
FINTECH FUTURES | REPORTS
PERCONA
The case for open source database software in finance: Key drivers and strategies for adoption

Introduction

Financial services firms face a multitude of challenges, with new ones coming on top of long-standing ones. Competition and customer expectations are increasing, at the same time as firms need to adhere to strict regulatory and security standards. In parallel, there is the need to modernise legacy systems and reduce infrastructure costs while accelerating time-to-market. It is essential to ensure 24/7 availability and optimal application performance.

Enterprise-grade open source software enables companies to freely scale without limits, licensing costs or compromises to security. Proprietary databases stifle innovation by limiting the ability to quickly scale - up, down or linearly - to support fluctuations in demand or deliver new products and features. Moving to fully supported, secure enterprise-grade open source replacements frees up companies to build, deploy and scale as they see fit.

The amount of data that needs to be held and analysed is growing but this shouldn't equate to a rise in infrastructure costs. In fact, the opposite is achievable with open source.

This paper looks at the business benefits and cost savings that can be achieved by migrating to open source. It dispels some of the myths around making such a move, and it describes some of the things that need to be considered when planning and carrying out the migration, then building and managing a high availability (HA) architecture.

Financial services – mounting challenges

Customer expectations are high, with a demand for financial services to match the user-friendly and resilient customer experience of other verticals. Downtime when executing a transaction or making a payment can undermine years of effort to build customer trust.

Indeed, trust in financial services firms has been eroded over the last few years. For instance, a 2022 survey carried out by NetApp found that only 78 per cent of UK customers felt their money was safe at their bank, and only 60 per cent felt the same when it came to their personal data. The survey found that while convenience is crucial, security concerns and a dislike of AI-based services are barriers to technology adoption in the financial services sector.

For financial services companies, any downtime can mean a loss of customer confidence. Customer loyalty can be lost in an instant and, particularly with the arrival of open banking across the globe, it is ever easier for customers to move their business. In the UK, by early 2023, there were seven million active open banking consumers and SMEs, with steep uptake as well in other countries.

78%

of UK customers felt their money was safe at their bank

60%

of UK customers felt their personal data was safe at their bank

The competition has been building for some time, from challenger banks, fintechs (typically seeking to excel in narrow areas of business, with leading edge technology and customer experience) and heavyweight players from other verticals, including high-tech and retail companies.

> *With enterprises always on the lookout for ways to reduce costs, eliminating millions in unnecessary spending on databases is a big win.*

The typically complex and legacy nature of most financial services firms' IT is putting a strain on resilience and up-time. Too many banks, in particular, have made headlines and incurred fines and loss of confidence because of high-profile outages.

The need for flexibility and breaking away from vendor lock-in

The ability to meet the challenges is hindered by a lack of flexibility and vendor lock-in.

The typical mix of legacy and proprietary databases that permeate financial services firms can restrict what companies can do. It can be difficult to innovate and move quickly to respond to customer demand or changes in the market.

As most financial services companies will know, vendor lock-in with proprietary databases such as Oracle also means exorbitant support costs and sky-high licensing fees. A single vendor owns the software, controls the code and can modify the licence.

Indeed, with enterprises always on the lookout for ways to reduce costs, eliminating millions in unnecessary spending on databases is a big win.

In fact, switching from a proprietary database to an open source architecture brings a multitude of benefits. In addition to avoiding costly licence fees, an open source database offers data portability; the freedom to deploy anywhere anytime; and great software designed by a community of contributors who prioritise innovation and quality.

The application uptime and availability that is a key factor in customer acquisition and retention in financial services can be achieved, in the form of highly available, resilient fintech products. In this, the back-end database architecture is a fundamental building block. Essentially, an HA database is an important enabler for high-performance, dependable apps that ensure customers have access to financial data whenever needed, with no downtime or outages.

Open source software brings the promise of freedom, combined with the types of 'premium' features that are associated with proprietary vendors. It allows users to seize new opportunities while controlling costs.

