Atoti
Under the Hood
Technical White Paper

ActiveViam's flagship product, Atoti, is a data analytics platform developed especially for financial services. This white paper provides an in-depth look at the technical characteristics that distinguish Atoti from general-purpose databases and mass-market business intelligence tools and explores the many benefits delivered by a true real-time analytics solution.

Table of Contents

Introduction:

Atoti is the data analytics platform developed by ActiveViam especially for financial services organizations. This unique technology meets the specific needs of the industry in ways general-purpose databases and mass-market business intelligence tools do not. Atoti offers:

- Real-time, self-service analysis of complex non-linear metrics, from top of the organization to trade level.
- Continuous recalculation of metrics on datasets that change in real time.
- Operational workflows with concurrent data updates, such as signoff and What-if analysis.
- The flexibility to handle both in-memory and in-database queries with DirectQuery, Atoti's tool to query external databases directly.

Atoti delivers speed and precision. Unlike mainstream data analytics technologies that pre-aggregate datasets, pre-calculate metrics and pre-load dashboards, Atoti works on the fly to give users fast access to the data they need along with complete analytical freedom.

This requires sub-second aggregation of complex metrics on large volumes of data, which Atoti delivers with database innovations such as in-memory column stores, bitmap indexes, work-stealing multithreading and distributed computing combined with user-defined functions that run at native speed.

Atoti is flexible, purpose-built, scalable and handles massive datasets in sub seconds. This paper examines the unique technical characteristics that make Atoti a true real-time analytics solution.

The need for speed

In a traditional stack, reference data, positions and numbers generated by risk engines are collected into a database. A batch process then aggregates metrics on this data and a BI tool visualizes the result. In this two-step architecture, the BI tool cannot model complex metrics or aggregate them quickly or efficiently enough for financial purposes.

Furthermore, users are unable to analyze data beyond what is predefined. Should they want to see a new metric or simply modify a report, they must request new development from IT (which can take weeks) or do it themselves in Excel.

This is the slow IT vs. shadow IT dilemma.

Atoti solves this issue by modeling all metrics and computing them on the fly from raw data. End users can slice and dice data any way they want. Therefore:

- Users don't just watch numbers — they understand and can explain them.
- Users aren't limited to detecting a risk anomaly — they can track it down to the single transaction that caused it.
- Users can freely explore the metrics and add new ones intraday.

Beyond pure performance, Atoti offers an unparalleled self-service experience — users never have to go back to IT to generate a new report.

A universal semantic layer for finance

Atoti is a semantic layer between data and user. It converts files and tables into business hierarchies and metrics that can be utilized and manipulated autonomously.

Atoti is designed to fit into any data analytics architecture — operating between data sources and visualization tools — delivering fast, precise and interactive analytics in any configuration, even for complex metrics that require long chains of calculations and non-linear aggregation.

One of Atoti's unique features is its data modeling capabilities. Once defined in the model, metrics become available as if they were simple sums in a pivot table. Users can include them in any query or report at the required level of granularity, using any of the hierarchies in the model for grouping or filtering without any further programming.

The result is that Atoti can be applied to any data analytics challenge faced by a financial institution. Atoti has been successfully deployed at banks, exchanges, hedge funds and asset managers around the globe, with uses ranging from real-time risk analysis for the front office to market risk management, regulatory reporting such as FRTB, xVA calculations, liquidity management, collateral optimization, credit risk, counterparty credit risk and limits management.

For any use case where other tools struggle to deliver interactive analytics at scale, Atoti is the go-to solution.

