



CLOUDERA


FEBRUARY 2026 | EXECUTIVE BRIEF

From Cloud Native On Premises to Hybrid

How Leading Banks Increase Agility and Efficiency by Bringing
Cloud Capabilities to On-Premises Data Estates With Cloudera

Joseph Montenaro, Senior Solutions Engineer, Cloudera

Matthew Michaelides, Principal, Value Management, Cloudera



Executive Summary

As enterprises adopt cloud-native capabilities on premises, they are positioning themselves at the forefront of modern data architecture. Yet, organizations who simply modernize on premises alone miss a significant opportunity to optimize infrastructure cost and drive accelerated time to value. Hybrid architectures enable customers to balance performance, compliance, and cost while enabling cutting-edge innovation with AI agents, next-generation AI/ML frameworks, and more.

Cloudera's banking and financial services customers are shifting rapidly toward hybrid data and AI operating models. As firms seek to realize the potential of generative AI (GenAI) and agentic AI capabilities, the need for hybrid agility has never been greater. According to Gartner, [data synchronization](#) across the hybrid cloud environment is the most urgent GenAI priority.

This whitepaper explores what hybrid means for organizations that are modernizing their on-premises data and AI estates, why hybrid matters to those organizations, and how Cloudera provides a practical path to hybrid modernization.



What Hybrid Means for Enterprise Banks

For enterprise banks, hybrid means bringing the best parts of the cloud, such as cost efficiency, scalability, and agility, to the data wherever it resides. To this end, Cloudera enables a significant and growing number of hybrid use cases, with six in particular that customers prioritize most frequently (see **Figure 1**).

The customer goals driving hybrid adoption vary, but there are patterns.

First, **infrastructure cost optimization** is a significant value driver for many hybrid use cases. Historically, large banks have struggled to maintain large on-premises infrastructure estates, often hamstrung by utilization of compute resources as low as 10 to 20%

while still struggling to meet service-level agreements (SLAs) at peak times. While on-premises modernization efforts have helped to mitigate this, hybrid architectures can be the next step to efficient compute resource allocation. By shifting workloads flexibly and dynamically to the cloud (known as cloud bursting), banks can significantly reduce infrastructure costs. For example, a recent internal Cloudera study found that enterprise banks operating real-time fraud detection models on Cloudera can save **31% by adopting a hybrid approach instead of using only the public cloud**. Customers can run predictable, consistent workloads like model inference and monitoring on premises, while more intermittent pipeline elements, including training and feature engineering, can run in the public cloud.

Top Use Cases for Hybrid in Banking and Financial Services

Use Case	Hybrid AI/ML	Unified Data Access	Efficient Data Pipelines	Cloud-Optimized DR	Future-Proof Flexibility	Regulation-Ready Data Control
Description	Build in cloud, run on premises, securely	Scale BI in the cloud when serving data anywhere (on premises, clouds)	Run “always-on” compute on premises, while bursting ad-hoc and “bursty” compute to the cloud	Cut DR costs while meeting uptime goals	Move workloads freely between on premises and cloud	Combine InfoSec and infra savings priorities with data and AI anywhere
Value	<p>Infra savings: Limit infra costs by running test/dev workloads in cloud and “always-on” workload on premises</p> <p>Ops savings: Reduce ops burdens when using cloud resources</p> <p>Time to value: Speed up dev/test cycles with cloud-native scaling</p> <p>Risk reduction: Keep sensitive data on premises for AI</p>	<p>Infra savings: Mitigate duplicate cloud storage costs with federated data access</p> <p>Time to value: Simplify and accelerate end user experience through unified AI/BI data access layer across on-premises and cloud data stores</p> <p>Risk reduction: Keep sensitive data on premises and serve from the cloud</p>	<p>Infra savings: Reduce infra costs by optimizing data pipelines by deployment style</p>	<p>Infra savings: Reduce infra costs by only paying for DR environment when active</p> <p>Risk reduction: Reduce downtime risk with different environments for prod vs. DR</p> <p>Time to value: Improve business agility by serving ad-hoc and/or bursty workloads with cloud-native scalability</p>	<p>Risk reduction: Improve ability to move workloads across deployment styles due to strategic, regulatory, or other shifts</p> <p>Vendor savings: Mitigate vendor lock-in concerns</p>	<p>Risk reduction: Meet stringent data security and data residency requirements</p> <p>Infra savings: Eliminate cost of idle compute on premises for “bursty” data serving workloads</p> <p>Time to value: Improve business agility by serving data in cloud</p>

