



WHITE PAPER

BY

Autumn Salama,
Vice President of Cloud
Operations, DataBank

The Reality of Cloud Repatriation

Why 86% of CIOs Are Rethinking Cloud Strategy – and How to Avoid Repatriation Pitfalls

Cloud repatriation has become increasingly common as organizations confront unpredictable costs, compliance complexity, and vendor lock-in. Yet, moving workloads back from the cloud can introduce substantial obstacles that may equal or exceed the original problems.

Strategic infrastructure planning and proactive decision-making prove far more effective than attempting to unwind poor cloud decisions after the fact.



DATABANK
Data Center Evolved™

Table of Contents

- Executive Summary** 1
- Introduction: The Next Phase of Cloud Maturity** 2
 - Inside this White Paper. 3
- Part 1: Why Cloud Repatriation is Happening** 4
 - Cost Overruns and Hidden Expenses 4
 - Security and Compliance Complexity 5
 - Support Limitations and Visibility Gaps 6
 - When Promise Meets Practice 6
 - Use Case: CarePoint Analytics: Part 1: The Wake-Up Call 7
- Part 2: The Repatriation Process: Challenges, Complications, and Considerations** 8
 - Understanding True Migration Costs..... 8
 - Capacity Planning and Performance Management 9
 - The Complexity Continues..... 10
 - Use Case: CarePoint Analytics: Part 2: The Complexity of Moving Back 10
- Part 3: Prevention and Strategic Planning** 11
 - Maintain Vendor Agnosticism from the Start..... 11
 - Choose Flexible Infrastructure Partners 12
 - Establish Clear Decision Criteria 13
 - Plan for Continuous Evolution 13
 - Use Case: CarePoint Analytics: Part 3: A Different Path Forward 14
- Conclusion: Strategic Infrastructure for an Evolving Landscape**..... 15
- About DataBank** 16

Executive Summary

Cloud repatriation – the process of moving workloads from public cloud back to private or hosted environments – has become increasingly common as organizations face unpredictable costs, complex compliance requirements, inadequate support, and other challenges. While it's true that public cloud offers significant benefits, many enterprises discover that certain workloads perform better and cost less in alternative infrastructure models.

However, repatriation itself presents substantial obstacles including data egress fees, proprietary service dependencies, and vendor lock-in that often require significant application modifications. Understanding these challenges is essential for making informed, proactive infrastructure decisions.

The most effective strategy is prevention: partner with infrastructure providers offering multiple deployment options while maintaining vendor agnosticism from the outset. This approach enables organizations to adapt as business requirements evolve without facing costly repatriation challenges or being trapped by a single provider's ecosystem.



Introduction: The Next Phase of Cloud Maturity

In the past decade, a common, often prevailing, view in enterprise IT was that the public cloud would completely displace on-premises infrastructure. Organizations rushed to embrace cloud-first and cloud-only strategies, drawn by promises of unlimited scalability, cost efficiency, and operational simplicity.

Not long after this migration wave began, cracks started to appear in the narrative. While cloud spending continues to surge – Gartner forecasts global cloud investments will reach approximately \$723 billion by the end of 2025 – a growing number of organizations are now re-evaluating their cloud strategies and making strategic decisions to move certain workloads back.¹

The numbers tell a compelling story:

- According to Flexera's 2025 State of the Cloud Report, roughly 21% of workloads and data have already been repatriated back to private clouds or on-premises environments.²
- More significantly, IDC research found that approximately 80% of respondents expected to repatriate some compute or storage workloads in the next 12 months.³
- Perhaps most striking, a Barclays survey revealed that 86% of enterprise CIOs planned to move at least some public cloud workloads back to private cloud or on-premises infrastructure – the highest percentage on record.⁴

It's important to note that any repatriation trend is not a complete rejection of cloud computing. Organizations aren't abandoning the cloud entirely. They're refining their strategies based on hard-earned experience. The real shift is from blanket "cloud-only" thinking to strategic workload placement – putting the right applications in the right environment based on actual business requirements rather than industry hype.

