

ClockWorkDB

DeMinimus Hardware

DeMaximus Performance

25 years in the making, developed to power some of the world's largest financial data warehouses, accelerate your time series, columnar, extensible workflows with our high-performance, developer-first platform. Don't trust us, see our [Stac Research](#) independant benchmarks! (Coming soon...)

1B+

Rows per second on commodity hardware.

1K+

Replication Nodes to reach the edge.

Peta+

Easily scale to multipetabyte datastores.

Slice+

Slice to Multiple Persistence Nodes.

25+

Years in production & development.



Storage Agnostic

Whether an RDBMS, NoSQL, or proprietary database like KDB, ClockWorkDB provides a consistent API for TimeSeries and Columnar data, for native types, structures, JSON, XML, Blobs and beyond.



Curated Proprietary Storage

QuantumChest and WarpDrive+ storage engines provide the ultimate in performance and extended features including HA, Replication, Petabyte+ storage, slicing and sharding.

TARGET: AI - RAG (Retrieval Augmented Generation)

A BRIEF HISTORY:

ClockWork DB was the answer for big data before the term had been coined. Initially founded to address the needs of American Century, now part of JP Morgan, and the sizeable financial data warehouse driving their modeling and investment decision systems split across IBM mainframes and Sun Solaris Servers. The need to have nightly models run to completion before market open, which was impossible at the time, until ClockWork DB came along.

This carried on to Putnam Investments, driving investment modeling, risk systems, and their nascent web portal. McKinsey & Co were brought in to do performance analysis against the then Gold Standard, FAME(Forecasting Analytic Modeling Environment: Purchased by Sunguard, and then FIS). ClockWork DB proved by every metric 3 magnitude faster.

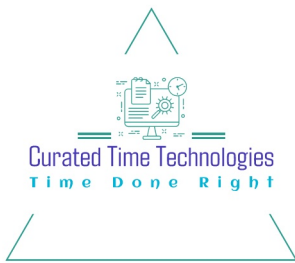
Shortly afterwards, ClockWork DB won the contract at Allianz in Germany to build their global financial data warehouse. At the time the biggest in the world, to service the data needs of Allianz, Deutsche Bank, Axa Insurance, PIMCO, and numerous others. Sunguard had spent 2 years architecting a solution. Upon production release, a use case for Risk Management, the query took 7 minutes. ClockWork DB was brought in, and the same query was reduced to less than 100 milliseconds **in the course of two weeks**.

When others were trying to force fit time series, vector/columnar data into RDBMS/SQL platforms, ClockWork DB was pioneering cutting edge technology, that moved on to topple systems at myriad top Banks, Broker Dealers, Central Banks including the ECB, the Bank of International Settlements, and more.

And development continues to this day, adding modern features, and sizeable performance improvements.



Curated Time Technology provides enterprise time series and columnar based storage and analytic solutions to Financial & Insurance Co's, Central Banks, BIS, and similar.



ClockWorkDB

DeMinimus Hardware

DeMaximus Performance

WHY RAG MATTERS:

1. OVERCOMES LLM KNOWLEDGE CUTOFF

- LLMs are trained on static datasets, meaning they can't "know" about events or facts after their last training date.
- RAG bridges this gap by retrieving the latest data from external sources (databases, APIs, document stores) at query time.

2. REDUCES HALLUCINATIONS

- LLMs sometimes generate plausible but incorrect content ("hallucinations") because they rely on learned patterns rather than explicit facts.
- With RAG, relevant documents or facts are retrieved and fed into the prompt, grounding the model's responses in verified information.

3. ENABLES DOMAIN-SPECIFIC KNOWLEDGE

- Training an LLM on niche or proprietary data from scratch is expensive.
- RAG allows you to keep a general-purpose LLM and enrich it dynamically with specialized knowledge (e.g., medical guidelines, legal documents) stored in a vector database.

4. KEEPS DATA PRIVATE & CONTROLLED

- Instead of fine-tuning with sensitive or proprietary data, RAG can retrieve it securely from an internal source at runtime, keeping the base model unchanged.

5. COST & EFFICIENCY

- Fine-tuning large models is resource-intensive.
- With RAG, you avoid retraining for every new dataset—just update the retrieval store.

ANALOGY:

Think of an LLM as a well-read expert with a great memory but no internet access. RAG is like handing that expert a fresh set of research papers right before they answer your question—so they can blend their deep reasoning with up-to-date facts. The faster the RAG data can be retrieved, the faster the response. The greater the depth of RAG data, the higher quality, higher relevancy the answer.

