

WHITE PAPER

YugabyteDB: The **Idea Database** **Choice for ISVs**

Discover Why YugabyteDB is Winning the ISV Database Battle

WHY ISVs NEED TO EXPLORE NEW DATA LAYER OPTIONS TO STAY AHEAD OF THE COMPETITION

Enterprises and users increasingly expect software-based services to be delivered “as a service” — delivered on demand, easily scalable, accessible from anywhere, and easy to use.

In response, modern independent software vendors (ISVs) have changed the way their products are architected. They have embraced a microservices-based architecture to facilitate rapid delivery cycles and optimized app deployment for flexible public cloud infrastructure. But, none of these updates touched the core data layer and underlying data that powers their apps and services.

Until now, ISVs have had to make heavy tradeoffs at the core transactional data layer underpinning their solutions, with three less-than-ideal options to choose from:

OPTION 1

Stick with traditional relational databases (RDBMS) to take advantage of SQL capabilities and consistency, but face having a monolithic ‘anchor’ at the heart of their cloud native solution stacks.

OPTION 2

Adopt a NoSQL database, but assume additional application complexity and manage the loss of strong data consistency.

OPTION 3

Choose the database provided by your cloud service vendor and find yourself locked into that proprietary (non-open source) platform and unavailable to customers using alternative cloud providers or their own on-prem hardware.

YugabyteDB offers ISVs another option, one that eliminates the forced trade-offs highlighted on the previous page.

YugabyteDB is a modern, cloud native [distributed SQL](#) database built for new and existing applications. YugabyteDB delivers:




Architectural Simplicity

YugabyteDB's advanced storage layer automates complex tasks, like sharding and re-balancing, to simplify app development and deployment. By delivering both rich relational and NoSQL capabilities, YugabyteDB also allows an ISV to eliminate the costs and complexity of requirements multiple different databases and/or relying on the application layer to managing this complexity.



Resilience and zero downtime

With automatic sharding and data distribution, YugabyteDB delivers resilient services with a highly available architecture that can survive node, zone, and region/DC failures. In addition, perform regular maintenance and upgrades with zero downtime.



Flexible Deployment Choices

With YugabyteDB, an ISV solution can be deployed anywhere: on-premises, public/private cloud, hybrid cloud, or fully-managed DBaaS. The flexibility allows ISVs to meet the needs of a broad range of customers.



100% Open Source

YugabyteDB is a [PostgreSQL-compatible](#) and completely open source database. It is committed to open standards and not locked into a single cloud or environment. This means ISVs can control their own destiny and avoid the risk of cloud lock-in.

This solution brief focuses on the long-term value that YugabyteDB provides to ISVs as they develop their data strategies and build a cloud native foundation to power their long-term success with flexibility and innovation.

WHAT'S HAPPENING IN THE WIDER MARKET?

When we look at today's ISV landscape we see newer ISVs predominantly adopting "cloud only" service delivery models like Snowflake, Monday.com, and Salesforce.com. These SaaS solutions enable the ISV to rapidly iterate and enhance their offerings unencumbered by the massive time and financial cost of taking their customers through painful upgrade cycles. End customers demand best-of-breed solutions, but they don't want to hire armies of operations folk to maintain these solutions.

More established ISVs are moving toward a "cloud-first" service delivery model. This means they prioritize their SaaS offerings, but continue to support end customers who demand a "self-managed" solution in their own environments (public cloud, private cloud, on-premises or hybrid). There are many underlying factors that determine whether an ISV has embraced this model, but it's clear that ISVs that don't embrace technological innovation will be disrupted by more nimble, cloud native competitors over the next 3-5 years.

MONOLITHIC DATABASES FAIL IN A CLOUD NATIVE WORLD

ISVs that fail to evolve will continue to spend time and resources dealing with many of the common challenges associated with monolithic databases, including:



Forced Tradeoffs

You have to pick between the strong consistency of RDBMS and the resilience and scale of NoSQL.



Complex Code

Your engineers have to build database scaling and resilience manually, increasing costs and fragility.



Operational Headaches

Limited automation, proprietary APIs and scaling woes increase time-to-value and lower responsiveness.

