure, we recommend that you implement the following multi-tenancy security recommendations:

- Use only a verified email address to authorize user access to a tenant based on domain match. Do not trust email addresses and phone numbers unless your app verifies them, or the external IdP gives a proof of verification. For more details on setting these permissions, see [Attribute Permissions and Scopes](#).
- Use immutable or mutable attributes for the user profile attributes that identify tenants. Administrators must be able to change these attributes. Also, give app clients read-only access to the attributes.
- Use 1:1 mapping between external IdP and application client to prevent unauthorized cross-tenant access. A user who has been authenticated by an external IdP, and who has a valid Amazon Cognito session cookie, can access other tenant apps that trust the same IdP.
- When you implement tenant-matching and authorization logic in your application, restrict users so that they can't modify the criteria that authorize user access to the tenants. Also, if an external IdP is being used for federation, restrict tenant identity provider administrators so that they can't modify user access.

Cross-service confused deputy prevention

The confused deputy problem is a multi-tenancy security issue where an entity that doesn't have permission to perform an action can coerce a more-privileged entity to perform the action. In AWS, cross-service impersonation can result in the confused deputy problem. Cross-service impersonation can occur when one service (the *calling service*) calls another service (the *called service*). The calling service can be manipulated to use its permissions to act on another customer's resources in a way it should not otherwise have permission to access. To prevent this, AWS provides tools that can help you protect your data for all services with service principals that have been given access to resources in your account. For more information which includes safeguards you should consider for addressing this security issue, see [Cross-service Confused Deputy Prevention](#) in the *Amazon Comprehend Developer Guide*.

Amazon Comprehend Medical

For guidance, see the previous [Amazon Comprehend](#) section.

Amazon Connect

Amazon Connect is a self-service, cloud-based contact center service that enables dynamic, personal, and natural customer engagement at any scale. Customers should not include any PHI in any fields associated with managing users, security profiles, and contact flows within Amazon Connect.

Amazon Connect Customer Profiles, a feature of Amazon Connect, equips contact center agents with a more unified view of a customer's profile with the most up to date information, to provide more personalized customer service. Customer Profiles is designed to automatically bring together customer information from multiple applications into a unified customer profile, delivering the profile directly to the agent as soon as the support call or interaction begins. Customers should refrain from naming domains or object keys with PHI data. The contents of Domains and Objects are encrypted and protected, but the key identifiers are not.

Amazon DocumentDB (with MongoDB compatibility)

Amazon DocumentDB (with MongoDB compatibility) (Amazon DocumentDB) offers encryption at-rest during cluster creation via AWS KMS, which allows customers to encrypt databases using AWS or customer-managed keys. On a database instance running with encryption enabled, data

stored at-rest is encrypted consistent with the Guidance in effect at the time of publication of this whitepaper, as are automated backups, read replicas, and snapshots. Because the Guidance might be updated, customers should continue to evaluate and determine whether Amazon DocumentDB encryption satisfies their compliance and regulatory requirements. For more information on encryption at-rest using Amazon DocumentDB, see [Encrypting Amazon DocumentDB Data at Rest](#).

Connections to Amazon DocumentDB containing PHI must use endpoints that accept encrypted transport (HTTPS). By default, a newly created Amazon DocumentDB cluster only accepts secure connections using Transport Layer Security (TLS). For more information, see [Encrypting Data in Transit](#). Amazon DocumentDB uses AWS CloudTrail to log all API calls. For more information, see [Logging and Monitoring in Amazon DocumentDB](#).

For certain management features, Amazon DocumentDB uses operational technology that is shared with Amazon RDS. Amazon DocumentDB console, AWS CLI, and API calls are logged as calls made to the Amazon RDS API.

Amazon DynamoDB

Connections to Amazon DynamoDB containing PHI must use endpoints that accept encrypted transport (HTTPS). For a list of regional endpoints, see [AWS service endpoints](#).

Amazon DynamoDB offers DynamoDB encryption, which allows customers to encrypt databases using keys that customers manage through AWS KMS. On a database instance running with Amazon DynamoDB encryption, data stored at-rest in the underlying storage is encrypted consistent with the Guidance in effect at the time of publication of this whitepaper, as are automated backups, read replicas, and snapshots.

Because the Guidance might be updated, customers should continue to evaluate and determine whether Amazon DynamoDB encryption satisfies their compliance and regulatory requirements. For more information on encryption at-rest using Amazon DynamoDB, see [DynamoDB Encryption at Rest](#).

Amazon Elastic Block Store

Amazon EBS encryption at-rest is consistent with the Guidance that is in effect at the time of publication of this whitepaper. Because the Guidance might be updated, customers should continue to evaluate and determine whether Amazon EBS encryption satisfies their compliance and regulatory requirements. With Amazon EBS encryption, a unique volume encryption key is generated for each EBS volume. Customers have the flexibility to choose which KMS key from

the AWS Key Management Service is used to encrypt each volume key. For more information, see [Amazon EBS encryption](#).

Amazon Elastic Compute Cloud

Amazon EC2 is a scalable, user-configurable compute service that supports multiple methods for encrypting data at rest. For example, customers might elect to perform application- or field-level encryption of PHI as it is processed within an application or database platform hosted in an Amazon EC2 instance. Approaches range from encrypting data using standard libraries in an application framework such as Java or .NET; leveraging Transparent Data Encryption features in Microsoft SQL or Oracle; or by integrating other third-party and software as a service (SaaS)-based solutions into their applications.

Customers can choose to integrate their applications running in Amazon EC2 with AWS KMS SDKs, simplifying the process of key management and storage. Customers can also implement encryption of data at rest using file-level or full disk encryption (FDE) by using third-party software from [AWS Marketplace Partners](#) or native file system encryption tools (such as dm-crypt, LUKS, etc.).

Network traffic containing PHI must encrypt data in transit. For traffic between external sources (such as the internet or a traditional IT environment) and Amazon EC2, customers should use open standard transport encryption mechanisms such as Transport Layer Security (TLS) or IPsec virtual private networks (VPNs), consistent with the [Guidance](#). Internal to an Amazon Virtual Private Cloud (VPC) for data traveling between Amazon EC2 instances, network traffic containing PHI must also be encrypted; most applications support TLS or other protocols providing in transit encryption that can be configured to be consistent with the Guidance. For applications and protocols that do not support encryption, sessions transmitting PHI can be sent through encrypted tunnels using IPsec or similar implementations between instances.

Amazon Elastic Container Registry

Amazon Elastic Container Registry (Amazon ECR) is integrated with Amazon Elastic Container Service (Amazon ECS) and allows customers to easily store, run, and manage container images for applications running on Amazon ECS. After customers specify the Amazon ECR repository in their Task Definition, Amazon ECS will retrieve the appropriate images for their applications.

No special steps are required to use Amazon ECR with container images that contain PHI. Container images are encrypted while in transit and stored encrypted while at-rest using Amazon S3 server-side encryption (SSE-S3).

Amazon Elastic Container Service

Amazon Elastic Container Service (Amazon ECS) is a highly scalable, high-performance container management service that supports Docker containers and allows customers to easily run applications on a managed cluster of Amazon EC2 instances. Amazon ECS eliminates the need for customers to install, operate, and scale their own cluster management infrastructure.

With simple API calls, customers can launch and stop Docker-enabled applications, query the complete state of their cluster, and access many familiar features like security groups, Elastic Load Balancing, EBS volumes, and IAM roles. Customers can use Amazon ECS to schedule the placement of containers across their cluster based on their resource needs and availability requirements.

Using ECS with workloads that process PHI requires no additional configuration. ECS acts as an orchestration service that coordinates the launch of containers (images for which are stored in S3) on EC2, and it does not operate with or upon data within the workload being orchestrated. Consistent with HIPAA regulations and the AWS Business Associate Addendum, PHI should be encrypted in transit and at-rest when accessed by containers launched with ECS. Various mechanisms for encrypting at-rest are available with each AWS storage option (for example, S3, EBS, and KMS). Ensuring complete encryption of PHI sent between containers may also lead customers to deploy an overlay network (such as VNS3, Weave Net or similar), in order to provide a redundant layer of encryption. Nevertheless, complete logging should also be enabled (for example, through CloudTrail), and all container instance logs should be directed to CloudWatch.

Using Firelens and AWS for Fluent Bit with workloads that process PHI requires no additional configuration, unless the logs contain PHI. If logs contain PHI, then they should not be emitted to log files, unless the disk encryption is enabled. Instead, configure your application to emit logs to standard out/error which will be automatically collected by FireLens. Similarly, do not enable file buffering for Fluent Bit, unless disk encryption is also enabled. Finally, the log destination must support encryption-in-transit; all of the AWS Service output plugins in AWS for Fluent Bit will always use TLS encryption to export logs.

Amazon Elastic File System (Amazon EFS)

Amazon Elastic File System (Amazon EFS) provides simple, scalable, elastic file storage for use with AWS Cloud services and on-premises resources. It is easy to use and offers a simple interface that allows customers to create and configure file systems quickly and easily. Amazon EFS is built to elastically scale on demand without disrupting applications, growing and shrinking automatically as customers add and remove files.

To satisfy the requirement that PHI be encrypted at-rest, two paths are available on EFS. EFS supports encryption at-rest when a new file system is created. During creation, the option for “Enable encryption of data at rest” should be selected. Selecting this option ensures that all data placed on the EFS file system will be encrypted using AES-256 encryption and AWS KMS-managed keys. Customers may alternatively choose to encrypt data before it is placed on EFS, but they are then responsible for managing the encryption process and key management.

PHI should not be used as all or part of any file name or folder name. Encryption of PHI while in transit for Amazon EFS is provided by Transport Layer Security (TLS) between the EFS service and the instance mounting the file system. EFS offers a mount helper to facilitate connecting to a file system using TLS. By default, TLS is not used and must be enabled when mounting the file system using the EFS mount helper. Ensure that the mount command contains the “-o tls” option to enable TLS encryption. Alternatively, customers who choose not to use the EFS mount helper can follow the instructions in the EFS documentation to configure their NFS clients to connect through a TLS tunnel.

Amazon Elastic Kubernetes Service (Amazon EKS)

Amazon Elastic Kubernetes Service (Amazon EKS) is a managed service that makes it easy for customers to run Kubernetes on AWS without needing to stand up or maintain their own Kubernetes control plane. Kubernetes is an open-source system for automating the deployment, scaling, and management of containerized applications. For additional Security and Compliance information, refer to the [Architecting for HIPAA Security and Compliance on Amazon EKS](#) whitepaper.

Amazon ElastiCache for Redis

Amazon ElastiCache for Redis is a Redis-compatible in-memory data structure service that can be used as a data store or cache. In order to store PHI, customers must ensure that they are running the latest HIPAA-eligible ElastiCache for Redis engine version and current generation node types. Amazon ElastiCache for Redis supports storing PHI for the following node types and Redis engine version:

- Node Types: current generation only (for example, as of the time of publication of this whitepaper, M4, M5, R4, R5, T2, T3)
- ElastiCache for Redis engine version: 3.2.6 and 4.0.10 onwards

For more information about choosing current generation nodes, see [Amazon ElastiCache pricing](#). For more information about choosing an ElastiCache for Redis engine, see [What Is Amazon ElastiCache for Redis?](#)

Customers must also ensure that the cluster and nodes within the cluster are configured to encrypt data at rest, enable transport encryption and enable authentication of Redis commands. In addition, customers must also ensure that their Redis clusters are updated with the latest 'Security' type service updates on or before the 'Recommended Apply by Date' (the date by which it is recommended the update be applied) at all times. For more information, see the sections below.

Topics

- [Encryption at Rest](#)
- [Transport Encryption](#)
- [Authentication](#)
- [Applying ElastiCache Service Updates](#)

Encryption at Rest

Amazon ElastiCache for Redis provides data encryption for its cluster to help protect the data at rest. When customers enable encryption at-rest for a cluster at the time of creation, Amazon ElastiCache for Redis encrypts data on disk and automated Redis backups. Customer data on disk is encrypted using hardware accelerated Advanced Encryption Standard (AES)-512 symmetric keys. Redis backups are encrypted through Amazon S3-managed encryption keys (SSE-S3). A S3 bucket with server-side encryption enabled will encrypt the data using hardware-accelerated Advanced Encryption Standard (AES)-256 symmetric keys before saving it in the bucket.

For more details on Amazon S3-managed encryption keys (SSE-S3), see [Protecting Data Using Server-Side Encryption with Amazon S3-Managed Encryption Keys \(SSE-S3\)](#). On an ElastiCache Redis cluster (single or multi-node) running with encryption, data stored at-rest is encrypted consistent with the Guidance in effect at the time of publication of this whitepaper. This includes data on disk and automated backups in S3 bucket. Because the Guidance might be updated, customers should continue to evaluate and determine whether Amazon ElastiCache for Redis encryption satisfies their compliance and regulatory requirements. For more information about encryption at-rest using Amazon ElastiCache for Redis, see [What Is Amazon ElastiCache for Redis?](#)

Transport Encryption

Amazon ElastiCache for Redis uses TLS to encrypt the data in transit. Connections to ElastiCache for Redis containing PHI must use transport encryption and evaluate the configuration for consistency with the Guidance. For more information, see [CreateReplicationGroup](#). For more information on enabling transport encryption, see [ElastiCache for Redis In-Transit Encryption \(TLS\)](#).

