



[www.ultipa.com](http://www.ultipa.com)

2342 Poppyview Avenue, San Ramon, CA 94582

Tel: +39 333 415 6248

Email: [support@ultipa.com](mailto:support@ultipa.com)

## Ultipa – the Powerhouse of the graph world

### The company

Ultipa is a relatively new graph database vendor. It was founded in 2019, and its flagship product (of the same name) was commercialised in 2021. Despite this, the company has offices across Europe, Asia, and the United States, as well as a substantial customer base in (retail) banking.

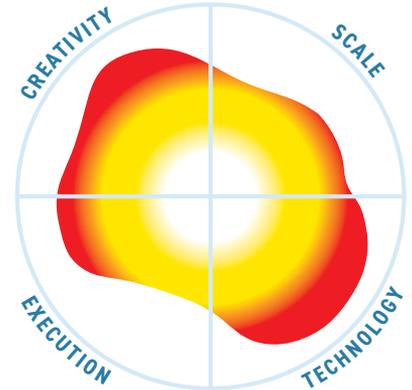
### What is it?

Ultipa is a highly performant property graph database built on an in-memory parallel graph computing engine, designed from the ground up to enable deep search and real-time analytics via an intuitive and powerful query interface. It is particularly well-suited for high-frequency transactional workloads and complex analytical processes, making it ideal for creating knowledge graphs to accelerate big data querying. To wit, some Ultipa customers are currently using it to process hundreds of millions of transactions per day, in real-time. Ultipa is also translytic: it offers true HTAP (Hybrid Transactional/Analytical Processing) functionality that does not need to separate your analytical and transactional workloads, enabling simultaneous data ingestion, updating, and real-time querying across both kinds of workload. In turn, this helps to ensure you have instantaneous access to up-to-date data for both transactional and analytical use cases.

Moreover, Ultipa is particularly notable for – and puts particular emphasis on – its ability to support AI, and especially XAI (“eXplainable AI”). For instance, the degree of performance it provides can make a very significant difference to the amount of time it takes to generate your AI models, thus allowing you to act on the predictions they make much more quickly. What is more, by analysing and discovering key features of your network and communicating them to downstream AI and modelling applications, it can be used to drive model explainability and improve model accuracy.

Ultipa’s latest major update includes an overhaul of the product’s architecture, referred to as Ultipa Powerhouse, that pushes the product’s strengths even further. Its two headline features are horizontal scalability and high performance, but as we will see, these are not the only benefits this new architecture offers. In addition, Ultipa is now natively available

This **Mutable Quadrant** is derived from 13 high level metrics, the more the image covers a section the better. **Execution** metrics relate to the company, **Technology** to the product, **Creativity** to both technical and business innovation and **Scale** covers the potential business and market impact.



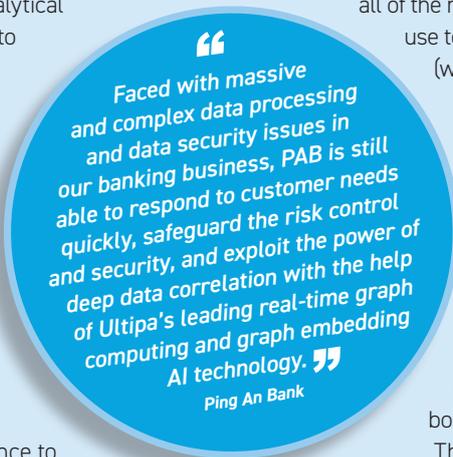
in the cloud (as well as on-prem) via AWS, with Azure and Google Cloud also supported. Various APIs are provided, as are Ultipa Deployer for “one-click deployment” and Ultipa Transporter for data import and export.

### What does it do?

Ultipa’s most significant differentiator is its performance, which (the company claims) outstrips or at the very least competes with all of the most popular and relevant graph databases in use today. This includes both open-source databases (which is frankly not very impressive) and proprietary offerings (which is).

There are a number of factors that contribute to the performance Ultipa offers. For example, it leverages high-density parallel computing to maximise the power of every CPU core available. It also utilises multi-level caching. Along with various other optimisations, these result in improved concurrency and significantly reduced latency. Moreover, Ultipa is highly scalable, and in fact boasts linear scalability.