¹ Gartner, [Gartner Forecasts Worldwide Public Cloud End-User Spending to Total \\$723 Billion in 2025](#), November 19, 2024.

² Flexera, [The Flexera 2025 State of the Cloud Report](#), March 19, 2025.

³ IDC, [Assessing the Scale of Workload Repatriation](#), June 2024.

⁴ Barclays, 2H24 CIO Survey: 2024 Outlook Sustained, 2025 Uplift, 2024.

» Inside This White Paper

This white paper examines cloud repatriation from three critical perspectives:

Part 1: Reasons why enterprises are moving workloads back, such as cost overruns, security complexities, and support limitations driving infrastructure reconsideration.

Part 2: Hidden repatriation challenges including data egress fees, vendor lock-in, and required architectural changes.

Part 3: Prevention strategies through flexible infrastructure planning, vendor agnosticism, and strategic workload placement.

The Goal: Help IT leaders make informed infrastructure decisions that maintain strategic flexibility as business needs evolve.



PART 1:

Why Cloud Repatriation is Happening

The shift toward cloud repatriation isn't driven by a single pain point but by a convergence of operational and financial challenges that emerge as organizations scale their cloud deployments.

What begins as an attractive proposition – flexible infrastructure, minimal upfront investment, and rapid deployment – often evolves into a more complex reality. As workloads grow and business requirements change, organizations frequently discover that public cloud economics and operational models don't always align with their long-term needs.

The following factors represent the most common drivers pushing enterprises to reconsider their cloud strategies.

COST OVERRUNS AND HIDDEN EXPENSES

Cloud providers market their services with straightforward pricing calculators that promise transparency and predictability. However, many organizations quickly discover that actual cloud costs far exceed initial projections. The challenge isn't just about compute and storage; it's the accumulation of variable charges that multiply as infrastructure scales.

Autoscaling is a prime example. While the ability to automatically add capacity during demand spikes may initially seem like pure efficiency, it can generate unexpected costs when not carefully configured. In many cases, charges for load balancer requests, monitoring and observability costs, data transfer between availability zones, and countless other micro-charges add up quickly. For example, what once appeared as a \$10,000 monthly infrastructure budget can easily balloon to \$15,000 or \$20,000 once these variables come into play.

The unpredictability becomes problematic for finance teams trying to forecast IT spending. Unlike fixed monthly costs for colocation or private cloud infrastructure, public cloud bills fluctuate based on usage patterns that may be difficult to anticipate or control. Organizations find themselves constantly adjusting budgets and implementing cost controls that can actually constrain business operations.

For companies operating at scale, these variable costs compound dramatically. What works economically for a startup or pilot project often becomes financially unsustainable at enterprise scale. Many organizations reach an inflection point where the total cost of ownership for public cloud exceeds what they would pay for dedicated infrastructure – sometimes by significant margins.

SECURITY AND COMPLIANCE COMPLEXITY

Achieving compliance in public cloud environments is possible, but it requires substantially more effort, tools, and ongoing management than many organizations first anticipate.

For example, the shared responsibility model sounds straightforward in theory: cloud providers secure the infrastructure, and customers secure their applications and data. Yet, in practice, this division creates complexity and gaps that organizations must actively address. The lines of responsibility aren't always clear, and misunderstandings about who controls what can leave organizations vulnerable or non-compliant until they invest significant resources into understanding and implementing the right controls.

Meeting regulatory requirements like HIPAA, PCI DSS, FedRAMP, or GDPR in public cloud often requires deploying multiple third-party security tools, implementing extensive logging and monitoring, and maintaining detailed documentation of security controls. Unlike private or hosted environments where infrastructure-level security controls are more directly accessible, public cloud customers must layer security solutions on top of the provider's baseline offerings.

Companies may also discover they inherit only limited controls from cloud providers. While providers maintain certifications and secure their underlying infrastructure, customers bear responsibility for configuration, access management, encryption, and data handling. This means security teams must become experts in cloud-specific security practices while managing tools and policies across potentially multiple cloud providers.

