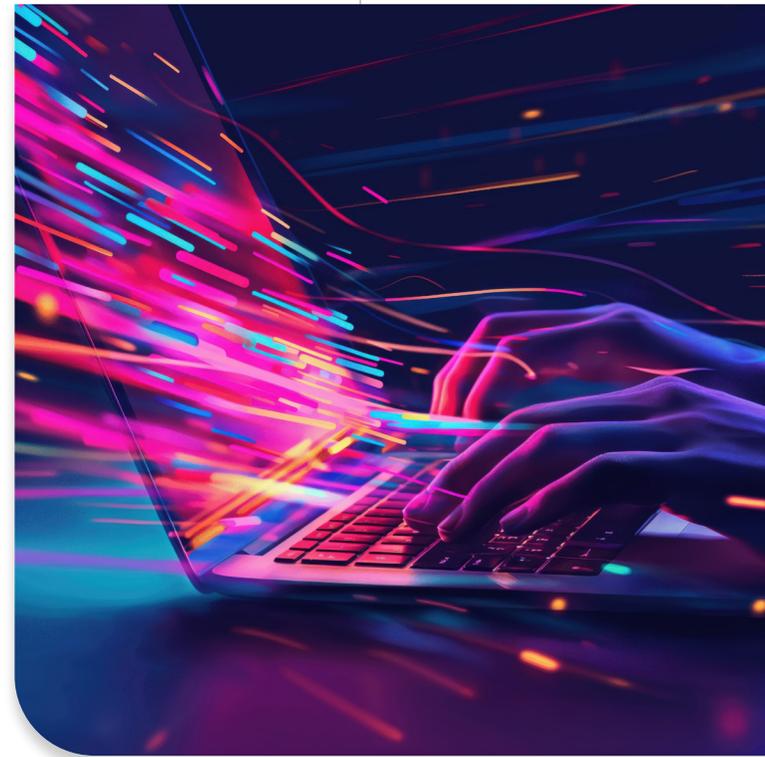


riverbed[®]