Figure 1: Top Use Cases for Hybrid in Banking and Financial Services

Second, hybrid architectures **support broader enterprise goals around agility and flexibility.**


By investing in a platform built on a foundation of open-source technologies, standards, and formats with feature parity across cloud and on-premises environments, companies can reduce risk across several dimensions:

- Security risks from siloed data that lacks a consistent governance framework
- Vendor or partner risks from any potential future changes to their technology strategy and/or downstream resilience impacts (for example, hyperscaler outages)
- Financial risks associated with large capital expenditure (CapEx) investments, such as hardware

In this way, investing in a truly hybrid platform is not a vendor swap; it's a purchase of leverage through flexibility and openness, mitigating punitive egress fees, vendor lock-in, and the costs of cross-environment coordination.

Above all else, Cloudera's most innovative customers stretch the bounds of a single use case, leveraging the flexibility of a unified hybrid cloud to address many use cases together. **Hybrid leaders often design for one use case but find that the flexibility they built into their platform enables them to capitalize on additional capabilities seamlessly.**

One Cloudera customer, a large Latin American bank, improved availability and optimized their on-premises infrastructure through this path. While they initially sought to build a resilient, cloud-optimized disaster recovery (DR) solution with Cloudera's hybrid data and AI platform, the bank ultimately found that they could expand the use of their cloud DR environment (including prepositioned cloud data synchronized via dual ingestion) to drive "burst-to-cloud" capabilities during peak times, reducing the need to expand its on-premises footprint and driving significant infrastructure cost reduction.



Defining Hybrid: The Three Keystones of Hybrid Architectures

Many organizations operate both on premises and in multiple public clouds, but this alone does not realize the full potential of a hybrid architecture. Instead, Cloudera's most successful hybrid customers implement a series of parallel and consistent capabilities across environments such that they collapse into a single, unified, hybrid and multi-cloud architecture.

This unified experience simplifies the data and AI technology landscape for all stakeholders—from data management professionals to platform managers and AI practitioners—across all deployments. The hybrid architecture is built on three keystones: the unified runtime, the hybrid control plane, and the hybrid experience (see **Figure 2**).

Keystone 1: Unified Runtime: Move your Data and Workloads Anywhere

A unified runtime, **a single codebase across on premises and cloud estates**, is the foundation of a hybrid architecture. The unified runtime dramatically reduces the refactoring and replatforming costs associated with moving between on-premises environments and public clouds, eliminating the need for replatforming, a time-consuming project that can cost **\$200–400K per application** (typically, \$10–50 million overall for larger enterprises) in refactoring alone.

Beyond simplifying migrations, a unified runtime ensures that features, updates, and workflows are consistent across infrastructures, meaning they work and feel the same across all environments, improving workload portability and user experience. This capability enables companies to run workloads where and when it makes sense, depending on potentially shifting goals, including but not limited to: data locality, infrastructure availability, and cost optimization.

The Three Keystones of Hybrid Architectures




Figure 2: How The Three Keystones Interoperate to Power Hybrid Architectures

Keystone 2: Hybrid Control Plane: View and Monitor your Data and Workloads Everywhere

Next is the hybrid control plane, a consolidated view of the entire data estate, providing centralized configuration and security management across hybrid and on-premises environments. It enables customers to monitor, visualize, and organize workloads and data assets no matter where they reside, dramatically simplifying efforts to audit and manage security policies across multiple environments and their constituent services. More tactically, a hybrid control plane is critical for maintaining unified data governance and data quality for AI/ML workloads.

