

WHITE PAPER

Financial Services: Banking on **Yugabyte** Delivers Business Value

Leverage deep experience with distributed
databases and microservices-based
architectures

“ By 2025, more than three-quarter of enterprises will have data spread across multiple cloud providers and on-premises data centres, requiring investment in data management products that span multiple locations.

Matt Aslett,
VP & Research
Director,
Ventana Research

Executive Overview

In this paper we focus on the value that YugabyteDB brings to financial services and banking organisations from the perspective of the **business**, and not solely in terms of technical features and functions.

We clearly identify the distinguishing capabilities of YugabyteDB that are of critical importance and inform the decision-making of distributed SQL adopters. We have based this paper on the demanding requirements identified as we developed our partnership with a top enterprise software provider for banks and financial services. But the resulting lessons learned are applicable to any banking or financial services organisation.

We have broken down the priorities of four key business areas: Risk, Operations, Engineering and overall Business Fit. Each section has vital and specific concerns that must be addressed, and requirements that must be fulfilled.

Risk



Concentration Risk

Regulators are increasingly concerned about concentration risk. Although CIOs are cognisant of the balance between lock-in and discount for scale, as banking and financial services regulators around the world begin to take a position, the issue becomes more pressing.

This is no longer just an issue of cloud-vendor independence. There is now an additional compliance factor added to the equation.

This reinforces the technology imperatives identified in the discussion of cloud-vendor independence: open-source software that can run on public clouds, private cloud or in multi-cloud. For simplicity, a managed cloud (DBaaS) offering adds further value. All these are imperatives YugabyteDB was designed to fulfill.



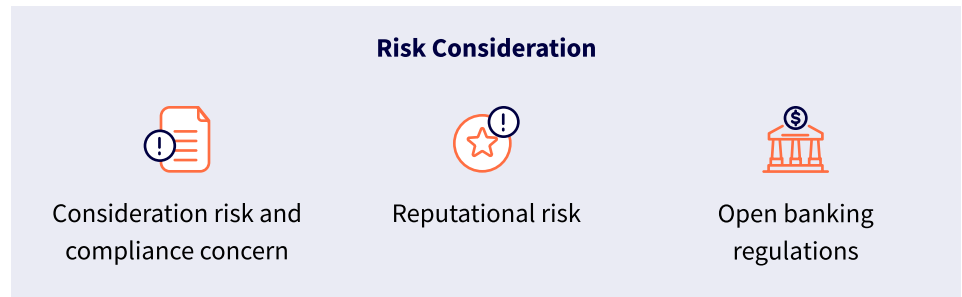
Open Regulatory Requirements

Publicly at least, the industry viewed the introduction of open regulatory requirements positively. But, there was slow adoption amongst some, even though they recognised the threats, and the need to innovate in response.

For example, although pioneered in the UK and EU, open banking now has global reach and there is widespread recognition of the opportunities and not just the threats. However, this is opportunity that comes with conditions.

New regulations do not just require API-based access to data and services, they often also mandate performance requirements, including availability guarantees and latency limits.

Traditional monolithic databases, deployed in international, trans-continental situations do not respond well to latency limits because no matter the data user's location, it will have to travel to and from a single, fixed database location.



A wise man once observed ‘ye cannae change the laws of physics,’ ([Lt. Commander Montgomery Scott, of the USS Enterprise](#)) This means built-in, minimum latency for continental and global scale data processing - unless the data can be replicated close to where it is needed.

YugabyteDB is a cloud-native database, built for exactly this scenario. With YugabyteDB, data can be located where it is needed and the database takes care of persistence and consistency. This delivers the transactional consistency that banking and financial services applications require, with the data proximity only a distributed database can deliver.

The other critical distributed database feature requirement is high availability. Once again monolithic databases can only deliver high availability through complex and expensive data replicas and failover processes. Even these may not meet regulatory requirements because failover can take hours and may lose data. This is never the case with YugabyteDB, which typically offers a three second failover for a lost node and zero data loss.