Authentication

Amazon ElastiCache for Redis clusters (single/multi node) that contain PHI must provide a Redis AUTH token to enable authentication of Redis commands. Redis AUTH is available when both encryption at-rest and encryption-in transit are enabled. Customers should provide a strong token for Redis AUTH with following constraints:

- Must be only printable ASCII characters
- Must be at least 16 characters and no more than 128 characters in length
- Cannot contain any of the following characters: '/', '"', or '@'

This token must be set from within the Request Parameter at the time of Redis replication group (single/multi node) creation and can be updated later with a new value. AWS encrypts this token using AWS Key Management Service (AWS KMS). For more information on Redis AUTH, see [ElastiCache for Redis In-Transit Encryption \(TLS\)](#).

Applying ElastiCache Service Updates

Amazon ElastiCache for Redis clusters (single/multi node) that contain PHI must be updated with the latest 'Security' type service updates on or before the 'Recommended Apply by Date.' ElastiCache offers this as a self-service feature that customers can use to apply the updates anytime on demand and in real time. Each service update comes with a 'Severity' and 'Recommended Apply by Date' and is available only for the applicable Redis replication groups.

The 'SLA Met' field in the service update feature will state whether the update was applied on or before the 'Recommended Apply by Date'. If customers choose to not apply the updates to the applicable Redis replication groups by the 'Recommended Apply by Date,' ElastiCache will not take any action to apply them. Customers can use the service updates history dashboard to review the application of updates to their Redis replication groups over time. For more information on how to use this feature, see [Self-Service Updates in Amazon ElastiCache](#).

Amazon OpenSearch Service

Amazon OpenSearch Service enables customers to run a managed OpenSearch or legacy Elasticsearch OSS cluster in a dedicated Amazon Virtual Private Cloud (Amazon VPC). When using OpenSearch Service with PHI, customers should use OpenSearch or Elasticsearch 6.0 or later. Customers should ensure PHI is encrypted at-rest and in-transit within Amazon OpenSearch Service. Customers may use AWS KMS key encryption to encrypt data at rest in their OpenSearch Service domains, which is only available for OpenSearch and Elasticsearch 5.1 or later. For more information about how to encrypt data at rest, see [Encryption of data at rest for Amazon OpenSearch Service](#).

Each OpenSearch Service domain runs in its own VPC. Customers should enable node-to-node encryption, which is available in all OpenSearch versions, and in Elasticsearch 6.0 or later. If customers send data to OpenSearch Service over HTTPS, node-to-node encryption helps ensure that their data remains encrypted as OpenSearch distributes (and redistributes) it throughout the cluster. If data arrives unencrypted over HTTP, OpenSearch Service encrypts the data after it reaches the cluster. Therefore, any PHI that enters an Amazon OpenSearch Service cluster should be sent over HTTPS. For more information, see [Node-to-node encryption for Amazon OpenSearch Service](#).

Logs from the OpenSearch Service configuration API can be captured in AWS CloudTrail. For more information, see [Monitoring Amazon OpenSearch Service API calls with AWS CloudTrail](#).

Amazon EMR

Amazon EMR deploys and manages a cluster of Amazon EC2 instances into a customer's account. For information on encryption with Amazon EMR, see [Encryption Options](#).

Amazon EventBridge

Amazon EventBridge (formerly Amazon CloudWatch Events) is a serverless event bus that enables you to create scalable event-driven applications. EventBridge delivers a stream of real-time data from event sources, such as Zendesk, Datadog, or Pagerduty, and routes that data to targets like AWS Lambda.

By default, EventBridge encrypts data using [256-bit Advanced Encryption Standard \(AES-256\)](#) under an AWS owned CMK, which helps secure customer data from unauthorized access. Customers

should ensure that any AWS resource emitting an event that is storing, processing, or transmitting PHI is configured in accordance with best practices.

Amazon EventBridge is integrated with AWS CloudTrail and customers can view the most recent events in the CloudTrail console in Event history. For more information, see [EventBridge Information in CloudTrail](#).

Amazon Forecast

Amazon Forecast is a fully managed service that uses machine learning to deliver highly accurate forecasts. Based on the same machine learning forecasting technology used by Amazon.com. Every interaction customers have with Amazon Forecast is protected by encryption. Any content processed by Amazon Forecast is encrypted with customer keys through Amazon Key Management Service, and encrypted at-rest in the AWS Region where customers are using the service.

Amazon Forecast is integrated with AWS CloudTrail, a service that provides a record of actions taken by a user, role, or an AWS service in Amazon Forecast. CloudTrail captures all API calls for Amazon Forecast as events. The calls captured include calls from the Amazon Forecast console and code calls to the Amazon Forecast API operations. If customers create a trail, customers can enable continuous delivery of CloudTrail events to an Amazon S3 bucket, including events for Amazon Forecast. For more information, see [Logging Forecast API Calls with AWS CloudTrail](#).

By default, the log files delivered by CloudTrail to their bucket are encrypted by Amazon [server-side encryption with Amazon S3-managed encryption keys \(SSE-S3\)](#). To provide a security layer that is directly manageable, customers can instead use [server-side encryption with AWS KMS-managed keys \(SSE-KMS\)](#) for their CloudTrail log files. Enabling server-side encryption encrypts the log files but not the digest files with SSE-KMS. Digest files are encrypted with [Amazon S3-managed encryption keys \(SSE-S3\)](#).

AWS Forecast imports and exports data to/from S3 buckets. When importing and exporting data from Amazon S3, customers should ensure S3 buckets are configured in a manner consistent with the guidance. For more information, see [Getting Started](#).

Amazon FSx

Amazon FSx is a fully-managed service providing feature-rich and highly-performant file systems. Amazon FSx for Windows File Server provides highly reliable and scalable file storage and is accessible over the Server Message Block (SMB) protocol. Amazon FSx for Lustre provides high-

performance storage for compute workloads and is powered by Lustre, the world's most popular high-performance file system.

Amazon FSx supports two forms of encryption for file systems, encryption of data in transit and encryption at rest. Amazon FSx for Windows File Server also supports logging of all API calls using AWS CloudTrail.

Encryption of data in transit is supported by Amazon FSx for Windows File Server on compute instances supporting SMB protocol 3.0 or newer, and by Amazon FSx for Lustre on Amazon EC2 instances that support encryption in transit. Alternatively, customers may encrypt data before storing on Amazon FSx but are then responsible for the encryption process and key management.

Encryption of data at rest is automatically enabled when creating an Amazon FSx file system, using AES-256 encryption algorithm and AWS KMS-managed keys. Data and metadata are automatically encrypted before being written to the file system, and automatically decrypted before being presented to the application. PHI should not be used in any file or folder name.

Amazon GuardDuty

Amazon GuardDuty is a managed threat detection service that continuously monitors for malicious or unauthorized behavior to help customers protect their AWS accounts and workloads. It monitors for activity such as unusual API calls or potentially unauthorized deployments that indicate a possible account compromise. Amazon GuardDuty also detects potentially compromised instances or reconnaissance by attackers.

Amazon GuardDuty continuously monitors and analyzes the following data sources: VPC Flow Logs, AWS CloudTrail event logs, and DNS logs. It uses threat intelligence feeds, such as lists of malicious IPs and domains, and machine learning to identify unexpected and potentially unauthorized and malicious activity within an AWS environment. As such, Amazon GuardDuty should not encounter any PHI as this data is not to be stored in any of the AWS based data sources listed above.

Amazon HealthLake

Amazon HealthLake enables customers in the healthcare and life sciences industries to store, transform, query, and analyze health data at petabyte scale. Customers can use Amazon HealthLake to transmit, process, and store PHI. Amazon HealthLake encrypts data at rest in customer's data stores by default. All service data and metadata is encrypted with a service owned KMS key. Per Fast Healthcare Interoperability Resources (FHIR) specifications, if a customer

deletes FHIR resource, it will only be hidden from retrieval, and will be retained by the service for versioning. When customers use `StartFHIRImportJob` API, Amazon HealthLake will enforce requirement to export data to an encrypted Amazon S3 bucket.

Amazon HealthLake encrypts data both in transit and at rest. For the encryption of data in transit, you can use AWS published API calls to access HealthLake through the network. Clients must support Transport Layer Security (TLS) 1.0 or later. We require TLS 1.2 and recommend TLS 1.3. Clients must also support cipher suites with perfect forward secrecy (PFS) such as Ephemeral Diffie-Hellman (DHE) or Elliptic Curve Ephemeral Diffie-Hellman (ECDHE). Most modern systems such as Java 7 and later support these modes. Additionally, requests must be signed by using an access key ID and a secret access key that is associated with an IAM principal. Alternatively, customers can use the AWS Security Token Service (AWS STS) to generate temporary security credentials to sign requests. For the encryption of data at rest, Amazon HealthLake encrypts data in customer's data stores with a customer-owned AWS KMS key or by a service-owned AWS KMS key by default. All service data and metadata is encrypted at rest with a service owned AWS KMS key.

Amazon HealthLake is integrated with AWS CloudTrail. CloudTrail captures all API calls to Amazon HealthLake as events, including calls made as result of interaction with AWS Management Console, command line interface (CLI), and programmatically using software development kit (SDK).

Amazon Inspector

Amazon Inspector is an automated security assessment service for customers seeking to improve their security and compliance of applications deployed on AWS. Amazon Inspector automatically assesses applications for vulnerabilities or deviations from best practices. After performing an assessment, Amazon Inspector produces a detailed list of security findings prioritized by level of severity. Customers may run Amazon Inspector on EC2 instances that contain PHI. Amazon Inspector encrypts all data transmitted over the network as well as all telemetry data stored at-rest.

Amazon Managed Service for Apache Flink

Amazon Managed Service for Apache Flink enables customers to quickly author SQL code that continuously reads, processes, and stores data in near real time. Using standard SQL queries on the streaming data, customers can construct applications that transform and provide insights into their data. Managed Service for Apache Flink supports inputs from Kinesis Data Streams and

Firehose delivery streams as sources for analytics application. If the stream is encrypted, Managed Service for Apache Flink accesses the data in the encrypted stream seamlessly with no further configuration needed. Managed Service for Apache Flink does not store unencrypted data read from Kinesis Data Streams. For more information, see [Configuring Application Input](#).

Managed Service for Apache Flink integrates with both AWS CloudTrail and Amazon CloudWatch Logs for application monitoring. For more information, see [Monitoring Tools](#) and [Working with Amazon CloudWatch Logs](#).

Amazon Data Firehose

When customers send data from their data producers to their Kinesis data stream, Amazon Kinesis Data Streams encrypts data using an AWS KMS key before storing it at-rest. When the Firehose delivery stream reads data from the Kinesis stream, Kinesis Data Streams first decrypts the data and then sends it to Firehose. Firehose buffers the data in memory based on the buffering hints specified by the customer.

It then delivers the data to the destinations without storing the unencrypted data at rest. For more information about encryption with Firehose, see [Data Protection in Amazon Data Firehose](#).

AWS provides various tools that customers can use to monitor Amazon Data Firehose, including Amazon CloudWatch metrics, Amazon CloudWatch Logs, Kinesis Agent and API logging and history. For more information, see [Monitoring Amazon Data Firehose](#).

Amazon Kinesis Streams

Amazon Kinesis Streams enables customers to build custom applications that process or analyze streaming data for specialized needs. The server-side encryption feature allows customers to encrypt data at rest. When server-side encryption is enabled, Kinesis Streams will use an AWS KMS key to encrypt the data before storing it on disks. For more information, see [Data Protection in Amazon Kinesis Data Streams](#). Connections to Amazon S3 containing PHI must use endpoints that accept encrypted transport (that is, HTTPS). For a list of regional endpoints, see [AWS service endpoints](#).

Amazon Kinesis Video Streams

Amazon Kinesis Video Streams is a fully managed AWS service that customers can use to stream live video from devices to the AWS Cloud, or build applications for real-time video processing or

batch-oriented video analytics. Server-side encryption is a feature in Kinesis Video Streams that automatically encrypts data at rest by using an AWS KMS key (formerly CMK) that is specified by the customer. Data is encrypted before it is written to the Kinesis Video Streams stream storage layer, and it is decrypted after it is retrieved from storage.

The Amazon Kinesis Video Streams SDK can be used to transmit streaming video data containing PHI. By default, the SDK uses TLS to encrypt frames and fragments generated by the hardware device on which it is installed. The SDK does not manage or affect data stored at-rest. Amazon Kinesis Video Streams uses AWS CloudTrail to log all API calls.

Amazon Lex

Amazon Lex is an AWS service for building conversational interfaces for applications using voice and text. With Amazon Lex, the same conversational engine that powers Amazon Alexa is now available to any developer, enabling customers to build sophisticated, natural language chatbots into their new and existing applications. Amazon Lex provides the deep functionality and flexibility of natural language understanding (NLU) and automatic speech recognition (ASR) so customers can build highly engaging user experiences with lifelike, conversational interactions, and create new categories of products.

Lex uses the HTTPS protocol to communicate both with clients as well as other AWS services. Access to Lex is API-driven, and appropriate IAM least privilege can be enforced. For more information, see [Data Protection in Amazon Lex](#).

Monitoring is important for maintaining the reliability, availability, and performance of customer's Amazon Lex chatbots. To track the health of Amazon Lex bots, use Amazon CloudWatch. With CloudWatch, customers can get metrics for individual Amazon Lex operations or for global Amazon Lex operations for their account. Customers can also set up CloudWatch alarms to be notified when one or more metrics exceeds a threshold that customers define. For example, customers can monitor the number of requests made to a bot over a particular time period, view the latency of successful requests, or raise an alarm when errors exceed a threshold. Lex is also integrated with AWS CloudTrail to log Lex API calls. For more information, see [Monitoring in Amazon Lex](#).