Atoti's Technical Characteristics

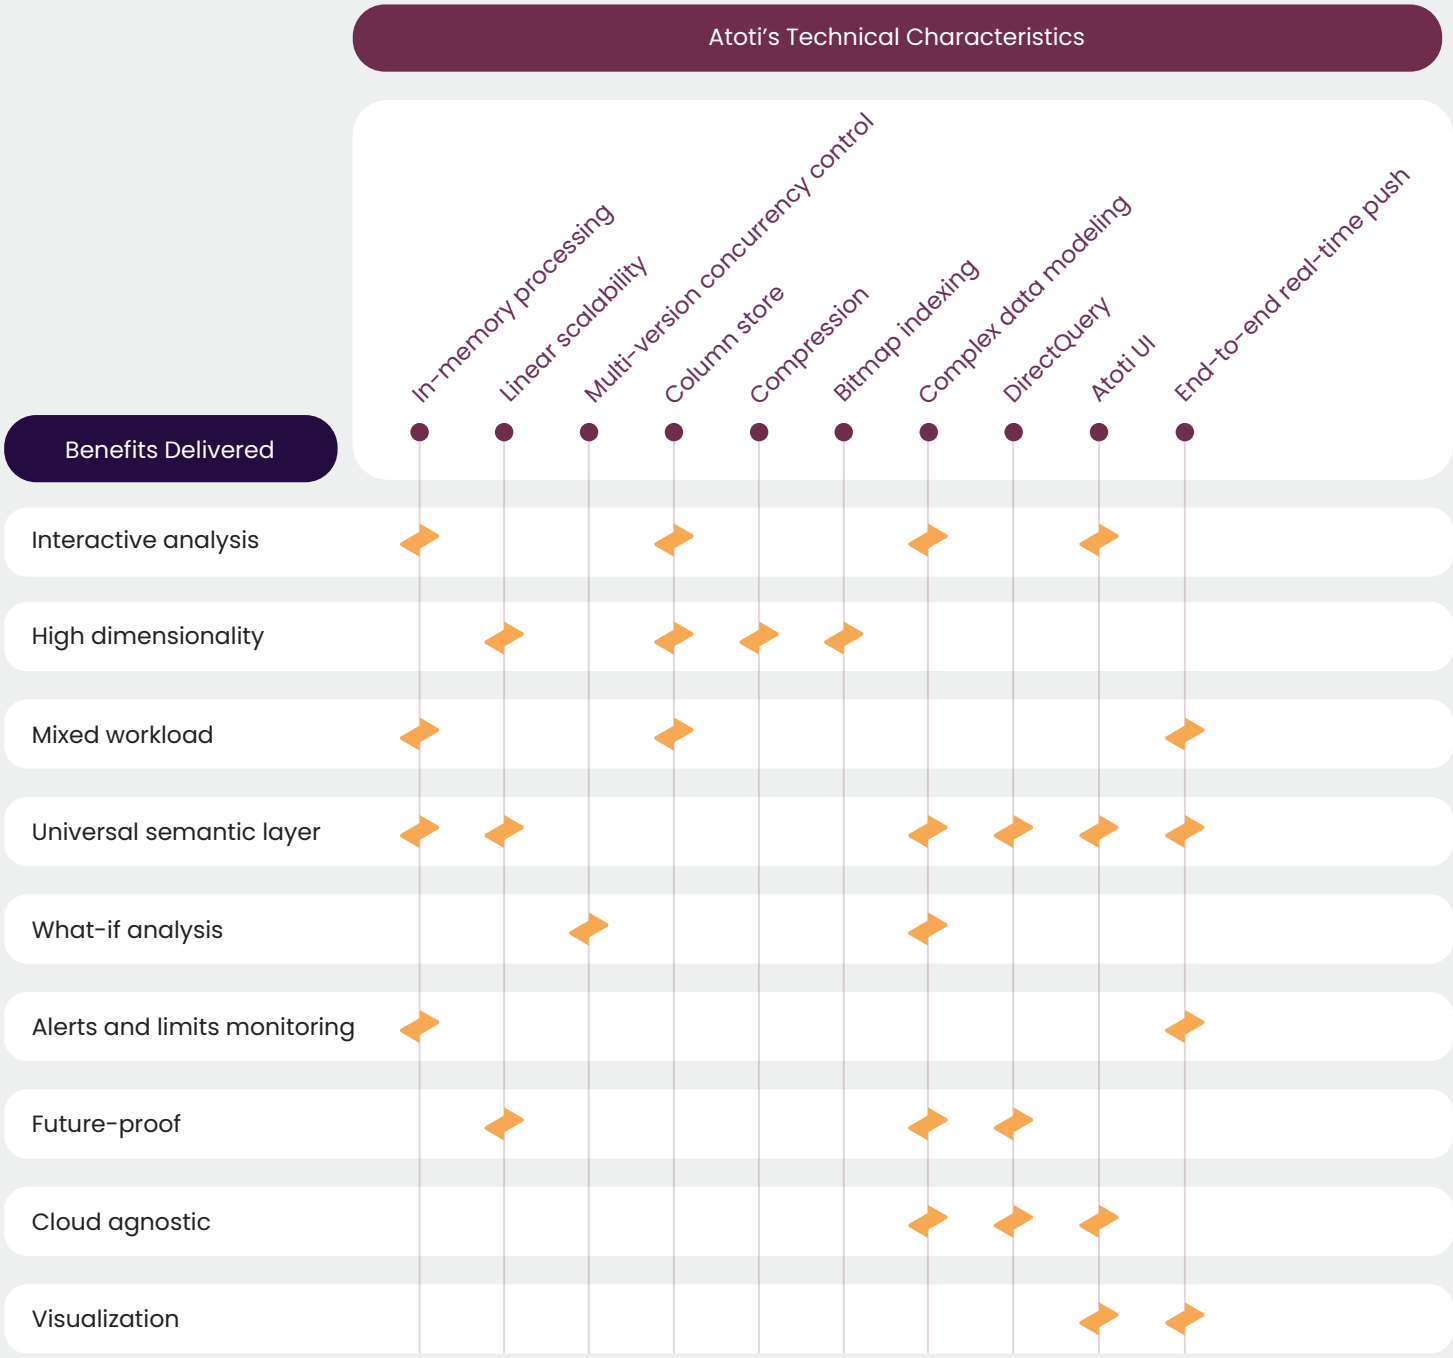

Atoti's Technical Characteristics

In-memory processing

The ability of in-memory computing to process data without interfacing with the data storage is one of the top prerequisites for rapid query performance. Atoti was natively designed to take advantage of in-memory computing to meet the extreme performance requirements of both transactional and analytical environments.

Atoti can process data directly from the computer's main memory. This means that transactions and queries never need to access a hard disk drive, network storage or any slower device which typically provides data access that is several orders of magnitude slower.

Another benefit of using memory as main storage is its random-access support. This eliminates fluctuations in speed and performance when reading and writing data to any part of memory. This is a tremendous advantage when compared to traditional disk storage that has significantly slower random access and requires highly optimized sequential access efficiency.

The overall impact of Atoti's architecture is a platform that outpaces traditional RDBMS and OLAP systems by several orders of magnitude in terms of aggregation velocity and modeling flexibility.

Linear scalability

Atoti is designed to handle massive volumes of data at interactive speed — scaling up to billions of rows on a single server — and can cluster servers together without any meaningful loss of performance.

Atoti can run on an old laptop or a tiny cloud VM, but it is dynamic enough to leverage "many-core" servers to the fullest, having been tested on machines of 1,000+ cores and 64 terabytes.

This is made possible through a combination of technologies, starting with state-of-the-art parallel computing code using work-stealing and fork-join techniques. Atoti implements a highly scalable memory allocator inspired by the Linux Kernel "SLAB" and benefits from the latest Java garbage collection improvements, developed in partnership between the Oracle Java team and Atoti R&D. On large servers with several processors, Atoti utilizes NUMA (non-uniform memory access) architectures to optimize parallel processing, employing 100% of the memory bandwidth of each individual processor to achieve massive performance gains.

When it comes to horizontal distribution, Atoti is built on a shared nothing architecture with three operating principles in mind.

First, the data nodes in the cluster are themselves cubes and can be started independently. Data nodes do not need to communicate with one another, and exchange data only with a lightweight "query node." If one server fails, cubes on the other servers remain unaffected and their data remains available for queries.