Breaking Barriers in Enterprise Data Movement

EBOOK

Table of Contents



Transforming Data Transfers into a Strategic Enabler

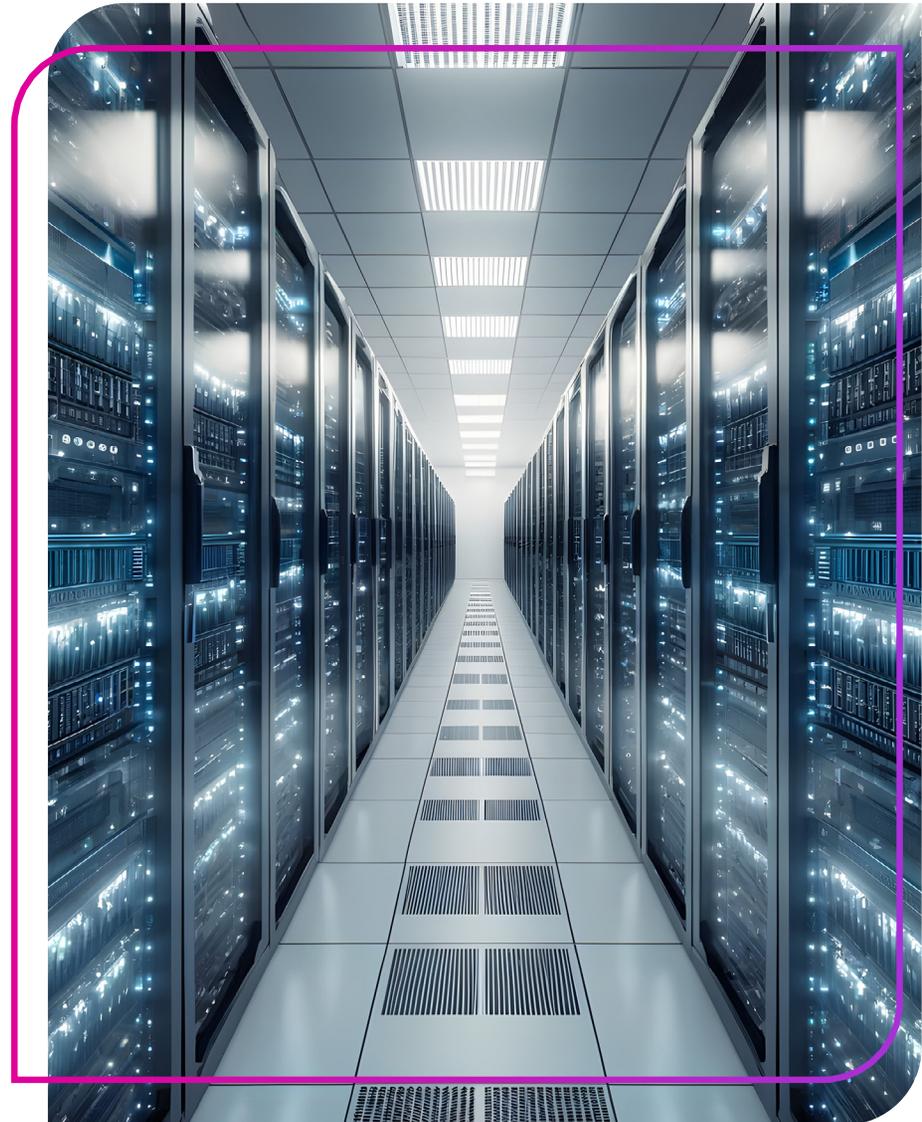
Executive Summary

As organizations consolidate data centers, go through M&A activity, transition to cloud or hybrid cloud models, or deploy resiliency capabilities, they face significant hurdles in data movement. Some of the key risks include

- Operational disruption
- Data loss
- Compliance violations
- Cost overruns

Data movement is a critical enabler, allowing organizations to migrate workloads, integrate applications, and support real-time operations across multiple environments. Without efficient data movement, workloads and applications remain tethered to legacy environments, limiting the ability to scale on demand or optimize costs.

This ebook explores these challenges in data movement and the strategic imperative of agile, secure, and cost-effective data movement approach. It then introduces Riverbed's Data Express Service that dramatically accelerates data transfer, reduces costs, and enhances observability across multiple use cases and environments.



Data Movement as a Business Imperative

The Strategic Imperative: Why Data Movement Matters More Than Ever

Data movement is not just a technical task; it is the foundation for unlocking the economic and operational benefits for an organization.

Across various use cases, data movement is never the headline—but it is always the outcome determinant. When data moves well, change feels deliberate and controlled. When it moves poorly, even the best architectural strategies unravel under latency, risk, and uncertainty. Here are four common data movement use cases.

1. Data Center Consolidation

In data center consolidation, data movement is the quiet force that determines whether consolidation delivers efficiency or chaos. Applications, datasets, and services rarely move together at the same pace, and the act of relocating data becomes the connective tissue that binds new architectures into something coherent. Successful consolidation depends not just on moving data, but on preserving performance, availability, and trust throughout the transition.

The role of data movement here is to decouple applications from physical locations, enabling workloads to be rehomed without disruption. As facilities are shuttered or reduced, data must traverse constrained networks, coexist temporarily in multiple locations, and remain accessible throughout prolonged transition windows. This movement is often incremental rather than singular, unfolding over months or years.

The central challenge lies in scale and simultaneity. Legacy systems, unpredictable traffic patterns, limited WAN capacity, and outdated protocols collide with modern consolidation goals. Latency sensitivity, incomplete data visibility, and the risk of extended outages can derail consolidation efforts, turning a cost saving initiative into an operational liability if data movement is not centrally governed and optimized.



Data Movement as a Business Imperative

2. Mergers and Acquisitions

In mergers and acquisitions, data movement becomes both a strategic enabler and a cultural fault line. Data is no longer just information—it is evidence, intellectual property, operational history and institutional memory. Moving it is essential to unify systems, standardize operations and realize the value promised by the deal.

Here, data movement plays the role of integration arbitrator, reconciling incompatible schemas, security policies and performance expectations. Systems built under different assumptions must suddenly interoperate, often before full architectural alignment is possible. Data movement allows organizations to operate in parallel states—integrated enough to function, separate enough to manage risk.

The challenges are deeply human and technical. Disparate governance models, regulatory constraints, data sensitivity and uneven infrastructure maturity create friction. Network constraints compound urgency, as value realization timelines pressure teams to move data faster than the environment was designed to tolerate. Without disciplined control, data movement can amplify security risks, propagate errors and surface hidden dependencies that delay integration and erode trust.

3. Backup, Replication, and Recovery

In backup, replication and recovery, data movement is the foundation of resilience. Every copy, snapshot, and replica exists because data is continuously in motion—quietly flowing away from primary systems to places where failure cannot reach it. The reliability of these movements determines how quickly systems recover and how confidently businesses operate under uncertainty.

The role of data movement in this context is insurance through redundancy. It must be predictable, efficient and invisible to production workloads. Whether replicating across sites or streaming to long term storage, data movement ensures continuity by maintaining temporal alignment between active systems and their recoverable counterparts.

