

Overview

The Modjoul software infrastructure is a highly scalable stack on top of IBM Bluemix that supports tera-byte scale data ingestion from Modjoul SmartBelts. The compute platform includes a large Spark cluster that is capable of performing activity classifications using machine learning for thousands of belts concurrently. The platform is built for scale and high availability using best-of-breed big data processing technologies on a robust Platform-as-a-Service.

Modjoul is built on top of the Watson Internet of Things Platform (WIoT) and uses a highly scalable and resilient Spark Streaming infrastructure to process the data from thousands of belts in near real-time and at scale. The Spark cluster, hosted on IBM BigInsights, provides a massive amount of compute power required for performing machine learning on a large scale.

Cloudant, an industry leading managed NoSQL storage technology by IBM, is used for powering the web and mobile apps and provides a highly scalable and performant storage capability.

Dashboards

The Modjoul dashboards can be accessed via web or mobile app. The web dashboard is an Angular-based responsive Single Page Application supported on major browsers. The Modjoul mobile app is a hybrid app which supports Android JellyBean and above and iOS 9 and above on phone and tablet form factors. The web dashboard and mobile apps are built using mobile web technologies and related standards.



Example of the web dashboard

Machine Learning at Scale

The Modjoul system processes sensor data from thousands of SmartBelts in the field in near real-time. The processing includes activity classification and computing metrics for display on the mobile and web dashboards. This is a very compute intensive task. Performing this at scale and within a short time window requires a distributed processing architecture built on top of Apache Hadoop and Apache Spark, both of which are hosted on IBM BigInsights.