Second, aggregate queries are distributed. The query node broadcasts the query to data nodes and partial results are merged automatically. This ensures that there is no degradation of performance as each server uses its full power to process the query.

Supporting polymorphic distribution, Atoti can be used for complex projects with heterogeneous analytics, such as PnL, Sensitivities and VaR. Each application runs on its own independent cube and Atoti federates them on the fly, joining on what they have in common and resolving any mismatches.

Finally, each data node cube loads data independently. Thus, if one cube can load data at a rate of 1,000,000 facts per second then a grid of four cubes will load 4,000,000 facts per second. This makes Atoti capable of reliably handling time-sensitive aggregations for the largest financial institutions. One of our large, real-time implementations is with Chicago Mercantile Exchange (CME), consolidating 21 servers' worth of data, 16 million trades per day, and 5,000 updates per second for the purpose of publishing margin calls in real time to all market members.

Multi-version concurrency control

The most common way to ensure concurrency control in a database is to prevent users from reading and writing the same piece of data at the same time. However, this method based on locking can prove expensive in mixed workload environments that involve many updates, queries and calculations simultaneously.

Multi-version concurrency control (MVCC) in Atoti works by storing multiple versions of the data in the database so that it can be modified by a user while others continue to work with their own consistent versions. This ability for every user to write-back to the database is the foundation of several key features of Atoti such as What-if analysis, adjustments, sign-off and validation processes.

Also, unlike the back-ups of traditional database designs, Atoti only stores variations between versions. In-memory architecture makes this type of recall highly efficient, resulting in easy and fast access to previous versions of the database for "as-of" and root-cause analysis.

In short, MVCC is a key enabler of operational applications over large amounts of data that change in real time.

Column store

Conventional database management systems use a row-based store for their tables. Each fact is stored in a compact row, and any access/update to a field requires reading/writing the entire row. While row-based storage is suitable for transactional processing, it is not optimized for multidimensional analysis or complex analytical workloads.

One of the key differentiators of Atoti is its column store. With Atoti the data is stored by column rather than by row. Each column consecutively stores the values of one given field, for all the facts in the table. When a new fact is inserted, all the declared fields are extracted from that fact and stored in the relevant columns.

The ability of column-based storage to access only the fields that are required for a particular query — instead of having to read/write the entire row — vastly accelerates query performance.

Additionally, the Atoti column store supports fact updates without requiring the use of an additional row-based store to serve data-refreshing purposes. Unlike most columnar databases whose compression algorithms prevent them from easily updating facts, Atoti's column store is fully updatable.

Consequently, Atoti does not require the combination of row-based tables with columnar tables to support mixed workload requirements. Atoti can process multidimensional queries at unparalleled speeds on data that is updated on the fly offering a true mixed workload DBMS.

Compression

Atoti's highly efficient memory storage takes advantage of advanced compression algorithms, enabling a low memory footprint and incremental updates.

The first type of compression is dictionary compression. In a database system, most fields only have a small number of possible values. Such fields are said to have low cardinality. For instance, a risk management system may store millions of facts with only a few possible values for each currency or date field.

Atoti keeps different currency values in a separate structure called a dictionary and assigns a unique number from 1 to N to each currency (where N is the number of different currencies).

This dictionary is small since it stores only a few hundred strings, and the date column only collates the integers that represent the date instead of the larger strings.

A second layer of compression uses Java primitive types. When storing floating-point numbers in a column, Atoti creates arrays of Java double primitives instead of highly inefficient double objects. Moreover, when storing integer values, if the range of values is known Atoti will only use the necessary number of bits instead of the full 4 bytes (32 bits) required by a Java int.

For example, if it is known that a column will only store values between 0 and 15, then only 4 bits will be used to store these numbers. This interacts nicely with dictionary compression; a dictionary typically has a low cardinality of N members and the associated integer column can therefore be compressed using only the number of bits required to encode an integer up to N while still supporting an unlimited number of members.

Bitmap indexing

The goal of an OLAP engine is to execute multidimensional analytical queries. Queries typically need to first select facts based on the value of their fields on some dimension of the analysis, and then aggregate the facts along these dimensions to return business indicators.