Amazon Managed Streaming for Apache Kafka (Amazon MSK)

Amazon MSK provides encryption features for data at rest and for data in-transit. For data at rest encryption, Amazon MSK cluster uses Amazon EBS server-side encryption and AWS KMS keys to

encrypt storage volumes. For data in-transit, Amazon MSK clusters have encryption enabled via TLS for inter-broker communication.

The encryption configuration setting is enabled when a cluster is created. Also, by default, in-transit encryption is set to TLS for clusters created from CLI or AWS Console. Additional configuration is required for clients to communicate with clusters using TLS encryption. Customers can change the default encryption setting by selecting the TLS/plaintext settings. For more information, see [Amazon MSK Encryption](#).

Customers can monitor the performance of customer's clusters using the Amazon MSK console, Amazon CloudWatch console, or customers can access JMX and host metrics using Open Monitoring with Prometheus, an open source monitoring solution.

Tools that are designed to read from [Prometheus](#) exporters are compatible with Open Monitoring, like: [Datadog](#), [Lenses](#), [New Relic](#), [Sumologic](#), or a Prometheus server. For details on Open Monitoring, see [Amazon MSK Open Monitoring documentation](#).

Please note that the default version of Apache Zookeeper bundled with Apache Kafka does not support encryption. However, it is important to note that communications between Apache Zookeeper and Apache Kafka brokers is limited to broker, topic, and partition state information. The only way data can be produced and consumed from an Amazon MSK cluster is over a private connection between their clients in their VPC and the Amazon MSK cluster. Amazon MSK does not support public endpoints.

Amazon MQ

Amazon MQ is a managed message broker service for Apache ActiveMQ that makes it easy to set up and operate message brokers in the cloud. Amazon MQ works with existing applications and services without the need for a customer to manage, operate, or maintain their own messaging system. To provide the encryption of PHI data while in transit, the following protocols with TLS enabled should be used to access brokers:

- AMQP
- MQTT
- MQTT over WebSocket
- OpenWire
- STOMP

- STOMP over WebSocket

Amazon MQ encrypts messages at-rest and in transit using encryption keys that it manages and stores securely. Amazon MQ uses CloudTrail to log all API calls.

Amazon Neptune

Amazon Neptune is a fast, reliable, fully managed graph database service that makes it easy to build and run applications that work with highly connected datasets. The core of Amazon Neptune is a purpose-built, high-performance graph database engine that is optimized for storing billions of relationships and querying the graph with milliseconds latency. Amazon Neptune supports the popular graph query languages Apache TinkerPop Gremlin and W3C's SPARQL.

Data containing PHI can now be retained in an encrypted instance of Amazon Neptune. An encrypted instance of Amazon Neptune can be specified only at the time of creation by choosing 'Enable Encryption' from the Amazon Neptune console. All logs, backups, and snapshots are encrypted for an Amazon Neptune encrypted instance. Key management for encrypted instances of Amazon Neptune is provided through the AWS KMS. Encryption of data in transit is provided through SSL/TLS. Amazon Neptune uses CloudTrail to log all API calls.

AWS Network Firewall

AWS Network Firewall is a managed firewall service that makes it easy to deploy essential network protections for all your Amazon Virtual Private Cloud (Amazon VPC). The service automatically scales with network traffic volume to provide high-availability protections without the need to set up or maintain the underlying infrastructure. Both customer rules and access logs may contain end user IP addresses, which are encrypted both at rest and in transit within the AWS architecture. Furthermore, AWS Network Firewall encrypts all data at rest and in transit between component AWS services (Amazon S3, Amazon DynamoDB, Amazon CloudWatch Logs, Amazon EBS). The service automatically encrypts data without requiring special configuration.

Amazon Pinpoint

Amazon Pinpoint offers developers a single API layer, CLI support, and client-side SDK support to extend application communication channels with users. The eligible channels include: email, SMS text messaging, mobile push notifications, and custom channels. Amazon Pinpoint also provides an analytics system that tracks app user behavior and user engagement. With this service, developers

can learn how each user prefers to engage and can personalize the user's experience to increase user satisfaction.

Amazon Pinpoint also helps developers address multiple messaging use cases, such as direct or transactional messaging, targeted or campaign messaging, and event-based messaging. By integrating and enabling all end-user engagement channels via Amazon Pinpoint, developers can create a 360-degree view of user engagement across all customer touch points. Amazon Pinpoint stores user, endpoint, and event data so customers can create segments, send messages to recipients, and capture engagement data.

Amazon Pinpoint encrypts data both at-rest and in-transit. For more information, see [Amazon Pinpoint FAQs](#). While Amazon Pinpoint encrypts all data at rest and in transit, the final channel, such as SMS or email, may not be encrypted, and customers should configure any channel in a manner consistent with their requirements.

Additionally, customers who need to send PHI through the SMS channel should use a dedicated short code (5-, 6- digit origination phone numbers) for the explicit purpose of sending PHI. For more information on how to request a short code, see [Requesting Dedicated Short Codes for SMS Messaging with Amazon Pinpoint](#). Customers may also choose not to send PHI through the final channel and instead provide a mechanism to securely access PHI over HTTPS.

API calls to Amazon Pinpoint can be captured using AWS CloudTrail. The captured calls include those from the Amazon Pinpoint console and code calls to Amazon Pinpoint API operations. If customers create a trail, customers can enable continuous delivery of AWS CloudTrail events to an Amazon S3 bucket, including events for Amazon Pinpoint. If customers don't configure a trail, they can still view the most recent events by using Event history on the AWS CloudTrail console. Using the information collected by AWS CloudTrail, customers can determine that the request was made to Amazon Pinpoint, the IP address of the request, who made the request, when the request was made, and additional details. For more information, see [Logging Amazon Pinpoint API Calls with AWS CloudTrail](#).

Amazon Polly

Amazon Polly is a cloud service that converts text into lifelike speech. Amazon Polly provides simple API operations that customers can easily integrate with existing applications. Amazon Polly uses the HTTPS protocol to communicate with clients. Access to Amazon Polly is API-driven, and appropriate IAM least privilege can be enforced. For more information, see [Data Protection](#). Some examples of use cases that include PHI:

- Caregiver converts a text report containing PHI into synthesized speech so they can listen to the report while walking or performing other duties.
- Visually impaired patient is given medical guidance and consumes the guidance in the form of synthesized speech.

The final delivery channel from Amazon Polly could result in playing audio with PHI in a public space and precautions should be taken that delivery takes this into consideration. The synthesized speech output can also be sent asynchronously to an Amazon S3 bucket with encryption enabled.

When supported event activity occurs in Amazon Polly, that activity is recorded in a AWS CloudTrail event along with other AWS service events in Event History. For an ongoing record of events in a customer AWS account, including events for Amazon Polly, create a trail. A trail enables CloudTrail to deliver log files to an Amazon S3 bucket. Using the information collected by CloudTrail, customers can determine the request that was made to Amazon Polly, the IP address from which the request was made, who made the request, when it was made, and additional details.

Amazon Quantum Ledger Database (Amazon QLDB)

Amazon QLDB is a fully managed ledger database that provides a transparent, immutable, and cryptographically verifiable transaction log owned by a central trusted authority. Amazon QLDB tracks each and every application data change and maintains a complete and verifiable history of changes over time. Data containing PHI can now be retained in a QLDB instance. By default, all Amazon QLDB data in transit and at rest is encrypted. Data in transit is encrypted using TLS and data at rest is encrypted using AWS managed keys. For data protection purposes, we recommend that customers protect AWS account credentials and set up individual user accounts with AWS Identity and Access Management (IAM), so that each user is given only the permissions necessary to fulfill their job duties. For more information, see [Data Protection in Amazon QLDB](#).

Amazon QLDB is integrated with AWS CloudTrail, a service that provides a record of actions taken by a user, role, or an AWS service in QLDB. CloudTrail captures all control plane API calls for QLDB as events. The calls that are captured include calls from the QLDB console and code calls to the QLDB API operations. If customers create a trail, customers can enable continuous delivery of CloudTrail events to an Amazon Simple Storage Service (Amazon S3) bucket, including events for QLDB. If customers don't configure a trail, customers can still view the most recent events on the CloudTrail console in Event history. Using the information collected by CloudTrail, customers can determine the request that was made to QLDB, the IP address from which the request was made, who made the request, when it was made, and additional details.

Amazon QuickSight

Amazon QuickSight is a business analytics service that customers can use to build visualizations, perform ad hoc analysis, and quickly get business insights from their data. Amazon QuickSight discovers AWS data sources, enables organizations to scale to hundreds of thousands of users, and delivers responsive performance by using a robust in-memory engine (SPICE).

Customers can only use the Enterprise edition of Amazon QuickSight to work with data containing PHI as it provides support for encryption of data stored at-rest in SPICE. Data encryption is performed using AWS managed keys.

Amazon RDS for MariaDB

Amazon RDS for MariaDB allows customers to encrypt MariaDB databases using keys that they manage through AWS KMS. On a database instance running with Amazon RDS encryption, data stored at-rest in the underlying storage is encrypted consistent with the Guidance in effect at the time of publication of this whitepaper, as are automated backups, read replicas, and snapshots.

Because the Guidance might be updated, customers should continue to evaluate and determine whether Amazon RDS for MariaDB encryption satisfies their compliance and regulatory requirements. For more information on encryption at-rest using Amazon RDS, see [Encrypting Amazon RDS Resources](#).

Connections to RDS for MariaDB containing PHI must use transport encryption. For more information on enabling encrypted connections, see [Using SSL/TLS to Encrypt a Connection to a DB Instance](#).

Amazon RDS for MySQL

Amazon RDS for MySQL allows customers to encrypt MySQL databases using keys that customers manage through AWS KMS. On a database instance running with Amazon RDS encryption, data stored at-rest in the underlying storage is encrypted consistent with the Guidance in effect at the time of publication of this whitepaper, as are automated backups, read replicas, and snapshots.

Because the Guidance might be updated, customers should continue to evaluate and determine whether Amazon RDS for MySQL encryption satisfies their compliance and regulatory requirements. For more information on encryption at-rest using Amazon RDS, see [Encrypting Amazon RDS Resources](#).

Connections to RDS for MySQL containing PHI must use transport encryption. For more information on enabling encrypted connections, see [Using SSL/TLS to Encrypt a Connection to a DB Instance](#).

Amazon RDS for Oracle

Customers have several options for encrypting PHI at-rest using Amazon RDS for Oracle. Customers can encrypt Oracle databases using keys that they manage through AWS KMS. On a database instance running with Amazon RDS encryption, data stored at-rest in the underlying storage is encrypted consistent with the Guidance in effect at the time of publication of this whitepaper, as are automated backups, read replicas, and snapshots.

Because the Guidance might be updated, customers should continue to evaluate and determine whether Amazon RDS for Oracle encryption satisfies their compliance and regulatory requirements. For more information on encryption at-rest using Amazon RDS, see [Encrypting Amazon RDS Resources](#).

Customers can also use Oracle Transparent Data Encryption (TDE), and they should evaluate the configuration for consistency with the Guidance. Oracle TDE is a feature of the Oracle Advanced Security option available in Oracle Enterprise Edition. This feature automatically encrypts data before it is written to storage and automatically decrypts data when the data is read from storage. Customers can also use AWS CloudHSM to store Amazon RDS Oracle TDE keys. For more information, see the following:

- Amazon RDS for Oracle Transparent Data Encryption: [Oracle Transparent Data Encryption](#).
- Using AWS CloudHSM to store Amazon RDS Oracle TDE keys: [What Is Amazon Relational Database Service \(Amazon RDS\)?](#)

Connections to Amazon RDS for Oracle containing PHI must use transport encryption and evaluate the configuration for consistency with the Guidance. This is accomplished using Oracle Native Network Encryption and enabled in Amazon RDS for Oracle option groups. For detailed information, see [Oracle Native Network Encryption](#).

Amazon RDS for PostgreSQL

Amazon RDS for PostgreSQL allows customers to encrypt PostgreSQL databases using keys that customers manage through AWS KMS. On a database instance running with Amazon RDS

encryption, data stored at-rest in the underlying storage is encrypted consistent with the Guidance in effect at the time of publication of this whitepaper, as are automated backups, read replicas, and snapshots.

Because the Guidance might be updated, customers should continue to evaluate and determine whether Amazon RDS for PostgreSQL encryption satisfies their compliance and regulatory requirements. For more information on encryption at-rest using Amazon RDS, see [Encrypting Amazon RDS Resources](#).

Connections to RDS for PostgreSQL containing PHI must use transport encryption. For more information on enabling encrypted connections, see [Using SSL/TLS to Encrypt a Connection to a DB Instance](#).

Amazon RDS for SQL Server

RDS for SQL Server supports storing PHI for the following version and edition combinations:

- 2008 R2 - Enterprise Edition only
- 2012, 2014 and 2016 - Web, Standard and Enterprise Editions

Important: SQL Server Express edition is not supported and should never be used for the storage of PHI.

In order to store PHI, customers must ensure that the instance is configured to encrypt data at rest, and enable transport encryption and auditing, as detailed below.

