PatientHealthStorageSystem:

Secure HIPAA-Complaint Collaboration — A Case Study Written by
Gregory Harman, Managing Partner, BigR.io, LLC

BigR.io
EMPOWERING DATA

PatientHealthStorageSystem:

Secure HIPAA-Complaint Collaboration — A Case Study Written by
Gregory Harman, Managing Partner, BigR.io, LLC

About the Author

Greg has an extensive background in data infrastructure and enterprise architecture. As a recognized industry leader, Greg's career has focused on data management platforms, enterprise architecture, integration, and data mobility. His past work has been utilized by organizations ranging from Fortune 100 companies to startups and has been commercialized for millions of end users.

Visit Greg on LinkedIn

Prior to starting BigR.io, Greg spent more than ten years in CTO roles, developing technology strategies and platforms to drive business and building and managing multifaceted teams to implement and support them. Greg believes strongly in agile development, open source technology, analytical approaches to problem solving, and data-driven solutions. Look for white papers and case studies from Greg, like this paper on Big Data Architecture: http://bigr.io/architecture/

Greg earned both his Bachelors and Masters Degrees in Electrical Engineering & Computer Science from the Massachusetts Institute of Technology (MIT).

Abstract

BigR.io architected and implemented a distributed, secure collaboration and document sharing platform that allows parents to exchange documents and conduct collaborative discussions with their children's healthcare and educational providers, even across provider organizations. This platform utilizes a distributed document storage system with multi-layered security that allows this sensitive information to be administered without a centralized point of access to PHI, relieving the accompanying responsibilities and risks of storing and utilizing this data. The platform is designed to support HIPAA and other privacy regulations. An anonymized data extraction pipeline supports advanced analytics without compromising user privacy and security.

About BigR.io

BigR.io is a technology consulting firm empowering data to drive analytics for revenue growth and operational efficiencies. Our teams deliver software solutions, data science strategies, enterprise infrastructure, and management consulting to the world's largest companies. We are an elite group with MIT roots, shining when tasked with complex missions: assembling mounds of data from a variety of sources, building high-volume, highly-available systems, and orchestrating analytics to transform technology into perceivable business value.

BigR.io
EMPOWERING DATA

About BigR.io (cont.)

With extensive domain knowledge, BigR.io has teams of architects and engineers that deliver best-in-class solutions across a variety of verticals. This diverse industry exposure and our constant run-in with the cutting edge, empowers us with invaluable tools, tricks, and techniques. We bring knowledge and horsepower that consistently delivers innovative, cost-conscious, and extensible results to complex software and data challenges. Learn more at www.bigr.io

The Situation

Despite privacy rules and standards, such as the Health Insurance Portability and Accountability Act (HIPAA) and supporting technologies within the internal IT and data systems at healthcare organizations, a secure communications gap exists between healthcare organizations and their patients. While some larger organizations do support document access portals, these portals are not ubiquitous and almost never span multiple organizations, leading patients or their caregivers to often resort to email and fax when coordinating care across a team of professionals at different organizations.

Even in the cases of organizations that offer more advanced portals, these portals and their underlying databases provide a single point of attack for malicious individuals as evidenced by a recent spate of hospital data being encrypted and ransomed. While this activity affects the organization's bottom line, it also exposes patients to privacy breaches and prevents them access to their own medical records.

A team of professionals from the MIT and Harvard communities have launched a new initiative to address these issues for coordinated healthcare and education of children. They are developing a cross-organization document sharing and collaboration portal that is built on a foundation of distributed control and privacy.

Objectives

BigR.io was engaged to architect and implement the secure, distributed document storage system that underpins this secure portal. Specific design objectives included:

- Distributed storage such that Protected Health Information (PHI) and other sensitive documents are stored in parent-controlled datastores. A breach of the core portal cannot expose PHI, and a breach of any specific datastore exposes PHI for only one patient.
- Flexible storage architecture allowing these parent-controlled datastores to exist on top of heterogeneous technologies such as AWS S3, Google Drive, and local file storage.
- Encryption of all PHI using a distributed key paradigm.
- Secure collaboration on PHI compatible with these storage and security constraints.
- Mechanism for removing Personally Identifiable Information (PII) and aggregating PHI data into (opt-in) data sets to support scientific research.

Engagement

This engagement was addressed in three phases:

Design & Architecture – The BigR.io team worked with stakeholders to understand specific use cases, constraints, and pain points. During this phase, candidate architectures were developed and presented with a description of relevant tradeoffs, and stakeholders were guided through a collaborative decision-making process culminating in single target architecture. An implementation plan was developed allowing the customer to divide the project into parallel paths, off shoring the low-risk portal front end development while continuing to engage BigR.io to develop the more complex data storage design.

Back-End Implementation – The back-end architecture was implemented by the BigR.io team, and subjected to third-party security testing and deployed to beta organizations. Additional roadmap features and future risks were identified, designed and backlogged for ongoing development.

Production Support and Transition – BigR.io has assisted the transition of technology and development to the customer's internal development team, providing both general knowledge transfers for ongoing development, as well as production support for early portal rollouts.

Results

The architecture that was selected is exposed as a REST API deployed on Amazon Web Services (AWS) and has the following characteristics:

- The REST API was implemented using AWS Lambda, API Gateway, and SNS.
- AWS Cognito and OAuth 2.0 were used to provide flexible authentication across data storage systems.
- Initial distributed storage implementations were created using Google Drive (primary) and AWS S3 (secondary)
- An asymmetric encryption scheme protects symmetric keys that support a mode of collaboration in which no unencrypted key is stored in any data system and every document is encrypted with a unique key for each user that has access to that document
- The secure nature of this architecture exceeds the security offered by any centralized document store, and carries the additional benefit of releasing the customer from the HIPAA requirement of Business Associate Agreements (BAAs) with infrastructure providers such as AWS and Google.