



# Make Business Intelligence Work on Big Data

**Speed. Scale. Simplicity.**

## Put the Power of Big Data in the Hands of Business Users

Connect your BI tools directly to your big data without compromising scale, performance, or control.

## Easily Access Big Data for Analytics

Make your big data directly accessible and interactive for BI and analytics – no more data movement. Describe complex data as simple measures and dimensions that anyone can understand and use.

## Get Performance Without Complexity

Let AtScale's query optimization engine take care of the complexity of creating and maintaining aggregations. Guarantee performance, at scale, for user-generated queries.

## Interactive Analysis on Big Data, from Your Favorite BI Tool

With AtScale, you can support current business processes, allowing analysts to use the BI tools they know and love.

AtScale allows data in Hadoop and other big data systems to be accessed by BI tools in a format users understand. IT gets the control and security they need, while business users get the interactivity and timeliness the business needs.

### ▲ BE MORE EFFICIENT

#### Analyze the Data Where it Lies

Scale-out BI with no data movement, marts or extracts

### ▲ LEVERAGE CHOICE

#### Use Your Favorite BI & Analytics App

Capitalize on investments in tools, people, skills and preferences

### ▲ GAIN AGILITY

#### Unify Data Definitions

Deliver consistent, fast updates across all models and reports

### ▲ DELIVER CONTROL

#### Empower Business Users + IT Alike

Enable self-service insights on big data with security + governance





*Adaptive cache is possibly one of the most meaningful breakthroughs in this space. We put the AtScale Adaptive Cache technology through a test on 57 billion rows of data. The results were 10–20 times faster.*

—Richard Langlois  
Director, Enterprise Data



## Intuitive Business Intelligence Design

Data modelers can interact directly with data in Hadoop and other big data systems via the AtScale Design Center. Modelers can design virtual BI cubes using familiar workflows and intuitive drag-and-drop interactions.

### ▲ Familiar BI Modeling Concepts

Define and visualize virtual relational models on top of data in your Hadoop and big data platforms using well-known BI and standard SQL concepts you already know. Create measures, dimensions, and relationships using drag-and-drop.

### ▲ Rich Multi-Dimensional Support

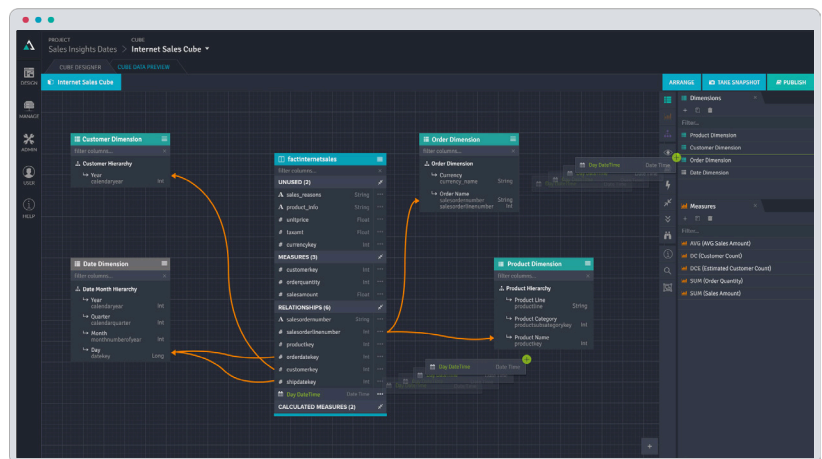
Design hierarchical dimensions for interactive analysis and exploration. Logically model dimensions and relationships of the underlying data format and schema in both Hadoop and big data platforms.

### ▲ Collaborative Virtual Cube Design

Collaborate on virtual BI models with other modelers. Take snapshots at any point in time, and restore previous versions if needed.

### ▲ ‘No Data Movement or ETL’ Virtual Model

Overlay a virtual cube on data in your Hadoop or big data cluster—no need to physically move or normalize your data up front. Analysts can create hierarchies, measures, metadata, and calculations that traditional and new BI tools can use to query data live in a data lake.





*A modern analytic platform standardizes on what's important; an interactive business view of big data that any tool or application can access. In this regard, AtScale is ahead of the pack.*

—Wayne Eckerson  
Principal Analyst



## Smart Query Optimization

AtScale takes the pain out of BI on big data by building and maintaining on-demand aggregate tables. It uses advanced machine-learning algorithms to optimize BI query workloads on-demand, and delivers the performance that users have come to expect from their legacy BI systems.

