Automated Safeguards for Kinesis Firehose

Amazon Kinesis Data Firehose is the easiest way to reliably load streaming data into data stores and analytics tools. It can capture, transform, and load streaming data into Amazon S3, Amazon Redshift, & Amazon Elasticsearch Service, enabling near real-time analytics with existing business intelligence tools and dashboards you're already using today. It is a fully managed service that automatically scales to match the throughput of your data and requires no ongoing administration. It can also batch, compress, transform, and encrypt the data before loading it, minimizing the amount of storage used at the destination and increasing security.

Pricing is based on volume of data ingested into Amazon Kinesis Data Firehose, which is calculated as the number of data records you send to the service, times the size of each record rounded up to the nearest 5KB. Please see Amazon Kinesis Firehose Pricing for more details.

Amazon Kinesis Data Firehose is the easiest way to load streaming data into data stores and analytics tools. Kinesis Data Firehose is a fully managed service that makes it easy to capture, transform, and load massive volumes of streaming data from hundreds of thousands of sources into Amazon S3, Amazon Redshift, and Amazon Elasticsearch Service enabling near real-time analytics and insights. Kinesis data delivery stream is the underlying entity of Kinesis Data Firehose. You use Kinesis Data Firehose by creating a Kinesis data delivery stream and then sending data to it. You can send data to the delivery stream by calling the Firehose API, or running the Linux agent we provide on the data source. Kinesis Data Firehose then continuously loads the data into Amazon S3, Amazon Redshift, and Amazon Elasticsearch Service. Please see Amazon Kinesis Firehose Features for more details.

ClearDATA Automated Safeguards evaluate the data source, ensuring that if the data is coming from a Kinesis Data Stream the stream is encrypted which ensures the data processed by Firehose is also encrypted. ClearDATA also ensures that if you choose to send data directly into Firehose, called a DIRECT PUT, that server-side encryption is enabled to ensure the data is encrypted. ClearDATA also limits the Firehose destinations to AWS services covered by ClearDATA Automated Safeguards, meaning that Splunk is not an available destination without a security exception.

Compliance Guidance

Encrypted Storage

HIPAA Technical Safeguard 45 CFR 164.312(a)(2)(iv) requires encryption and decryption addressable standard strongly suggests that you implement a mechanism to encrypt and decrypt electronic protected health information (ePHI). ClearDATA's interpretation of this regulation is that all storage must be encrypted, and as a result the Automated Safeguards for Kinesis Firehose ensure that all data sources, including DIRECT PUT, are encrypted.

Remediation

ClearDATA automatically, and transparently, enables encryption storage for all Firehoses.

Customer Managed Keys

If customers wish to use a Customer Master Key instead of the default key, that can be done via the console or SDK. See https://docs.aws.amazon.com/AmazonS3/latest/dev/UsingKMSEncryption.html for details.

Audit Logging

HIPAA Technical Safeguard 45 CFR § 164.312(b) requires a Covered Entity to "Implement hardware, software, and/or procedural mechanisms that record and examine activity in information systems that contain or use electronic protected health information." Amazon Kinesis Firehose is integrated with CloudTrail, a service that logs API calls made by or on behalf of Amazon Kinesis Firehose in your AWS account and delivers the log files to the specified Amazon S3 bucket. CloudTrail captures API calls made from the Amazon Kinesis Firehose console or from the Amazon Kinesis Firehose API.

Shared Responsibility

ClearDATA will enforce encryption on all data sources, and ensure that the data is sent to a compliant data source such as S3 or Redshift. Customers are not allowed to use Splunk as an output source and any Firehose that is configured to send data to Splunk will be deleted.

If there is an existing Kinesis Data Stream that was created prior to the deployment of Automated Safeguards, and it is not encrypted, when the Data Stream is connected to the Firehose ClearDATA's Automated Safeguards will not remediate the data stream. ClearDATA will only send an alert regarding the unencrypted data stream and it is the customer responsibility to remediate the Kinesis Data Stream.

Please contact your ClearDATA team for a copy of the Responsibilities Matrix.

Exclusion

Disabling automated remediation at done the Firehose level. Please contact ClearDATA Support to request that an exclusion be placed to allow for the instance to be created.

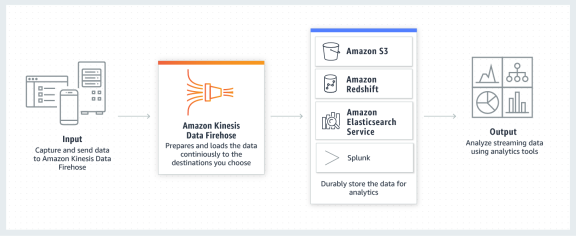
Input
Capture and send data to Amazon Kinesis Data Firehose
Amazon Kinesis Data Firehose
Prepares and loads the data continiously to the destinations you choose
Amazon S3
Amazon Redshift
Amazon Elasticsearch Service
Splunk
Durably store the data for analytics
Output
Analyze streaming data using analytics tools

Users can be added to the **Safeguards-Firehose** group to gain access to create and manage RDS instances.

Item	**ClearDATA**	**Customer**
Enforcement of Automated Safeguards	RA	IC
Remediate any existing unencrypted Kinesis Data Stream that is connected to a Firehose.	IC	RA
Alert when an unencrypted Kinesis Data Stream is detected when Automated Safeguards are deployed	RA	IC
Ensure any service excluded from automated remediation does not contain any PHI/PII	IC	RA
Creation, configuration, and management of all Firehose streams	C	RA