WHY ORGANIZATIONS ARE RECONSIDERING CLOUD

- **Cost unpredictability:** Variable charges for autoscaling, data transfer, and micro-services can cause monthly bills to balloon by 50-100% beyond initial projections.
- **Compliance complexity:** Shared responsibility models require extensive third-party tools, specialized staff, and ongoing management that often exceeds the cost of managing secure on-premises environments.
- **Support limitations:** Mid-market customers frequently lack access to responsive, personalized support during critical incidents, creating real business risk for mission-critical applications.

For organizations in highly regulated industries, these challenges can become critical, tipping-point factors in repatriation decisions. The resource investment required to maintain compliant cloud environments – specialized staff, security tools, audit preparation – can often exceed the total operational costs of managing on-premises or colocation infrastructure where security controls are more directly managed.

SUPPORT LIMITATIONS AND VISIBILITY GAPS

Public cloud providers excel at scale, but that scale often comes at the expense of personalized support, especially for customers who aren't spending millions annually. Organizations discover that support tiers matter significantly, and the level of assistance available to mid-market companies often falls short during critical incidents.

When infrastructure problems arise, many customers find themselves dependent on provider status pages, community forums, and support ticket systems with limited visibility into underlying issues. Unlike working with a colocation or managed services provider where direct communication with engineers is standard, public cloud support can feel impersonal and difficult to navigate. During outages or performance problems, customers may have little insight into what's actually wrong, or when resolution will occur.

This visibility gap extends beyond support interactions. Public cloud abstracts infrastructure details that IT teams may need for troubleshooting, optimization, or capacity planning. While this abstraction simplifies some tasks, it can frustrate teams accustomed to deeper infrastructure control and visibility.

For mission-critical applications, these support limitations represent real business risk. When revenue-generating systems experience problems, waiting hours or days for cloud provider support to escalate and resolve issues isn't acceptable. Organizations with demanding uptime requirements often find that the support and visibility they need justifies moving certain workloads back to environments where they have direct control and responsive support partnerships.

WHEN PROMISE MEETS PRACTICE

The initial appeal of public cloud – easy setup, self-service provisioning, and minimal upfront investment – often gives way to operational realities as deployments mature. Organizations that once rushed to embrace cloud-first strategies often find themselves reassessing their decision once they encounter unpredictable costs, complex compliance requirements, and support limitations.

These aren't failures of cloud technology itself, but mismatches between cloud operational models and specific business requirements. Understanding why organizations want to leave the cloud is only half the story. The repatriation process itself introduces a new set of challenges that can be just as costly and complex as the problems organizations are trying to solve.

› Use Case:

CarePoint Analytics: Part 1: The Wake-Up Call

CarePoint Analytics seemed like the perfect cloud success story. The 250-person healthcare analytics company migrated to a public cloud provider three years ago, attracted by scalability, reduced capital expenses, and freedom from infrastructure management.

Initially, everything worked. The team loved self-service provisioning, finance appreciated predictable bills, and leadership avoided a costly data center upgrade. Within 18 months, CarePoint Analytics was processing patient data for over 200 hospital systems. So far, so good.

Then the costs started climbing. Monthly bills jumped from \$85,000 to \$180,000. API calls, data transfers, and autoscaling created expenses finance couldn't forecast. The VP of Engineering spent more time explaining overruns than discussing products.

Compliance issues added complexity. HIPAA certification required multiple security tools, detailed logging, and specialized consultants. When a critical job failed during a holiday weekend, their support tier meant waiting 18 hours for help while clients demanded answers.

By year three, leadership faced an uncomfortable question: Was cloud really the right decision?

In Part 2, CarePoint Analytics discovers that moving back introduces unexpected complications.

PART 2:

The Repatriation Process: Challenges, Complications, and Considerations

Recognizing the need to move workloads back from the public cloud is one thing, but executing that migration successfully is quite another. Organizations often approach repatriation expecting a straightforward reversal of their cloud migration, only to discover a new landscape of technical, operational, and financial obstacles.

The very features that made the cloud attractive, such as platform services, elastic scaling, and proprietary tools, become barriers when trying to leave. Understanding these challenges upfront helps organizations make realistic assessments of repatriation costs and complexity.

