



Atoti+

performance
& scalability

Reference handbook



Atoti+ is a one-of-a-kind unified analytics platform that features a unique combination of advanced technologies (in-memory computing, cloud-native architecture, multi-dimensional analysis) and delivers precise, flexible, interactive analytics even on the largest datasets without losing performance.

It scales up and scales out seamlessly and ensures the most cost-efficient use of hardware and reduces pause times to mere milliseconds. Atoti+ features comprehensive Python and Java APIs, offering the best possible combination of accessibility, speed-of-development and control.

Why is Atoti+ a leading in-memory scale-up database?

Atoti+ has become the industry standard for the ability to scale up and exploit large-memory servers to their highest capacity and has been used by technology leaders such as Oracle, HP Enterprise and Azure to demonstrate the capabilities of their own top-of-the-class hardware.

1. On heap / off heap for optimum memory management

ActiveViam's R&D department has developed and refined best-in-class off-heap memory management, partnering with Oracle to help with the development of Java 11's own memory management capabilities.

2. Lockless updates, MVCC: state of the art mixed workload architecture

Atoti+'s Multiversion Concurrency Control (MVCC) capability enables lockless updates by a virtually unlimited number of users. It is also the foundation of several other key features of Atoti+, such as what-if analysis, adjustments and validation processes. Unlike the back-ups of traditional database designs, Atoti+ only stores variations between versions. Accessing data in old versions has no overhead, thanks to the in-memory architecture, resulting in easy and fast access to previous versions of the database for "as-of" and root-cause analysis.

3. Column store, dictionary world class compression

- One of the key attributes of Atoti+ is its column store. In Atoti+, all incoming facts are stored in the column store, where each row represents a fact, and each column stores the value of a given field in the fact. 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, greatly improves query performance. Unlike most columnar databases, whose compression algorithms prevent them from easily updating facts, Atoti+'s column store is fully updateable.

Offering a true mixed-workload DBMS, Atoti+ is able to process multi-dimensional queries at unparalleled speeds on data that is updated on the fly.

- Atoti+'s advanced compression algorithms enable a low memory footprint and incremental updates. Atoti+ leverages dictionary compression: in a database system, most fields have a low cardinality (which means that they only have a small number of possible values).

For these fields, Atoti+ keeps all the different possible values in a separate structure, called a dictionary, and assigns a unique number from 1 to N to each possible value. This dictionary will be very small, since it only needs to store a few hundred strings (the possible values). The destination column will then only store the integers that represent the destinations, instead of the large destination strings.

4. Modern partition based multi-threading

- Writing single-threaded algorithms with the lowest possible runtime complexity is no longer the best way to optimize the use of computer resources. Instead, resource optimization is achieved by writing parallel, multi-threaded algorithms that take advantage of all the available processing cores.
- Atoti+ has been designed from the ground up to leverage the computing power of a multi-CPU machine, from dual-core Java VMs to the modern so-called "many-core" machines that contain several hundred cores. Atoti+ takes full advantage of NUMA (Non-Uniform Memory Access) architectures. On large servers with several processors, Atoti+ can access 100% of the memory bandwidth of each individual processor, resulting in massive performance gains.

The data structures used are thread-safe and all algorithms involved in transaction and

query processing are heavily multithreaded – creating a design that lets Atoti+ exploit as much performance as possible from multi-core hardware.

Client Case: Streaming Front Office PnL & Risk

- 8 cores 256 GB memory servers
- 3 days data online, 24 hours by 6 days continuous operation
- 25 million facts per CoB date
- Periodic full revaluation
- Continuous streaming risk for new and amended trades
- 100 dimensions, 300 concurrent continuous queries
- 200 users worldwide in three global regions

Client Case: CVA at one of the largest global banks

- **1 million trades**, 32 Cores, 3 TB RAM.
- CVA full calculations per counterparty across 600k trades (EE, EEPE, RWA, PFE)
- 136 time steps, 1,000 scenarios

+ Fully processed in 2 minutes

Client Case: CME Market Risk

- 5 days of historical data
- 20 million prices kept at any given time
- 8-16 million trades per day on average
- 21 source systems
- **1 TB dataset**
- Up to 5K inserts/updates per second