A complex part of the query is the selection of facts based on multidimensional criteria.

An Atoti instance can contain billions of facts, so the ability to efficiently retrieve and filter is crucial. As an in-memory engine, Atoti has a huge advantage over its disk-based competitors. Even in cases where a query needs to retrieve all the facts and aggregate them, reading the data directly from memory provides a significant speed advantage.

However, when a query needs to filter facts based on conditions by specifying the values of fields, an OLAP engine must use one or more indexes to efficiently retrieve them. Common database indexes (such as B-tree) are only capable of indexing a single attribute. For example, a currency index quickly locates records in a particular currency, or a year index locates records for a given year. Traditional databases therefore perform multidimensional queries by using several simple indexes. This process becomes extremely slow, even when applying only a few variables on a sizable dataset.

Conventional databases have tried to resolve this issue using composite indexes. For example, an index dedicated to "year + quarter + territory". However, to answer all queries quickly, one would also need to build multiple composite indexes such as "year + quarter" and a "year + territory", etc. Assuming the cube has more than a few dimensions, the resources required for such indexes are unmanageable (N-factorial number of dimensions).

Atoti resolves this issue with a proprietary multi-dimensional index called Bitmap index. The Bitmap index a name that reflects its underlying binary arithmetic and compression schemes selects records for any random combination of predicates on any number of dimensions. It is a lightweight data structure that filters over billions of records at high speeds and easily supports more than 100 dimensions.

Complex data modeling

In financial institutions, complex metrics are common and often require long chains of calculations and non-additive aggregations. They cannot be modeled in BI tools and demand bespoke development that result in batch processing and canned reports.

With Atoti, key metrics can be modeled precisely using Java or Python APIs. Once defined in the model, metrics become available to end users as simple sums in a pivot table. They can be included in any query or report at the level of granularity required, using any of the hierarchies in the model for grouping or filtering. Typical examples include VaR, expected shortfall, potential exposure, CVA and xVA, weighted averages, bucketing, real-time PnL estimation and liquidity ratios.

DirectQuery

DirectQuery is Atoti's tool to query external databases directly such as Snowflake, BigQuery, Redshift, Databricks, ClickHouse and others without loading the data in-memory beforehand.

This means that Atoti can target cloud-scale datasets that would not fit in memory. This gives users the ability to analyze large historical datasets. With DirectQuery, Atoti deploys plug-and-play atop a data warehouse, with cubes starting in seconds since they do not need to load the data. This simplifies projects and allows for on-demand analytics and elastic provisioning.

DirectQuery also works in a hybrid mode. Atoti can combine in-memory and database queries in the same model, dramatically expanding its flexibility and modeling capabilities. Users can cache hot data in memory and combine real-time datasets with cloud-scale datasets that do not fit in-memory, such as historical data, to receive a complete picture within the same environment.

This capability delivers three main benefits:

- End users have access to granular, detailed data with unlimited historical perspective. All analyses are performed in a single environment, ensuring consistency, auditability and ease of access.
- IT managers can precisely optimize which parts of the dataset to load in memory and which parts to leave on a persistent, on-cloud or on-premises database. Hot data queried often and by many users can be loaded in memory for instant, fast access, while cold data, accessed less often, can remain where it is for cost-effective storage.
- Projects launch faster because it is all plug-and-play. Cubes start in minutes — or even seconds — because data does not need to be moved, opening new possibilities for deploying analytics on demand and dynamic provisioning.

Because the data models in Atoti are separate from the queries, IT developers only need to define a data model once to apply it to any part of a dataset, regardless of where it is stored.

DirectQuery offers a consistent deployment experience across all cloud providers as well as on-premises databases.

The same models, dashboards, and data access rules can be deployed on Snowflake in AWS, on BigQuery in GCP or on ClickHouse on-premises. Applications may be lifted and shifted between environments.

