

Unleashing Your Data to Deliver Transformational Business Outcomes



Couchbase



Red Hat

Data is the lifeblood of transformative business outcomes

Digital transformation starts with data. New and existing digital processes are only as good as the data that supports them. Access to accurate and timely information is essential for these digital processes. As organizations begin to create processes to look forward, not back, they require an underlying data layer that supports such efforts.

However, legacy data management tools and databases aren't a good match for the new demands from digital transformation and new processes. A Couchbase study found that 72% of data architects in larger organizations said legacy databases limit an organization's ability to complete digital transformation projects. That is not to disparage the older products; they are well suited to the jobs they were designed for. But they were not designed to meet the new demands of digital business.

A big problem is that, with legacy systems, too much data remains unavailable, and that can have an adverse impact on the organization. In the Couchbase study, 86% of data architects said factors that have prevented them from pursuing a new digital service or other transformation project include reliance on legacy technology, the complexity of implementing such technologies, and a lack of resources and skills.

Digital transformation requires more than merely the legacy application data. New data sources are essential for transformative apps. Critical information found in social and news media and various new repositories have unstructured and semi-structured data that provides much more insight than that found in organizations' legacy data inventory. Ensuring that this information can be easily incorporated into a comprehensive data layer is essential to supporting transformative outcomes.

For both data architects and IT teams in general, the ability to leverage all types of data effectively is necessary to deliver new and updated apps. Without the ability to use all of the organization's data, there will be blind spots. Without all the information, management may make the wrong decisions as incomplete data leads them astray. Legacy databases are poorly positioned to deliver the comprehensive data layer necessary to prevent uninformed decision-making. Some key limitations include:

- **Data silos.** A good example of the data silo problem is the difficulty of building an integrated, 360-degree customer perspective. With legacy data siloed in different systems, not only is integration complex, but delivering this perspective in real or near real time is impossible using legacy technology.
- **Older relational technologies.** Such technologies are hard to use because their strict schemas and data models are difficult to modify and integrate. Instead, data must become truly fluid so that it can be easily reused in many different apps or services. It is time to eliminate older approaches such as long-running batch processes and repetitive manual ETL methods to support data movement.

The time is now for a new approach designed to support a fluid data layer from the start.





The next-generation data layer will support transformative outcomes

It is abundantly clear that legacy database solutions are not up to the task of enabling digital transformation initiatives. Moving into the future, data architects, DevOps practitioners and other technical professionals will require new capabilities to meet the organization's demand for new and innovative digital systems.

Fluidity is among the most important capabilities of next-generation data layer solutions. Data silos are no longer allowed, and data must be independent of any application. True data fluidity is enabled and enhanced using a cloud-native platform that supports mobility, not only among different apps or software stacks, but with different infrastructure instances, regardless of physical location. The historical design and development limitations of finding, moving and joining data are eliminated using this modern approach, enabling DevOps to focus more on delivering solutions and less on being data “janitors.”

The second fundamental difference between legacy and modern data layers is the inclusion of numerous new data types—and including both unstructured and semi-structured data is essential. Four Cornerstone, a data center services provider, estimates that only 20% of corporate data is now contained in structured data sets. Unstructured and semi-structured data sets are not only treasure troves of critical information but often provide insight on aspects of the business operation that is not gathered anywhere else. For example, video usage is increasing exponentially and may provide the key data for transformational new apps about customer behavior.

Another key capability of the modern data layer is the ability to incorporate incoming data at scale. Current business processes such as marketing and engineering are creating huge amounts of data, and as IoT and the billions of sensors supporting it come online, the rate of incoming data will grow in geometric fashion. A key element of incorporating data at scale is the ability to utilize data streamed from the edge.

The last point to note is that the new data layer begins life as a cloud-centric solution and includes native integration with public cloud services and hybrid cloud deployments. The data layer simplifies dealing with issues such as data sovereignty without encumbering developers with those details. It also makes it simple to move and shift workloads dynamically depending on changing requirements.



NoSQL technology delivers this functionality as the data layer for cloud native applications. The technology provides support for a dynamic schema that delivers the agility needed to support the regular modifications and upgrades that are common with modern transformative applications. NoSQL also makes inclusion of new data types simple during the DevOps application lifecycle. This will become much more important as new data sets emerge and need to be incorporated into an increasing number of applications.

In addition, NoSQL provides performance across a range of parameters that ensure transformative apps deliver necessary service levels. To begin with, NoSQL delivers the scalability that modern, highly dynamic workloads require. The ability to scale both up and down in real time is essential to maintaining fluidity and delivering efficiency. NoSQL also provides far greater intrinsic performance compared with legacy products. Again, as the amount of data being used continues to grow, ensuring that data layer platforms can provide the necessary service levels is key. One of the most important benefits of best in-class products like Couchbase is functionality that makes NoSQL a good match for database administrators' current skills, leveraging their expertise and knowledge.