Encryption at Rest

Customers can encrypt SQL Server databases using keys that they manage through AWS KMS. On a database instance running with Amazon RDS encryption, data stored at-rest in the underlying storage is encrypted consistent with the Guidance in effect at the time of publication of this whitepaper, as are automated backups, and snapshots. Because the Guidance might be updated, customers should continue to evaluate and determine whether Amazon RDS for SQL Server encryption satisfies their compliance and regulatory requirements. For more information about encryption at-rest using Amazon RDS, see [Encrypting Amazon RDS Resources](#).

If customers use SQL Server Enterprise Edition, they can use Server Transparent Data Encryption (TDE) as an alternative. This feature automatically encrypts data before it is written to storage and

automatically decrypts data when the data is read from storage. For more information on RDS for SQL Server Transparent Data Encryption, see [Support for Transparent Data Encryption in SQL Server](#).

Transport Encryption

Connections to Amazon RDS for SQL Server containing PHI must use transport encryption provided by SQL Server Forced SSL. Forced SSL is enabled from within the parameter group for Amazon RDS SQL Server. For more information on RDS for SQL Server Forced SSL, see [Using SSL with a Microsoft SQL Server DB Instance](#).

Auditing

RDS for SQL Server instances that contain PHI must have auditing enabled. Auditing is enabled from within the parameter group for Amazon RDS SQL Server. For more information on RDS for SQL Server auditing, see [Compliance Program Support for Microsoft SQL Server DB Instances](#).

Amazon Redshift

Amazon Redshift provides database encryption for its clusters to help protect data at rest. When customers enable encryption for a cluster, Amazon Redshift encrypts all data, including backups, by using hardware-accelerated Advanced Encryption Standard (AES)-256 symmetric keys. Amazon Redshift uses a four-tier, key-based architecture for encryption. These keys consist of data encryption keys, a database key, a cluster key, and a KMS key.

The cluster key encrypts the database key for the Amazon Redshift cluster. Customers can use either AWS KMS or an AWS CloudHSM (Hardware Security Module) to manage the cluster key. Amazon Redshift encryption at-rest is consistent with the Guidance that is in effect at the time of publication of this whitepaper. Because the Guidance might be updated, customers should continue to evaluate and determine whether Amazon Redshift encryption satisfies their compliance and regulatory requirements. For more information, see [Amazon Redshift database encryption](#).

Connections to Amazon Redshift containing PHI must use transport encryption and customers should evaluate the configuration for consistency with the Guidance. For more information, see [Configuring security options for connections](#). Amazon Redshift Spectrum enables customers to run Amazon Redshift SQL queries against exabytes of data in Amazon S3. Redshift Spectrum is a feature of Amazon Redshift, and thus is also in scope for the HIPAA BAA.

Amazon Rekognition

Amazon Rekognition makes it easy to add image and video analysis to customer applications. A customer only needs to provide an image or video to the Amazon Rekognition API, and the service can identify the objects, people, text, scenes, and activities, as well as detect any inappropriate content. Amazon Rekognition also provides highly accurate facial analysis and facial recognition.

Amazon Rekognition is eligible to operate with images or video containing PHI. Amazon Rekognition operates as a managed service and does not present any configurable options for the handling of data. Amazon Rekognition only uses, discloses, and maintains PHI as permitted by the terms of the AWS BAA. All data is encrypted at-rest and in transit with Amazon Rekognition. Amazon Rekognition uses AWS CloudTrail to log all API calls.

Amazon Route 53

Amazon Route 53 is a managed DNS service that provides customers the ability to register domain names, route internet traffic customer domain resources and check the health of those resources. While Amazon Route 53 is a HIPAA Eligible Service, no PHI should be stored in any resource names or tags within Amazon Route 53 as there is no support for encrypting such data. Instead, Amazon Route 53 can be used to provide access to customer domain resources that transmit or store PHI such as web servers running on Amazon EC2 or storage such as Amazon S3.

Amazon S3 Glacier

Amazon S3 Glacier automatically encrypts data at rest using AES 256-bit symmetric keys and supports secure transfer of customer data over secure protocols. Connections to Amazon S3 Glacier containing PHI must use endpoints that accept encrypted transport (HTTPS). For a list of regional endpoints, see [AWS service endpoints](#).

Do not use PHI in archive and vault names or metadata because this data is not encrypted using Amazon S3 Glacier server-side encryption and is not generally encrypted in client-side encryption architectures.

Amazon S3 Transfer Acceleration

Amazon S3 Transfer Acceleration (S3TA) enables fast, easy, and secure transfers of files over long distances between a customer's client and an S3 bucket. Transfer Acceleration takes advantage of Amazon CloudFront's globally distributed edge locations. As the data arrives at an edge location,

data is routed to Amazon S3 over an optimized network path. Customers should ensure that any data containing PHI transferred using AWS S3TA is encrypted in transit and at-rest. Refer to the Guidance for Amazon S3 to understand the available encryption options.

Amazon SageMaker

Amazon SageMaker is a fully managed machine learning service. With Amazon SageMaker, data scientists and developers can quickly and easily build and train machine learning models, and then directly deploy them into a production-ready hosted environment. It provides an integrated Jupyter authoring notebook instance for easy access to data sources for exploration and analysis. Amazon SageMaker also provides common machine learning algorithms that are optimized to run efficiently against extremely large data in a distributed environment.

With native support for bring-your-own-algorithms and frameworks, Amazon SageMaker offers flexible distributed training options that adjust to a customer's specific workflows. Amazon SageMaker is eligible to operate with data containing PHI. Encryption of data in transit is provided by SSL/TLS and is used when communicating both with the front-end interface of Amazon SageMaker (to the Notebook) and whenever Amazon SageMaker interacts with any other AWS service (for example, pulling data from Amazon S3).

To satisfy the requirement that PHI be encrypted at-rest, encryption of data stored with the instance running models with Amazon SageMaker is enabled using AWS Key Management Service (KMS) when setting up the endpoint (`DescribeEndpointConfig:KmsKeyID`). Encryption of model training results (artifacts) is enabled using AWS KMS and keys should be specified using the `KmsKeyID` in the `OutputDataConfig` description. If a KMS Key ID isn't provided, the default Amazon S3 KMS Key for the role's account will be used. Amazon SageMaker uses AWS CloudTrail to log all API calls.

Amazon Simple Notification Service (Amazon SNS)

Customers should understand the following key encryption requirement in order to use Amazon Simple Notification Service (SNS) with Protected Health Information (PHI). Customers must use the HTTPS API endpoint that SNS provides in each AWS Region. The HTTPS endpoint leverages encrypted connections, and protects the privacy and integrity of the data sent to AWS. For a list of all HTTPS API endpoints, see [AWS service endpoints](#).

Additionally, Amazon SNS uses CloudTrail, a service that captures API calls made by or on behalf of Amazon SNS in the customer's AWS account and delivers the log files to an Amazon S3 bucket

that they specify. CloudTrail captures API calls made from the Amazon SNS console or from the Amazon SNS API. Using the information collected by CloudTrail, customers can determine what request was made to Amazon SNS, the source IP address from which the request was made, who made the request, and when it was made. For more information on logging SNS operations, see [Logging Amazon SNS API calls using CloudTrail](#).

Amazon Simple Email Service (Amazon SES)

Amazon Simple Email Service (Amazon SES) is a flexible and highly scalable email sending and receiving service. It supports both S/MIME and PGP protocols to encrypt messages for full end-to-end encryption, and all communication with Amazon SES is secured using SSL (TLS 1.2). Customers have the option to store messages encrypted at-rest by configuring Amazon SES to receive and encrypt messages before storing them in an Amazon S3 bucket. For more information, see [How Amazon Simple Email Service \(Amazon SES\) uses AWS KMS](#) to find out more information about encrypting messages for storage. Messages are secured in transit to Amazon SES either through an HTTPS endpoint or encrypted SMTP connection.

For messages sent from Amazon SES to a receiver, Amazon SES will first attempt to make a secure connection to the receiving mail server, but if a secure connection cannot be established, it will send the message unencrypted. To require encryption for delivery to a receiver, customers must create a configuration set in Amazon SES and use the AWS CLI to set the `TlsPolicy` property to `Require`. For more information, see [Amazon SES and Security Protocols](#). Amazon SES integrates with AWS CloudTrail to monitor all API calls. Using the information collected by AWS CloudTrail, customers can determine that the request was made to Amazon SES, the IP address of the request, who made the request, when the request was made, and additional details. For more information, see [Logging Amazon SES API Calls with AWS CloudTrail](#). Amazon SES also provides methods to monitor sending activity such as sends, rejects, bounce rates, deliveries, opens, and clicks. For more information, see [Monitoring Your Amazon SES Sending Activity](#).

Amazon Simple Queue Service (Amazon SQS)

Customers should understand the following key encryption requirements in order to use Amazon SQS with PHI.

- Communication with the Amazon SQS Queue via the Query Request must be encrypted with HTTPS. For more information on making SQS requests, see [Making Query API requests](#).

- Amazon SQS supports server-side encryption integrated with the AWS KMS to protect data at rest. The addition of server-side encryption allows customers to transmit and receive sensitive data with the increased security of using encrypted queues. Amazon SQS server-side encryption uses the 256-bit Advanced Encryption Standard (AES-256 GCM algorithm) to encrypt the body of each message. The integration with AWS KMS allows customers to centrally manage the keys that protect Amazon SQS messages along with keys that protect their other AWS resources. AWS KMS logs every use of encryption keys to AWS CloudTrail to help meet regulatory and compliance needs. For more information, and to check Region for the availability for SSE for Amazon SQS, see [Encryption at Rest](#).
- If server-side encryption is not used, the message payload itself must be encrypted before being sent to SQS. One way to encrypt the message payload is by using the Amazon SQS Extended Client along with the Amazon S3 encryption client. For more information about using client-side encryption, see [Encrypting Message Payloads Using the Amazon SQS Extended Client and the Amazon S3 Encryption Client](#).

Amazon SQS uses CloudTrail, a service that logs API calls made by or on behalf of Amazon SQS in a customer's AWS account and delivers the log files to the specified Amazon S3 bucket. CloudTrail captures API calls made from the Amazon SQS console or from the Amazon SQS API. Customers can use the information collected by CloudTrail to determine which requests are made to Amazon SQS, the source IP address from which the request is made, who made the request, when it is made, and so on. For more information about logging SQS operations, see [Logging Amazon SQS API calls using AWS CloudTrail](#).

Amazon Simple Storage Service (Amazon S3)

Customers have several options for encryption of data at rest when using Amazon S3, including both server-side and client-side encryption, and several methods of managing keys. For more information, see [Protecting data using encryption](#).

Connections to Amazon S3 containing PHI must use endpoints that accept encrypted transport (HTTPS). For a list of regional endpoints, see [AWS service endpoints](#).

Do not use PHI in bucket names, object names, or metadata because this data is not encrypted using S3 server-side encryption and is not generally encrypted in client-side encryption architectures.

Amazon Simple Workflow Service

Amazon Simple Workflow Service (Amazon SWF) helps developers build, run, and scale background jobs that have parallel or sequential steps. Amazon SWF can be thought of as a fully managed state tracker and task coordinator in the Cloud.

The Amazon Simple Workflow Service is used to orchestrate workflows and is not able to store or transmit data. PHI should not be placed in metadata for Amazon SWF or within any task description. Amazon SWF uses AWS CloudTrail to log all API calls.

Amazon Textract

Amazon Textract uses machine learning technologies to automatically extract text and data from scanned documents that goes beyond simple optical character recognition (OCR) to identify, understand, and extract data from forms and tables. For example, customers can use Amazon Textract to automatically extract data and process forms with protected health information (PHI) without human intervention to fulfill medical claims.

Amazon Textract can also be used to maintain compliance in document archives. For example, customers can use Amazon Textract to extract data from insurance claims or medical prescriptions, and automatically recognize key-value pairs in those documents so that sensitive ones can be redacted.

Amazon Textract supports server-side encryption (SSE-S3 and SSE-KMS) for input documents and TLS encryption for data in transit between the service and agent. Customers can use Amazon CloudWatch to track resource usage metrics and AWS CloudTrail to capture API calls to Amazon Textract.

Amazon Transcribe

Amazon Transcribe uses advanced machine learning technologies to recognize speech in audio files and transcribe them into text. For example, customers can use Amazon Transcribe to convert US English and Mexican Spanish audio to text and to create applications that incorporate the content of audio files. Amazon Transcribe can be used with data containing PHI. Amazon Transcribe does not retain or store any data and all calls to the API are encrypted with SSL/TLS. Amazon Transcribe uses CloudTrail to log all API calls.

Amazon Translate

Amazon Translate uses advanced machine learning technologies to provide high-quality translation on demand. Customers can use Amazon Translate to translate unstructured text documents or to build applications that work in multiple languages. Documents containing PHI can be processed with Amazon Translate. No additional configuration is required when translating documents that contain PHI. Encryption of data while in transit is provided by SSL/TLS and no data remains at-rest with Amazon Translate. Amazon Translate uses CloudTrail to log all API calls.

Amazon Virtual Private Cloud

Amazon Virtual Private Cloud (Amazon VPC) offers a set of network security features well-aligned to architecting for HIPAA regulated workloads. Features such as stateless network access control lists and dynamic reassignment of instances into stateful security groups afford flexibility in protecting the instances from unauthorized network access.

Amazon VPC also allows customers to extend their own network address space into AWS, as well as providing a number of ways to connect their data centers to AWS. VPC Flow Logs provide an audit trail of accepted and rejected connections to instances processing, transmitting, or storing PHI.