The new Powerhouse architecture (see *Figure 1*) is designed to further improve Ultipa’s performance. For instance, it provides greater horizontal scaling via features like automated sharding, has a minimised memory footprint (also minimising any associated costs, especially the cost of running in-memory), and reduces the need for and scope of data migration by placing storage near compute. Elastic graph computing is available through the use of HDC (High Density Computing) nodes that are extremely performant and enable on-the-fly provisioning, including on-demand scaling and synchronisation. More traditional storage-compute coupled sharding setups are also available, and you can use both styles together by utilising data sync and similar features.



“  
 Ultipa graph database has been the backbone in empowering SARS' next-gen knowledge graph system in scenarios like KYC, Customer-360, Tax-Evasion and more.”  
 South African Revenue Service

For a concrete example of what this performance does for you, consider the concept of a "supernode". Supernodes are graph nodes that connect to an abnormally large (read: massive) number of other nodes, and therefore have a large number of edges connected to them. Many real-world systems are centralised around supernodes, and due to their nature supernodes tend to have a very substantial presence in whatever system they are found in. This makes them prime targets for analysis. Unfortunately, it is this very thing – their size – that makes them hard to deal with, because you need to have both the capacity to process every edge attached to the supernode as well as the performance to run queries involving them in a reasonable timeframe. Most analytics platforms, and even graph databases, fail at one of these hurdles, diminishing the efficacy of analysing any system which contains one of these supernodes and leading to fundamentally incomplete analyses. Ultipa, by contrast, has both the capacity and the performance to allow you to include supernodes in your analytics.

In addition, Ultipa provides an extensible and hot-pluggable collection of semi-/unsupervised graph algorithms, as well as graph embedding algorithms that convert high-dimensional, sparse graphs into low-dimensional, dense, continuous vector

spaces while preserving graph structural properties. Graph embeddings in particular can be very beneficial for performance, especially for machine learning-based applications such as link prediction and node classification. This makes Ultipa well-suited for addressing such use cases.

The other major point about Ultipa is that it is easy to use. Although it provides a bespoke query language (which has the inherent problem of making it harder to onboard), the "Ultipa Query Language" (UQL) is genuinely simpler and easier to use than much of what else is available. The fact that it hews close to Cypher certainly helps matters, and there is even an adaptor for moving code from Cypher to UQL. The product is also a leading supporter of the ISO GQL standard, if UQL is not to your taste. In addition, Ultipa provides easy multi-graph capabilities (which is to say, it supports multiple edges between nodes without intermediaries) and comes with over 120 graph algorithms built-in, roughly half of which are generically useful (the other half are bespoke to specific customers). Ultipa is also "demi-schematic", meaning that it can be used both with and without schemas (which are equivalent to labels in Ultipa). Last – but certainly not least – it comes with an integrated command line interface as well as Ultipa Manager, a web application you can use to access