The benefits of open source databases

While the business benefits of moving to open source are clear from many case studies across many sectors, there can still be concerns about a lack of maturity of open source offerings, a lack of inbuilt HA tools and other potential obstacles to making the move.

In fact, these issues have been addressed and there is a strong body of evidence from many institutions across the globe that shows how to go about the migration and demonstrates the strong business benefits that can be achieved as a result.

Open source databases are now mature, functionally rich and scalable. PostgreSQL, for instance, has benefitted from more than 23 years of active development and is known for its reliability, robust features and performance.

Open source frees users from proprietary ecosystems, allowing them to use the systems of their choice. They are portable, giving the option to deploy anywhere and move at any time. They are highly extensible, supporting different data types and sources (including Oracle), as well as various programming languages, indexing methods and non-relational workloads. And they offer a strong procedural language and a wide range of special functions, making it easy to expand to fit any environment or need.

This brings the need to carefully review the various extensions and tools available. These extensions and tools are excellent for enriching their respective databases but, as the environment scales, may not be able to keep up with evolving, more complex requirements.

Of course, in financial services, scalability is an important consideration. The amount of data that firms need to process continues to grow exponentially, in part for regulatory reasons.

> *Open source encourages technological agility by removing obstacles to innovation, including burdensome contracts, lengthy procurement processes and prohibitive costs.*

The emerging area of environmental, social and governance (ESG) reporting, for instance, is adding to firms' data collection, management and reporting requirements.

Open source databases can scale to support heavy workloads through capabilities such as efficient handling of concurrent connections, query parallelism and logical replication.

That maturity extends to another area that financial services firms need to get right: **security**. Database security is evaluated within the authentication, authorisation and accounting (AAA) framework. To meet this standard, a database needs comprehensive features for authenticating users, authorising what they can do with the database and auditing (accounting for) what they did while using it. Open source databases provide robust AAA capabilities, such as authentication based on users and roles and encryption at multiple levels, from passwords to specific data columns to file systems, plus a comprehensive audit logging system.

Open source databases also allow you to own your data and dictate where it is stored (on-premises/cloud/hybrid cloud), which may be necessary to comply with more stringent regulations. In other words, a proprietary vendor may own your data in a way that violates a regulation/standard. Basically, you can mold open source databases to fit needed standards rather than hoping to find a vendor that has the capability.

In addition, migrating to open source means financial services firms can tap into a **collective pool of expertise**. This was

highlighted by The Linux Foundation in 2022's "A Guide to Enterprise Open Source". By participating in open source projects, enterprise developers can build on the work of others. They can experiment with new features and optimisations much faster within the community than in-house, leading to breakthroughs in technology. There is the ability to focus their efforts on differentiation in higher stack levels and improve their unique value to their consumers. In doing so, they contribute to an organisation's overall innovation pipeline.

This can also be one factor in creating a more **vibrant development environment**, thereby addressing another headache for many financial services firms – how to attract and retain good developers. Working in open source, rather than with proprietary databases, can be more interesting and collaborative.

There is not only a larger pool of development resources, but developers can form different groups that simultaneously work on various projects along multiple paths.

These concurrent efforts accelerate feature creation and enhancement in areas such as availability, performance, scalability and security.

Open source also encourages **technological agility** by removing obstacles to innovation, including burdensome contracts, lengthy procurement processes and prohibitive costs. By cutting 'red tape', it enables users to customise the source code, try new applications and services with lower upfront investment and deliver features to market faster.