+ Both IT and Market risk teams report little or no impact on query performance



Large regional APAC bank

- Market Risk use case (Sensitivities, VaR)
- Loading and aggregating time: **1 min 54**
- Records: **140 millions**
- Direct memory used: **37 GB** / Heap: **8 GB**
- Data source: **1,500 files**, 16 different types of CSV files
- Size on disk (zipped): **20 GB**
- CSV avg throughput: **347 MB/s**
- Loading throughput (20 GB / 1 min 54): **175 MB/s**
- **Sub-second** query time

Client Case: Large international bank

- Market risk use case (Sensitivities, VaR)
- **1.2 million trades** in a full business day
- **79 GB** total memory
- **41 minutes** loading and aggregating time
- **Sub-second** query time

FRTB & CVA Benchmarks in partnership with HPE: FRTB

- 1 TB of data per day
- ES calculations on-the-fly in **5 seconds**

CVA

- 5,000 Monte-Carlo simulations, each with
- 200 time points and **1 million trades**

+ 8 TB of data

- PFE and ES calculated on-the-fly in 15 seconds
- 99% of pauses under 12 ms.

Client Case: Top 10 global reinsurer/asset manager

- 150 GB dataset
- 2 TB system in total
- 300,000 positions by dataset
- 10,000 Monte Carlo vector size
- **VaR** processed in **less than 10 seconds**

Client Case: Top Hedge Fund

Risk and PnL Dashboard use case (VaR, Stress Test, Flash PnL)

- Max limit: 200 GB
- Average: ~120 GB
- Direct Limit: 500 GB
- Average: 90 G Positions Initial Load days: 30
- **1.2M records per day**

Why is Atoti+ a leading in-memory scale-out database?

Scale-out architectures are a good fit for applications that require a great deal of flexibility, for instance to accompany the growth of historical data or to provide for fluctuations in the usage of memory and computing power.

Atoti+ is designed to ensure smooth scale-out capabilities with a linear increase in performance with the addition of any new VMs, thanks to the integration of several key technologies.

1. Speed


In an Atoti+ cluster, the different nodes share absolutely no hardware, ensuring that there is never a bottleneck or a decrease in performance because one node may be hogging a hardware resource. This also ensures that the platform does not use any more memory than is strictly necessary.

2. Incremental Updates

Atoti+ clusters are composed of a majority of data nodes, which hold the actual data, supported by a handful of dedicated query nodes, which are tasked with coordinating data updates between the other nodes and routing queries from end-users to the correct destinations. This construction is the most optimized for multi-dimensional aggregation, which is the core functionality of Atoti+. It minimizes the redundancy between nodes as well as the network and memory used by nodes to exchange messages.

We also leverage ActiveViam's off-heap memory management technology as well as the Netty library to further optimize messaging.

Furthermore, when used in a scale-out configuration, Atoti+ remains permanently elastic and will automatically add or remove nodes, depending on usage.



Advanced Data Modelling Example #1

Atoti+ running on a cluster of **128 VMs of 140 GB each** (19 TB total). We are pleased to say that query times remain constant when scaling up the size of the dataset in parallel with the number of VMs:

- 5-6 seconds for full scan queries
- 800-1,000 ms when using ActiveViam's high performance bitmap index

Advanced Data Modelling Example #2

We doubled the size of the dataset as well as the number of cores for this second test.

128 nodes, 4,096 cores, 40 TB

- 400 days of historical data, 40 TB of data
- Loaded and aggregated in memory in 15 minutes
- Queries takes between 5 and 10 seconds

About ActiveViam

ActiveViam provides precision data analytics tools to help organizations make better decisions faster.

ActiveViam started in 2005 with the vision of leveraging in-memory technology to create an analytics platform where businesses could leverage the largest data sets without restrictions, keep them up-to-date in real time and use them to empower their decision makers.

Our goal at ActiveViam, is to let organizations not only make decisions faster, but better; to not only reach their data, but their potential; to not only see their data, but find their way into the future.

ActiveViam is a privately owned company with offices in London, New York, Paris, Singapore and Hong Kong.

For more information please visit: www.activeviam.com

LONDON

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

NEW YORK

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

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

HONG KONG

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