### ▲ Automated Aggregate Creation and Tuning

With AtScale's virtual BI cube definitions and end-user query patterns your aggregates are dynamically created and tuned to optimize performance without the need for manual intervention.

### ▲ Easy Aggregate Maintenance

You have complete visibility into aggregate usage and performance. View aggregate hit rates, configure the number of aggregates, and configure incremental updates so data never gets stale.

### ▲ Manual Aggregate Overrides

Create aggregates ahead of time if needed—for example, to support known business-critical dashboards, KPIs, and reports.

### ▲ Support across BI Tool and Excel Queries

With AtScale both queries used by most BI tools (SQL format) as well as queries used by Excel, MicroStrategy, Business Objects and Tableau (MDX) are covered. AtScale creates an optimized plan across both types to deliver sub-second query response times.

### ▲ Real Distinct Counts on Big Data

Create distinct count queries to drive BI tool controls for users, including filter and drop downs, with built-in support for both 'exact' and 'estimated' distinct counts that perform.

Definition Status	Project	Cube	Type
Active	Project Clickstream Analytics	Cube Clickstream Health Cube	Type Demand-Defined
Latest Instance Status	Last Successful Build Started Apr 4, 2018 11:43 AM PDT	Last Successful Build Ended Apr 4, 2018 11:43 AM PDT	Last Successful Build Duration 4.6 seconds
Active	Project Customer Service	Cube myServiceDrill	Type User-Defined
Latest Instance Status	Last Successful Build Started Apr 4, 2018 11:42 AM PDT	Last Successful Build Ended Apr 4, 2018 11:42 AM PDT	Last Successful Build Duration 6.8 seconds
Active	Project Customer Service	Cube myServiceDrill	Type User-Defined
Latest Instance Status	Last Successful Build Started Apr 4, 2018 11:38 AM PDT	Last Successful Build Ended Apr 4, 2018 11:38 AM PDT	Last Successful Build Duration 6.5 seconds
Active	Project Brent Sales Insights	Cube Internet Sales Cube	Type User-Defined
Latest Instance Status	Last Successful Build Started Feb 16, 2018 10:50 AM PST	Last Successful Build Ended Feb 16, 2018 10:50 AM PST	Last Successful Build Duration 11.8 seconds



*With AtScale, Cloudera customers can easily and securely connect their favorite BI tools to their enterprise data hub. Users leverage the power of Cloudera Impala and AtScale to maximize the value of Hadoop data in real-time and with optimal speed.*

—Amr Awadallah  
CTO

**cloudera**



## Big BI Solution

AtScale is purpose-built for BI on big data. It leverages the latest advancements in the big data ecosystem to support existing and new BI workloads, using Hadoop and other data lakes as the modern platform for data storage, discovery, optimization, and processing.

### ▲ Support for Leading SQL-on-Hadoop Engines

AtScale works out-of-the-box with the leading SQL-on-Hadoop engines, such as Impala, SparkSQL, or Hive, and allows them to function as an analytics engine.

### ▲ In the Cloud and On-Premises

AtScale works with on-premises and cloud big data systems including Hadoop (Apache, Cloudera, Hortonworks, MapR), Microsoft Azure HDInsight, Google BigQuery and Amazon Redshift.

### ▲ Query the Data Where it Lies

Query the data directly in Hadoop—no data movement, no data silos. Use advanced statistics and schema-on-read to optimize queries on the data in its native format.

### ▲ Built-in Support for Complex Data Types

AtScale is designed for modern data; structured and unstructured—including clickstream data, network logs, and IoT data—with built-in support for complex data such as arrays.

### ▲ Single Gateway Node Drop-in Deployment

AtScale deploys on a single gateway node in your Hadoop environment—no additional cluster to maintain, no software footprint on your big data nodes.



*AtScale's no-ETL and no-data movement approach is simply a game-changer. This application should be required for anyone who wants to do BI on Hadoop.*

—Kevin Johnson  
CEO

EBATES



## Enterprise Security & Control

AtScale works with big data deployments to support enterprise data governance and security requirements. Administrators have complete control over who can access which data across all your clusters.

### ▲ Role-Based Access Control

Manage users across departments and organizations using role-based access control. Use AtScale cubes as a metadata layer to control which data in Hadoop is available to which BI users.

### ▲ Pluggable Security Authentication

Connect AtScale to secure Hadoop and big data services. We support Kerberos, username/password, delegated authorization, Active Directory, LDAP authentication, SASL, TLS and more.