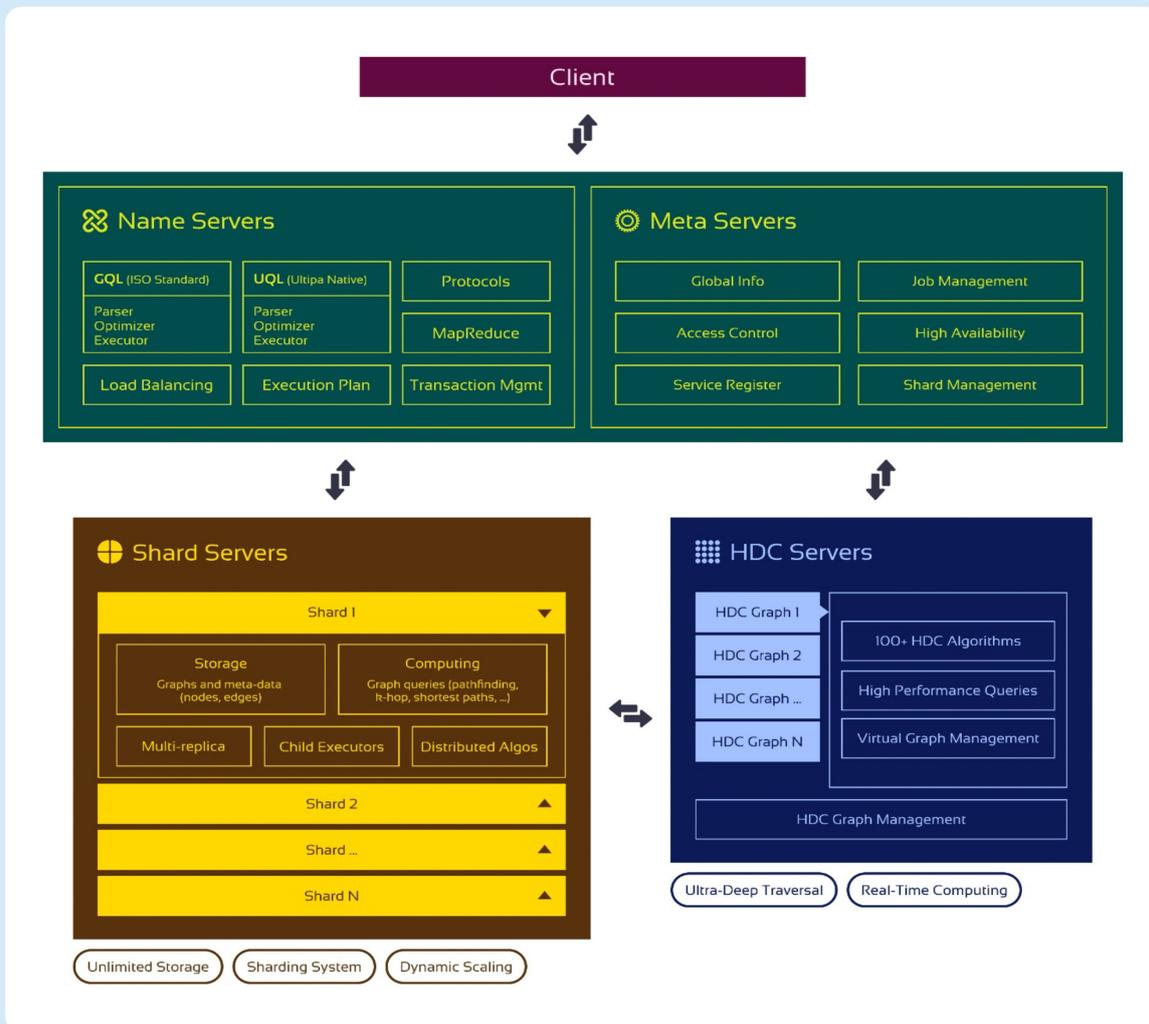


Figure 1 – Ultipa Powerhouse architecture

“  
 Ultipa system on average is 20 to over 1,000 times faster [than Neo4j] – in one scenario, we tried to explore an enterprise’s investment network which is 32-layer deep, Neo4j takes 2-hour to penetrate, while Ultipa takes only 0.2 second.”  
 Shenzhen Stock Exchange

all of Ultipa's functionality while also being a way to visualise and interact with your graph (this is shown in **Figure 2**). All of this contributes to the product's ease of use.

### Why should you care?

There are several reasons to care about Ultipa, but the two most notable are performance and ease of use. It is clear that these are the driving forces behind the product, and the Powerhouse update offers significant enhancements to both. That said, it is worth bearing in mind that high performance can improve query accuracy as well as speed, as it allows you to extract more information in less time, including details that might otherwise be impossible to extract in a manageable time frame. Essentially, Ultipa offers deep graph traversal, at scale, while maintaining high performance. It should go without saying that this is a valuable capability. The product's application to AI use cases is an additional factor that may also be worth your consideration.

We have already described several of the features that allow Ultipa to perform very competitively with its peers. This is only appropriate, as sheer speed and performance is Ultipa's main selling point. But in this case, discussing the technology only goes so far: what's important are the results. Luckily, the company provides some rather impressive benchmarks to this effect on its website, and although we generally feel it pays to be sceptical of these kinds of performance benchmarks, they are nevertheless impressive.

In terms of ease of use, the fact that Ultipa have positioned this as a major selling point is in itself a good sign: too many graph vendors do not seem to overly care how complex or

difficult their chosen query language is, so Ultipa's approach is, at the very least, refreshing. Ultipa Manager, in particular, is very welcome in how it centralises (and visualises) graph interactions. Bespoke query languages like UQL can be a hard sell if your users are already familiar with an alternative, but compatibility with GQL in addition to UQL makes this irrelevant: new users can start with a standardised query language in GQL, and incorporate UQL as and when it is appropriate to do so.

