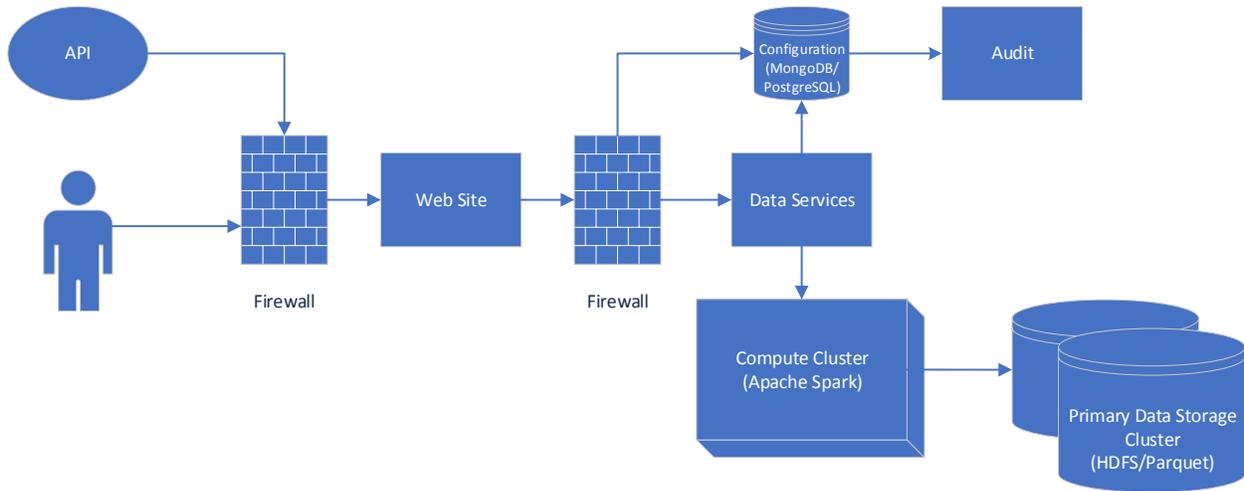


Ki™ Technology Overview

Overview

Ki is architected as a multi-tier highly scalable microservice application, built to address the unique data, funding, analytics, deal structuring, and reporting needs of the structured finance market. The platform is built upon leading edge “big-data” technologies, and is capable of handling data workloads of all sizes in a very efficient and cost-effective fashion.

Ki has been architected using best of breed technologies across all layers of the technology stack. Our design principles ensure that the solution provides high performance, options for high availability, and a deployment model which enables scalability in all dimensions.



The front-end of Ki consists of a single-page application built upon technologies such as React, Javascript, and Node.js. This approach provides a highly responsive user experience.

Data operations are executed via a collection of microservices that are written in Scala, and utilize Apache Spark as a data processing engine. Apache Spark is a scalable analytics engine, capable of handling datasets of all sizes. Solver operations, such as asset selection for funding activities, is run using a high-performance mixed-integer solver.

The primary data layer of Ki is housed in a scalable HDFS filesystem, storing data in the Apache Parquet format. Apache Parquet is a column-oriented storage format that is ideal for storage of structured data, and the preferred format for big-data technologies, such as Apache Spark. By using Parquet for storage rather than traditional row-based relational database technologies, data is stored in an extremely efficient fashion, in a structure that is perfectly aligned with Ki to

deliver optimal performance. Additional SQL and NoSQL technologies such as PostgreSQL and MongoDB are also used within Ki for storing lightweight metadata and configuration.

Ki is deployed using container technology, which allows the platform to be operated in a secure fashion across a multi-node cluster environment. This model enables Ki to be highly scalable, and accommodates the option of deployment across multiple zones in active/passive or active/active configurations.

Security

Authentication and Authorization

Web forms-based authentication and role-based authorization are provided as standard functionality in Ki. Organizations can directly manage users and entitlements users via an administration console, which is available to users who are assigned to the “User Admin” role in Ki. Support for SAML 2.0 is also supported for organizations that desire single sign-on authentication integration with enterprise identity providers.

Hosting

Moody's hosts Ki client sites within a production Virtual Private Cloud (VPC) in Amazon Web Services. This VPC is segregated into separate Virtual LANs (VLANs), with firewalls between them to ensure access is restricted to authorized personnel. The outermost VLAN, which provides external access to web users of the application, permits access only to TCP web traffic operating on port 443, secured using TLS version 1.2.

A second tier, separated from the front-end tier via an additional firewall, provides all data services for the application.

All client data is stored in dedicated volumes on dedicated server instances in the AWS environment. When requested, these volumes may also be encrypted.

Installation, maintenance, and operational access to the environment is restricted to authorized IT personnel only. In order to provide support for Ki, technical support personnel are assigned read-only access to the machines that host Ki for each client environment.

All servers within the environment are hardened based on Moody's Information Security standards. We also perform periodic vulnerability assessments across the application.

Data Exchange

Data can be submitted into Ki either manually via the application, or programmatically using Ki's RESTful API. In addition to Client-driven custom programmatic interaction with the Ki API, third-party data pipelining tools such as Talend can be used for data submission. Talend, in particular, is integrated with Ki and when necessary, Moody's works closely with its clients to define automated jobs for data conversion, scheduling and other relevant data exchange procedures.

Data within Ki can be accessed outside of the system via Ki's report functionality. Reports can be retrieved via standard HTTPS requests. Report output is available in multiple formats, including PDF, XLS, CSV, and JSON. The JSON representation of the data in these reports can be easily consumed using a number of standard applications, including Excel via PowerQuery, JasperReports, Tableau, etc. Natively, Ki interacts with JasperReports for pixel-perfect report creation and Excel for traditional structured finance related report templates (e.g. settlement statements & investor reports).

System Availability

Currently, Moody's has AWS production datacenters in a number of geographic regions, including US East (Virginia), US West (Oregon), Ireland, and Australia. Additional regions will be added as options over time. Ki can be hosted in any or all of these regions, depending on client requirements. Hosted systems are monitored 24x7 by Moody's IT, providing mission critical-level uptime.

Accommodation for disaster recovery is accomplished via nightly system snapshots.

In addition to the standard system monitoring and disaster recovery, high availability both within and across availability zones, is offered as an option.