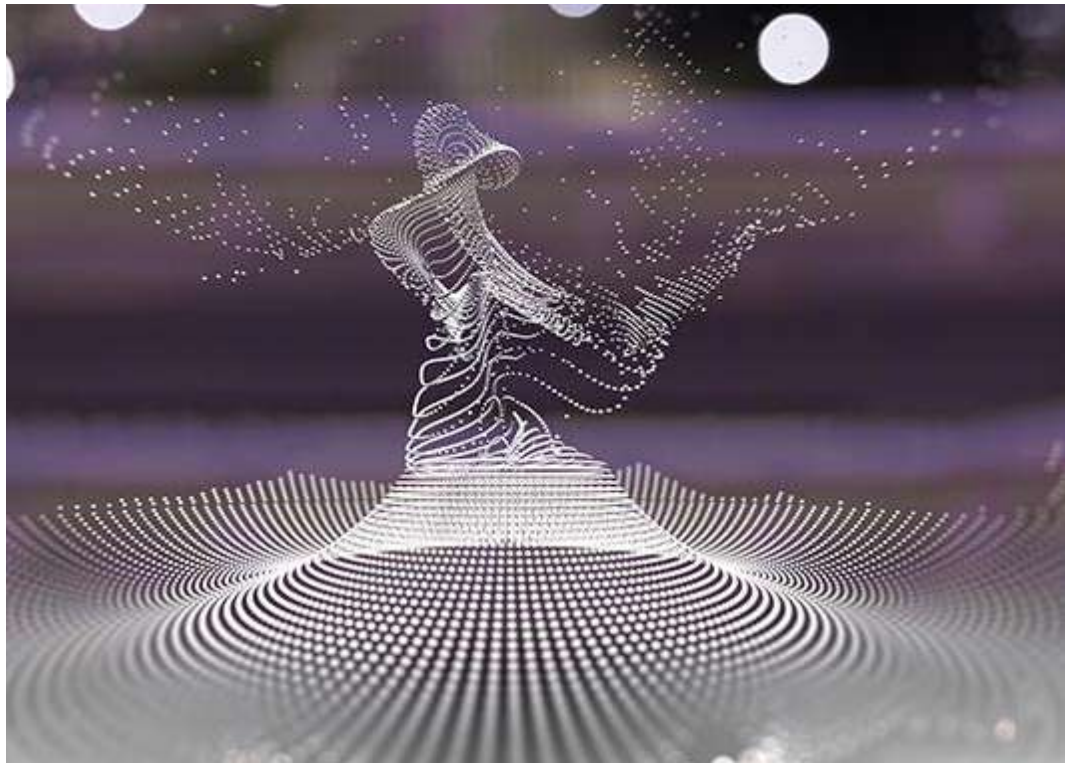

Migration challenges

As with any migration journey, planning is vital. One of the common mistakes with respect to database migration is underestimating how difficult and time consuming it is. Many companies do not develop a thorough plan and instead struggle with the process, bringing undesirable outcomes.

Also critical to success is to manage the likely dearth of internal technical skills, often combined with an unclear understanding of business requirements and project scope. This can apply even within large financial services companies with massed ranks of IT staff. For most, this is new territory, requiring new skills and expertise.

There are other considerations, such as dealing with unpredictable downtime, preventing data loss, handling syntax and programming language differences, and mapping different data types.

A successful, cost-effective migration will be built on a solid strategy and deep expertise to assure that critical workloads are moved with minimal disruption and unforeseen issues.

This requires assessing the existing environment to identify whether the target platform is appropriate and to analyse the level of complexity and potential obstacles.

More specifically it involves:

- Identifying changes that must be made to applications, including whether or how they need to be changed. It is also important to estimate the time it will take to convert the embedded standard SQL code to a new syntax.
- Identifying incompatible features. These can create the need for application rewrites, which increases development time. Addressing incompatibilities in advance allows users to identify, apply and document patches to existing extensions, thereby avoiding delays.
- Listing objects and application code to be converted. By listing the database schema(s) for complex objects and application code, firms are able to rank their complexity and convert them ahead of time.
- Determining whether new extensions are needed. This means identifying where a database has unique features for which there is no compatible solution in the new environment. This gives users the chance to write and document new extensions prior to the migration.
- Completing the assessment allows users to document their database schema in detail so they can create a proposed setup as a guide to a smooth migration.