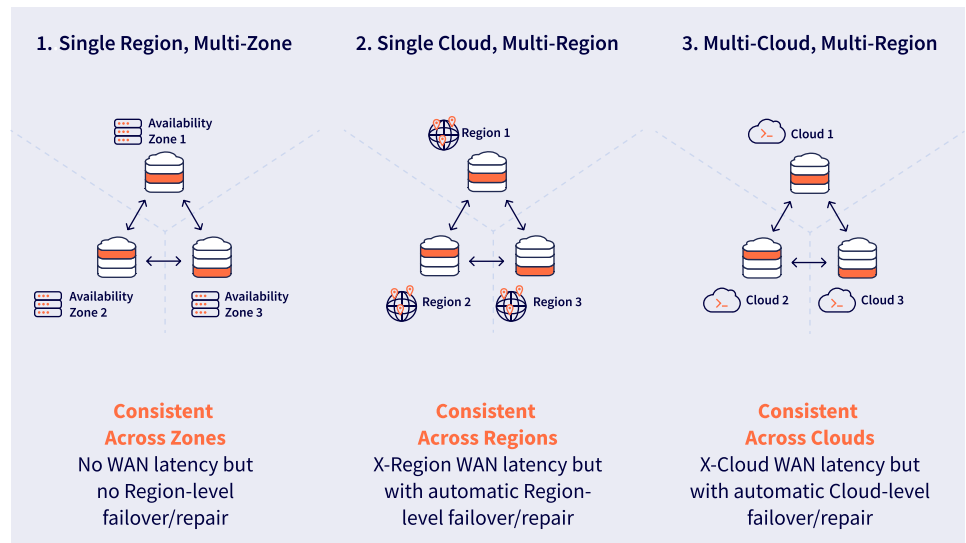
Reputational Risk

Leaving aside the high cost of service failure with respect to regulatory compliance and loss of business, even the briefest outage or the smallest data loss carries a real risk of reputational damage.

[According to a study reported in Forbes magazine](#), the third most cited reason for switching banks, after fees and rates, is “better online services and apps.” There may be lots of competition, but the service that is unavailable is certainly worse than all the options that are.

YugabyteDB’s ability to withstand infrastructure failures without loss of data enables it to deliver always-on, everywhere.





Three availability zones in a single cloud region provide resilience, except in the event of regional failure. At the cost of additional latency a multi-region deployment resolves that issue. **Deployment to multiple clouds**, or a hybrid cloud, protects against even total failure of any one provider. This last option is increasingly being adopted as the default enterprise modus operandi.

✓ Security

Securing the enterprise and its customer-facing services is a much bigger challenge than any single component like the database. But, conversely, overall security is contingent on being able to secure each component; a chain is only as strong as its weakest link.

YugabyteDB offers a powerful set of tools and capabilities to allow DBAs to secure data at rest and in motion. It supports a wide range of user authentication options, including password, LDAP and Host. This rich capability is managed via comprehensive, **simple-to-use tooling**.

Operations

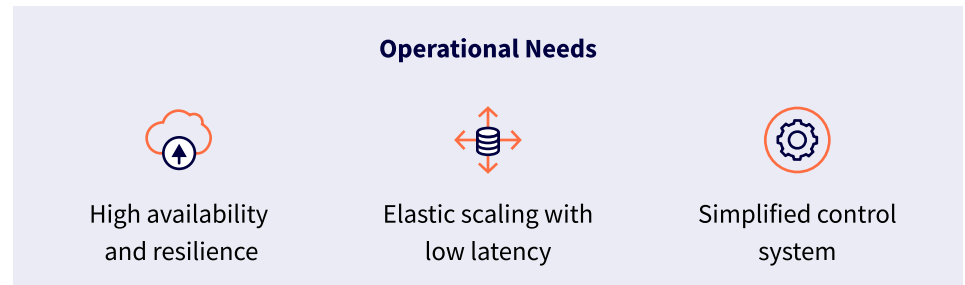
☁️ High Availability

When online availability is business critical, avoiding downtime is a key priority.

As discussed, there may be regulatory restrictions on down time. There is also reputational damage risk in downtime. And, there is just plain loss of business.

These three risks set a very high bar for availability. A bar that monolithic databases struggle to clear. Not only do they demand costly and complex Disaster Recovery (DR) strategies, but failure can still often result in significant failover down-time and even data loss.

YugabyteDB's arbitrary levels of resilience (see above) give you the resilience you need, with only three second failover and no data loss. Importantly, for operations professionals, this capability is configurable using native tooling.



For **YugabyteDB Managed** users (fully managed YugabyteDB-as-a-Service) and for **YugabyteDB Anywhere** (YugabyteDB delivered as a self-managed private, public, or multi-cloud service) mature and well-documented tooling exists to simplify operations and integrate into operations process automation. With YugabyteDB there are no third-party products to integrate and manage in complex and brittle processes, just a focus on built-in resilience.