Large enterprises have already taken significant steps to build a unified enterprise architecture across public cloud and on-premises IT operations workflows, delivering as much as 50% reduction in cybersecurity staff costs. The hybrid control plane extends these capabilities, taking advantage of the same operational logic to drive cost reduction and innovation.

Keystone 3: Hybrid Experience: Action your Data, Workloads, and Apps Anywhere

The culmination of the journey is the hybrid experience, providing data management professionals, platform teams, and practitioners alike with a singular experience that supports managing, governing, and extracting value from data regardless of where it lives, across three capabilities tied to their core professional activities:

- **Hybrid data management:** Discover, map, and move your data anywhere
- **Hybrid platform experience:** Provision compute & storage with resources anywhere - permanently and dynamically
- **Hybrid practitioner experience:** Query, visualize, and operationalize data and AI anywhere

Building out your organization's hybrid experience is the pinnacle of the hybrid journey, encompassing unique capabilities required to power hybrid activities performed by all critical stakeholders, including data security and governance professionals, platform managers, and data and AI practitioners (see **Figure 3**). With these capabilities in place, your organization can undertake each aspect of the data management lifecycle *anywhere*.

The Hybrid Experience Capabilities Map




Figure 3: The Hybrid Experience Capabilities Map

Getting Started On Your Hybrid Journey

Building a hybrid architecture is a significant undertaking for enterprise banks and financial services institutions (FSIs) that typically figures into a broader set of strategic initiatives including on-premises modernization, enterprise AI, and the broader data strategy (see **Figure 4**).

Considering these interdependencies, the emerging model for integrating these capabilities relies on parallel development of on-premises and cloud capabilities through iterative transformation. In executing in this fashion, companies make incremental measurable progress toward sophisticated hybrid use cases while also acting on long-term strategic priorities. The key parallel and overlapping phases of this coordinated, iterative transformation program include:

Phase 0: Set Your Hybrid Vision

For most enterprises, their hybrid architecture is not the result of strategic decision, but rather a default. Many organizations end up supporting disjointed on-premises and cloud deployments by way of myopic,

reactive architectural decision making focused on solving specific technology challenges, as opposed to a well-considered, strategic approach. This accidental hybrid architecture often limits the potential of organizational data, which is difficult to discover, access, integrate, and leverage for analytics and AI.

Companies that implement hybrid as a conscious, strategic decision build on two key components:

- First, they bet on specific, strategic hybrid use cases. They understand that to get value from a hybrid model, they must ensure that the capabilities they deliver align to broader business priorities.
- Second, they invest in innovative enterprise capabilities that they believe can drive long-term value delivery over and above specific use cases.

As an example, a global bank is working with Cloudera to define a programmatic approach to implementing their hybrid architecture. While they are working to support a list of their top 10 most critical hybrid use cases, they are also ensuring that architectural decisions preserve flexibility and interoperability, so they are equipped to implement additional hybrid use cases years into the future.

The Model for Delivering Hybrid Data and AI/ML Capabilities




Figure 4: The Model for Delivering Hybrid Data and AI/ML Capabilities

Phase 1: Hybrid Readiness On Premises

Once a clear hybrid vision is in place, companies must modernize their on-premises estate to drive hybrid readiness. This process involves updating existing data centers, servers, and legacy systems to improve efficiency, performance, and self-service, and ensure compatibility between on-premises environments and the public cloud, reducing on-premises technical debt.

The most critical modernization efforts include:

- Separating compute and storage to drive independent resource scaling
- Densifying hardware by leveraging modern storage solutions like Cloudera Object Store, powered by Apache Ozone
- Adopting container-based compute
- Shifting to open architectures that reduce data and compute silos

This effort is often lengthy, complex, and requires a significant investment in time and resources, but it is critical for seamless integration with cloud services.