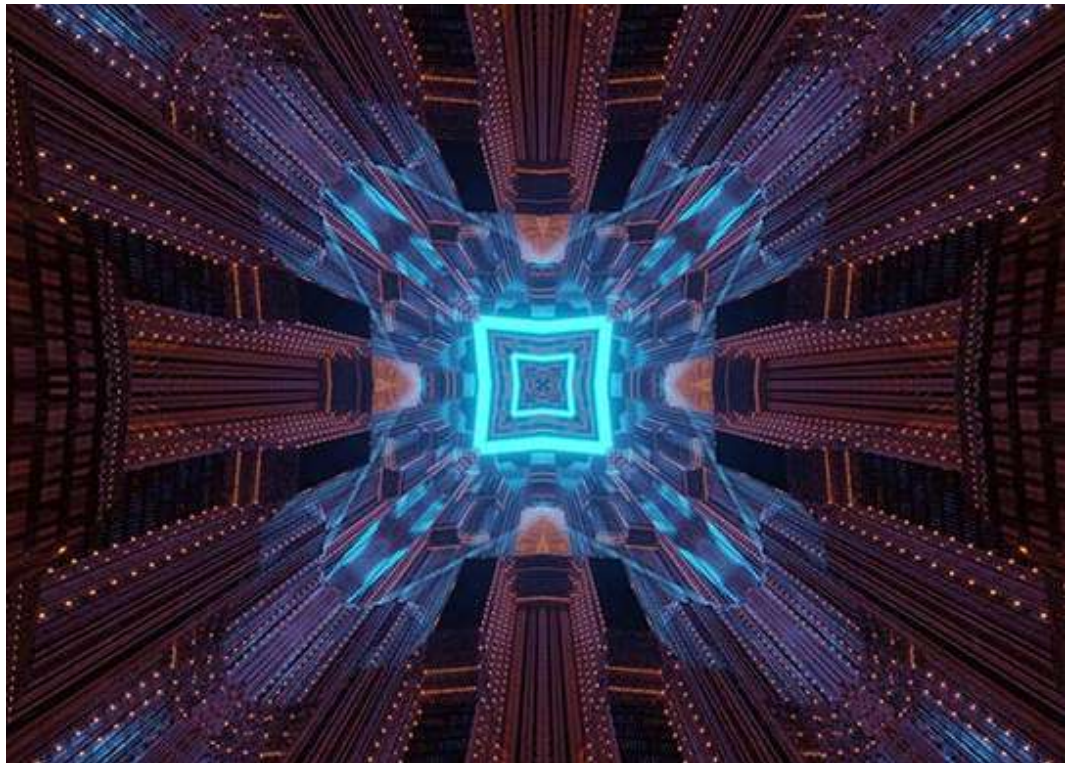

Ensuring high availability with open source databases

All of the perceived obstacles can be overcome but it requires planning and expertise.

While there are key business benefits from moving to open source, there are key considerations when looking to build an architecture for high availability/uptime. It requires expertise to design, configure, deploy and maintain highly available open source databases. There needs to be a level of expertise and commitment to ensuring databases continue to run smoothly and continuously.

Greater complexity requires certain skills and knowledge. For example, how to write application connectors, integrate multiple systems and align business requirements with the HA solution. There is also a need to understand open source databases, applications and infrastructure. Consider the possibility that outside expertise may be necessary to help manage the HA architecture.

An HA database system provides failover (preferably automatic) from a primary database node to redundant nodes within a cluster. An HA system provides quick recovery of all system components to minimise downtime. Some disruption may occur but will be minimal.

This is different from 'fault tolerance', which aims for zero downtime and data loss. Fault tolerance is much more expensive to implement because it requires dedicated infrastructure that completely mirrors the primary system. It also demands a lot of resources to maintain it.

