

Helping students learn with Course Hero, powered by Amazon SageMaker

[Course Hero](#) is an online learning platform that provides students access to over 25 million course-specific study materials, including study guides, class notes, and practice problems for numerous subjects. The platform, which runs on AWS, is designed to enable every student to take on their courses feeling confident and prepared. To make that possible, Course Hero is equipped to do some learning of its own, using [Amazon Machine Learning](#) (Amazon ML), which powers Course Hero and serves as its primary artificial intelligence and ML platform.

The artificial intelligence group at Course Hero is tasked with building the company's semantic knowledge graph. This constantly expanding graph enables students to access personalized learning experiences and gives educators tools to create unique course content.

Most aspects of Course Hero's offerings rely on AWS in some form or another (either compute or ML). For example, [Amazon OpenSearch Service](#) (Amazon ES) powers the search function that students and educators use to search for materials. The Amazon OpenSearch platform allows the Course Hero team to write custom implementations through its [API extension plugin](#). The plugin gives them the flexibility to create relevant user experiences, even for more esoteric searches that require locally dense semantic search capability.

Students and educators search within Course Hero's document library (which is freely accessible) in exchange for uploading one's own content. Course Hero does not accept all documents as publishable library material; documents gain acceptance to the library after going through a cloud-driven vetting process. When new documents are uploaded, an artificial intelligence platform running on [Amazon EMR](#) and [Amazon SageMaker Inference Pipelines](#) checks and validates the documents for fraud, honor code violations, copyright infringements, and spam.

The documents that pass quality review then move to further processing and tagging using ML models that are built on the label data that [Amazon SageMaker Ground Truth](#) has collected. This document labeling enables Course Hero to learn what kind of materials are used by a given student, then predict what else might be useful for them.

By personalizing the experience in this way, Course Hero provides each user with relevant content for their studying needs. With the right content in hand, students gain a deeper understanding and meet their learning objectives more efficiently.

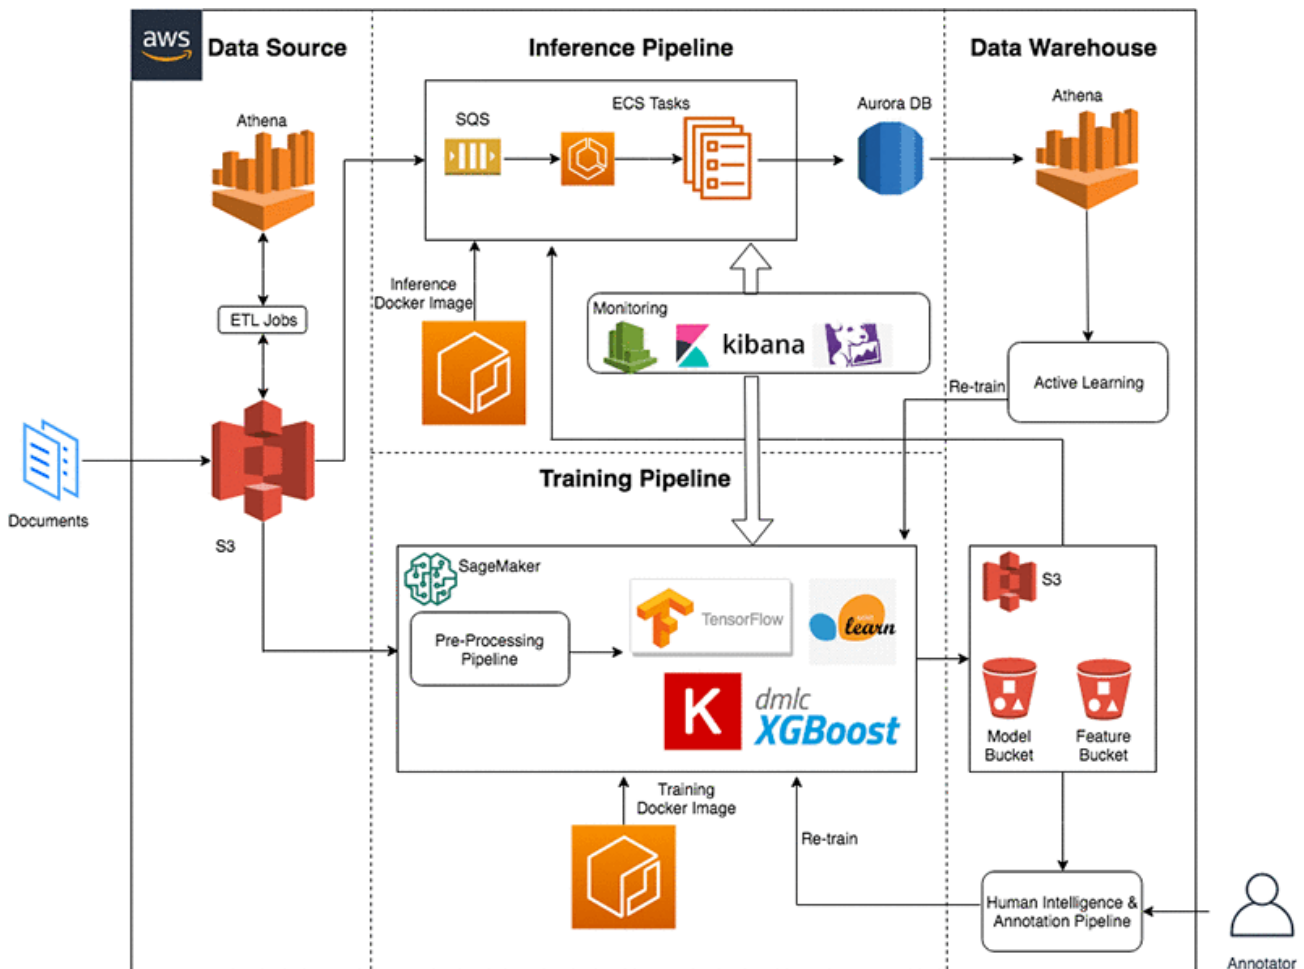
AWS is a comprehensive platform for Course Hero. In addition to the student-facing use cases described above, Course Hero uses AWS services for ad hoc analyses, data exploration, trend discovery, real-time analytics, fraud detection, and more. Course Hero constructs its data platform using key AWS services, including the following:

- [Amazon S3](#)
- [Amazon Athena](#)
- [AWS Glue](#)
- [Amazon EMR](#)
- [Amazon QuickSight](#)
- [Amazon RDS](#)



Course Hero's planning, tracking, and monitoring platforms also use [Kibana](#), [Logstash](#), and [Amazon CloudWatch](#) to keep all monitoring and service centers running smoothly.

The following diagram shows how all of these components work together.



To further augment the existing AWS technology that powers Course Hero, the team is exploring additional Amazon services, including [Amazon Forecast](#), for time series and financial forecasting. It is also looking at possibilities using [Amazon Echo](#) that will allow users to ask questions via Alexa,

Course Hero's Saurabh Khanwalkar, the Director of Machine Learning & Search Sciences, says, "The entire machine learning, engineering, and artificial intelligence stack runs on AWS. From our CI/CD pipelines to our code workbench to our end-to-end model development to our staging and production inferences, we're on AWS."