The challenge lies in balancing speed, cost and consistency. Volumes grow faster than networks evolve, recovery expectations shrink while compliance demands expand, and replication traffic competes with business critical workloads. Latency, bandwidth saturation, and incomplete recovery testing undermine confidence. When backup data moves poorly, recovery becomes a theory rather than a certainty.

Data Movement as a Business Imperative

4. Data Center-to-Cloud Migration

In data center-to-cloud migration, data movement is the bridge between intent and reality. Cloud promises elasticity and innovation, but none of it materializes until data arrives intact, performant, and secure. Migration is not a single transfer, but a sustained choreography of synchronization, validation, and cutover.

Here, data movement plays the role of transformation catalyst, enabling organizations to run hybrid operations while progressively shifting gravity toward the cloud. Data must move repeatedly, often bi-directionally, as applications are refactored, dependencies discovered, and operating models rewritten in real time.

Data movement is essential for achieving the affordability, flexibility and scalability promised by cloud adoption. By moving data into the cloud, businesses can take advantage of elastic storage and compute resources, paying only for what they use rather than maintaining expensive on-premises infrastructure. Without the right data movement strategy, the promise of affordability, flexibility, and scalability from cloud remains out of reach. Yet the challenge of transferring large, complex datasets across cloud and data center platforms is made more complicated when an organization has a multi-cloud strategy.

Hybrid and multicloud strategies require data to flow seamlessly between multiple providers, each with unique APIs, storage formats, and security models. With different architectures, protocols, and

90%
of enterprises are multicloud

compliance requirements, this all becomes exponentially more complex. This lack of standardization creates interoperability issues that complicate integration and governance. The challenge is that cloud magnifies imperfections. Latency becomes visible, egress costs accumulate quietly, and inconsistencies between premise and cloud datasets create subtle application failures. Without disciplined movement strategies, migrations stall under the weight of prolonged dual environments, performance degradation, and escalating costs. The difficulty is not moving data once but sustaining high fidelity data movement long enough for transformation to complete.

Despite these differences, all face shared challenges. Performance major concerns. Moving terabytes or petabytes of data across networks can take weeks to months, strain bandwidth, introduce latency, and disrupt critical applications. Cost is another common challenge since cloud providers typically charge egress fees for data leaving their service, making frequent or inefficient transfers expensive. Security and compliance are top concerns, as data in transit is vulnerable to breaches and cross-border transfers can violate regulations. When data moves across the internet, between on-premises systems and cloud platforms, it becomes vulnerable to security threats if encryption and monitoring are not enforced. Regulatory compliance adds complexity, as cross-border transfers can violate data residency laws and industry regulations.

Top Four Challenges in AI Data Movement

Time

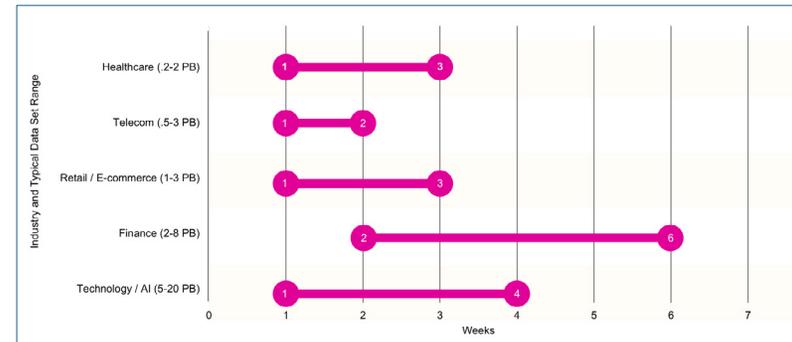
Moving the petabytes of data needed to feed enterprise applications present significant time-related challenges. Even with high-speed connections, transferring this volume can take days or weeks, creating delays that are unacceptable for time-sensitive enterprise application and cloud migration use cases. For example, over a perfectly-tuned 10 Gbps connection, it would take roughly 9 days to move 1 PB of data continuously.

However, different industries need to regularly move different amounts of data, and their different infrastructure investment can greatly impact how fast they can move 1 PB of data. (See “Typical Time to Move Data Sets By Industry” graphic.)

Before data can be moved, it often needs to be cleaned, transformed, and validated. These preprocessing steps add significant time overhead, especially when dealing with unstructured or inconsistent data sources.

Distance, inadequate network infrastructure and protocol overhead can slow down data movement, especially when data must be moved from remote geographic locations of data centers that can provide operational fail over, improved performance and scale. These bottlenecks erode the core benefits of cloud—agility, scalability, and cost efficiency while introducing operational risks that can impact competitiveness and customer experience.