Some organizations mistakenly hold back on developing hybrid capabilities until completion of their on-premises modernization efforts. This approach leaves significant value on the table by delaying innovation and realization of infrastructure and operational efficiencies with cloud-native capabilities. Additionally, lengthy security review processes, which are typical in most highly regulated organizations, can delay hybrid cloud efforts unless they begin well in advance of executing on-premises modernization efforts.



Phase 2: Hybrid Cloud Acceleration

Hybrid cloud acceleration enables self-service and agile platform operations, both in the data center and cloud.

From an execution perspective, it involves deploying the most foundational hybrid capabilities to build an integrated view of a hybrid ecosystem as well as accelerating migration efforts from on-premises to cloud via a unified runtime.

While they are developing the foundational capabilities, companies must refine their initial goals through rigorous use case-specific prioritization exercises. After identifying potential use cases with specific tagged business priorities, they should size the value of those opportunities, and estimate the cost and timeline required to achieve them. This will enable them to prioritize and justify their hybrid roadmap.

While these prioritization analyses vary by company, we typically see two primary hybrid use case archetypes (see **Table 1**):

- **“Big win” hybrid use cases (e.g., hybrid AI/ML, unified data access, future-proof flexibility):**
These are the use cases that drive an enterprise strategy toward a hybrid architecture and are most likely the motivators behind the overall hybrid vision (see Phase 0). By prioritizing these use cases, companies can justify their hybrid capability investments.
- **“Quick win” hybrid use cases (e.g., cloud-optimized DR, efficient data pipelines):**
These use cases typically have a shorter timeline to implementation, demonstrate value from hybrid capability investment, and make incremental progress toward more strategic wins later.

Cloud-optimized DR and other related data back-up or offload use cases are a common start to the hybrid journey. First, these initiatives are usually driven by IT, requiring less stakeholder management to execute. Second, they typically offer immediate reprieve for customers contending with data center capacity issues. And finally, they enable data to be staged and accessible across both the public cloud and on-premises data centers, providing an environment for ad-hoc analytics, R&D, and other use cases that run well in the public cloud.

Typical Hybrid Use Case Prioritization (Note: Varies by Company Context)

	Hybrid AI/ML	Unified Data Access	Efficient Data Pipelines	Cloud-Optimized DR	Future-Proof Flexibility	Regulation-Ready Data Control
Business Value Impact: Growth, business P&L, risk reduction	High Impact Single net new AI/ML use case can drive significant multi-million dollar value realization	High Impact Critical for democratizing data across the organization to pursue business goals	Low Impact Use case focused on 'tech value' TCO optimization	Moderate Impact Improved resilience via environment diversification	High Impact Avoid (1) vendor lock-in & (2) being locked out of other CSP-specific or on-prem only AI/ML capabilities (Note: ROI is not immediate)	High Impact Typically drives ability to support new use cases not otherwise achievable
Tech Value Impact: TCO reduction	High Impact Massive AI/ML compute burden with inference reduced	High Impact Significant opportunity to de-duplicate data & ETL costs	High Impact Significant opportunity to properly allocate always-on & bursty compute	High Impact Large infra savings; on-prem DR typically up to 40% of total environment & low utilization	High Impact Future ability to optimize pipelines, avoid risks (regulatory, vendor lock-in, etc.) (Note: ROI is not immediate)	Moderate Impact Can drive ability to offload on-premises infra
Capabilities Cost: Upfront cost/effort to develop capabilities	High Cost Significant GPU cost/investment to materialize capability	High Cost Significant custom effort to unify the data estate, with unified governance	Moderate Cost Control Plane configuration required to direct workloads & burst accordingly	Moderate Cost Largely out-of-the-box capabilities; however, must adapt to custom DR req's	High Cost Significant complexity in establishing end-to-end hybrid	High Cost Significant complexity in establishing custom unified enterprise governance layer
Use Case Cost: Incremental cost/effort to deploy a new hybrid app/use case	Low Cost Repeatable pattern design; however, staffing & execution costs are significant	Moderate Cost Key dependency is validating SLA performance & optimizing hybrid/multi-cloud networking	Moderate Cost Key dependency is validating SLA performance & optimizing hybrid/multi-cloud networking	Low Cost Repeatable pattern design	Low Cost Repeatable pattern design	High Cost Significant, highly custom pipeline design efforts required to meet custom regulatory requirements