KEY REASONS ISVs SHOULD ASSESS THEIR DATA LAYER

How ISVs choose to build and manage their data layer will have major implications on their operational efficiency and effectiveness. It will also impact how quickly they can introduce new product features and new data-centric services, and the choices they can offer to existing and prospective customers.

For example, how a solution is architected for elastic scale, resilience, availability, or latency (the list goes on) has a big impact on the end customer experience (outbound), as well as the complexity, management overhead, cost and agility (inbound) that will result.

Let's explore this further.

SIMPLIFY YOUR ARCHITECTURE AND OPERATIONS

Many ISV solutions require multiple databases to deliver all the functionality they need.

Typically we see both NoSQL and traditional SQL databases deployed side-by-side to overcome each other's deficiencies. NoSQL databases rose to prominence to cover the shortcomings of traditional monolithic SQL databases, namely [geo-distribution](#), high availability, and horizontal scalability. Unfortunately, as NoSQL databases cannot guarantee multi-row ACID transactions, they sacrifice data accuracy and consistency.

ISVs have had to overcome these tradeoffs by either implementing expensive and often manual workarounds in the application layer, or by accepting the complexity of deploying both database types within their solution.

One Fortune 500 customer attempted to build and maintain transactional consistency on a NoSQL database. Despite their best efforts they had to make major changes to their application code, and still experienced repeated production data inconsistencies that necessitated expensive remediation. By adopting YugabyteDB, they completely removed that challenge.

Justuno

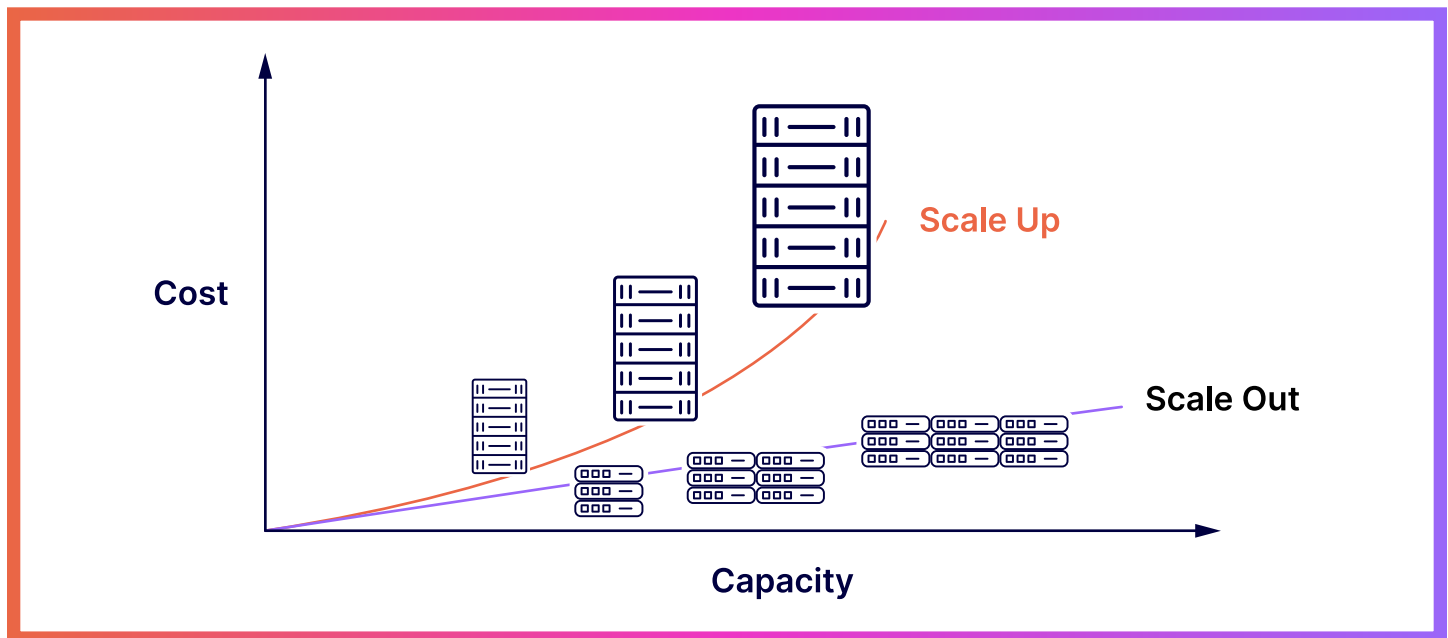
[Justuno](#) provides Conversion Rate Optimization for websites to help retailers increase revenue by enabling personalized onsite messaging, intelligent cross-selling, and upselling to visitors. Justuno originally relied on a variety of databases including CockroachDB and Microsoft SQL Server to support their critical applications. A mixed database environment led to a complex and brittle architecture that didn't meet their performance expectations. YugabyteDB enabled Justuno to consolidate their database sprawl, reduce latency, and dramatically simplify their data infrastructure operations.