To achieve an HA database solution, three key principles must be put into practice:

- **Single point of failure (SPOF)** – Eliminating any single point of failure in the database environment, including the physical or virtual hardware the database system relies on and which would cause it to fail.
- **Redundancy** – Ensuring sufficient redundancy of all components within the database environment and reliable crossover to these components in the event of failure.
- **Failure detection** – Monitoring the entire database environment for failures.

Expert consultants can work with fintech companies to build robust architectures, including providing on-demand support or full-time staff to maintain them, as well as supplying architectures - essentially blueprints – as a solid starting point.

HA architectures pull together tools and components to enhance and harden interoperability, ensuring applications always have access to the data they need.

An option for **small and start-up fintechs** is to use open source tools to achieve HA without an exorbitant price tag or the need for an overly complex environment. The HA infrastructure should be built within a single data centre on a local node. It is intended to keep the database available should the primary node go down, either from an automatic failover or planned switchover.

For **medium and large financial services businesses**, such as established banks and insurance companies, one route is for availability to be spread across data centres to add more layers of availability to the cluster. It is intended to keep the infrastructure available and data safe and consistent, should a problem arise in the primary data centre.

The infrastructure will typically feature a disaster recovery site and an external layer of strong, consistent, distributed key-value store that provides a reliable way to store data that needs to be accessed by a distributed system or cluster of machines. Should communication between the sites be lost, the node layer will act as a 'source of truth' and decide what replica will be promoted as a primary, thereby keeping the cluster healthy and the infrastructure highly available.

“High availability architectures pull together tools and components to enhance and harden interoperability, ensuring applications always have access to the data they need.”

For **enterprise and globally distributed businesses**, such as tier one and two international banks, the infrastructure might feature two disaster recovery sites, which ensure HA. Based on tightly coupled database clusters spread across data centres and availability zones, it offers an HA level of 99.999% when using synchronous streaming replication, having the same hardware configuration in all nodes, and a fast internode connection.

Managing the open source database infrastructure

After the migration, there is the need to ensure that the new environment, whether in the cloud, in a managed Database as a Service (DBaaS), in a hybrid situation or on-premise, stays agile, available and highly scalable as it grows.

Consider finding a vendor that can provide a comprehensive and cost-effective plan to support you post-migration. This needs to be someone who can be available around the clock to help ensure uptime, efficiently restore service, implement performance improvements and navigate the complexity of modern deployments using industry best practices. They should be experts when it comes to tuning the new database environment for scaling and performance.

Databases need to be optimally managed so it is easier for developers to deliver new applications and features, to keep pace with the disruption. You can't have your database holding you back. The new competitors typically have a much faster time-to-market than established financial services firms.

There is the need to monitor the health of the database infrastructure, explore new patterns in behaviour and improve the performance of the databases no matter where they are located.

Database performance monitoring can be done with customisable dashboards and real-time alerting. Critical performance issues can be detected faster, with users able to better understand the root cause of incidents and troubleshoot them more efficiently. Optimisation encapsulates zooming in and drilling down into database performance from node to single query levels, to allow in-depth troubleshooting and performance optimisation. It should be possible to create and configure database clusters no matter where the infrastructure is deployed.

> *The tools and expertise are available in the market to support the migration and that the business benefits can be relatively swiftly unlocked.*

Also important is the ability to back up critical data with zero downtime and minimal performance impact. Users should be able to schedule various types of back-ups (hot, incremental, physical) and restore databases up to a specific moment with Point-in-Time-Recovery. And on the security front, built-in advisors can run regular checks of the databases to identify and alert users of potential security threats, performance degradation, data loss and data corruption.

Conclusion

Ever more companies – within financial services and across many other sectors – are making the shift to an open source future. The financial services converts include fintechs, payment providers, exchanges and retail, corporate and universal banks.

