

Transforming Healthcare with near Real-Time Analytics



Customer Challenge

The goal of Modernizing Medicine® is to increase efficiency and improve outcomes by transforming creation and consumption of healthcare information. Its data-driven electronic health records (EHR) and practice management (PM) systems transform the clinical, financial and operational aspects of medical practices.

Why did the customer choose AWS and Cloudwick as an AWS partner?

Modernizing Medicine turned to Cloudwick to support this goal. Cloudwick provided the company with its AWS Modern Analytics QuickStart services, aiding with data ingestion from MySQL to Cloudera CDH, integrating the company's applications for predictive analytics, and providing professional services to ensure Health Insurance Portability and Accountability Act (HIPAA) and Protected Health Information (PHI) compliance.

Cloudwick Solution

Modernizing Medicine has developed a specialty-specific electronic medical records (EMR) platform. The solution incorporates features like digital record tracking, revenue cycle management, along with advanced analytics enabling managers to make informed decisions about patients and their practices. The EMR application uses MySQL databases running on Amazon Web Services. A mix of Hadoop, Microsoft SQL and MySQL data-marts and various visualization tools enables analytics. To support visualization and reporting, Modernizing Medicine developed a Sqoop based data ingestion process that took data from the MySQL sharded environments into a Cloudera-based Hadoop cluster running on AWS EC2 instances. There were about 800 tables to import daily with a total data size of roughly 4.7 TB. Many of these tables were small and a Cloudera CDH cluster was running with a fixed number of nodes regardless of the ingestion workload. These factors resulted in a load time exceeding 48 hours, impacting their ability to use this data to manage their practices. To eliminate the delay and introduce near-real time predictive analytics, Modernizing Medicine moved to transient AWS EMR Spark clusters to handle the data ingestion, providing flexibility, performance, cost savings and reliability improvements. The new data ingestion process runs within 3 hours, enabling Modernizing Medicine to make crucial healthcare and business decisions with near real-time data.

About Modernizing Medicine



Modernizing Medicine offers electronic health records, practice management, revenue cycle management and data analytics for doctors, and has software for dermatologists, ophthalmologists, orthopedic surgeons, gastroenterologists, plastic surgeons, ENTs (or otolaryngologists), urologists and pain management physicians. As part of its business model, Modernizing Medicine uses practicing doctors to help program the software for each specialty.



Solution Architecture

The MySQL environment for Modernizing Medicine comprised 12 different shards. The redesigned ingest process leverages a transient, HIPAA compliant EMR Spark cluster which connects to the MySQL pods in parallel with different Spark executors using Spark JDBC. The loaded Spark Dataframes save the data to disk before sending it to S3 using 'S3Distcp'. During a later step, the Cloudera cluster ingests the data using Distcp from S3. The entire process takes around 3 hours, compared to the earlier Sqoop-based ingestion time of 48hrs. Modernizing Medicine turned to Cloudwick to support this effort. Cloudwick provided the company with its AWS Modern Analytics QuickStart services, aiding with data ingestion from MySQL to Cloudera CDH, integrating the company's applications for predictive analytics, and providing professional services to ensure Health Insurance Portability and Accountability Act (HIPAA) and Protected Health Information (PHI) compliance.

Results and Benefits

Cloudwick chose Amazon Simple Storage Service (Amazon S3) for storing the PHI information and AWS Elastic Map Reduce (EMR) for resizable compute capacity. To achieve HIPAA compliance, Cloudwick set up encryption-at-rest and encryption-in-transit for both Cloudera Data Hub and EMR. We secure the EMR transient cluster and the Cloudera Data Hub using AWS VPC, Subnet ACLs, Routing tables, Sentry, IAM role, LUKS encryption, HDFS transparent encryption, Kerberos authentication and S3 bucket policies. Modernizing Medicine now enjoys significant improvements in ingest performance, thus enabling analytics with near real-time data which is a crucial capability in medicine. This makes it easier to diagnose patients more accurately and enables the company to make crucial healthcare and business decisions on a day-to-day basis. With the new solution, the company enjoys a cloud-based modernized business intelligence and analytics infrastructure, managed by Cloudwick. With near real-time predictive analytics, it gives doctors the information they need to treat patients and improve healthcare.

About Cloudwick

Cloudwick is the leading provider of enterprise business and technology modernization services and solutions to the Global 1000 and helps enterprises gain competitive advantage from open source, data lake, big data, cloud and advanced analytics. Cloudwick is an advanced AWS partner with Big data and machine learning competency.

Cloudwick