Typical Time to Move AI Data Sets By Industry



When data transfer between environments is delayed, applications that rely on timely access to information suffer from performance degradation. This can disrupt real-time analytics, delay decision-making, and impair customer-facing services that depend on up-to-date data.

From an operational standpoint, slow data movement creates bottlenecks in workflows such as backup, disaster recovery, and workload migration. These delays can extend maintenance windows, increase downtime, and compromise business continuity. In multicloud scenarios, where workloads often span multiple providers, sluggish data transfers undermine the flexibility and scalability that cloud adoption promises, limiting the ability to shift workloads dynamically for cost or performance optimization.

Top Four Challenges in AI Data Movement

Scale

The scale of data movement in all four use cases presents a unique set of challenges because of the sheer volume, velocity and diversity of data involved. Modern enterprises generate massive amounts of data across multiple clouds, on-premise data centers and edge locations. Moving this data efficiently is not a simple task—it requires high-bandwidth networks, optimized transfer protocols, and sophisticated orchestration to avoid bottlenecks.

As the scale grows, so do the risks. Large-scale transfers can overwhelm network capacity, leading to latency and degraded application performance. They also increase operational complexity, as organizations must coordinate movement across different providers with varying APIs, security models and compliance requirements. Additionally, maintaining data consistency and integrity across distributed environments becomes more difficult at scale, raising the risk of synchronization errors and compliance violations.

In short, the scale of data movement amplifies every challenge—performance, cost, security and governance—making it a strategic priority for organizations to adopt automation, intelligent routing and policy-driven controls to manage data efficiently across multicloud and hybrid ecosystems.

Common Industry-Specific Data Set Sizes and Use Cases

Industry	Daily Data Volume	Key Cloud Use Cases
Financial Services	1-100+ TB	Transactional data, fraud detection, risk modeling, real-time trading analytics
Healthcare	1-50 TB	Electronic Health Records, medical imaging, genomic sequencing, telemedicine data
Retail & E-commerce	10-100+ TB	Customer purchase history, inventory management, personalization engines, demand forecasting
Manufacturing	1-10 TB	IoT sensor data from machines, predictive maintenance, supply chain optimization
Media & Entertainment	10-100+ TB	Video streaming libraries, digital content distribution, rendering and post-production workflows

Not One and Done

Keeping data fresh for enterprise applications is essential to maintain accuracy, but the process of constantly ingesting and transferring data introduces delays that can undermine time-sensitive applications.

This places a heavy burden on resources, both human and infrastructure, consuming staff time, bandwidth, compute and storage resources—especially when dealing with real-time or high-volume data. Managing these transfers adds operational complexity, requiring robust orchestration and governance, often requiring new, dedicated teams to properly perform these functions, as the risk to the business of not doing it correctly is too great to do otherwise.

Top Four Challenges in AI Data Movement

Complexity

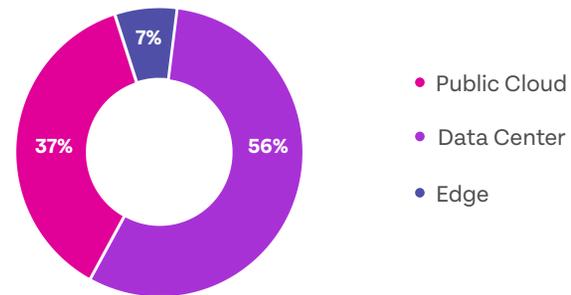
Beyond the sheer scale required to move data, there is increased complexity. As data becomes more distributed across clouds, edge, and on-premise environments, keeping it all in sync introduces the first challenge: data consistency—ensuring that every copy of the data reflects the same state across multiple environments. This becomes difficult when updates occur simultaneously in different locations, requiring sophisticated conflict resolution and synchronization mechanisms.

Latency and slow performance add to the complexity of data movement. Synchronizing large volumes of data across geographically dispersed systems can lead to delays, impacting real-time applications, analytics and a “single source of truth.” Network bandwidth limitations further complicate timely synchronization.

Additionally, transferring data between data center, cloud and edge environments raises compliance complexity risks, as data may cross jurisdictional boundaries or violate privacy regulations. The number of different and sometimes competing regulations creates a complex matrix of what can be moved when and to where. Just keeping up with the ever-changing compliance regulations of different jurisdictions is challenge onto itself.

Finally, operational complexity grows as organizations manage multiple platforms with different APIs, protocols, and storage formats. Automation and orchestration tools are essential but implementing them at scale requires advanced expertise and careful planning.

Where Enterprise Data Resides



56% of all enterprise data, structured and unstructured, is stored across multiple public clouds.