Atoti UI

General purpose BI and visualization products are not always the best fit for financial services analysis. Because they must accommodate hundreds of use cases across many industries, they are not specialized. Therefore, some tasks that are extremely common, yet very specific to financial analysis can be unreasonably burdensome.

In contrast, Atoti UI has been designed by engineers who understand the unique needs of business users in finance including traders, risk analysts, heads of trading desks and senior managers. The interface is tailored to their needs and habits based on years of experience and feedback from finance professionals.

Atoti UI is a full-fledged, web-based, customizable UI platform designed to enable users to explore data freely, spot inconsistencies, and explain any figure quickly and easily. Atoti UI is for live, collaborative data exploration. It comes with a wide range of out-of-the-box widgets while offering an extensive API for creating custom widgets and actions. Alternatively, Atoti UI widgets may also be used to integrate seamlessly into existing web-based applications.

Built in JavaScript with the React framework, Atoti UI complements its powerful backend, Atoti Server. It ensures visualizations are rendered as quickly as they are calculated, handles on-demand calculations and seamlessly performs What-if analysis, and edits and writes back the data source as well as incoming real-time dataset. The result is a responsive, interactive UI where each component may be updated by multiple sources.

Native compatibility with other front ends

In addition to Atoti UI, Atoti can use third-party BI tools to serve as front end. Atoti uses MDX, the standard language for OLAP systems to describe queries, which makes it natively compatible with front ends such as Tableau and Excel.

Using MDX, a user can:

- Start with high-level aggregates, then navigate into the details
- Slice and dice
- Filter and group
- Crossjoin

End-to-end real-time push

The result of Atoti's state-of-the-art architecture is an unprecedented end-to-end real-time push that empowers users to register continuous queries and to receive updates to the query results when the database is updated.

When new data is streamed into Atoti, such as a new record in the cube itself or a contextual update like a stock price tick, the engine precisely computes the impact on the registered continuous queries and optimizes performance by recomputing only what has potentially changed. The cells that have changed (and only those cells) are pushed to real-time clients or downstream applications.

Contrary to naïve implementations that fully recompute queries and thus fail after a few dozen concurrent subscriptions, Atoti optimizes processing power to minimize the bandwidth for transporting updates. Consequently, a single Atoti node can scale up to hundreds of concurrent, real-time continuous queries with Atoti processing tens of thousands of updates per second.

Atoti's end-to-end, real-time push consists of:

- **Real-time data sources** — Atoti naturally accepts multiple heterogeneous real-time data sources such as Apache Kafka or AWS Kinesis.
- **Continuous update to the cube** — Atoti updates a cube continually in real-time. While some users may work with ad-hoc queries, others need to see live reports that are refreshed whenever the cube data changes.
- **Streaming API** — Atoti notifies registered clients through a proprietary streaming API as soon as the data is updated.

Benefits delivered

Interactive analysis of large datasets

Because of the speed and multidimensional modeling capabilities of Atoti, end users can freely explore complex metrics over large datasets, from top-of-the-house level to single records. They may group metrics using any combination of hierarchies, filter out elements, create their own dashboards on the fly — all independently.

- Users don't just watch numbers — they understand and can explain them.
- Users don't just detect a risk figure issue — they can track the issue down to the single record that caused it.
- Users may freely explore the metrics and add new metrics intraday.

High dimensionality

One of the key challenges of OLAP systems is quickly aggregating data across different dimensions. Unlike most OLAP engines that only support a small number of analysis dimensions, Atoti supports multidimensional queries via:

- **A multidimensional bitmap index** that filters billions of records and supports hundreds of dimensions.
- **A column store indexer** that enables the processing of multidimensional queries at unparalleled speeds for data that is updated on the fly.
- **An efficient data compression mechanism** using a dictionary and Java primitive types, which, when combined, enable efficient data storage and access in memory.
- **Multi-core architecture** designed from the ground up, with parallel algorithms running on multiple threads that leverage the computing power of multi-CPU machines.
- **Polymorphic distribution** that enables developing simple independent applications for heterogeneous data, which are federated on the fly.