UNDERSTANDING TRUE MIGRATION COSTS

The most immediate surprise in cloud repatriation is often the cost of extracting data when it's time to part ways. Cloud providers design their pricing to encourage adoption, to the point where data ingress is typically free or minimal. Data egress, however, carries substantial fees that organizations may have already experienced during regular operations like analytics processing or backups. When repatriation requires moving terabytes or petabytes of data out permanently, these charges can represent a significant barrier to exit.

Beyond egress charges, a comprehensive exit cost analysis must account for application downtime during migration, staff time dedicated to execution, performance testing and validation, and the risk of business disruption if transitions don't go smoothly.

For mission-critical applications, organizations may need to maintain parallel environments during transition periods, effectively paying for both cloud and alternative infrastructure simultaneously. The timeline itself becomes a cost factor. For example, unlike faster cloud migrations through lift-and-shift approaches, repatriation requires careful execution to avoid service disruptions, extending both the duration and expense of the process.

CAPACITY PLANNING AND PERFORMANCE MANAGEMENT


Private environments replace the cloud's elastic scaling with the need for proactive resource forecasting. Organizations must determine appropriate infrastructure sizing, forecast growth requirements, and maintain sufficient capacity buffers without wastefully overprovisioning. This shift from reactive scaling to predictive planning requires different skills and processes than cloud management, where capacity decisions previously happened automatically or on-demand.

There are real risks associated with getting this wrong. Underprovisioning leads to user dissatisfaction, performance problems, and potential business disruptions. Overprovisioning wastes capital and operating expenses on unused capacity. Finding the right balance requires skills and processes that may have atrophied during long periods of cloud reliance on automatic scaling.

Performance expectations also require recalibration. While properly designed private infrastructure often delivers better and more consistent performance than public cloud, the transition period can introduce challenges as teams optimize configurations and tune systems for specific workloads.

HIDDEN REPATRIATION CHALLENGES

- **Data egress fees:** Moving large datasets out of public cloud can cost tens of thousands of dollars before touching a single server – a barrier many organizations underestimate.
- **Vendor lock-in:** Applications built on proprietary cloud services require significant refactoring, often taking 6-12 months of engineering work to achieve portability.
- **Lost institutional knowledge:** Years of relying on automatic scaling can leave teams unable to confidently plan capacity, risking either expensive overprovisioning or performance problems.
- **Extended timelines:** Unlike cloud migrations that use lift-and-shift approaches, repatriation requires careful execution to avoid disruptions – often doubling planned budgets and schedules.



THE COMPLEXITY CONTINUES

It's important to note that repatriation isn't a simple undo button for cloud adoption. The process introduces substantial technical, operational, and financial challenges that organizations must navigate carefully.

Success requires realistic planning, adequate resources, and often significant system and application modifications. These complexities explain why prevention and making better infrastructure decisions from the start proves to be far more effective than attempting to unwind poor cloud strategies after the fact.

› Use Case:

CarePoint Analytics: Part 2: The Complexity of Moving Back

CarePoint Analytics' leadership approved the repatriation plan, confident they could reverse their cloud migration within six months. Reality proved far more challenging.

The first shock came from data egress fees. Moving 500 terabytes out of the public cloud would cost nearly \$45,000 before touching a single server. Their finance team, already frustrated by unpredictable cloud bills, now faced a massive one-time charge just to leave.

Technical complications multiplied. CarePoint Analytics' analytics platform relied heavily on public cloud-native services for machine learning pipelines, managed databases, and automated scaling. The development team estimated six months of engineering work to refactor applications for their new colocation environment. Critical features built on proprietary APIs had no direct equivalents anywhere else.

Meanwhile, the infrastructure team realized they'd lost institutional knowledge about capacity planning. After three years of automatic scaling, no one could confidently size servers for peak processing loads without risking expensive overprovisioning or performance problems.

What started as a six-month project stretched to a full year, with costs doubling as the team encountered one unexpected complexity after another.

In Part 3, CarePoint Analytics learns how they could have avoided these challenges entirely.