Lastly, Ultipa is intent on taking full advantage of AI to drive its graph (and conversely, make full use of its graph to facilitate AI), leaning on the basic conceit that graph and AI analytics go well together. This provides several benefits for Ultipa, accelerating its graph computing and AI model-building capabilities alike, while also adding explainability to AI via the heightened understanding that graph provides. For example, we are told that one of Ultipa's customers reduced the time it took to generate predictive models from around 2 weeks to a matter of hours, while simultaneously improving their accuracy and gaining insight into their inner workings.

### The Bottom Line

Ultipa is a strong addition to the graph database space, differentiated first and foremost by the level of real-time performance it can provide. Its HTAP capabilities and support for ISO GQL are also very much worth noting, as are its newly redeveloped architecture and AI-ready design. Put all this together, and you have a very compelling solution for high-speed, high-scalability graph analytics.

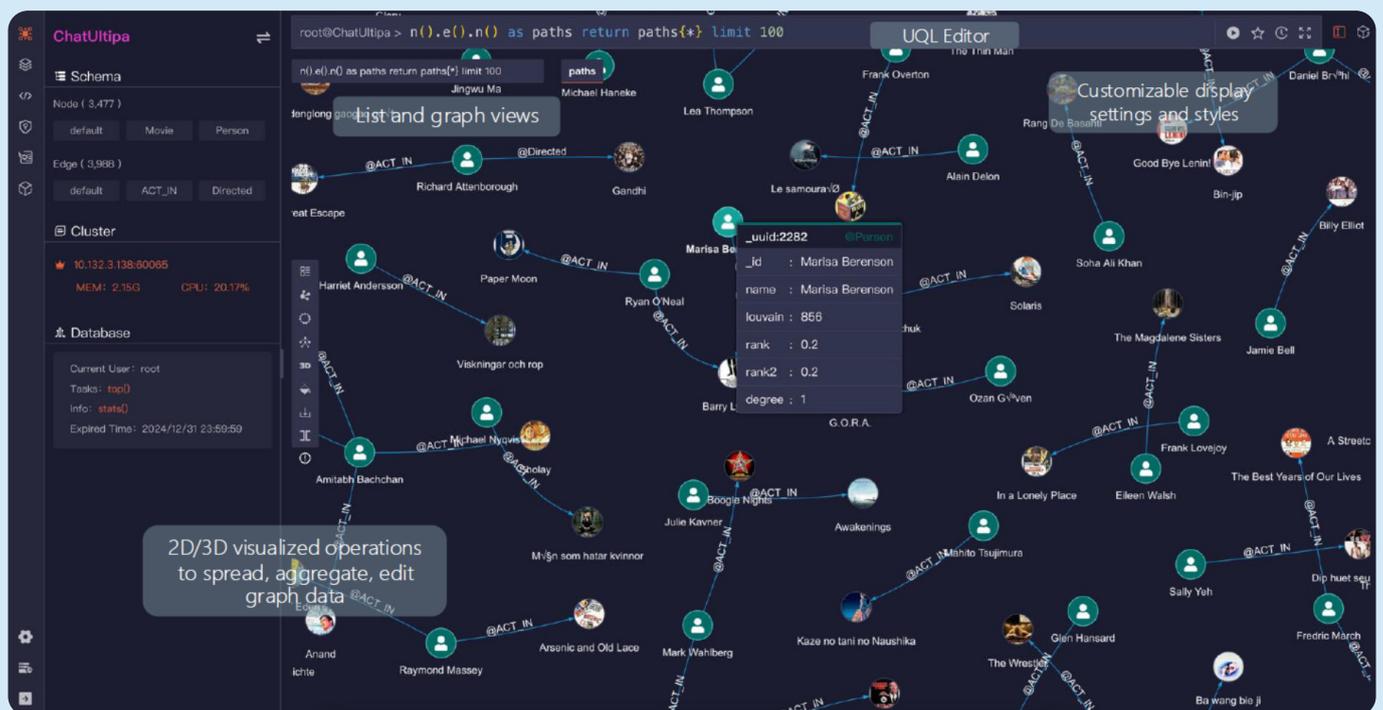


Figure 2 – Graph visualisation in Ultipa Manager