Mixed workload

Typically, users are forced to select either an OLTP or OLAP system, depending on the workload. Atoti offers a mixed workload solution that supports both. The same database engine performs data updates and complex analytical queries so that analytical work is always done on up-to-date data reflecting the current state of the business. Atoti's mixed workload is enabled via:

- **Multi-version concurrency control** that supports multiple queries to the database (both big and small) as well as write-backs simultaneously.
- **An in-memory database engine** offering ultra—fast data access and the ability to insert/update large amounts of data while concurrently performing queries and analytics.
- **Multi-threaded algorithms** that maximize the usage of all computer cores and process transactions and queries simultaneously to deliver the best possible throughput and response time.

- **An updatable column store** that preserves the underlying facts so they are easily updatable and quickly accessible for multidimensional queries.
- **A continuous real-time push engine** that ensures the data used for analytical queries is always the most recent and propagated in real time to end users.

Universal semantic layer

Atoti integrates with many data sources and risk engines, commercial or in-house. Many financial institutions have migrated to open architectures that decouple analytics, allow federated data from several systems and blend in external data. Atoti supports this open architecture as the semantic layer turning data into actionable business metrics.

Atoti provides end users with a global picture that consolidates data from position — keeping systems, market data providers, and metrics generated by risk engines such as prices, sensitivities and simulations. Data can be read from open format files such as CSV, CRIF, and Parquet, fetched from relational databases (including Snowflake, BigQuery, and RedShift), datalakes or streamed from messaging queues such as Kafka.

Atoti integrates with any in-house or commercial risk engine that exports its results, with a track record of successful implementations alongside MUREX, Calypso, RiskMetrics, Numerix and others.

Measures calculated in Atoti may be visualized in any MDX-compatible front-end, including Excel, Tableau, Python notebooks, and of course Atoti UI.

As a result, Atoti has been deployed successfully in production for use cases as diverse as real-time front-office risk and PnL, FRTB-compliant risk management, xVA calculations, liquidity management, collateral optimization, credit risk, counterparty credit risk and limits management.

Atoti brings consistency to workflows that are often spread across several analytical environments. With Atoti, a metric (e.g., VaR and expected shortfall) is defined concretely. Every user has the same understanding, avoiding common errors that arise from different definitions across systems. Atoti resolves operational risk and misalignment between teams working on the same data, but with different tools. By providing a single global solution, Atoti offers the best possible return on investment.

Cloud-agnostic

Atoti runs seamlessly on Amazon Web Services (AWS), Microsoft Azure and Google Cloud Platform (GCP). It supports all major cloud-based data warehouses (Snowflake, BigQuery, Databricks, and RedShift) as well as their deployment and orchestration systems with high-performance connectors to facilitate integration.

The deployment experience is consistent across all public clouds as well as on-premises. Customers do not need to reinvent data models, dashboards or retrain their teams. They can lift and shift from one environment to another, and even create hybrid solutions that are part in the cloud and part on-premises.

While public cloud stacks are comprehensive, well-managed and state-of-the-art, they also come with risk of vendor lock in. Atoti makes it easy to switch from one to another or to support a mix of public and private clouds for different functions. Users do not need to reinvent processes for data modeling, publishing dashboards and set up data access rights for each environment they support. They can lift and shift their models and visualizations from one platform to another.

What-if analysis

With Atoti, you can perform instant simulations and What-if analysis on huge datasets to evaluate alternative scenarios, make projections, and reach informed business decisions. What-if analysis leverages Atoti's ability to perform real-time transactions and to execute queries immediately on the updated data. Thanks to MVCC, users create their own splinter environments on the fly where they can perform personalized What-If simulations and queries with no impact on other users.

Atoti's What-if capability is enabled via:

- **User-defined aggregate functions,** which allow for complex and non-trivial aggregation mechanisms that can rely on both internal and external data.
- **An updatable column store** that stores the underlying facts that are easily updatable and quickly accessible for multidimensional queries.
- **An in-memory database engine** that supports real-time transactions (such as inserting or updating data) with queries performed immediately on the new data.
- **Multi-version concurrency control** enables users to create private environments on the fly, adjust any parameters, analyze, and share the results, and ultimately validate or discard the changes.

