



Sqrri Enterprise

Powering Secure, Real-Time Applications

Sqrri Enterprise is a secure, massively scalable Big Data platform that powers real-time applications and is built with Apache Accumulo™ and Hadoop.

Sqrri Enterprise extends the capabilities of Accumulo with additional data ingest, security, and analytical features that help unlock the power of Big Data. *Sqrri Enterprise* excels in environments that require multitenancy, strong security controls, petabytes of data, flexible schemas, and real-time analytics.

Sqrri Enterprise features:

- Rapidly develop real-time dashboards and applications for Big Data
- Utilizes a variety of open source projects, such as Hadoop and Accumulo
- Multi-petabyte scalability and stability at scale
- Works with static or streaming data feeds
- Automated secondary indexing and JSON document support
- Fine-grained access controls for unstructured, semi-structured, and structured data
- Encryption of data-at-rest and data-in-motion
- Easy-to-use analytic APIs for search, statistics, graph analysis, and SQL
- Integrations with business intelligence tools
- Compatibility with open source or commercial distributions of Hadoop
- Run on bare metal or in public or private clouds
- Utilize low cost commodity hardware

Sqrri Enterprise

Sqrri Enterprise is a secure, multitenant Big Data platform for building real-time applications that use massive amounts of structured, semi-structured, and unstructured data. It scales elastically to tens of petabytes of data and enables organizations to eliminate their internal data silos through fine-grained security controls. *Sqrri Enterprise* supports development of both exploratory and operational real-time applications and helps companies find hidden value in their data via dynamic and interactive queries. *Sqrri Enterprise's* architecture consists of four layers, as depicted below:

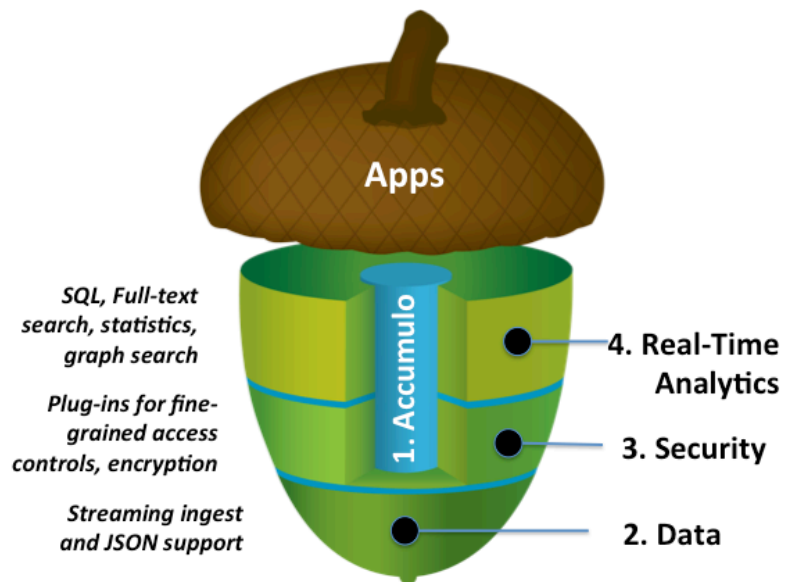


Figure 1. Sqrri Enterprise Technology Architecture

1. Apache Accumulo. Apache Accumulo powers the core of *Sqrri Enterprise*. Accumulo is a sorted, distributed key/value store that integrates natively with Apache Hadoop. It can process petabytes of structured, semi-structured, and unstructured data as a robust, scalable, and real-time data storage and retrieval system. Accumulo possesses fine-grained access controls (i.e., “cell-level” security) that enable granular tagging of key/value pairs with visibility labels. Accumulo’s innovative iterator framework distributes computation across your cluster to achieve real-time analytics.

What People Are Saying

"Wikibon practitioners tell us that increasingly, value from their big data initiatives will come from making applications work with multiple data sources in real-time" said David Vellante, Chief Analyst at Wikibon.org. "The challenges of delivering on this vision are substantial, however, and include constraints around performance, application architecture, scale, data access and skill sets. Apache Accumulo provides a core architectural foundation to solve these problems, and when combined with *Sqrrl Enterprise*, has the potential to dramatically simplify Big Data application development and unlock data sets that are sealed off in silos."

