Healthcare Interoperability Stack on AWS

This reference architecture describes how to build a modular interoperability platform to ingest, parse, and store healthcare data of any shape, size, and format. Ingestion of both standard and non-standard data is achieved through connectors with syntactic and semantic parsers responsible for conversion and standardization of data.

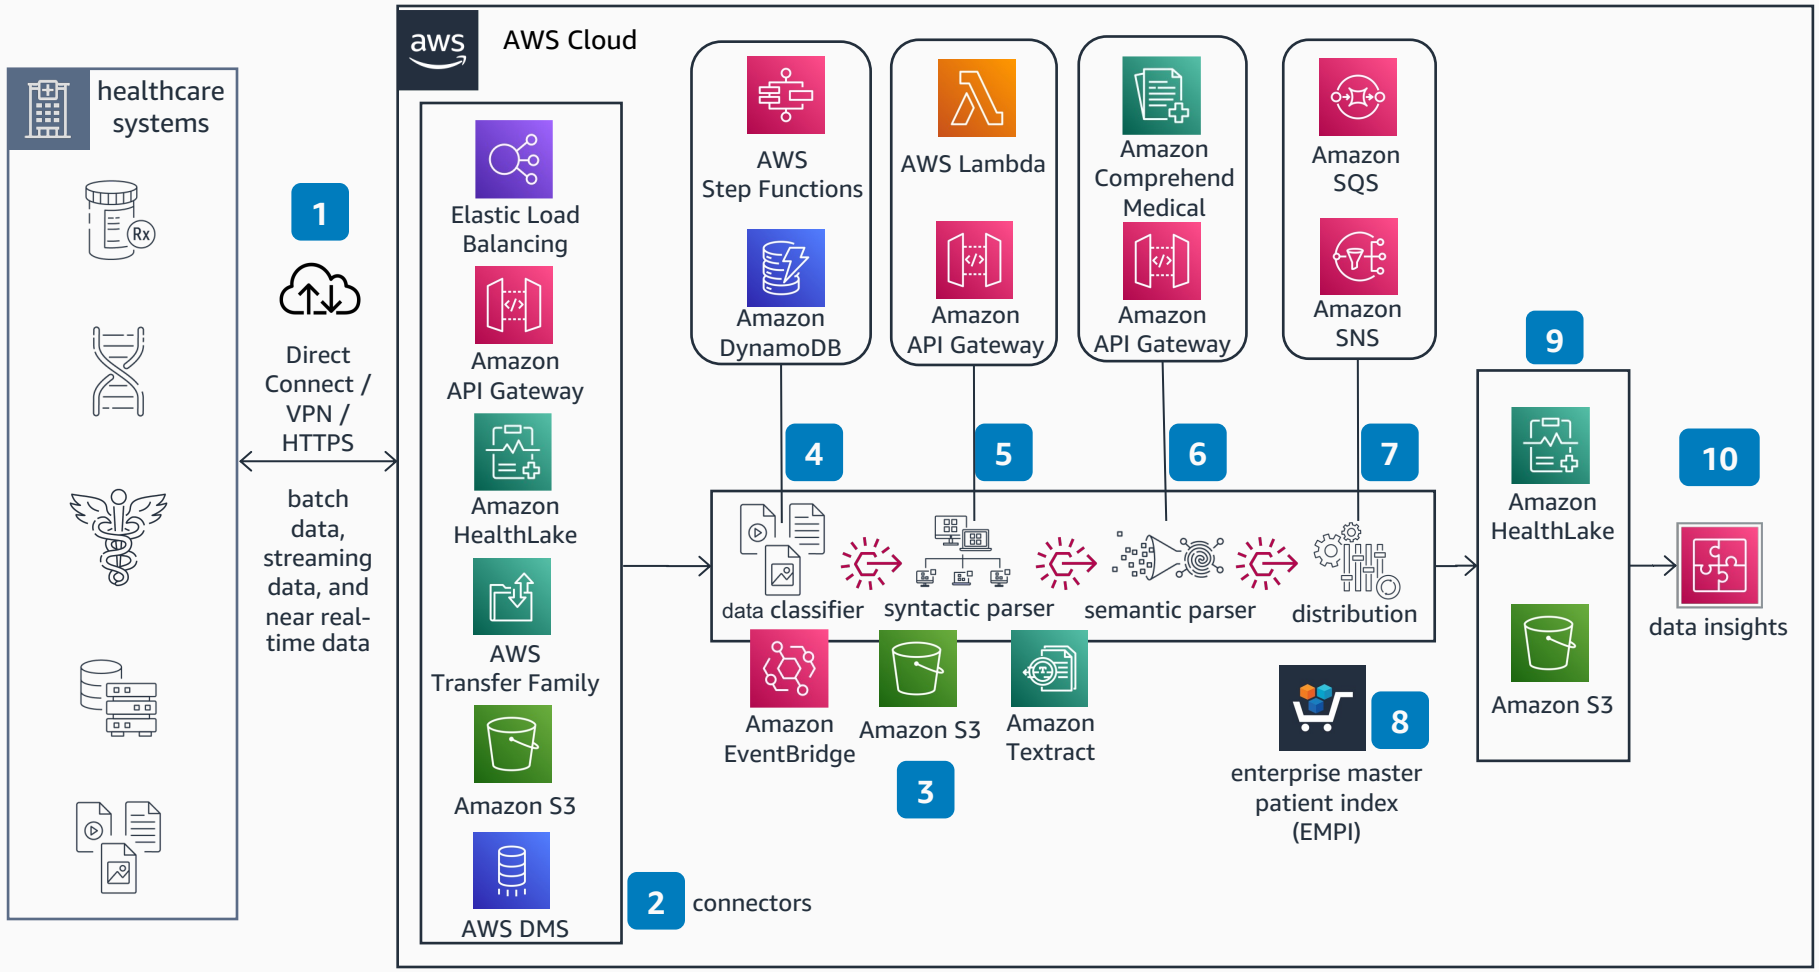

1. On-premises connectivity is required for data ingestion. A private channel such as an **AWS Site-to-Site VPN** or **AWS Direct Connect** or HTTPS connectivity is recommended.
2. Depending on the incoming message type, different connectors will be used. For public-facing **Amazon API Gateway** instances, properly configured **AWS WAF**, a web application firewall is required.
3. **Amazon EventBridge** handles the data processing pipeline orchestration. Intermediate data is stored in **Amazon Simple Storage Service** (Amazon S3) and **Amazon Textract** enables extraction of text from files where required (such as PDFs).
4. A data classifier determines the incoming message type and provides information to the appropriate parser using event rules.
5. Using a set of configurable templates, the syntactic parser transforms incoming data formats into HL7 FHIR if required.
6. The semantic parser standardizes data values into common ontologies.
7. The distribution module is responsible for the fanout of converted messages to appropriate destinations using **Amazon Simple Queue Service** (Amazon SQS) and **Amazon Simple Notification Service** (Amazon SNS).
8. Customers can bring their own identity solutions to normalize resources such as patients and providers.
9. **Amazon HealthLake** serves as a durable storage destination for messages transformed into HL7 FHIR. **Amazon S3** is used to store other resources.
10. Customers can benefit from standardized data to gain insights and to improve care.

Reviewed for technical accuracy January 31, 2023

AWS Reference Architecture