Alerts and limits monitoring

The key requirement for monitoring business indicators and alerting users is to be informed of a change in underlying data. Atoti's alert and monitoring capabilities immediately notify users when any change to specified data is made.

Using advanced measures, users may also be alerted to any business logic defined. For instance, an alert may be issued if a KPI increases/decreases by more than the average of competitors during a specific timeframe. Contrary to conventional databases that use triggers or polling to monitor data, Atoti's alerts and monitoring capabilities are enabled via:

- **An updatable column store** that stores the underlying facts so they are easily updatable and quickly accessible for multidimensional queries.

- **An in-memory database engine** offering fast data access and the ability to insert/update large amounts of data while concurrently performing queries and analytics.
- **A best-in-class aggregation engine** that recalculates aggregated data and limits on the fly.
- **A continuous, real-time push engine** that alerts users when data of interest has changed.
- **Powerful What-if capabilities** that enable users at all levels to test potential decisions and assess their impact on limit breaches.
- **Tools for reporting** that ensure auditability, enable root-cause analysis and assist with incident tracking.
- **An optional extension module, Atoti Limits,** that allows customers to centralize all limits into a dedicated cube to facilitate management, monitoring and the investigation of potential breaches.

Future-proof

Designed to scale-up and scale-out, Atoti supports expansion of business data and the need to continuously increase the scope and depth of data analysis over time.

This is achieved via:

- **Multi-threaded parallel algorithms,** which enable Atoti to maximize its use of computer cores. Adding more cores allows Atoti to scale-up on a single box.
- **Efficient data-compression algorithms,** which allow for the storage of large amounts of data within each Atoti instance. This enables the application to grow while using a minimal amount of memory.
- **Horizontal distribution**, which automatically adds more instances to a distributed Atoti deployment, resulting in increased computing power and available storage.
- **Polymorphic distribution,** which simplifies the data architecture and enables the efficient management of multiple cubes for different business areas.

Data visualization

Atoti provides intuitive visualization and data analysis in near real time. Data visualization is enabled via:

- **Multiple UIs like Excel, Python notebooks, Tableau and Atoti UI** which provide a graphical interface to the Atoti engine in MDX, the standard query language for multidimensional analysis. Users may select the front-end they prefer to perform work.
- **A continuous, real-time push engine** which ensures users view the most recent data to facilitate optimal business decisions.
- **A multidimensional index** that supports hundreds of dimensions.
- **User-defined aggregation functions and post-processors** which allow users to define the exact datasets to be visualized for their business goals.

ActiveViam

NEW YORK

550 Seventh Avenue
19th Floor
New York, NY 10018
Tel: +1 646 688 4442

LONDON

6th floor, Shaftesbury House
151 Shaftesbury Avenue
London WC2H 8AL, UK
Tel: +44 20 7836 8820

PARIS

46 rue de l'Arbre Sec
75001 Paris, France
Tel: +33 1 40 13 91 00

SINGAPORE

80 Amoy Street #02-01
Singapore 069899
Tel: +65 6816 4988

SYDNEY

Level 9, 64 York St
Sydney NSW 2000
Australia
Tel: +61 400 043111

HONG KONG

21/F, On Hing Building
1 On Hing Terrace Central
Hong Kong

FRANKFURT

Germany,
Junghofstrasse 16, 8th Floor
60311 Frankfurt

About ActiveViam

We provide the world's leading financial institutions with a fully flexible data aggregation and analytics platform. Our technology and solutions continuously calculate risk and performance metrics on large volumes of fast—moving data, revealing greater insights that empower our clients to make better business decisions, optimize their profits, and manage their risk.

ActiveViam operates in the global financial market, and is a trusted partner of HSBC, Société Générale, CME, Danske Bank, and many more.

For more information, please visit www.activeviam.com or follow us on Twitter and LinkedIn.

For more information please visit:

www.activeviam.com