AWS Transit Gateway acts as a network hub and simplifies the connectivity between Amazon VPCs as well as on-premises networks. AWS Transit Gateway also provides inter-region peering capabilities to other Transit Gateways to establish a global network using the AWS backbone. For more information on Amazon VPC, refer to [Amazon Virtual Private Cloud](#).

Amazon WorkDocs

Amazon WorkDocs is a fully managed, secure enterprise file storage and sharing service with strong administrative controls and feedback capabilities that improve user productivity. Amazon WorkDocs files are encrypted at-rest using keys that customers manage through AWS Key Management Service (AWS KMS). All data in transit is encrypted using SSL/TLS. AWS web and mobile applications, and desktop sync clients, transmit files directly to Amazon WorkDocs using SSL/TLS.

Using the Amazon WorkDocs Management Console, WorkDocs administrators can view audit logs to track file and user activity by time, and choose whether to allow users to share files with others outside their organization. Amazon WorkDocs is also integrated with CloudTrail (a service that

captures API calls made by or on behalf of Amazon WorkDocs in customer's AWS account), and delivers CloudTrail log files to an Amazon S3 bucket that customers specify.

Multi-factor authentication (MFA) using a RADIUS server is available and can provide customers with an additional layer of security during the authentication process. Users log in by entering their user name and password followed by an OTP (One-Time Passcode) supplied by a hardware or a software token.

For more information, see:

- [Amazon WorkDocs feature](#)
- [Logging Amazon WorkDocs API calls using AWS CloudTrail](#)

Customers should not store PHI in file names or directory names.

Amazon WorkSpaces

Amazon WorkSpaces is a fully managed, secure Desktop-as-a-Service (DaaS) solution that runs on AWS. With Amazon WorkSpaces, customers can easily provision virtual, cloud-based Microsoft Windows desktops for their users, providing them access to the documents, applications, and resources they need, anywhere, anytime, from any supported device.

Amazon WorkSpaces stores data in Amazon Elastic Block Store volumes. Customers can encrypt customer's WorkSpaces storage volumes using keys that customers manage through AWS Key Management Service. When encryption is enabled on a Workspace, both the data stored at-rest in the underlying storage and the automated backups (EBS Snapshots) of the disk storage are encrypted consistent with the Guidance. Communication from the Workspace clients to Workspace is secured using SSL/TLS. For more information on encryption at-rest using Amazon WorkSpaces, see [Encrypted WorkSpaces](#).

AWS App Mesh

AWS App Mesh is a service mesh that provides application-level networking to make it easy for your services to communicate with each other across multiple types of compute infrastructure, like Amazon ECS, Amazon EKS, or Amazon EC2 services. App Mesh configures Envoy proxies to collect and transmit observability data to the monitoring services that you configure, to give you end-to-end visibility. It can route traffic based on routing and traffic policies configured to ensure high-

availability of your applications. Traffic between applications can be configured to use TLS. App Mesh can be used using AWS SDK or App Mesh controller for Kubernetes. While AWS App Mesh is a HIPAA Eligible Service, no PHI should be stored in any resource names/attributes within AWS App Mesh as there is no support for protecting such data. Instead, AWS App Mesh can be used to monitor, control and secure customer domain resources that transmit or store PHI.

AWS Application Migration Service

AWS Application Migration Service (AWS MGN) allows you to quickly migrate your servers and applications to AWS, without changes and with minimal downtime. AWS MGN is the primary migration service recommended for lift and shift migrations to AWS.

AWS MGN uses block level data replication to copy source disks directly to EBS volumes in the customer account - the data is never transmitted through an AWS MGN controlled cloud environment. Replicated data is encrypted in transit by default. Data in the customer's EBS volumes is encrypted by default using a customer's own keys.

AWS Auto Scaling

AWS Auto Scaling enables customers to configure automatic scaling for the AWS resources that are part of a customer's application in a matter of minutes. Customers can use AWS Auto Scaling for a number of services that involve PHI, such as Amazon DynamoDB, Amazon ECS, Amazon RDS Aurora replicas, and Amazon EC2 instances in an Auto Scaling Group.

AWS Auto Scaling is an orchestration service that does not directly process, store or transmit customer content; for that reason, customers can use this service with encrypted content. The [AWS shared responsibility model](#) applies to data protection in AWS Auto Scaling: AWS is responsible for the AWS network security procedures, whereas the customer is responsible for maintaining control over a customer's content that is hosted on this infrastructure. This content includes the security configuration and management tasks for the AWS services that customers use. For data protection purposes, we recommend that customers protect AWS account credentials and set up individual user accounts with AWS Identity and Access Management (IAM). That way each user is given only the permissions necessary to fulfill their job duties.

We strongly recommend that customers never put sensitive identifying information, such as customers' account numbers, into free-form fields such as a Name field. This includes when customers work with AWS Auto Scaling or other AWS services using the AWS Management Console, API, AWS CLI, or AWS SDKs.

Any data that customers enter into AWS Auto Scaling or other services might get picked up for inclusion in diagnostic logs. When customers provide a URL to an external server, they should not include credentials information in the URL to validate their request to that server. AWS also recommends that customers secure their data in the following ways:

- Use multi-factor authentication (MFA) with each account.
- Use SSL/TLS to communicate with AWS resources. We recommend TLS 1.2 or later
- Set up API and user activity logging with AWS CloudTrail.
- Use AWS encryption solutions, along with all default security controls within AWS services.
- Use advanced managed security services such as Amazon Macie, which assists in discovering and securing personal data that is stored in Amazon S3.

AWS Backup

AWS Backup offers a centralized, fully-managed, and policy-based service to protect customer data and ensure compliance across AWS services for business continuity purposes. With AWS Backup, customers can centrally configure data protection (backup) policies and monitor backup activity across customer AWS resources, including Amazon EBS volumes, Amazon Relational Database Service (Amazon RDS) databases (including Aurora clusters), Amazon DynamoDB tables, Amazon Elastic File System (Amazon EFS), Amazon FSx file systems, Amazon EC2 instances, and AWS Storage Gateway volumes.

AWS Backup encrypts customer data in transit and at rest. Backups from services with existing snapshot capabilities are encrypted using the source service's snapshot encryption methodology. For example, EBS snapshots are encrypted using the encryption key of the volume that the snapshot was created from.

Backups from newer AWS services that introduce backup functionality built on AWS Backup, such as Amazon EFS, are encrypted in-transit and at-rest independently from the source services, giving customer backups an additional layer of protection. Encryption is configured at the Backup Vault level. Default Vault is encrypted. When customers create a new vault, an Encryption Key must be selected.

AWS Batch

AWS Batch enables developers, scientists, and engineers to easily and efficiently run hundreds of thousands of batch computing jobs on AWS. AWS Batch dynamically provisions the optimal

quantity and type of compute resources (such as CPU or memory-optimized instances) based on the volume and specific resource requirements of the batch jobs submitted. AWS Batch plans, schedules, and executes batch computing workloads across the full range of AWS compute services and features.

Similar to guidance for Amazon ECS, PHI should not be placed directly into the job definition, the job queue or the tags for AWS Batch. Instead, jobs scheduled and executed with AWS Batch may operate on encrypted PHI. Any information returned by stages of a job to AWS Batch should also not contain any PHI. Whenever jobs being executed by AWS Batch must transmit or receive PHI, that connection should be encrypted using HTTPS or SSL/TLS.

AWS Certificate Manager

AWS Certificate Manager is a service that lets customers easily provision, manage, and deploy public and private SSL/TLS certificates for use with AWS services and their internal connected resources. AWS Certificate Manager uses CloudTrail to log all API calls.

Users need programmatic access if they want to interact with AWS outside of the AWS Management Console. The way to grant programmatic access depends on the type of user that's accessing AWS.

To grant users programmatic access, choose one of the following options.

Which user needs programmatic access?	To	By
Workforce identity (Users managed in IAM Identity Center)	Use temporary credentials to sign programmatic requests to the AWS CLI, AWS SDKs, or AWS APIs.	Following the instructions for the interface that you want to use. <ul style="list-style-type: none">• For the AWS CLI, see Configuring the AWS CLI to use AWS IAM Identity Center in the <i>AWS Command Line Interface User Guide</i>.• For AWS SDKs, tools, and AWS APIs, see IAM Identity

Which user needs programmatic access?	To	By
		Center authentication in the <i>AWS SDKs and Tools Reference Guide</i> .
IAM	Use temporary credentials to sign programmatic requests to the AWS CLI, AWS SDKs, or AWS APIs.	Following the instructions in Using temporary credentials with AWS resources in the <i>IAM User Guide</i> .
IAM	(Not recommended) Use long-term credentials to sign programmatic requests to the AWS CLI, AWS SDKs, or AWS APIs.	<p>Following the instructions for the interface that you want to use.</p> <ul style="list-style-type: none"> For the AWS CLI, see Authenticating using IAM user credentials in the <i>AWS Command Line Interface User Guide</i>. For AWS SDKs and tools, see Authenticate using long-term credentials in the <i>AWS SDKs and Tools Reference Guide</i>. For AWS APIs, see Managing access keys for IAM users in the <i>IAM User Guide</i>.

AWS Cloud Map

AWS Cloud Map is a cloud resource discovery service. With AWS Cloud Map, customers can define custom names for application resources, such as Amazon ECS tasks, Amazon EC2 instances, Amazon S3 buckets, Amazon DynamoDB tables, Amazon SQS queues, or any other cloud resource. Customers can then use these custom names to discover the location

and metadata of cloud resources from their applications using AWS SDK and authenticated API queries. While AWS Cloud Map is a HIPAA Eligible Service, no PHI should be stored in any resource names/attributes within AWS Cloud Map as there is no support for protecting such data. Instead, AWS Cloud Map can be used to discover customer domain resources that transmit or store PHI.

AWS CloudFormation

AWS CloudFormation enables customers to create and provision AWS infrastructure deployments predictably and repeatedly. It helps customers leverage AWS products such as Amazon EC2, Amazon Elastic Block Store, Amazon SNS, Elastic Load Balancing, and Auto Scaling to build highly reliable, highly scalable, cost-effective applications in the cloud without worrying about creating and configuring the underlying AWS infrastructure. AWS CloudFormation enables customers to use a template file to create and delete a collection of resources together as a single unit (a stack).

AWS CloudFormation does not itself store, transmit, or process PHI. Instead, it is used to build and deploy architectures that use other AWS services that might store, transmit, and/or process PHI. Only HIPAA Eligible Services should be used with PHI. Please refer to the entries for those services in this Whitepaper for guidance on use of PHI with those services. AWS CloudFormation uses AWS CloudTrail to log all API calls.

AWS CloudHSM

AWS CloudHSM is a cloud-based hardware security module (HSM) that enables customers to easily generate and use their own encryption keys on the AWS Cloud. With CloudHSM, customers can manage their own encryption keys using FIPS 140-2 Level 3 validated HSMs. CloudHSM offers customers the flexibility to integrate with their applications using open standard APIs, such as PKCS#11, Java Cryptography Extensions (JCE), and Microsoft CryptoNG (CNG) libraries.

CloudHSM is also standards-compliant and enables customers to export all of their keys to most other commercially available HSMs. As AWS CloudHSM is a hardware appliance key management service, it is unable to store or transmit PHI. Customers should not store PHI in Tags (metadata). No other special guidance is required.

AWS CloudTrail

AWS CloudTrail is a service that enables governance, compliance, operational auditing, and risk auditing of AWS accounts. With CloudTrail, customers can log, continuously monitor, and retain

account activity related to actions across their AWS infrastructure. CloudTrail provides event history of their AWS account activity, including actions taken through the AWS Management Console, AWS SDKs, command line tools, and other AWS services. This event history simplifies security analysis, resource change tracking, and troubleshooting.

AWS CloudTrail is enabled for use with all AWS accounts and can be used for audit logging, as required by the AWS BAA. Specific Trails should be created using the CloudTrail console or the AWS Command Line Interface. CloudTrail encrypts all traffic while in transit and at-rest when an encrypted Trail is created. An encrypted trail should be created when the potential exists to log PHI.

By default, an encrypted Trail stores entries in Amazon S3 using Server-Side Encryption with Amazon S3 (SSE-S3) managed keys. If an additional management over keys is desired, it can also be configured with AWS KMS-managed keys (SSE-KMS). As CloudTrail is the final destination for AWS log entries, and thus, a critical component of any architecture that handles PHI, CloudTrail log file integrity validation should be enabled and the associated CloudTrail digest files should be periodically reviewed. Once enabled, a positive assertion that the log files have not been changed or altered can be established.

AWS CodeBuild

AWS CodeBuild is a fully managed build service in the cloud. AWS CodeBuild compiles source code, runs unit tests, and produces artifacts that are ready to deploy. AWS CodeBuild uses an AWS KMS key to encrypt build output artifacts. A KMS key should be created and configured before building artifacts that contain PHI, secrets/passwords, certificates, etc. AWS CodeBuild uses AWS CloudTrail to log all API calls.

AWS CodeDeploy

AWS CodeDeploy is a fully managed deployment service that automates software deployments to a variety of compute services including Amazon EC2, AWS Fargate, AWS Lambda, and on-premises servers. Customers use AWS CodeDeploy to rapidly release new features of containerized workload and handles the complexity of updating applications.