YugabyteDB is a distributed SQL database that supports full ACID compliance for geo-distributed transactions. It offers the best of SQL and NoSQL in one database, and it automates some of the painful processes that are managed either manually or through application-level changes. For example, the distributed storage layer automatically handles sharding across clusters and the built-in, intelligent load balancer automatically and dynamically distributes data when nodes come on or off line. ISVs can radically simplify their solution architectures and can reduce the risks, resources, and costs that are tied up in non-differentiating activities.

xignite®

Xignite provides financial market data, including real-time and reference market data to 700+ customers. As their business continues to grow, their data layer must seamlessly scale while still ensuring data consistency. Xignite looked at scaling SQL Server, adopting MySQL, or implementing a caching layer, but each alternative had drawbacks including lack of scalability, added complexity, and operational overheads. Instead they chose YugabyteDB, which provided multi AZs and region geo distribution and scale, while eliminating the need for a separate caching tier.

SCALE YOUR BUSINESS WITH CONFIDENCE THANKS TO RAPID, EFFORTLESS SCALABILITY



There is no doubt that traditional SQL databases can scale. Some of the most data intensive banking and retail ISV applications have traditionally run on monolithic SQL databases such as Oracle or DB2. However, to grow, ISVs had to scale UP, buying and managing expensive,

customized hardware to meet their needs—destroying the cloud dream of rapid scaling on commodity hardware.

As an application scales, the upgrade path for the database is to run on bigger and bigger machines, with costs rising exponentially. As an ISV (or end customer that uses an application with a traditional SQL database) grows, they are punished with non-linear cost increases.

In contrast, YugabyteDB is a scale OUT database. It can scale horizontally without limits by leveraging commodity hardware and near-linear increases in performance. The scale out architecture means YugabyteDB can massively scale while infrastructure costs grow linearly (each additional server is a linear cost). Reduced costs are not the only benefit. ISVs can also avoid expensive and risky migrations caused by the scaling limitations of monolithic SQL databases.

What happens when workloads are unpredictable, spike without warning, and/or have only seasonal demands?

Capacity planning with a monolithic SQL database means anticipating and building for maximum peak requirements and then adding disaster recovery. Sharding and other replication strategies can introduce some elasticity to reduce costs, but this adds significant complexity and can introduce stability concerns that need repeated re-tuning and refactoring.

YugabyteDB is engineered from the ground up to operate in a cloud environment, which means it can respond quickly to meet capacity requirements and flex around peaks in demand. The database can automatically scale up and down, with all [sharding](#) included out-of-the-box. This simplifies capacity planning, dramatically reduces risk, removes complexity, avoids stability tradeoffs, and can have a big impact on costs.



Turvo is a company that provides real-time logistics information. They needed to rapidly scale their solution and underlying database, while maintaining low latency and a favorable total cost of ownership. With YugabyteDB, they were able to point, click, and reshard the cluster to double the size of the database. *“YugabyteDB has the ability to just naturally expand the cluster and everything just works out without any hiccups,”* said Nick Hristov, Sr. Software Architect at Turvo.