Introducing Anthropic's Claude Sonnet 3.5 to ClockWork DB - RAG

`/getTimeSeries SPY`

I'll retrieve the current time series data for SPY:

> Ran `getTimeSeries - tsdb (MCP Server)` ✓

Here's the SPY data Formatted in a table with daily changes:

Date	Price (\$)	Daily Change	% Change
2025-Jul-25	637.10	--	--
2025-Jul-28	636.94	-0.16	-0.03%
2025-Jul-29	635.26	-1.68	-0.26%
2025-Jul-30	634.46	-0.80	-0.13%
2025-Jul-31	632.08	-2.38	-0.38%
2025-Aug-01	621.72	-10.36	-1.64%
2025-Aug-04	631.17	+9.45	+1.52%
2025-Aug-05	627.97	-3.20	-0.51%
2025-Aug-06	632.78	+4.81	+0.77%
2025-Aug-07	632.25	-0.53	-0.08%
2025-Aug-08	637.18	+4.93	+0.78%

Key Statistics:

1. Total Return: +0.08 points (+0.01%)
2. Highest Close: \$637.18 (Aug-08)
3. Lowest Close: \$621.72 (Aug-01)
4. Trading Range: \$15.46
5. Largest Daily Gain: +\$9.45 (Aug-04)
6. Largest Daily Loss: -\$10.36 (Aug-01)
7. Average Daily Move: ±3.83 points

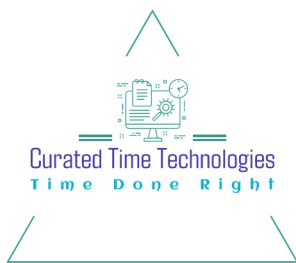
Would you like any specific analysis of this data or comparison with other securities?

Once Claude is introduced to the `/getTimeSeries` resource, it never need be mentioned again.

Here the ClockWork DB - RAG retrieves the entire history of SPY (8485 days) in just under 1 millionth of a second.



Curated Time Technology provides enterprise time series and columnar based storage and analytic solutions to Financial & Insurance Co's, Central Banks, BIS, and similar.



ClockWorkDB

DeMinimus Hardware

DeMaximus Performance

RAG / TRAINING DATA SIZES ACROSS VARIOUS LLM MODEL SIZES

Model Size	Parameters	Training Data (Tokens)	Approx. Data Size (Cleaned Text)	Typical Capabilities
Small	100M–500M	1–20B	~10–50 GB	Basic text completion, keyword extraction, classification, domain-limited Q&A.
Medium	1B–7B	50–500B	~200 GB–2 TB	General-purpose Q&A, summarization, translation, basic reasoning, good in one language.
Large	13B–30B	0.8–2T	~3–8 TB	Strong reasoning, multilingual support, good coding skills, broad general knowledge.
Very Large	65B–175B+	2–10T+	~8–40 TB+	Advanced reasoning, complex problem-solving, creativity, high factual accuracy (within training cutoff).
Frontier	500B+ (speculative)	10T–100T+	40 TB–400 TB+	Near-human conversational depth, cross-domain mastery, highly adaptable reasoning.

COMMENTARY:

These performance metrics were run against ~2 trillion tokens. A small machine is capable without tuning to achieve ~9 billion tokens/second. They were performed on a server two generations behind current Intel XEON Processors. They were performed with a small cache (128GiB) relative to the overall dataset size (8.56 TiB). No slicing or partitioning is in effect, which does not reflect true production use cases. Zlib compression is in use, resulting in a datastore size of 2.4 TiB. Snappy and LZ4 compression are also available.

CLOCKWORK DB: 1 YEAR OF REALTIME ORDER BOOK EVENTS

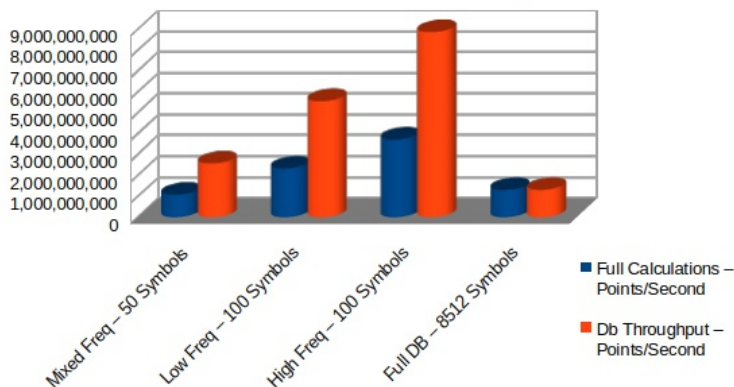
Instruments: 8,512

Trading Days: 251 (261 calendar, 10 holidays)