AWS CodeDeploy supports server-side encryption (SSE-S3) for deployment artifacts and TLS encryption for data in transit between the service and agent. Customers can use Amazon CloudWatch Events to track deployments and AWS CloudTrail to capture API calls to AWS CodeDeploy.

AWS CodeCommit

AWS CodeCommit is a secure, highly scalable, managed source control service that hosts private Git repositories. AWS CodeCommit eliminates the need for customers to manage their own source control system or worry about scaling its infrastructure.

AWS CodeCommit encrypts all traffic and stored information while in transit and at-rest. By default, when a repository is created within AWS CodeCommit, an AWS managed key is created with AWS KMS and is used only by that repository to encrypt all data stored at-rest. AWS CodeCommit uses AWS CloudTrail to log all API calls.

AWS CodePipeline

AWS CodePipeline is a fully managed [continuous delivery](#) service that helps customers automate customer release pipelines for fast and reliable application and infrastructure updates. Customers use AWS CodePipeline to allow researchers to automatically process clinical trial data, lab results and genomic data are few examples of workflow pipeline used by customers.

AWS CodePipeline supports server-side encryption (SSE-S3 and SSE-KMS) for Code artifacts and TLS encryption for data in transit between the service and agent. Customers can use Amazon CloudWatch Events to track pipeline changes and AWS CloudTrail to capture API calls to AWS CodePipeline.

AWS Config

AWS Config provides a detailed view of the resources associated with a customer's AWS account, including how they are configured, how they are related to one another, and how the configurations and their relationships have changed over time.

AWS Config cannot itself be used to store or transmit PHI.

Instead, it can be leveraged to monitor and evaluate architectures built with other AWS services, including architectures that handle PHI, to help determine whether they remain compliant with their intended design goal. Architectures that handle PHI should only be built with HIPAA Eligible Services. AWS Config uses AWS CloudTrail to log all results.

AWS Data Exchange

AWS Data Exchange makes it easy to find, subscribe to, and use third-party data in the cloud. Once subscribed to a data product, customers can use the AWS Data Exchange API to load data directly into [Amazon S3](#) and then analyze it with a wide variety of AWS [analytics](#) and [machine learning](#) services. For data providers, AWS Data Exchange makes it easy to reach the millions of AWS customers migrating to the cloud by removing the need to build and maintain infrastructure for data storage, delivery, billing, and entitlement.

AWS Data Exchange always encrypts all data products stored in the service at-rest without requiring any additional configuration. This encryption is automatically done via a service managed KMS key. AWS Data Exchange uses Transport Layer Security (TLS) and client-side encryption for encryption in transit. Communication with AWS Data Exchange is always done over HTTPS so customer's data is always encrypted in transit. This encryption is configured by default when customers use AWS Data Exchange. For more information, see [Data Protection in AWS Data Exchange](#).

AWS Data Exchange is integrated with AWS CloudTrail. AWS CloudTrail captures all calls to AWS Data Exchange APIs as events, including calls from the AWS Data Exchange console and from code calls to the AWS Data Exchange API operations. Some actions customers can take are console-only actions. There is no corresponding API in the AWS SDK or AWS CLI. These are actions that rely on AWS Marketplace functionality, such as publishing or subscribing to a product. AWS Data Exchange provides CloudTrail logs for a subset of these console-only actions. For more information, see [Logging AWS Data Exchange API Calls with AWS CloudTrail](#).

Please note that all listings using AWS Data Exchange must adhere to AWS Data Exchange's [Publishing Guidelines](#) and [AWS Data Exchange FAQs](#) for AWS Marketplace Providers, which restrict certain categories of data. For more information, see [AWS Data Exchange FAQs](#).

AWS Database Migration Service

AWS Database Migration Service (AWS DMS) helps customers migrate databases to AWS easily and securely. Customers can migrate their data to and from most widely used commercial and open-source databases, such as Oracle, MySQL, and PostgreSQL. The service supports homogeneous migrations such as Oracle to Oracle, and also heterogeneous migrations between different database platforms, such as Oracle to PostgreSQL or MySQL to Oracle.

Databases running on-premises and being migrated to the cloud with AWS DMS can contain PHI data. AWS DMS encrypts data while in transit and when data is being staged for final migration into the target database on AWS. AWS DMS encrypts the storage used by a replication instance and the endpoint connection information. To encrypt the storage used by a replication instance, AWS DMS uses an AWS KMS key that is unique to the AWS account. Refer to the Guidance for the appropriate target database to ensure that data remains encrypted once migration is complete. AWS DMS uses CloudTrail to log all API calls.

AWS DataSync

AWS DataSync is an online transfer service that simplifies, automates, and accelerates moving data between on-premises storage and AWS. Customers can use AWS DataSync to connect their data sources to either Amazon S3 or Amazon EFS. Customers should ensure that Amazon S3 and Amazon EFS are configured in a manner consistent with the Guidance. By default, customer data is encrypted in transit using TLS 1.2. For more information about encryption and AWS DataSync, see [AWS DataSync features](#). Customers can monitor DataSync activity using AWS CloudTrail. For more information on logging with CloudTrail, see [Logging AWS DataSync API Calls with AWS CloudTrail](#).

AWS Directory Service

AWS Directory Service for Microsoft AD

AWS Directory Service for Microsoft Active Directory (Enterprise Edition), also known as AWS Microsoft AD, enables directory-aware workloads and AWS resources to use managed Active Directory in the AWS Cloud. AWS Microsoft AD stores directory content (including content containing PHI) in encrypted Amazon Elastic Block Store volumes using encryption keys that AWS manages. For more information, see [Amazon EBS Encryption](#).

Data in transit to and from Active Directory clients is encrypted when it travels through Lightweight Directory Access Protocol (LDAP) over customer's Amazon Virtual Private Cloud (VPC) network. If an Active Directory client resides in an on-premises network, the traffic travels to customer's VPC by a virtual private network link or an AWS Direct Connect link.

Amazon Cloud Directory

Amazon Cloud Directory enables customers to build flexible cloud-native directories for organizing hierarchies of data along multiple dimensions. Customers also can create directories for a variety

of use cases, such as organizational charts, course catalogs, and device registries. For example, customers can create an organizational chart that can be navigated through separate hierarchies for reporting structure, location, and cost center. Amazon Cloud Directory automatically encrypts data at rest and in transit by using 256-bit encryption keys that are managed by the AWS Key Management Service (AWS KMS).

AWS Elastic Beanstalk

With AWS Elastic Beanstalk, customers can quickly deploy and manage applications in the AWS Cloud without having to learn about the infrastructure that runs those applications. Customers can simply upload code and AWS Elastic Beanstalk automatically handles the deployment, from capacity provisioning, load balancing, automatic scaling to application health monitoring. At the same time, customers retain full control over the AWS resources powering their application and can access the underlying resources at any time.

AWS Elastic Beanstalk does not itself store, transmit, or process PHI. Instead, customers can use it to build and deploy architectures with other AWS services that might store, transmit, and/or process PHI. Customers should ensure that when picking the services that are deployed by AWS Elastic Beanstalk to only use HIPAA Eligible Services with PHI. Refer to the entries for those services in this whitepaper for guidance on use of PHI with those services.

Customers should not include PHI in any free-form fields within AWS Elastic Beanstalk such as the **Name** field. AWS Elastic Beanstalk uses AWS CloudTrail to log all API calls.

AWS Elastic Disaster Recovery

AWS Elastic Disaster Recovery (AWS DRS) minimizes downtime and data loss with fast, reliable recovery of on-premises and cloud-based applications using affordable storage, minimal compute, and point-in-time recovery.

Customers can set up AWS Elastic Disaster Recovery on their source servers to initiate secure data replication. Their data is replicated to a staging area subnet in your AWS account, in the AWS Region they select. The staging area design reduces costs by using affordable storage and minimal compute resources to maintain ongoing replication. Customer data replicated by AWS Elastic Disaster Recovery is encrypted in transit using TLS 1.2, and is transferred directly from their source servers to their VPC. Customers can leverage private connectivity such as AWS Direct Connect or VPN to configure the replication route. Customer data can also be [encrypted at rest](#) on AWS using Amazon EBS encryption.

Customers can perform non-disruptive tests to confirm that implementation is complete. During normal operation, maintain readiness by monitoring replication and periodically performing non-disruptive recovery and failback drills. If customers need to recover applications, they can launch recovery instances on AWS within minutes, using the most up-to-date server state or a previous point in time. After customer applications are running on AWS, they can choose to keep them there, or they can initiate data replication back to their primary site when the issue is resolved. Customers can fail back to their primary site whenever they're ready.

AWS Fargate

AWS Fargate is a technology that allows customer to run containers without having to manage servers or clusters. With AWS Fargate, customers no longer have to provision, configure, and scale clusters of virtual machines to run containers. This removes the need to choose server types, decide when to scale clusters, or optimize cluster packing. AWS Fargate removes the need for customers to interact with or think about servers or clusters. With Fargate, customers focus on designing and building their applications instead of managing the infrastructure that runs them.

Fargate does not require any additional configuration in order to work with workloads that process PHI. Customers can run container workloads on Fargate using container orchestration services like Amazon ECS. Fargate only manages the underlying infrastructure and does not operate with or upon data within the workload being orchestrated. In keeping with the requirements for HIPAA, PHI should still be encrypted whenever in transit or at-rest when accessed by containers launched with Fargate. Various mechanisms for encrypting at-rest are available with each AWS storage option described in this paper. For additional HIPAA security and configuration information, see the [Architecting for HIPAA Security and Compliance on Amazon EKS](#) whitepaper.

AWS Firewall Manager

AWS Firewall Manager is a security management service which allows customers to centrally configure and manage firewall rules across customer accounts and applications in AWS Organizations. As new applications are created, Firewall Manager makes it easy to bring new applications and resources into compliance by enforcing a common set of security rules. Now customers have a single service to build firewall rules, create security policies, and enforce them in a consistent, hierarchical manner across their entire infrastructure, from a central administrator account.

AWS Firewall Manager is an orchestration service that does not directly process, store or transmit user data. The service does not encrypt customer content, but underlying services that AWS Firewall Manager uses, such as DynamoDB, encrypts user data.

AWS Global Accelerator

AWS Global Accelerator is a global load balancing service that improves the availability and latency of multi-region applications. To ensure that PHI remains encrypted in transit and at-rest while using AWS Global Accelerator, architectures being load balanced by Global Accelerator should use an encrypted protocol, such as HTTPS or SSL/TLS. Refer to the guidance for Amazon EC2, Elastic Load Balancing, and other AWS services to better understand the available encryption options for backend resources. AWS Global Accelerator uses AWS CloudTrail to log all API calls.

AWS Glue

AWS Glue is a fully managed ETL (extract, transform, and load) service that makes it simple and cost-effective for customers to categorize their data, clean it, enrich it, and move it reliably between various data stores. In order to ensure the encryption of data containing PHI while in transit, AWS Glue should be configured to use JDBC connections to data stores with SSL/TLS. Additionally, to maintain encryption while in-transit, the setting for server-side encryption (SSE-S3) should be passed as a parameter to ETL jobs run with AWS Glue. All data stored at-rest within the Data Catalog of AWS Glue is encrypted using keys managed by AWS KMS when encryption is enabled upon creation of a Data Catalog object. AWS Glue uses CloudTrail to log all API calls.

AWS Glue DataBrew

AWS Glue DataBrew is a fully managed visual data preparation service that makes it easy for data analysts and data scientists to clean and normalize data to prepare it for analytics and machine learning. In order to ensure the encryption of data containing PHI while in transit, DataBrew should be configured to use JDBC connections to data stores with SSL/TLS. When connecting to JDBC data sources, DataBrew uses the settings on your AWS Glue connection, including the “Require SSL connection” option. Additionally, to maintain encryption while at rest in S3 buckets, the setting for server-side encryption (SSE-S3 or SSE-KMS) should be passed as a parameter to DataBrew jobs.

AWS IoT Core and AWS IoT Device Management

AWS IoT Core and AWS IoT Device Management provide secure, bidirectional communication between internet-connected devices, such as sensors, actuators, embedded micro-controllers, or smart appliances, and the AWS Cloud. AWS IoT Core and AWS IoT Device Management can now accommodate devices that transmit data containing PHI. All communication with AWS IoT Core and AWS IoT Device Management is encrypted using TLS. AWS IoT Core and AWS IoT Device Management use AWS CloudTrail to log all API calls.

AWS IoT Greengrass

AWS IoT Greengrass lets customers run local compute, messaging, data caching, sync, and ML inference capabilities for connected devices in a secure way. AWS IoT Greengrass uses X.509 certificates, managed subscriptions, AWS IoT policies, and IAM policies and roles to ensure that customer's Greengrass applications are secure. AWS IoT Greengrass uses the AWS IoT transport security model to encrypt communication with the cloud using TLS. In addition, AWS IoT Greengrass data is encrypted when at-rest (in the cloud). For more information on Greengrass security, see [Overview of AWS IoT Greengrass Security](#).

Customers can log AWS IoT Greengrass API actions using AWS CloudTrail. For more information, see [Logging AWS IoT Greengrass API Calls with AWS CloudTrail](#).

AWS Lambda

AWS Lambda lets customers run code without provisioning or managing servers on their own. AWS Lambda uses a compute fleet of Amazon Elastic Compute Cloud (Amazon EC2) instances across multiple Availability Zones in a Region, which provides the high availability, security, performance, and scalability of the AWS infrastructure.