Admiral is the leading Visitor Relationship Management (VRM) platform, serving thousands of publishers worldwide. Admiral must handle relationship data instantly and in a highly configurable way based on user and publisher preferences. *“YugabyteDB is a solid technical choice that meets all of our key requirements, including massive scale, high performance, low latency reads, and cloud native out-of-the-box.”* said James Hartig, Co-Founder at Admiral.

DELIVER ALWAYS-ON SERVICES WITH HIGH AVAILABILITY AND ZERO-DOWNTIME UPGRADES

ISVs know that they must meet their end customer expectations of always-on, any-device, any-place access to services.

Service availability commitments (planned downtime) are often a key factor of whether an end customer will consider an ISV solution at all. SLAs are becoming more onerous, compounded by stricter regulatory demands. However, it's unplanned downtime that can really hurt an ISV, as it comes with costly penalties, reputational damage, erosion of customer confidence, and loss of business.

Monolithic databases try to achieve high availability through complex and expensive data replicas and failover processes. At best this disaster recovery can take hours and often results in loss of data. For a global bank, retailer, telco provider or any number of other industries that deliver always-on services, every minute of downtime can represent six figures of lost revenue.

YugabyteDB is a geographically distributed, resilient database that can survive loss of an entire failure zone and offers a three second failover for a lost node and zero data loss. The geo-distributed design allows YugabyteDB to withstand infrastructure failures and upgrades without downtime and with zero data loss. This is because nodes of the database cluster can be taken offline individually, upgraded, and then restarted. An ISV can take advantage of new features and releases without impacting the delivery of their service to end customers.



Temenos, a leading ISV that supports over 80% of the world's Top 50 banks, is constantly investing in innovation and looking for new ways to help its base of 3,000 global banking clients. As their clients shift to cloud services and reduce the overhead of on-premises infrastructure management, Temenos has a goal to rapidly expand the value and performance of their managed cloud services. However, the monolithic transactional databases traditionally used in the Temenos Banking Cloud posed a major obstacle to delivering on that goal. By switching to YugabyteDB for their mission-critical transactional workload, which demands strong consistency across sites, Temenos successfully delivered over 100k business transactions per second, while also delivering a cloud native and cloud agnostic solution.

MAXIMIZE YOUR REACH BY PROVIDING DEPLOYMENT CHOICE AND FLEXIBILITY

For many ISVs, the public cloud infrastructure requirement for their hosted solutions is typically one of, if not the, largest operational costs—adding up to millions (or tens of millions) a year. Some ISVs find themselves locked into just one cloud provider because they chose to use a cloud service provider's proprietary database offering. Over time, this reduces their ability to take advantage of both technical and commercial offerings from other providers.

YugabyteDB runs in AWS, GCP, and Azure, but is not committed to any specific provider. This means an ISV can avoid being locked into a single cloud provider, giving them leverage to negotiate effective rates with each.

ISVs want to serve the largest possible market, while maintaining a single tech stack for supportability. Many large enterprises still want to operate mission critical transactional workloads in a hybrid topology as they gradually move to the public cloud.

Additionally, enterprises want the choice of running workloads on VMs, containers, or bare metal depending on use case and their maturity. ISVs are finding that their customers mandate where their application will run - some demand AWS, some GCP, others want a portion on prem, another portion on AWS, and a backup on Azure with the option of moving the on-premises components to GCP over time!

YugabyteDB allows an ISV to build one solution stack that runs in the cloud, on prem, hybrid, and with deployment options. This means an ISV can address the widest possible customer base. And because YugabyteDB is 100% open source and leverage open standards, you have added transparency into the solution and can avoid vendor "lock in." As an ISV, you can control your destiny. If you want, you can contribute to the open source project, get involved with the community, and benefit from the rich and extensive open source PostgreSQL and YugabyteDB ecosystem and the rich tooling that comes with it.



Broadleaf Commerce provides a flexible, extensible, customizable, and scalable microservices-oriented eCommerce framework including 26+ services. When Broadleaf is deployed with YugabyteDB, a retailer can achieve a true "always-on" geographically distributed, multi-master eCommerce environment that approaches five nines of availability. Broadleaf's solution is designed to be deployed using Docker and Kubernetes in multiple cloud environments or a private cloud or data center, and YugabyteDB supports all these deployment options.