One of the most important design points of NoSQL is that it is cloud-native. Rather than try to retrofit cloud capabilities or functionality in products that were conceived before cloud became the way IT is done, some NoSQL products were designed to support cloud from the outset.

Red Hat OpenShift and Couchbase deliver an optimal data layer foundation for transformative solutions

The joint Red Hat-Couchbase solution delivers native integration of Couchbase with Red Hat OpenShift. This gives organizations an optimal solution for cloud-native apps that require database support. IT can transparently run hybrid and multicloud infrastructure without impacting the data layer. Not only is this solution easy to deploy, but modifications and additions to apps are greatly simplified as well. Couchbase integrates memory-first performance with caching and a persistent document store that can both sync and manage data on edge devices. To summarize, Couchbase provides the data layer management and containers, and Red Hat OpenShift manages the container platform across cloud and on-premise infrastructure.



The solution checks all the boxes necessary to provide organizations with an agile, comprehensive and high-performance modern data layer:

- Container-friendly, with native support for microservices
- Ability to use structured, unstructured and semi-structured data
- Cloud-native design
- Full and consistent capabilities from cloud to edge
- Performance at scale for both data intake and manipulation
- Enterprise-class security, scalability and reliability
- Reduced complexity, with seamless native integration of Couchbase and Red Hat OpenShift

The benefits of the Red Hat-Couchbase collaboration are many, but one of the most striking is that the solution greatly simplifies daily IT operations and reduces the demands for scarce staff resources. IT and DevOps teams can provide more transformational apps and services, more quickly. The Couchbase Autonomous Operator simplifies using Kubernetes platforms to create or modernize applications so that they have lower resource requirements. Also, increased levels of automation reduce the need for manual intervention.

Additionally, it is possible to manage more than one Couchbase deployment more simply and with less interaction. The solution reduces operational complexity, and simpler operations always reduce the need for human resources. Finally, with full cloud portability, there is no lock-in to a specific cloud service, which enables simpler workload migration and eliminates the need to modify a workload to move between clouds.

This solution provides a great deal of value in the retail industry, for example, where numerous problems must be solved. For one, structured data is often stored in many different apps and in inventory, order entry, website activity and other tracking tools. In addition, emails with customer service agents have insight on specific transactions and social media may have critical perceptual data, but for many brands, this information cannot be viewed or analyzed holistically. With the Couchbase and Red Hat solution, it is now possible to integrate all this disparate data into a usable form.



Extended architectures with Couchbase Capella

With Couchbase Capella, organizations can extend their cloud architectures to broaden and diversify their deployment options and footprint. Couchbase Capella is the fully managed deployment of Couchbase, providing all the power of Couchbase Server as a database-as-a-service (DBaaS) available on AWS, Azure and GCP. It provides the ability to replicate data between database clusters leveraging Cross Data Center Replication (XDCR), a feature of Couchbase available in both self-managed deployments as well as the Capella DBaaS. This capability enables extended architectures where self-managed deployments of Couchbase running on Red Hat OpenShift can replicate data to Couchbase Capella.

This extended architecture expands the options for hybrid and multi-cloud deployments, allowing organizations to meet any regional coverage or availability requirements. It also enables edge computing architectures, where Red Hat OpenShift deployments running Couchbase at the edge can replicate data to Capella in the cloud.

Key takeaways

For any organization, the ability to harness all its data and use it effectively is the single most important aspect of successful digital transformation. The new generation of truly transformative applications depends on far more and varied data than legacy apps. IT management must ensure that the business has the tools it needs to compete. The modern data layer that supports transformative applications cannot be supported with legacy technologies. Deploying new platforms such as NoSQL and a cloud-native, enterprise-class application platform that provides support for containers and microservices is the way forward. Red Hat and Couchbase are partnering to make it easier to deploy a seamless and simplified platform for deploying a modern, enterprise-class data layer. For more information on this game-changing solution, [click here](#).





Modern customer experiences need a flexible database platform that can power applications spanning from cloud to edge and everything in between. Couchbase's mission is to simplify how developers and architects develop, deploy and consume modern applications wherever they are. We have reimaged the database with our fast, flexible and affordable cloud database platform Capella, allowing organizations to quickly build applications that deliver premium experiences to their customers—all with best-in-class price performance. More than 30% of the Fortune 100 trust Couchbase to power their modern applications.

For more information, visit www.couchbase.com and follow us on Twitter.

© 2023 Couchbase. All rights reserved.