To ensure that PHI remains encrypted while using AWS Lambda, connections to external resources should use an encrypted protocol such as HTTPS or SSL/TLS. For example, when S3 is accessed from a Lambda procedure, it should be addressed with `https://bucket.s3-aws-region.amazonaws.com`.

If any PHI is placed at-rest or idled within a running procedure, it should be encrypted client-side or server-side with keys obtained from AWS KMS or AWS CloudHSM. Follow the related guidance for Amazon API Gateway when triggering AWS Lambda functions through the service. When using

events from other AWS services to trigger AWS Lambda functions, the event data should not contain (in and of itself) PHI. For example, when a Lambda procedure is triggered from an S3 event, such as the arrival of an object in S3, the object name that is relayed to Lambda should not have any PHI, although the object itself can contain such data.

AWS Managed Services

AWS Managed Services provides ongoing management of AWS infrastructures. By implementing best practices to maintain a customer's infrastructure, AWS Managed Services helps to reduce their operational overhead and risk. AWS Managed Services automates common activities such as change requests, monitoring, patch management, security, and backup services, and provides full-lifecycle services to provision, run, and support infrastructures.

Customers can use AWS Managed Services to manage AWS workloads that that operate with data containing PHI. Usage of AWS Managed Services does not alter the AWS Services eligible for the use with PHI. Tooling and automation provided by AWS Managed Services cannot be used for the storage or transmission of PHI.

AWS OpsWorks for Chef Automate

AWS OpsWorks for Chef Automate is a fully managed configuration management service that hosts Chef Automate, a set of automation tools from Chef for infrastructure and application management. The service itself does not contain, transmit, or handle any PHI or sensitive information, but customers should ensure that any resources configured by OpsWorks for Chef Automate is configured consistent with the Guidance. API calls are captured with AWS CloudTrail. For more information, see [Logging AWS OpsWorks Stacks API Calls with AWS CloudTrail](#).

AWS OpsWorks for Puppet Enterprise

AWS OpsWorks for Puppet Enterprise is a fully managed configuration management service that hosts Puppet Enterprise, a set of automation tools from Puppet for infrastructure and application management. The service itself does not contain, transmit, or handle any PHI or sensitive information, but customers should ensure that any resource configured by OpsWorks for Puppet Enterprise is configured consistent with the Guidance. API calls are captured with AWS CloudTrail. For more information, see [Logging AWS OpsWorks Stacks API Calls with AWS CloudTrail](#).

AWS OpsWorks Stack

AWS OpsWorks Stacks provides a simple and flexible way to create and manage stacks and applications. Customers can use AWS OpsWorks Stacks to deploy and monitor applications in their stacks.

AWS OpsWorks Stacks encrypts all traffic while in transit. However, encrypted data bags (a Chef data storage mechanism) are not available and any assets that must be stored securely, such as PHI, secrets/passwords, certificates, etc., should be stored in an encrypted bucket in Amazon S3. AWS OpsWorks Stack uses AWS CloudTrail to log all API calls.

AWS Organizations

AWS Organizations helps customers centrally manage and govern their environment as they grow and scale their AWS resources. Using AWS Organizations, they can programmatically create new AWS accounts and allocate resources, group accounts to organize their workflows, apply policies to accounts or groups for governance, and simplify billing by using a single payment method for all of their accounts.

In addition, AWS Organizations is integrated with other AWS services so customers can define central configurations, security mechanisms, audit requirements, and resource sharing across accounts in their organization. AWS Organizations is available to all AWS customers at no additional charge.

AWS Organizations is an orchestration service that does not directly process, store or transmit user data. The service does not encrypt customer content, but underlying services that are launched within AWS Organizations, do encrypt user data. AWS Organizations is integrated with AWS CloudTrail, a service that provides a record of actions taken by a user, role, or an AWS service in AWS Organizations.

AWS RoboMaker

AWS RoboMaker enables customers to execute code in the cloud for application development and provides a robotics simulation service to accelerate application testing. AWS RoboMaker also provides a robotics fleet management service for remote application deployment, update, and management.

Network traffic containing PHI must encrypt data in transit. All management communication with the simulation server is over TLS, and customers should use open standard transport

encryption mechanisms for connections to other AWS services. AWS RoboMaker also integrates with CloudTrail to log all API calls to a specific Amazon S3 bucket.

AWS RoboMaker logs do not contain PHI, and the EBS volumes used by the simulation server are encrypted. When transferring data that may contain PHI to other services, such as Amazon S3, customers must follow the receiving service's guidance for storing PHI. For deployments to robots, customers must ensure that encryption of data in transit and at-rest is consistent with their interpretation of the Guidance.

AWS SDK Metrics

Enterprise customers can use the AWS CloudWatch agent with AWS SDK Metrics for Enterprise Support (SDK Metrics) to collect metrics from AWS SDKs on their hosts and clients. These metrics are shared with AWS Enterprise Support. SDK Metrics can help customers collect relevant metrics and diagnostic data about their application's connections to AWS services without adding custom instrumentation to their code, and reduces the manual work necessary to share logs and data with AWS Support.

Please note that SDK Metrics is **only** available to AWS customers with an Enterprise Support subscription. Customers can use SDK Metrics with any application that directly calls AWS services and that was built using an AWS SDK that is one of the versions listed in the [AWS Metrics documentation](#).

SDK Metrics monitors calls that are made by the AWS SDK and uses the CloudWatch agent running in the same environment as a client application.

The CloudWatch agent encrypts the data in-transit from the local machine to delivery in the destination log group. The log group can be configured to be encrypted following the directions at [Encrypt Log Data in CloudWatch Logs Using AWS KMS](#).

AWS Secrets Manager

AWS Secrets Manager is an AWS service that makes it easier for customers to manage "secrets." Secrets can be database credentials, passwords, third-party API keys, and even arbitrary text. AWS Secrets Manager might be used to store PHI if such information is contained within "secrets". All secrets stored by AWS Secrets Manager are encrypted at-rest using the AWS Key Management System (KMS). Users can select the AWS KMS key used when creating a new secret. If no key is

selected, the default key for the account will be used. AWS Secrets Manager uses AWS CloudTrail to log all API calls.

AWS Security Hub

AWS Security Hub collects and consolidates findings from AWS security services enabled in a customer's environment, such as intrusion detection findings from Amazon GuardDuty, vulnerability scans from Amazon Inspector, Amazon S3 bucket policy findings from Amazon Macie, publicly accessible and cross-account resources from IAM Access Analyzer, and resources lacking WAF coverage from AWS Firewall Manager. AWS Security Hub also consolidates findings from integrated AWS Partner Network (APN) security solutions.

AWS Security Hub integrates with Amazon CloudWatch Events, enabling customers to create custom response and remediation workflows. Customers can easily send findings to SIEMs, chat tools, ticketing systems, Security Orchestration Automation and Response (SOAR) tools, and on-call management platforms. Response and remediation actions can be fully automated or they can be triggered manually in the console. Customers can also use AWS Systems Manager Automation documents, AWS Step Functions, and AWS Lambda functions to build automated remediation workflows that can be initiated from AWS Security Hub.

To ensure data protection, AWS Security Hub encrypts data at rest and data in transit between component services. Third-party auditors assess the security and compliance of AWS Security Hub as part of multiple AWS compliance programs. AWS Security Hub is part of AWS's SOC, ISO, PCI, and HIPAA compliance programs.

AWS Server Migration Service

AWS Server Migration Service (AWS SMS) automates the migration of on-premises VMware vSphere or Microsoft Hyper-V/SCVMM virtual machines to the AWS Cloud. AWS SMS incrementally replicates server VMs as cloud-hosted Amazon Machine Images (AMIs) ready for deployment on Amazon EC2.

Servers running on-premises and being migrated to the cloud with (AWS SMS) can contain PHI data. AWS SMS encrypts data while in transit and when server VM images are being staged for final placement onto EC2. Refer to the guidance for EC2 and setting up encrypted storage volumes when migrating a server VM containing PHI with AWS SMS. AWS SMS uses CloudTrail to log all API calls.

AWS Serverless Application Repository

The AWS Serverless Application Repository (SAR) is a managed repository for serverless applications. It enables teams, organizations, and individual developers to store and share reusable applications, and easily assemble and deploy serverless architectures in powerful new ways. The applications are AWS CloudFormation templates, which contain definitions of the application infrastructure and compiled binaries of application AWS Lambda function code.

Although it is possible for applications that are in the AWS Serverless Application Repository to process PHI, they would only do this after being deployed to a customer's account and not as part of the SAR itself. The AWS Serverless Application Repository encrypts files that customers upload, including deployment packages and layer archives. For data in transit, the AWS Serverless Application Repository uses TLS to encrypt data between the service and the agent. AWS Serverless Application Repository is integrated with AWS CloudTrail, which is a service that provides a record of actions taken by a user, role, or an AWS service in the AWS Serverless Application Repository.

Service Catalog

Service Catalog allows IT administrators to create, manage, and distribute portfolios of approved products to end users, who can then access the products they need in a personalized portal. Service Catalog is used to catalog, share, and deploy self-service solutions on AWS and cannot be used to store, transmit, or process PHI. PHI should not be placed in any metadata for Service Catalog items or within any item description. Service Catalog uses AWS CloudTrail to log all API calls.

AWS Shield

AWS Shield is a managed Distributed Denial of Service (DDoS) protection service that safeguards web applications running on AWS. AWS Shield provides always-on detection and automatic inline mitigations that minimize application downtime and latency, so there is no need to engage AWS Support to benefit from DDoS protection.

AWS Shield cannot be used to store or transmit PHI, but instead can be used to safeguard web applications that do operate with PHI. As such, no special configuration is needed when engaging AWS Shield.

All AWS customers benefit from the automatic protections of AWS Shield Standard, at no additional charge. AWS Shield Standard defends against most common, frequently occurring

network and transport layer DDoS attacks that target their website or applications. For higher levels of protection against attacks targeting their web applications running on Elastic Load Balancing (ELB), Amazon CloudFront, and Amazon Route 53 resources, customers can subscribe to AWS Shield Advanced.

AWS Snowball

With AWS Snowball (Snowball), customers can transfer hundreds of terabytes or petabytes of data between their on-premises data centers and Amazon Simple Storage Service (Amazon S3). PHI stored in AWS Snowball must be encrypted at-rest consistent with the Guidance. When creating an import job, customers must specify the ARN for the AWS KMS key to be used to protect data within the Snowball. In addition, during the creation of the import job, customers should choose a destination S3 bucket that meets the encryption standards set by the Guidance.

While Snowball does not currently support server-side encryption with AWS KMS-managed keys (SSE-KMS) or server-side encryption with customer provided keys (SSE-C), Snowball does support server-side encryption with Amazon S3-managed encryption keys (SSE-S3). For more information, see [Protecting Data Using Server-Side Encryption with Amazon S3-Managed Encryption Keys \(SSE-S3\)](#).

Alternatively, customers can use the encryption methodology of their choice to encrypt PHI before storing the data in AWS Snowball.

Currently, customers may use the standard AWS Snowball appliance as part of our BAA.

AWS Snowball Edge

AWS Snowball Edge connects to existing customer applications and infrastructure using standard storage interfaces, streamlining the data transfer process and minimizing setup and integration. Snowball Edge can cluster together to form a local storage tier and process customer data on-site, helping customers ensure that their applications continue to run even when they are not able to access the cloud.

To ensure that PHI remains encrypted while using Snowball Edge, customers should make sure to use an encrypted connection protocol such as HTTPS or SSL/TLS when using AWS Lambda procedures powered by AWS IoT Greengrass to transmit PHI to/from resources external to Snowball Edge. Additionally, PHI should be encrypted while stored on the local volumes of Snowball Edge, either through local access or via NFS. Encryption is automatically applied to data placed into Snowball Edge using the Snowball Management Console and API for bulk transport

into S3. For more information on data transport into S3, see the related guidance for [the section called “AWS Snowball”](#).

AWS Step Functions

AWS Step Functions makes it easy to coordinate the components of distributed applications and microservices using visual workflows. AWS Step Functions is not able to store, transmit, or process PHI. PHI should not be placed within the metadata for AWS Step Functions or within any task or state machine definition. AWS Step Functions uses AWS CloudTrail to log all API calls.

AWS Storage Gateway

AWS Storage Gateway is a hybrid storage service that enables customers' on-premises applications to seamlessly use AWS Cloud storage. The gateway uses open standard storage protocols to connect existing storage applications and workflows to AWS Cloud storage services for minimal process disruption.

File Gateway

File gateway is a type of AWS Storage Gateway that supports a file interface into Amazon S3 and that adds to the current block-based volume and VTL storage. File gateway uses HTTPS to communicate with S3 and stores all objects encrypted while on S3 using SSE-S3, by default, or using client-side encryption with keys stored in AWS KMS. File metadata, such as file names, remains unencrypted and should not contain any PHI.

Volume Gateway

Volume gateway provides cloud-backed storage volumes that customers can mount as internet Small Computer System Interface (iSCSI) devices from on-premises application servers. Customers should attach local disks as Upload buffers and Cache to the Volume Gateway VM in accordance with their internal compliance and regulatory requirements. It is recommended that, for PHI, these disks should be capable of providing encryption at-rest. Communication between the Volume Gateway VM and AWS is encrypted using TLS 1.2 to secure PHI in transport.

