BigRio LLC
Harvard Square,
One Mifflin Place, Suite 400,
Cambridge, MA 02138
(617) 500-5093 | info@bigr.io
www.BigR.io

PATIENT HEALTH DATA-STORAGE SYSTEM: SECURE HIPAA-COMPLAINT COLLABORATION

By Greg Harman

ABSTRACT

BigRio architected and implemented a distributed and secure, document sharing platform. This platform allows parents to exchange documents and have collaborative discussions with their children's healthcare and educational providers across provider organizations. This platform utilizes a distributed document-storage system with multi-layered security which allows sensitive information to be administered without a centralized point of access to private health information (PHI). This platform relieves providers of the risks and responsibilities of storing and utilizing PHI. The platform is designed to support HIPAA and other privacy regulations—including an anonymized data extraction pipeline that supports advanced analytics without compromising user privacy and security.

PROBLEM

Despite privacy rules and standards such as HIPAA and other supporting technologies within the internal IT and data systems of healthcare organizations, a secure communications gap exists between healthcare organizations and their patients. While some larger organizations do support document access portals, they are not ubiquitous and almost never span multiple organizations. This makes patients and their caregivers resort to emailing and faxing when coordinating care across a team of professionals at different organizations.

Even in the cases where organizations offer more advanced portals, they often have underlying databases which provide a single point-of-attack as evident 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 from accessing 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 the education of children. This team developed a cross-organization document sharing and collaboration portal that is built on a foundation of distributed control and privacy.

OBJECTIVES

BigRio was asked to architect and implement the secure, distributed document storage system that underpins this secure portal. **Specific design objectives included**:

- Distributed storage such that PHI and other sensitive documents are stored in parent-controlled datastores; wherein 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 the 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.
- Mechanisms for removing Personally Identifiable Information (PII) and aggregating PHI data into (opt-in) data sets to support scientific research.

METHODS

This engagement was addressed in three phases:

DESIGN & ARCHITECTURE

The BigRio 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 a 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 BigRio to develop a more complex data-storage design.

BACK-END IMPLEMENTATION

The back-end architecture was implemented by the BigRio 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

BigRio has assisted in the transition of technology and development to the customer's internal development team. This provides 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 and supports a mode of collaboration in which no unencrypted key is stored in any data system. Additionally, every document is encrypted with a unique key for each user who 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). This is done through infrastructure providers such as AWS and Google.

BigRio LLC, Harvard Square, One Mifflin Place, Suite 400, Cambridge, MA 02138
(617) 500-5093 | info@bigr.io | www.BigR.io