ACCELERATE INNOVATION BY TURBO CHARGING YOUR DEVELOPERS' PRODUCTIVITY

ISVs feel the pressure to constantly push new capabilities into production. Blink and your competition has released five new features. Many compelling technologies don't gain traction because engineering and development teams don't have the bandwidth or desire to learn a new technology. The same is true for databases. Developers cannot dedicate cycles to learning entirely new, custom-built database APIs as this is a massive hit to feature velocity and difficult due to talent sparsity.

YugabyteDB provides all the benefits described above, while offering a traditional PostgreSQL API to developers. YugabyteDB is not just wire compatible, but also virtually fully code-compatible with PostgreSQL, supporting advanced features like triggers and stored procedures along with the ecosystem of PG extensions. As PostgreSQL expertise is widespread, this means developers can be up and running in minutes.

Meanwhile the complexity behind a geo-distributed, resilient, elastically scalable database requires no additional activity from the developer—it just works. Sometimes the release of a new application feature can become an unexpected success, requiring the database to scale quickly and elastically. YugabyteDB takes this in its stride, ensuring that success doesn't turn into a headache for operators and developers.



Narvar enables over 650 retailers to deliver personalized customer experience based on their own data. As their business grows, Narvar wanted a database that would allow them to avoid vendor lock-in, stay GDPR compliant and be cost effective, but most of all, they wanted to maintain and improve developer productivity. Narvar's Chief Architect, Uday Shanmugam, commented; *"It was easy for our developers to get started with YugabyteDB using the distributed SQL APIs. YugabyteDB eliminated the need for a separate cache and multiple databases, allowing us to roll-out new features a lot quicker."*

DELIVER ENTERPRISE-GRADE SECURITY WITH NATIVE CAPABILITIES

Ensuring the security of customer-facing services is a far bigger challenge than any single component like a database. That said, data security should be at the forefront of an ISV's technology strategy, and for good reason. ISVs are being put through increasingly rigorous security assessments by their customers and prospects. Any security lapses can affect many end customers, with major reputational and financial consequences. ISVs need the confidence that every aspect of their solution stack provides world-class security capabilities.

YugabyteDB is built for sensitive, mission-critical transactional data, and as expected, security features front and center in everything we do. Aligned to our security mission and objectives, Yugabyte has established a security program based on the ISO 27001 security framework and is audited by an independent accounting firm using the SSAE18 SOC 2 standards. This includes testing against the Trust Services Criteria related to security, availability, and confidentiality. Our security protocols fall into 3 separate areas:

- **Product Security:** security features of our products, including how we help customers meet security and privacy compliance requirements
- **Product Development Security:** how we build our product securely
- **Corporate Security:** how we secure our company and comply with privacy obligations

For more information on Yugabyte security, visit our Security & Trust Center [here](#).



Netskope is a leading provider of cloud native security solutions. Netskope selected YugabyteDB to store critical data as it is processed and analyzed to generate alerts and insights. Serving a mission-critical role in their service, YugabyteDB was selected thanks to its ability to meet key SLAs related to transactions per second, average latency, and P99 query latency.

CONCLUSION

This solution brief identifies how YugabyteDB can positively impact the services and applications that an ISV provides to its customers. It also evidences the way that YugabyteDB can improve how an ISV builds, manages, maintains, and delivers those services and applications.

Cloud, containerisation, microservices, and orchestration represent capabilities that have powered the digital transformation of stateless application development. Up until now, the data layer has lagged behind as first monolithic databases, then NoSQL databases, failed to deliver all the capabilities of a true cloud-native data layer.

YugabyteDB was built specifically to deliver these strong, future-proof capabilities. It has been successfully adopted and endorsed by digital transformation leaders in multiple industries across the world.

Yugabyte counts major ISVs servicing retail, financial services, telco, and others as customers—each running mission-critical services to ensure the satisfaction and success of their end customers.

Get started today by trying a free version of YugabyteDB Managed in the cloud, or explore just some of the YugabyteDB customers and ISVs that have already made the switch to distributed SQL.



Get in Touch

yugabyte.com | contact@yugabyte.com