They are finding that the tools and expertise are available in the market to support the migration and that the business benefits can be relatively swiftly unlocked. There is the ability to ensure that applications remain secure, compliant and highly available, allowing firms to deliver the frictionless experiences that their customers now expect, all while lowering infrastructure costs.

This allows firms to be competitive and compliant in ever more fluid markets and turbulent times, creating a technology base that allows resilience, flexibility and future growth.

About Percona

Percona is widely recognised as a world-class open source database software, support and services company for MySQL, MongoDB and PostgreSQL databases in on-premises and cloud environments. Percona is dedicated to helping make your databases and applications run better through a unique combination of expertise and open source software.

Your data is the foundation of your business, and Percona believes it can be simpler, faster and easier to manage. The company's services and software accommodate rapid growth and application development, help you develop and deploy agile solutions and improve your ability to keep up with customer needs. Percona has progressively built a suite of open source database solutions that empower our customers with greater flexibility.

Percona is on a mission to provide a unified approach to open source database software, support and services. At Percona, we make your database – and the applications that depend on it – better.

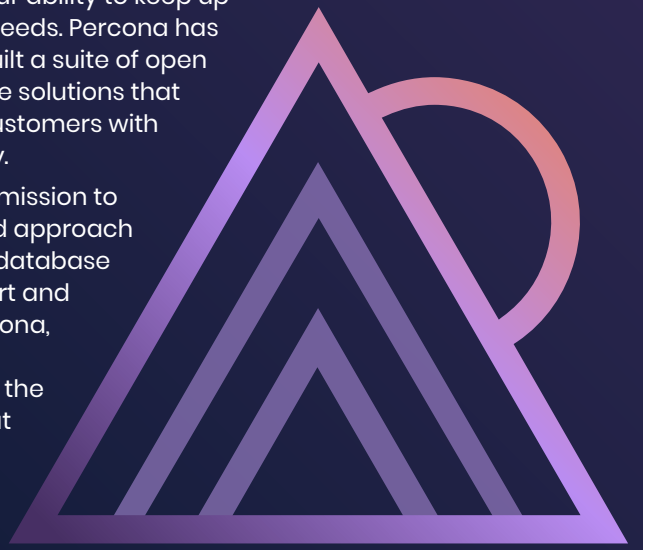

About
FINTECH FUTURES

Worldwide fintech news, intelligence & analysis

FF 40YEARS
Leading fintech media coverage

Celebrating **40 years of excellence**, FinTech Futures stands as the leading source of cutting-edge insights and resources for the global fintech sector. Our reputable platform serves industry professionals across various areas, including **FinTech**, **BankingTech**, **PayTech**, **RegTech**, **WealthTech** and **LendTech**.

Adopting a digital-first approach, we provide comprehensive coverage to keep readers informed and empowered to make knowledgeable decisions. Our extensive portfolio comprises Banking Technology magazine, daily newsletters, an intelligence library, **Banking Tech Awards**, **Banking Tech Awards USA**, **PayTech Awards** and the popular **What the Fintech? podcast** with over 1,000 listeners per episode.

Join our dedicated global community of **100k+ monthly visitors** and **25k newsletter subscribers** as we shape the future of banking, payments, and fintech together. Access a wealth of resources, including **reports**, **white papers**, **e-books**, **industry surveys**, **webinars** and more – all available for free. Our refreshed **website**, enhanced user experience and expanded digital channels offer an integrated platform for multimedia content and in-depth intelligence resources.

Explore advertising, brand awareness, thought leadership, and lead generation opportunities through our suite of products and services. Experience our commitment to digital excellence and adaptability in the ever-changing landscape. Trust FinTech Futures as your catalyst for change and growth in the banking, payments, and fintech industries, delivering well-researched, quality content to our global community.

Contact us

Sam Hutton
Head of Sales
sam.hutton@fintechfutures.com
+44 208 052 0434

Kate Stevenson
Business Development Manager
kate.stevenson@fintechfutures.com
+44 782 593 0099

Download our media kit here