### ▲ Query Audit Trails

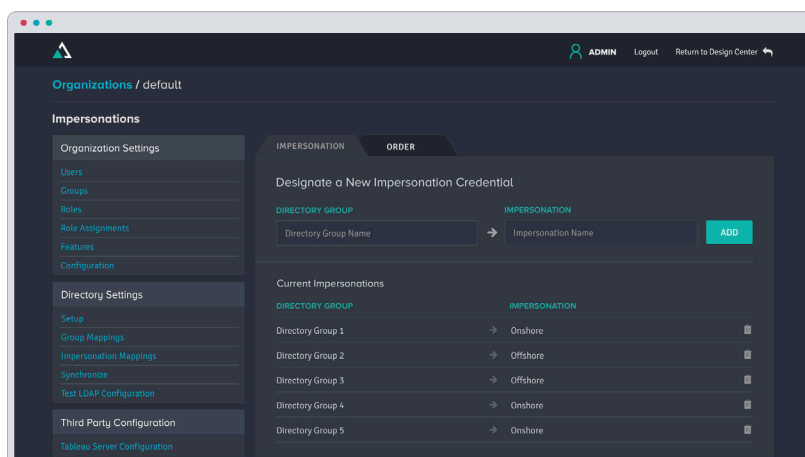
Track cube access and query metrics for every query executed by AtScale. Know who is accessing what data, and the data size and response times of all results.

### ▲ Multi-Cluster Support

Create separate AtScale execution environments to connect to data from different physical or virtual Hadoop and other big data clusters. Easily move published virtual cubes from one environment to another without any downtime.

### ▲ Production Ready High-Availability (HA)

AtScale's customers depend on our ability to always be available to satisfy queries from customer-facing applications that are powered on the back end by AtScale. With AtScale you can deploy in a "hot-hot" mode, avoiding downtime and maximizing uptime.





## AtScale Works With:

### Business Intelligence Applications

- Microsoft Excel
- Power BI
- Tableau
- Qlik
- MicroStrategy
- Business Objects

### Languages & Protocols

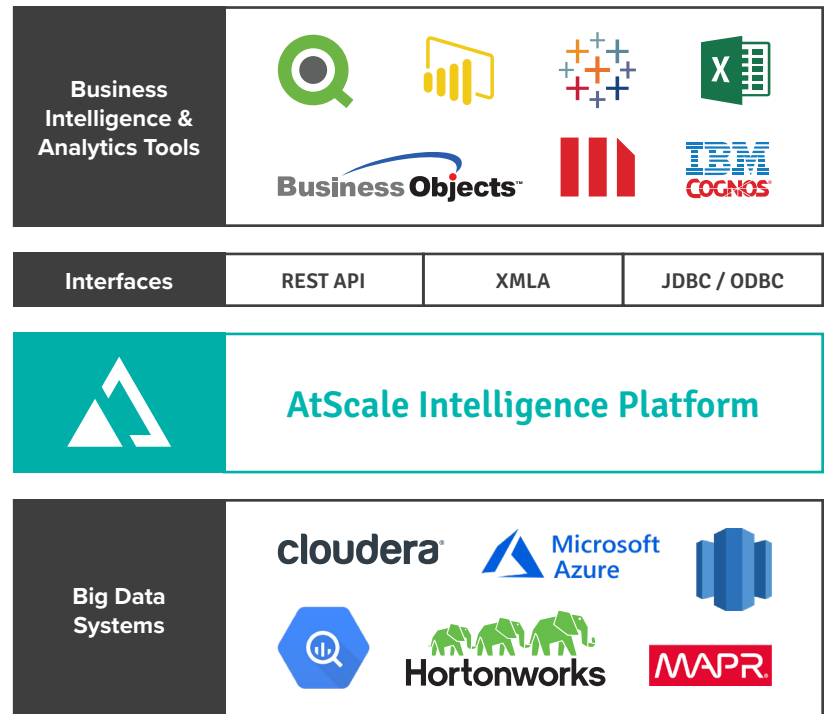
- SQL & MDX
- ODBC, JDBC
- REST API
- Sentry
- Ranger
- LDAP
- Active Directory
- Kerberos

### Big Data Systems

- Cloudera
- HortonWorks
- MapR
- Microsoft Azure HDInsight
- Google BigQuery
- Amazon RedShift

## The AtScale Architecture

AtScale bridges the gap between your BI applications and your big data platform both on-prem and in the cloud. AtScale runs a number of services that interact with BI tools using standard interfaces including ODBC and JDBC. We also use various standard services to optimize and execute BI queries directly on your data lake.



## Contact

[atscale.com](https://atscale.com)  
[info@atscale.com](mailto:info@atscale.com)

400 S. El Camino Real, Suite 800  
San Mateo, CA 94402