Capacity Planning And Performance

There is no cap on the scale of operations a client might demand, but the time-windows available for certain processes, for example close-of-day processing, have strict constraints. Even though operational transaction rates are linked to a bank or financial services' customer base they are inevitably peaky and seasonal.

Capacity planning with a monolithic database means costing for peak requirements and then adding in DR. Sharding and other replication strategies can help reduce costs, but will add further complexity that will normally involve application changes. Even then they may not remain stable and might require repeated re-tuning and refactoring.

YugabyteDB Managed and YugabyteDB Anywhere deliver unlimited, horizontal, elastic scaling with a much more rapid response. This allows granular, timely investment, that is appropriate to the gain the investment produces. It provides responsiveness in times of peak demand and makes capacity planning for the database much simpler and more agile. It also takes risk and complexity out of big-ticket capacity planning around seasonal peaks and long term trends.

Database Operations

As we say about YugabyteDB Managed, "You bring the apps. We handle the database." The rationale is that the user focuses on building business value and Yugabyte focuses on delivering data-layer services to simplify the entire process.

Of course, users still have control of configuration, scaling, resilience levels, and backup frequency. But, they do not need to be concerned with how those requirements are delivered.

Delegating the operational minutiae of your database is hugely advantageous for many situations and workloads. But there are situations where an enterprise wants to retain management responsibility.

For these workloads, YugabyteDB Anywhere places responsibility for fine-grained management of the database on the user, but gives them a rich set of tools and monitors to achieve that, and to continuously manage and control the database.

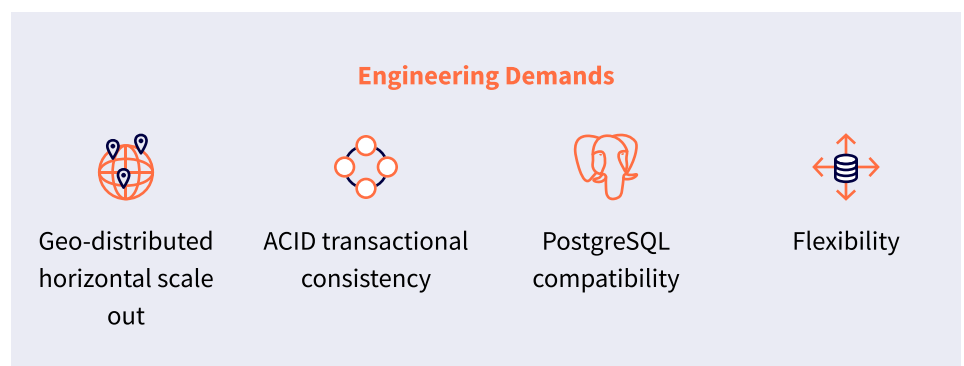
One major feature of the YugabyteDB geo-distributed, resilient database is that upgrades require no downtime. This is because nodes of the database cluster can be taken offline individually, upgraded and then restarted. This has the same effect on the database as a node failing. Being a highly resilient database, that means it just keeps working - albeit at a lower maximum capacity.

With YugabyteDB Managed, users set windows during which upgrades are permitted - to make sure maximum capacity is available at peak business periods. With YugabyteDB Anywhere, the user controls the upgrade process directly and can schedule it for the most convenient time.

Whichever solution you choose, no downtime is required to keep your database up to date; important patches and desirable new features for systems operations do not have to inconvenience or damage everyday business operations.

Engineering

Agility, familiarity, and flexibility are the three watchwords of Yugabyte's pledge to applications engineers. Our users are almost universally high-velocity, agile teams building microservices. They are driven by the demands of their businesses and resent time spent on non-business-value activities.





Database Operations

YugabyteDB's horizontal scaling and resilience means no designing and maintaining shards, read replicas or other app-level scaling and availability features. There are no complex multi-tool or even multi-vendor DR architectures.

With YugabyteDB we believe scalability and availability are mandatory database capabilities, not application engineering requirements. So, they are built-in rather than imposed as a diversionary load on application developers.



ACID Transactional Consistency

YugabyteDB supports full ACID compliance, even for geo-distributed transactions. Not every use case needs transactional consistency, but many do, especially in financial services.

One of Yugabyte's Fortune 500 customers adopted YugabyteDB principally because they had spent far too much time and energy attempting to build transactions into applications on a non-ACID compliant database. The result had been complexity in application code and repeated production data inconsistencies that necessitated expensive remediation.