Tape Gateway

Tape gateway provides a VTL (virtual tape library) interface to third-party backup applications running on-premises. Customers should enable encryption for PHI within the third-party backup

application when setting up a tape backup job. Communication between the Tape Gateway VM and AWS is encrypted using TLS 1.2 to secure PHI in transport. Customers using any of the Storage Gateway configurations with PHI should enable full logging. For more information, see [What Is AWS Storage Gateway?](#)

AWS Systems Manager

AWS Systems Manager is a unified interface that allows customers to easily centralize operational data, automate tasks across their AWS resources, and shortens the time to detect and resolve operational problems in their infrastructure. Systems Manager provides a complete view of a customer's infrastructure performance and configuration, simplifies resource and application management, and makes it easy to operate and manage their infrastructure at scale.

When outputting data that may contain PHI to other services, such as Amazon S3, customers must follow the receiving service's guidance for storing PHI. Customers should not include PHI in metadata or identifiers, such as document names and parameter names.

AWS Transfer for SFTP

AWS Transfer for SFTP provides Secure File Transfer Protocol (SFTP) access to a customer's S3 resources. Customers are presented with a virtual server, which is accessed using the standard SFTP protocol at a regional service endpoint. From the point of view of the AWS customer and the SFTP client, the SFTP gateway looks like a standard, highly available SFTP server. Although the service itself does not store, process, or transmit PHI, the resources that the customer is accessing on Amazon S3 should be configured in a manner that is consistent with the Guidance. Customers can also use AWS CloudTrail to log API calls made to AWS Transfer for SFTP.

AWS WAF – Web Application Firewall

AWS WAF is a web application firewall that helps protect customer web applications from common web exploits that could affect application availability, compromise security, or consume excessive resources. Customers may place AWS WAF between their web applications hosted on AWS that operate with or exchange PHI, and their end users. As with the transmission of any PHI while on AWS, data containing PHI must be encrypted while in transit. Refer to the guidance for Amazon EC2 to better understand the available encryption options.

AWS X-Ray

AWS X-Ray is a service that collects data about requests that a customer's application serves, and provides tools that they can use to view, filter, and gain insights into that data to identify issues and opportunities for optimization. For any traced request to a customer's application, they can see detailed information not only about the request and response, but also about calls that their application makes to downstream AWS resources, microservices, databases, and HTTP web APIs. AWS X-Ray should not be used to store or process PHI. Information transmitted to and from AWS X-Ray is encrypted by default. When using AWS X-Ray, do not place any PHI within segment annotations or segment metadata.

Elastic Load Balancing

Customers can use Elastic Load Balancing to terminate and process sessions containing PHI. Customers can choose either the Classic Load Balancer or the Application Load Balancer. Because all network traffic containing PHI must be encrypted in transit end-to-end, customers have the flexibility to implement two different architectures:

Customers can terminate HTTPS, HTTP/2 over TLS (for Application), or SSL/TLS on Elastic Load Balancing by creating a load balancer that uses an encrypted protocol for connections. This feature enables traffic encryption between the load balancer and the clients that initiate HTTPS, HTTP/2 over TLS, or SSL/TLS sessions, and for connections between the load balancer and customer backend instances. Sessions containing PHI must encrypt both front-end and backend listeners for transport encryption. Customers should evaluate their certificates and session negotiation policies and maintain them consistent to the Guidance. For more information, see [HTTPS Listeners for Your Classic Load Balancer](#).

Alternatively, customers can configure Amazon ELB in basic TCP-mode (for Classic) or over WebSockets (for Application) and pass-through encrypted sessions to backend instances where the encrypted session is terminated. In this architecture, customers manage their own certificates and TLS negotiation policies in applications running in their own instances. For more information, see [Listeners for Your Classic Load Balancer](#). In both architectures, customers should implement a level of logging which they determine to be consistent with HIPAA and HITECH requirements.

FreeRTOS

FreeRTOS is an operating system for microcontrollers that makes small, low-power edge devices easy to program, deploy, secure, connect, and manage. FreeRTOS is based on the FreeRTOS kernel,

a popular open source operating system for microcontrollers, and extends it with software libraries that make it easy to securely connect small, low-power devices to AWS Cloud services like AWS IoT Core or to more powerful edge devices running AWS IoT Greengrass.

Data containing PHI can now be encrypted in transit and while at-rest when using a qualified device running FreeRTOS. FreeRTOS provides two libraries to provide platform security: TLS and PKCS#11. The TLS API should be used to encrypt and authenticate all network traffic that contains PHI. PKCS#11 provides a standard interface for software cryptographic operations and should be used to encrypt any PHI stored on a qualified device running FreeRTOS.

Using AWS KMS for Encryption of PHI

KMS keys can be used to encrypt/decrypt data encryption keys used to encrypt PHI in a customer's applications or in AWS services that use AWS KMS. AWS KMS can be used in conjunction with a HIPAA account, but PHI can only be processed, stored, or transmitted in HIPAA Eligible Services. AWS KMS is normally used to generate and manage keys for applications running in other HIPAA Eligible Services.

For example, an application processing PHI in Amazon EC2 could use the `GenerateDataKey` API call to generate data encryption keys for encrypting and decrypting PHI in the application. The data encryption keys would be protected by a customer's KMS keys stored in AWS KMS, creating a highly auditable key hierarchy as API calls to AWS KMS are logged in AWS CloudTrail. PHI should not be stored in the Tags (metadata) for any keys stored in AWS KMS.

VM Import/Export

VM Import/Export enables customers to easily import virtual machine images from existing environment to Amazon EC2 instances and export them back to your on-premises environment. This offering allows customers to leverage existing investments in the virtual machines that you have built to meet their IT security, their configuration management, and their compliance requirements by bringing those virtual machines into Amazon EC2 as ready-to-use instances. Customers can also export imported instances back to their on-premises virtualization infrastructure, allowing them to deploy workloads across your IT infrastructure.

VM Import/Export is available at no additional charge beyond standard usage charges for Amazon EC2 and Amazon S3.

To import customer images, customers can use the AWS CLI or other developer tools to import a virtual machine (VM) image from their VMware environment. If customers use the VMware vSphere

virtualization platform, they can also use the AWS Management Portal for vCenter to import their VM. As part of the import process, VM Import will convert customer VM into an Amazon EC2 AMI, which they can use to run Amazon EC2 instances. Once their VM has been imported, they can take advantage of Amazon's elasticity, scalability and monitoring via offerings like Auto Scaling, Elastic Load Balancing and CloudWatch to support their imported images.

Customers can export previously imported Amazon EC2 instances using the Amazon EC2 API tools. Simply specify the target instance, virtual machine file format, and a destination Amazon S3 bucket, and VM Import/Export will automatically export the instance to the Amazon S3 bucket along with encryption options to secure the transmission and storage of their VM images. Customers can then download and launch the exported VM within their on-premises virtualization infrastructure.

Customers can import Windows and Linux VMs that use VMware ESX or Workstation, Microsoft Hyper-V, and Citrix Xen virtualization formats. And customers can export previously imported Amazon EC2 instances to VMware ESX, Microsoft Hyper-V or Citrix Xen formats. For a full list of supported operating systems, versions, and formats, see [VM Import/Export Requirements](#). AWS plans to add support for additional operating systems, versions and formats in the future.

Auditing, backups, and disaster recovery

HIPAA's Security Rule has detailed requirements related to in-depth auditing capabilities, data back-up procedures, and disaster recovery mechanisms. The services in AWS contain many features that help customers address their requirements. For example, customers should consider establishing auditing capabilities to allow security analysts to examine 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.

HIPAA also has detailed requirements related to maintaining 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 EBS offers persistent storage for Amazon EC2 virtual server instances. These volumes can be exposed as standard block devices, and they offer off-instance storage that persists independently from the life of an instance. To align with HIPAA guidelines, customers can create point-in-time snapshots of Amazon EBS volumes that are stored automatically in Amazon S3 and are replicated across multiple Availability Zones, which are distinct locations engineered to be insulated from failures in other 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.

Moreover, AWS inherently offers a variety of disaster recovery mechanisms. Disaster recovery, the process of protecting an organization's data and IT infrastructure in times of disaster, involves maintaining highly available systems, keeping both the data and system replicated off-site, and enabling continuous access to both.

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 graceful failover 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 designed to provide 99.99% availability.

Using [AWS Elastic Disaster Recovery](#) (AWS DRS), customers can quickly recover applications on AWS, either at the applications's most up-to-date state, or from an earlier point in time.

Document revisions

To be notified about updates to this whitepaper, subscribe to the RSS feed.

Change	Description	Date
Minor update	Minor update	May 12, 2023
Minor update	Updated whitepaper to expanded the available content on services.	September 28, 2022
Minor update	Fix non-inclusive language.	April 6, 2022
Whitepaper updated	Added information about AWS Application Migration Service, and updated information for Amazon ECS	December 6, 2021
Whitepaper updated	Updated information in Amazon Healthlake and Amazon VPC sections	November 9, 2021
Whitepaper updated	Added information about AWS Network Firewall	September 9, 2021
Whitepaper updated	Updated information about Amazon Connect Customer Profiles	August 26, 2021
Whitepaper updated	Added sections Amazon AppFlow and AWS Glue DataBrew	July 22, 2021
Whitepaper updated	Updated navigation and organization.	April 26, 2021
Whitepaper updated	Added the following sections: AWS CodeDeploy, AWS	March 31, 2021

CodePipeline, Amazon Aurora, Aurora PostgreSQL, Amazon Textract, Amazon Polly, Amazon FSx, AWS Auto Scaling, AWS Backup, AWS Elastic Beanstalk, AWS Firewall Manager, AWS Organizations, AWS Security Hub, AWS Serverless Application Repository, VM Import/Export, Amazon HealthLake, Amazon EventBridge. Updated Amazon Aurora section.

[Whitepaper updated](#)

Added section on AWS App Mesh, and updated AWS System Manager content

August 25, 2020

[Whitepaper updated](#)

Added sections Amazon Appstream 2.0, AWS SDK Metrics, AWS Data Exchange, Amazon MSK, Amazon Pinpoint, Amazon Lex, Amazon SES, and Amazon Forecast, Amazon Quantum Ledger Database (QLDB), AWS Cloud Map.

May 7, 2020

Whitepaper updated	Added sections on Amazon CloudWatch, Amazon CloudWatch Events, Amazon Data Firehose, Amazon Managed Service for Apache Flink, Amazon OpenSearch Service, Amazon DocumentDB (with MongoDB compatibility), AWS Mobile Hub, AWS IoT Greengrass, AWS OpsWorks for Chef Automate, AWS OpsWorks for Puppet Enterprise, AWS Transfer for SFTP, AWS DataSync, AWS Global Accelerator, Amazon Comprehend Medical, and AWS RoboMaker.	January 1, 2020
Whitepaper updated	Added sections on Amazon Comprehend, Amazon Transcribe, Amazon Translate, and AWS Certificate Manager.	January 1, 2019
Whitepaper updated	Added sections on Amazon Athena, Amazon EKS, AWS IoT Core and AWS IoT Device Management, Amazon FreeRTOS, Amazon GuardDuty, Amazon Neptune, AWS Server Migration Service, AWS Database Migration Service, Amazon MQ, and AWS Glue.	November 1, 2018

[Whitepaper updated](#)

Added sections on Amazon Elastic File System (EFS), Amazon Kinesis Video Streams, Amazon Rekognition, Amazon SageMaker, Amazon Simple Workflow, AWS Secrets Manager, Service Catalog, and AWS Step Functions.

June 1, 2018

[Whitepaper updated](#)

Added sections on AWS CloudFormation, AWS X-Ray, AWS CloudTrail, AWS CodeBuild, AWS CodeCommit, AWS Config, and AWS OpsWorks Stack.

April 1, 2018

[Whitepaper updated](#)

Added section on AWS Fargate.

January 1, 2018

Updates made prior to 2018:

Date	Description
November 2017	Added sections on Amazon EC2 Container Registry, Amazon Macie, Amazon QuickSight, and AWS Managed Services.
November 2017	Added sections on Amazon ElastiCache for Redis and Amazon CloudWatch.
October 2017	Added sections on Amazon SNS, Amazon Route 53, AWS Storage Gateway, and AWS CloudHSM. Updated section on AWS Key Management Service.

Date	Description
September 2017	Added sections on Amazon Connect, Amazon Kinesis Streams, Amazon RDS (Maria) DB, Amazon RDS SQL Server, AWS Batch, AWS Lambda, AWS Snowball Edge, and the Lambda@Edge feature of Amazon CloudFront.
August 2017	Added sections on Amazon EC2 Systems Manager and Amazon Inspector.
July 2017	Added sections on Amazon WorkSpaces, Amazon WorkDocs, AWS Directory Service, and Amazon ECS.
June 2017	Added sections on Amazon CloudFront, AWS WAF, AWS Shield, and Amazon S3 Transfer Acceleration.
May 2017	Removed requirement for Dedicated Instances or Dedicated Hosts for processing PHI in EC2 and EMR.
March 2017	Updated list of services to point to AWS Services in Scope by Compliance Program page. Added description for Amazon API Gateway.
January 2017	Updated to newest template.
October 2016	First publication

Notices

Customers are responsible for making their own independent assessment of the information in this document. This document: (a) is for informational purposes only, (b) represents current AWS product offerings and practices, which are subject to change without notice, and (c) does not create any commitments or assurances from AWS and its affiliates, suppliers or licensors. AWS products or services are provided “as is” without warranties, representations, or conditions of any kind, whether express or implied. 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.

© 2023 Amazon Web Services, Inc. or its affiliates. All rights reserved.