Data Size: 8.56 Terabytes

	Points/Second	Points/Second
ID	Full Calculations	Db Throughput
Mixed Freq – 50 Symbols	1,084,651,558	2,582,503,710
Low Freq – 100 Symbols	2,338,635,200	5,568,179,048
High Freq – 100 Symbols	3,727,043,878	8,873,913,994
Full DB – 8512 Symbols	1,314,529,516	1,314,529,516

Stac A3 and CTT DB Throughput

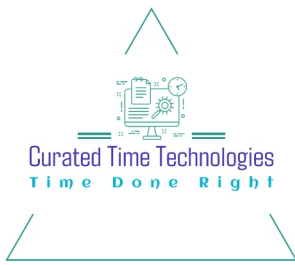


HARDWARE / SOFTWARE CONFIGURATION

Data Set	Terrabytes
CSV Drive Space	8.56
CTT Drive Space	2.40
Savings	357%
Setting	Value
Compressor	zlib
Chunk Size	4,096
Page Size	65,536
Partitioned	NO
Sliced	NO
DB Worker Threads	96
Data Points/Record	6
Cache Size	128GB
Cache Segments	8
ACID Semantics	YES
Server Attributes	Value
AWS Id	m7i.metal-48xl
Processor	Intel Xeon Platinum 8475B
Base Clock	2.7GHz
All Core Turbo	3.2GHz
Single Core Turbo	3.8GHz
Architecture	Sapphire Rapids
Memory	750 GB
Physical Cores/Die	48
Dies	2
AWS gp3 (OS/raw data)	10 TB
AWS Provisioned (DB)	10 TB
IOPs (DB)	100,000



Curated Time Technology provides enterprise time series and columnar based storage and analytic solutions to Financial & Insurance Co's, Central Banks, BIS, and similar.



ClockWorkDB

DeMinimus Hardware

DeMaximus Performance

ABOUT CLOCKWORK DB:

• ON RELIABILITY:

- Five 9's uptime in large scale production use for over 20 years.
- Serverless architecture requires little support from Sys Admins or Database Admins.
- Low/No maintenance serverless system reduces complexity, downtime and associated costs.
- Embedded Database technology supporting multiprocess, free threaded, fine grained locking for low contention scalability, for both depth and breath.

• ON DATA:

- Support for Time Series, Ordinal Series (Vector/Columnar), Sparse Series, and Scalars.
- All series and scalar objects can store integers, floats, strings, JSON, XML, BLOBs (binary large objects), documents (XLS,CSV,PDF,...), as well as plugable custom structures.
- Key objects on canonical names while including aliases for quick reverse mapping between token ids and names.
- Attach unlimited meta-data attributes to each object. All objects have description and documentation meta-data by default.

• ON TECH:

- Support for plugable compression with native snappy, zlib, and lz4 compression.
- Replication to ~1000 nodes.
- High Availability, Resilience and Robust recovery configurations with automatic quorum election during network segmentation or process fallover.
- Exabyte Ready. Support for multiple repository, multiple datastore configurations. Each datastore supports up to 256 TiB storage.
- Database Slicing allows extreme data-thruput with isolated WAL logging, locking and caching per slice. It also provides for datastores that can store sizes to N-slices * 256 TiB.
- Highly tunable to match data structures, for chunking, and page size of objects, with tools to analyze and recomend page size.
- Fine grained locking on pages (defined by the page size tuning attribute) means massive parallelism can be achieved with DeMinimus contention.

Prompt: compare spy qqq nvda msft aapl and orcl

CHAT

> Ran getTimeSeries - tsdb (MCP Server)

Now I'll create a comprehensive comparison of all securities. I'll calculate the percentage change from the starting point (Jul-25) for each security to show relative performance:

Date	SPY	QQQ	NVDA	MSFT	AAPL	ORCL
2025-Jul-25	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
2025-Jul-28	-0.03%	+0.31%	+1.87%	-0.24%	+0.08%	+1.06%
2025-Jul-29	-0.29%	+0.16%	+1.16%	-0.22%	-1.22%	+1.98%
2025-Jul-30	-0.41%	+0.29%	+3.33%	-0.09%	-2.26%	+2.24%
2025-Jul-31	-0.79%	-0.24%	+2.52%	+3.85%	-2.95%	+3.53%
2025-Aug-01	-2.42%	-2.21%	+0.13%	+2.02%	-5.38%	-0.29%
2025-Aug-04	-0.93%	-0.40%	+3.75%	+4.27%	-4.92%	+3.02%
2025-Aug-05	-1.43%	-1.08%	+2.74%	+2.73%	-5.12%	+4.31%
2025-Aug-06	-0.68%	+0.17%	+3.41%	+2.19%	-0.29%	+4.61%
2025-Aug-07	-0.76%	+0.51%	+4.19%	+1.39%	+2.88%	+1.74%
2025-Aug-08	+0.01%	+1.44%	+5.30%	+1.62%	+7.23%	+2.01%