PART 3:

Prevention and Strategic Planning

Cloud repatriation challenges aren't inevitable. While moving workloads back has become common, the pain and expense can largely be avoided through strategic planning from the outset.

The key is designing infrastructure strategies that prioritize flexibility and vendor agnosticism rather than committing wholesale to any single approach. Organizations that think strategically about workload placement, partner selection, and architectural decisions position themselves to adapt as business requirements evolve without facing costly migrations in either direction.

MAINTAIN VENDOR AGNOSTICISM FROM THE START

The most effective prevention strategy is avoiding deep dependencies on proprietary services and platforms from day one. This doesn't mean rejecting cloud providers or their offerings entirely, but rather making deliberate, informed choices about where and how to use vendor-specific features.

Whenever possible, applications should be built using open standards and portable technologies. Containerization with Kubernetes, for example, provides far more flexibility than building directly on a cloud provider's proprietary container orchestration. Open-source databases offer portability that managed database services don't. Additionally, infrastructure-as-code tools work across multiple providers, while cloud-specific tools lock organizations into specific ecosystems, some owned by the cloud provider.

Application architecture matters enormously. Designing with portability in mind means avoiding tight coupling to proprietary APIs, using abstraction layers that can be swapped out, and documenting dependencies clearly. The upfront investment in portable architecture pays dividends when business needs to change or when vendor pricing and policies shift in unfavorable directions.

Admittedly, this type of approach requires discipline. Cloud providers intentionally make their proprietary services attractive and easy to adopt. The short-term convenience of using managed services can create long-term constraints. Organizations need to weigh the immediate development speed against future flexibility and make conscious trade-offs rather than defaulting to whatever seems easiest at the moment.

CHOOSE FLEXIBLE INFRASTRUCTURE PARTNERS

Partner selection proves just as important as technical architecture. Organizations seeking to avoid repatriation challenges should work with infrastructure providers that offer multiple deployment options rather than forcing them into a single model.

The ideal partner provides a full spectrum including colocation, private cloud, bare metal, managed services, and seamless public cloud connectivity. This diversity ensures that as workload requirements change, organizations can shift approaches without changing partners or starting over.

Contract terms and pricing transparency require proactive investigation. Cloud providers' pricing isn't necessarily deceptive, but it's often incomplete. For example, they may not volunteer information about expensive scenarios like high-volume data transfers or complex scaling patterns.

Organizations should ask detailed questions upfront and expect candid answers about potential cost drivers before committing. Clear pricing without hidden fees, straightforward terms, and honest conversations about worst-case scenarios reduce the risk of unpleasant surprises. Look for partners that make it easy to move between deployment models or leave entirely if business needs demand it.

STRATEGIC PREVENTION FRAMEWORK

- **Build with portability:** Use open standards, containerization with Kubernetes, and vendor-agnostic tools rather than proprietary services that create lock-in.
- **Choose flexible partners:** Work with providers offering multiple deployment options (colocation, private cloud, bare metal, managed services) so workload changes don't require changing partners.
- **Ask the right questions:** Cloud pricing isn't deceptive, just incomplete – demand candid answers upfront about expensive scenarios like data transfers and complex scaling patterns.
- **Apply clear criteria:** Match workloads to optimal environments based on characteristics like predictability, compliance requirements, performance needs, and true total cost of ownership.
- **Plan for evolution:** Conduct regular infrastructure assessments and maintain the flexibility to adapt as business requirements change.

ESTABLISH CLEAR DECISION CRITERIA

Strategic infrastructure planning requires frameworks for making consistent, rational decisions about where to place different workloads. Without clear criteria, organizations risk making ad hoc choices that create complexity and technical debt.

Certain workload characteristics make public cloud an excellent fit. Applications with highly variable demand that benefit from elastic scaling, development, and testing environments that need fast provisioning, or workloads that genuinely require global geographic distribution often work well in public cloud. Understanding these use cases helps organizations use cloud deployments strategically rather than reflexively.

Conversely, stable production workloads with predictable resource requirements, applications handling sensitive data with strict compliance requirements, High-Performance Computing that demands consistent low latency, and large-scale data processing that incurs substantial egress fees often perform better and cost less in colocation or private cloud environments.