When NoSQL databases overcame the scalability and resilience challenges of monolithic relation databases they did so at the expense of transactional consistency. They have spent the decade-plus since trying to add it back in, largely at the cost of latency lags that are unacceptable for most workloads. With the benefit of hindsight, YugabyteDB was able to be uncompromising in delivering scale-out and transactional consistency. YugabyteDB's commitment to applications engineers is again evident as we believe ACID compliance (where it is needed) should be a core database feature. With YugabyteDB, it is.



PostgreSQL Compatibility

The other established feature NoSQL databases abandoned, and there's a clue in the title, was SQL. SQL was the data manipulation language (DML) of choice for a generation of application developers.

SQL is a rich, powerful, and mature DML, with many stable implementations. Indeed the plurality of SQL implementations stands as a measure of its success and a challenge, since every dialect is subtly, but significantly, different.

Fortunately, in cloud-native data layers there is one SQL that increasingly rules the roost - PostgreSQL. This is now the most widely chosen SQL in microservices architectures. Yugabyte recognised the importance of PostgreSQL early and took the decision to fork and maintain in parallel the Postgres query layer as the front end to its (unlike Postgres) distributed database backend. This gives YugabyteDB a major advantage. Other databases position compatibility as useful because it confers developer familiarity. That is enough for new applications using a familiar, but modestly, different dialect.

The YugabyteDB aspiration is that compatibility should facilitate the simple migration of the many, many existing applications and services. YugabyteDB customers do not just adopt YugabyteDB for new applications, they use it to unify their whole SQL services portfolio on a single database. This enables them to deliver elastic scale-out, geo-distribution, and resilience in a single vehicle.

Flexibility

The YugabyteDB's mission is to make it easier for developers to adopt YugabyteDB for more workloads, beyond just PostgreSQL. [YugabyteDB also supports Apache Cassandra CQL](#) to cover a wider set of workloads in a single database and provide a comparable familiarity to NoSQL developers.

But, perhaps an even more important facet of the YugabyteDB commitment to ease of adoption, is YugabyteDB's support for all major cloud providers, for private cloud, for hybrid clouds and even for on-prem deployment. YugabyteDB can be deployed on these platforms in VMs, in containers, and orchestrators (i.e., Kubernetes).

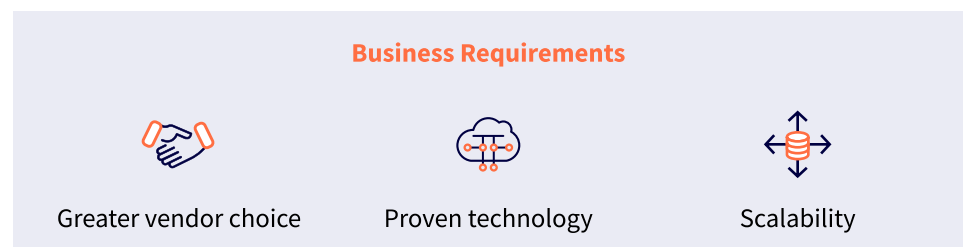
Business Fit

Choice

Over 60% of the world market for cloud services provision is held by just three providers.

However, there are many other choices available, and even between the big three players there are pluses and minuses that affect the evaluation of different workloads in different geographies. This is also all in flux, as new and improved offerings appear and competitors strive to wrench market share from incumbents.

The net outcome is that large enterprises no longer discuss which cloud supplier to select; multi-cloud is now a given. A single cloud provider may deliver initial simplicity, but the multi-cloud trend advances because enterprises rightly fear swapping on-premises lock-in for cloud lock-in.



Businesses that rely on proprietary solutions from specific cloud providers reduce their ability to adopt another cloud provider, as they cannot deliver the same solution. This is a key reason why third-party, open-source solutions are gaining such traction.

YugabyteDB is 100% open source and available on all the major clouds, in private clouds and on premises. Significantly, it is also available as a DBaaS, improving operational simplicity, but without the lock-in. The same data and applications will run unchanged on any other YugabyteDB installation or service.

✓ Confidence

Choice is important, but organisations are right to be wary of the unproven attractions of innovative new technology.