Example Use Case

Sqrrl Enterprise can be utilized to rapidly develop cybersecurity analysis applications across diverse datasets, such as logs, configuration info, system and audit trails, identity context, network flows and anomalies, external intelligence feeds, full packet and DNS capture, email /social activity, and customer transactions. These datasets are analyzed using to detect suspicious behavior in real-time and analysts are alerted when thresholds are triggered. Access to sensitive data, such as PII, can be restricted based on organizational roles.

2. Data. *Sqrrl Enterprise* can ingest up to millions of records per second as either static or streaming data feeds. These records are converted into hierarchical JSON documents, giving *Sqrrl Enterprise* document store functionality. *Sqrrl Enterprise* uses the Hadoop Distributed File System (HDFS) for storage and natively connects with Hadoop MapReduce and other Hadoop projects. *Sqrrl Enterprise* can plug into commercial or open source distributions of Hadoop.

3. Security. *Sqrrl Enterprise* builds on Apache Accumulo's fine-grained access controls and provides seamless integration to Enterprise Identity and Access Management (IAM) Systems. This allows data access to be controlled at the cell-level. *Sqrrl Enterprise* can also support encryption of data-in-motion and data-at-rest, database auditing and monitoring tools, and has been thoroughly tested for software vulnerabilities.

4. Real-Time Analytics. *Sqrrl Enterprise's* analytics layer is designed to make it easier for developers to create powerful Big Data analytic applications. *Sqrrl Enterprise* APIs support real-time full-text search, statistics, graph/network analysis, and SQL. These APIs are compatible with popular programming languages, such as Python, Ruby, and Java, and they allow web developers to develop operational or exploratory applications for Big Data using *Sqrrl Enterprise's* extensible secondary indexing techniques.

Applications. Applications developed on top of *Sqrrl Enterprise* can be created in a variety of ways. *Sqrrl* works directly with clients to develop Big Data applications that solve specific business problems. *Sqrrl Enterprise's* APIs are also simple enough for clients to rapidly develop their own applications. *Sqrrl Enterprise* can also integrate with Tableau through Tableau's Data Engine API.

Benefits

Sqrrl Enterprise is most beneficial for organizations that want to build real-time applications on top of Big Data but are constrained in doing so due to limited resources. The most important requirement for building real-time applications across distributed computing systems is high-end Big Data engineering skills. *Sqrrl Enterprise* greatly simplifies Big Data application development and bridges the gap between traditional application developers and distributed computing experts.

Sqrrl Enterprise also greatly improves information sharing within and across organizations that have security- or privacy-sensitive information. By utilizing Accumulo's fine-grained access controls, *Sqrrl Enterprise* helps organizations bring together datasets that were previously kept segregated for security reasons. When these security controls are combined with Accumulo's petabyte scalability, organizations can bring their diverse datasets together for real-time analysis on a single platform.

Who We Are For

- CTOs and CIOs that desire to integrate datasets across an organization and provide developers with a platform to build analytical tools on top of the data.
- Application developers and business unit executives that wish to develop real-time applications on top of Big Data to solve critical business problems, such as risk management, cybersecurity analysis, and deep customer analysis
- Infrastructure managers who seek low cost and low administrative overhead solutions for storing Big Data and deploying Big Data applications
- Business users who want real-time applications that leverage Big Data

ABOUT Sqrrl

Sqrrl was founded in 2012 by creators of Apache Accumulo. With their roots in the U.S. Intelligence Community, Sqrrl's founders have deep experience working with and building applications for complex petabyte-scale datasets. Sqrrl is headquartered in Cambridge, MA and is a venture-backed company with investors from Matrix Partners and Atlas Venture.



275 Third Street
Cambridge, MA 02142

P: (617) 902-0784
E: info@sqrrl.com

@sqrrl_inc
sqrrl.com