Table 1: Typical Hybrid Use Case Prioritization (Note: Varies by Company Context)

Key

- Favorable
- Neutral
- Least Favorable

This approach—starting with “quick wins” before moving toward “big win” use cases—is how many companies begin their hybrid journey. For example, as part of a broad strategy to optimize their infrastructure, a large Brazilian bank has partnered with Cloudera to build out foundational hybrid capabilities and prioritize both (1) high-level use cases, including cloud-optimized DR and unified data access, as well as (2) specific priority applications for these new patterns based on impact potential and key requirements. For this bank, beginning their hybrid journey with cloud-optimized DR is the first step toward a broader array of hybrid use cases.

Phase 3: Hybrid Use Case Rollout and Continuous Improvement

The final step is to work with application teams to roll out top-priority hybrid use cases, including specific hybrid AIOps and MLOps applications. This step is both the culmination of the hybrid journey and the start of long-term, production-scale value delivery from hybrid.

Typically, organizations embrace a “pilot & scale” approach to launching these use cases. Depending on the roadmap of prioritized hybrid use cases, they build the hybrid experience capabilities needed. Once the pilots demonstrate business value, they can be scaled as individual applications and more broadly as organization-wide capabilities.

Even as companies make progress launching hybrid models, ongoing continuous improvement efforts are critical. Hybrid architectures continue to evolve, and the use cases grow. Bleeding-edge companies are monitoring innovation and preparing to adapt, backed by existing, highly flexible capabilities.



Conclusion: Power the Hybrid Enterprise

Several years ago, one of the largest global oil and gas companies began its journey from the data center to the public cloud on Cloudera, with the belief that a unified, flexible enterprise data platform would serve an immediate high-priority use case, as well as their broader enterprise goals. Today, that hybrid use case and dozens of others power \$400M+ per year in internally validated value delivery by reducing the costs of data egress, centralizing query logic, and reducing compute redundancy over hundreds of data pipelines across geographies, on-premises data centers, and public clouds. However, their current hybrid platform is so much more—encompassing an industry-leading data ingestion and integration workflow, unified security and governance, and cutting-edge hybrid AIOps and MLOps. At each stage of their journey, they chose to invest in the power of their hybrid ecosystem and drive usability across their data estate—and it is paying off.


This story showcases how companies win with a hybrid approach. They start with specific, high-priority use cases and a clear enterprise vision, and they do so while building maximum flexibility into their architecture at the outset, because they do not want infrastructure decisions they make today to limit innovation in the future. These companies understand and commit to their goals, while acting transformationally to improve data access and coordination across deployments.

Hybrid architecture with Cloudera is not a compromise—it is the strategy for resilience, agility, innovation, and openness. By enabling workload portability, reducing the risk of vendor lock-in, and accelerating AI/ML, Cloudera’s hybrid platform ensures the right workload runs in the right place at the right time.

About Cloudera

Cloudera is the only data and AI platform company that large organizations trust to bring AI to their data anywhere it lives. Unlike other providers, Cloudera delivers a consistent cloud experience that converges public clouds, data centers, and the edge, leveraging a proven open-source foundation. As the pioneer in big data, Cloudera empowers businesses to apply AI and assert control over 100% of their data, in all forms, delivering unified security, governance, and real-time predictive insights. The world's largest organizations across all industries rely on Cloudera to transform decision-making and ultimately boost bottom lines, safeguard against threats, and save lives.

To learn more, visit [Cloudera.com](https://cloudera.com) and follow us on [LinkedIn](#) and [X](#).



CLouDERA

Cloudera, Inc. | 6220 America Center Dr, 5th Fl, San Jose, CA 95002 USA | cloudera.com