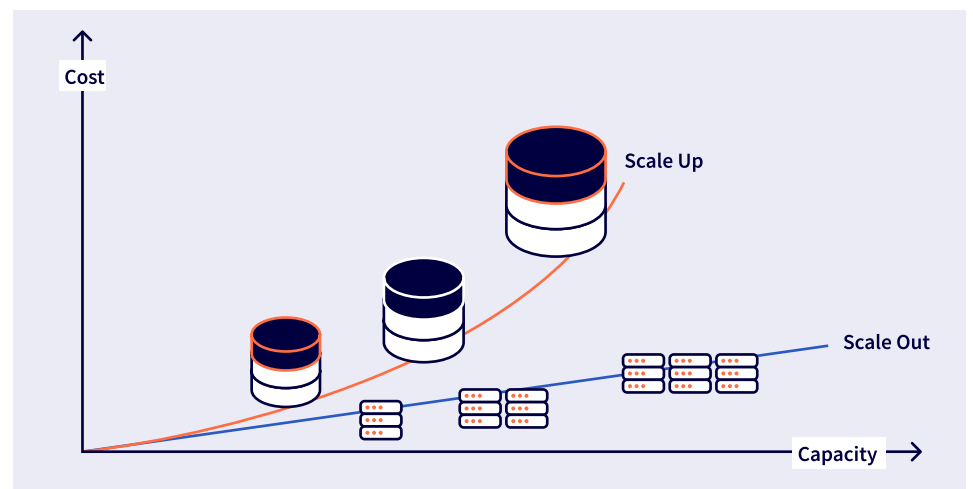
There are countless war stories about predicted grand technological leaps that history proved to be delusions. The cloud-native distributed transactional database capability that YugabyteDB offers has been a data management grail for decades. The last half decade has seen this finally become reality and many banking and financial services organisations are now delivering value to their customers from YugabyteDB including Wells Fargo, Charles Schwab, US Bank Mastercard (Ciphertrace), Xignite, Hudson River Trading, Fiserv, and Temenos.

↕ Scalability

For on-premises data centres, capacity planning is a complex business. The challenge is to balance expensive upgrade investment against the likelihood and timing of the business expansion that justifies it.

This was one of the initial drivers for cloud migration: turning “lumpy” cap-ex into “smooth” op-ex. But, data processing has been dominated for decades by monolithic databases whose only upgrade path was to run them on a bigger single machine.

In the cloud, huge proprietary hardware boxes are replaced by myriad smaller commodity boxes (typically virtual) to share the load. Lift and shift of legacy databases is not that simple in those circumstances; it entails diverting valuable engineering capacity away from adding business value, to managing and maintaining the complex and brittle application changes required.



“ At Temenos, we believe that the banking industry is moving inexorably towards open banking as it spreads across the world and proprietary data and proprietary relationships get relegated to the past.

Kanika Hope,
Temenos Chief
Strategy Officer

An ideal solution is to adopt a cloud-native database engineered from the ground up to operate in a cloud environment. One that scales out to as many machines as needed. One that can elastically expand and contract as capacity requirements flex around peaks in demand. YugabyteDB is such a database.



Summary and Conclusions

This white paper specifically focused on how the capabilities or limitations of the data layer can impact different aspects of your business. By viewing the data layer from these different perspectives, certain features of the database assume greater or lesser significance, but a clearer picture emerges of the critical capabilities required.

How rapidly, securely and effectively an organisation, in any position in the banking or financial services supply chain, moves to embrace and exploit this new world will be one of the most critical measures of its success.

However, the ability to embrace this new paradigm necessitates an API-driven, microservices, cloud-native database architecture. One that maps a demanding combination of database capabilities: distributed data, unlimited scale-out, always-on resilience, global transactional consistency, and more.

YugabyteDB was built to deliver these strong and future-proof capabilities and is being successfully adopted and endorsed by digital transformation leaders across the world for banking and financial services.



Yugabyte is the company behind YugabyteDB, the open source, high-performance **distributed SQL database** for building global, cloud-native applications. YugabyteDB serves business-critical applications with SQL query flexibility, high performance and cloud-native agility, thus allowing enterprises to focus on business growth instead of complex data infrastructure management. It is trusted by global companies in cybersecurity, financial markets, IoT, retail, e-commerce, and other verticals. Founded in 2016 by former Facebook and Oracle engineers, Yugabyte is backed by Lightspeed Venture Partners, 8VC, Dell Technologies Capital, Sapphire Ventures, and others.

Get in Touch

www.yugabyte.com | contact@yugabyte.com

