

EOS Platform Topology

Data Storage, Processing, Endpoints, and Flow

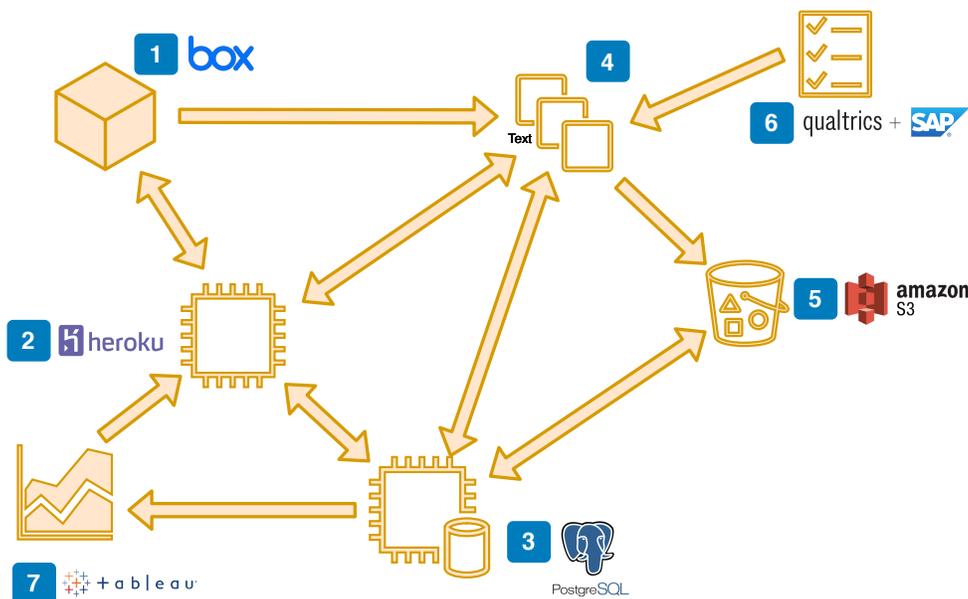
EOS maintains a proprietary Software as a Service (SAS) portal that is used to perform the following functions:

- Store school partnership information and data.
- Document interactions between school resources and students.
- Store and interpret survey information collected from school staff and students.
- Provide reporting about individual students and also collective progress as it relates to the opportunity gaps in access to curriculum.

The portals architecture is a combination of custom development using the latest web technologies. These features and functions are maintained by internal resources and at times are supplemented by outside contractors who agree to be bound by EOS' non-disclosure and follow FERPA requirements. The portal is not all custom, as we do partner with external industry leading package technologies and platforms to ensure we have the latest advances in technologies as It relates to:

- Web site hosting
- Surveying
- Reporting
- Cloud based file storage

The following diagram shows the current technologies that are used in portal and is updated on a yearly basis based on any significant changes to technology. The focus of our architecture is summarized in three areas: Security, Scalability and Functionality.



1 The portal is integrated with BOX.com, a FERPA compliant storage service. Access to it is controlled via portal permissions, even for EOS staff. However, EOS staff have the option of accessing the Box folders directly.

Partnership Directors and Partners share office documents, which may contain PII, through this (and only this) interface. It is also the location where student photos are dropped and later processed by the portal. Data is encrypted in flight and at rest.

2 The portal hosted on AWS EC2 through Heroku. It scales dynamically based on server loads. No data is stored in the EC2 instance, only code. Communication between the portal and client browsers or services is restricted to SSL encrypted ports.

3 The RDMS is comprised of a primary Postgres database and a secondary follower database hosted on AWS EC2 through Heroku. It is the main repository of all the data we collect. Access to the database is restricted to the Portal, Tableau, and a few staff who perform either maintenance or analysis. Data is encrypted in flight and at rest. Although the follower mirrors the primary's data, all data is backed up nightly. Full and partial rollbacks are possible within minutes.

4 The portal is supplemented by a number of workers, independent processes that operate in the same EC2 instance and inherit its security. Some are created on-demand while others are persistent. Workers are used for maintenance and processing, especially CSV uploads, student photos, and survey responses. No data is kept in the workers, and no user or programming interfaces are available. Workers may store and retrieve to/from Postgres (3) or store to S3 (5).

5 Hard assets, such as partner-supplied CSV files, student photos, and portal generated PDFs (e.g. Student Insight Cards), are securely stored in AWS S3. Access to the bucket contents is restricted to the portal's domain. The portal can serve out the contents using secure, self-expiring links (time based and single-use depending on the purpose). Each district's data is individually sequestered to prevent data bleed between entities. Data is encrypted in flight and at rest.

6 Qualtrics, a FERPA compliant survey platform, administers staff and student surveys. The responders interact directly with Qualtrics, although EOS takes care to brand it for the individual schools. All survey data, which may contain PII, is stored in Qualtrics for one year. It is retrieve from Qualtrics in near-realtime by portal workers (4) where it is processed and analyzed. Results are stored in the RDMS (5) for as long as the Partnership Data Agreement requires.

7 Tableau, a FERPA compliant data visualization platform, accesses aggregate student and staff information for reporting purposes. Because data is presented in aggregate, Tableau does not generally expose PII.

Access to the visualizations and reports is restricted via portal permissions. Data is pulled from the RDBMS (3) and rendered in the portal (2) to the client browser. EOS self-hosts a federated Tableau server on AWS EC2.