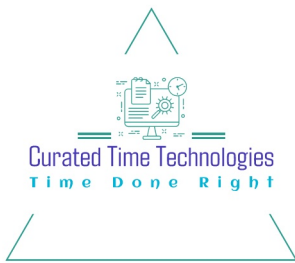
Key Observations:

1. Best to Worst Performance (Jul-25 to Aug-08):
 - AAPL: +7.23%
 - NVDA: +5.30%
 - ORCL: +2.01%
 - MSFT: +1.62%
 - QQQ: +1.44%
 - SPY: +0.01%
2. Volatility Analysis:
 - Most volatile: AAPL (12.61% range from low to high)
 - Least volatile: QQQ (3.73% range from low to high)
 - Tech stocks showed higher volatility than indices
3. Notable Patterns:
 - Aug-01 market-wide dip affected all securities
 - Tech stocks (AAPL, NVDA) showed strongest recovery
 - Indices (SPY, QQQ) showed more stable performance
 - MSFT had strong performance mid-period but faded
 - ORCL maintained consistent positive performance
4. Sector Trends:
 - Technology sector outperformed the broader market
 - Large-cap tech (AAPL, NVDA) showed strongest momentum
 - QQQ (tech-heavy) outperformed SPY (broader market)
5. Correlation:
 - High correlation during market decline (Aug-01)
 - Divergent performance during recovery phase
 - Individual stocks showed more independence from market movements

Would you like any specific aspect of this comparison analyzed in more detail?

curated time
technology

Curated Time Technology provides enterprise time series and columnar based storage and analytic solutions to Financial & Insurance Co's, Central Banks, BIS, and similar.



ClockWorkDB

DeMinimus Hardware

DeMaximus Performance

ABOUT CLOCKWORK DB:

- ON PERFORMANCE:
 - Arguably the fastest time series, vector/columnar database tech in the world.
 - Arguably the highest scaling database for absolute data size and breath of concurrent user load.
 - Benchmarked against Traditional RDBMS/SQL servers (Oracle, Sybase, MS SQL Server, PostgreSQL, MySQL,...)
 - Benchmarked against NoSQL servers (Cassandra, Scylla, MongoDB, ...)
 - Benchmarked against BigData servers (Google BigTable, Amazon Timestream, InfluxDB, ...)
 - Benchmarked against Proprietary Vector Servers (KX/KDB+, FAME, ...)
 - ClockWork DB outperformed every alternative on comparable hardware.
- ON COST & EFFICIENCY
 - Fewer clock cycles per retrieval means lower costs.
 - Dense storage saves on storage requirements.
 - Maximum scale with greatly reduced hardware requirements means lower capital costs, lower rack space costs, lower energy cost.
 - Low maintenance, reduced SA/DBA human capital costs.
- ON DEPLOYMENT
 - Small footprint allows use on smart devices (Phones, Tablets, ...) but scales to super computer clusters.
 - Supports Windows, Apple OS, Linux (favored), iOS, Android.
 - Supports x86_64 and ARM64
 - Cloud friendly with Docker Images available.
 - Easy to deploy on development workstations or blades.

Prompt: compare prices of banking stocks

CHAT

compare prices of banking stocks

I'll help you compare major banking stocks. Let me retrieve data for some of the largest banks: JPMorgan Chase (JPM), Bank of America (BAC), Wells Fargo (WFC), Citigroup (C), and Goldman Sachs (GS):

Let me analyze the banking sector data we've collected from July 25 to August 8, 2025:

JPMorgan Chase (JPM):

- Started at \$298.62
- Ended at \$288.76
- Overall decline: -3.30%

Bank of America (BAC):

- Started at \$48.45
- Ended at \$46.01
- Overall decline: -5.04%

Wells Fargo (WFC):

- Started at \$84.32
- Ended at \$77.85
- Overall decline: -7.67%

Goldman Sachs (GS):

- Started at \$728.98
- Ended at \$721.27
- Overall decline: -1.06%

Citigroup (C):

- Started at \$96.07
- Ended at \$92.73
- Overall decline: -3.48%

Key observations:

Claude is aware of what "Banking Stocks" are...

Claude picks 5 major banking stocks, knows the symbols, and comes to ClockWork DB to get current prices to make the response accurate. No additional training, just RAG context injection!



Curated Time Technology provides enterprise time series and columnar based storage and analytic solutions to Financial & Insurance Co's, Central Banks, BIS, and similar.