Finally, total cost of ownership (TCO) analyses should extend beyond simple monthly infrastructure bills. Realistic TCO includes data transfer costs, management overhead, compliance implementation, support quality, and the hidden costs of vendor lock-in. Organizations that conduct honest TCO comparisons across deployment models often discover that public cloud economics work differently at scale than they do for pilot projects.

PLAN FOR CONTINUOUS EVOLUTION

Infrastructure strategies shouldn't be set-in-stone decisions. Business requirements change, technology evolves, and vendor landscapes shift. Organizations need processes for regularly reassessing infrastructure choices and maintaining the flexibility to adapt.

Regular infrastructure assessments create opportunities to identify workloads that no longer fit their current environments, evaluate new deployment options, and optimize spending. These reviews should examine both technical performance and financial efficiency, looking for mismatches between workload characteristics and hosting environments.

Understanding how business requirements evolve allows organizations to anticipate infrastructure needs rather than reacting to them. A startup's infrastructure needs differ dramatically from those of a mature enterprise. Growth patterns, customer requirements, compliance obligations, and competitive pressures all influence optimal infrastructure strategies.

Maintaining flexibility requires ongoing investment. This means preserving in-house infrastructure expertise even when using managed services, keeping architectural options open rather than over-optimizing for current deployment models, and building relationships with multiple infrastructure partners rather than depending entirely on a single vendor.

› Use Case:

CarePoint Analytics: Part 3: A Different Path Forward

CarePoint Analytics' painful repatriation experience taught expensive lessons. The company eventually succeeded in moving critical workloads to a colocation partner, but the 18-month process and doubled budget left leadership wondering how they could have avoided the entire ordeal.

The answer became clear in hindsight: strategic planning from the start. Had CarePoint Analytics chosen an infrastructure partner offering multiple deployment options – colocation, private cloud, and public cloud connectivity – they could have placed workloads strategically from day one rather than defaulting everything to their chosen public cloud provider.

More critically, had they built their analytics platform using portable technologies and open standards instead of public cloud-native services, moving workloads would have been a business decision rather than an engineering crisis. Container orchestration with Kubernetes, open-source databases, and vendor-agnostic APIs would have preserved their options.

Today, CarePoint Analytics maintains a hybrid approach. The company runs development in public cloud and production analytics in colocation, with the flexibility to shift workloads based on business needs rather than technical constraints. Prevention proved far less expensive than cure.

Conclusion: Strategic Infrastructure for an Evolving Landscape

The rise of cloud repatriation tells an important story about infrastructure maturity. While cloud spending continues its dramatic growth trajectory – as evidenced by Gartner’s \$723 billion forecast – the 86% of CIOs planning selective workload repatriation demonstrates something equally significant. Organizations aren’t rejecting cloud computing. They’re rejecting the notion that any single approach works for every workload, every business requirement, or every stage of growth.

The pattern is clear: unpredictable costs, compliance complexity, and support limitations drive organizations to reconsider their strategies. Yet the repatriation process itself introduces substantial obstacles including data egress fees, vendor lock-in, and architectural challenges that can equal or exceed the original problems.

Prevention remains far more effective than reaction. By designing systems and applications with portability in mind, choosing partners offering genuine flexibility, and making deliberate workload placement decisions, organizations avoid painful choices between untenable cloud costs and expensive repatriation projects. Strategic infrastructure planning delivers reduced risk, better cost management, and freedom to evolve as business needs demand.

About DataBank

DataBank enables the world's largest enterprises, technology, and content providers to consistently deploy and manage their infrastructure, applications, and data on the right platform, at the right time, in the right place. Our colocation and edge infrastructure footprint consists of 65+ data centers and 20 interconnection hubs in 27+ markets, on-ramps to an ecosystem of cloud providers, and a modular edge data center platform with virtually unlimited reach.

We combine these platforms with contract portability, managed security, compliance enablement, hands-on support, and a guarantee of 100% uptime availability, to give our customers absolute confidence in their IT infrastructure and the power to create a limitless digital future for their business.