Top Four Challenges in AI Data Movement

Cost

The bigger the data footprint, the higher the cost implications, with egress fees and storage replication expenses escalating quickly.

Financially, inefficiencies in data movement can lead to higher costs. Extended transfer times consume more network resources and may trigger additional egress fees from cloud providers. Furthermore, slow synchronization between environments can result in inconsistent datasets, increasing the risk of errors, compliance violations, and security vulnerabilities.

Perhaps biggest, most important cost is the opportunity cost in delaying delivery of superior customer experience or timely business analytics. When data isn't delivered to applications in a timely manner, it can lead to outdated insights, slower decision-making, and missed chances to respond to emerging trends or threats. In customer-facing applications, stale data can reduce personalization accuracy, leading to lower engagement or conversion rates.

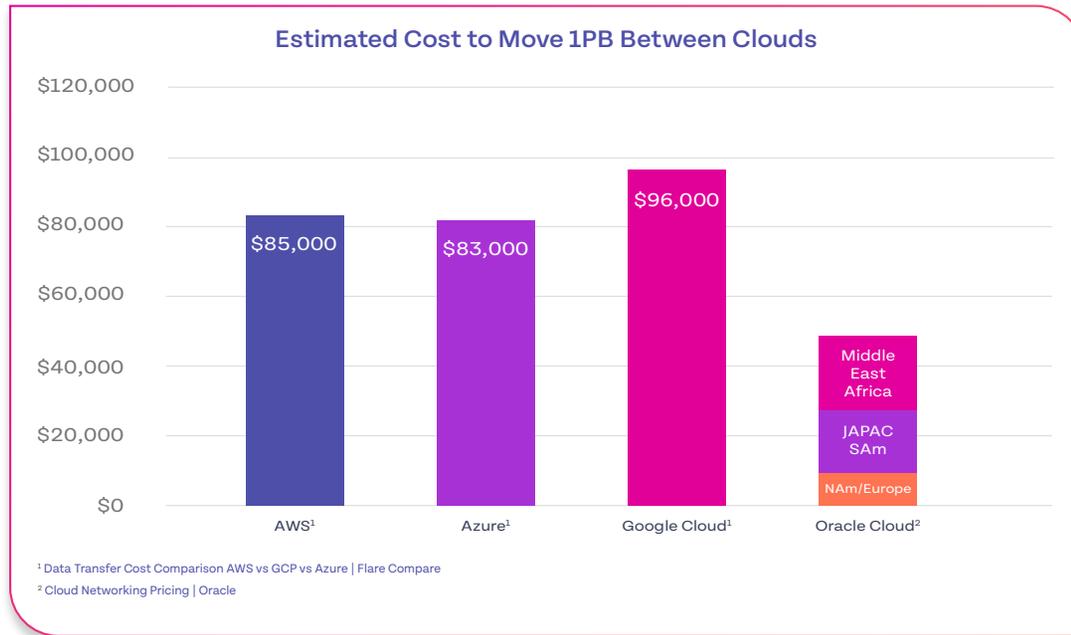
Hard Costs, Too

Beyond lost opportunity costs, there are actual real costs in moving data. One of the most significant is cloud egress fees, which are charged when data is transferred out of cloud environments—between different cloud providers or different regions within the same cloud provider, or to on-premises systems. These fees can be substantial when dealing with petabytes of data.

It can easily cost \$80,000 in egress charges to move 1 PB of data out of a cloud provider, although each of the four major hyperscalers have different rates depending on various factors including where the data resides, so actual costs do vary. (See “Estimated Cost to Move 1PB Between Clouds” graphic).

Beyond egress fees, moving data out of a cloud provider incurs several other hard and hidden costs that can significantly impact enterprise budgets. These may include network transfer and bandwidth costs, especially for high-throughput or cross-region transfers; storage tiering and duplication to prepare data for transfer; compute costs for data preparation which consume compute resources and can drive up cloud bills, particularly when done at scale, various licensing, API access fees, orchestration tools and monitoring expenses and of course labor costs.

Top Four Challenges in AI Data Movement



- **Fast:** Data must be moved quickly so it can go to work sooner
- **Easy:** Scattered data is complicated so moving it has to be simple to manage
- **Secure:** The risk involved in highly sensitive data being put in motion increases cyber threats
- **Insightful:** IT and business leaders will want to know how things are performing and what further network changes may be required

Organizations must treat it as a strategic priority by implementing secure transfer protocols, enforcing governance, and optimizing for performance and cost. Hence, there is a strategic imperative for organizations to develop robust migration planning, enforce consistent governance, implement secure sharing protocols, and optimize for performance and cost

Since data feeding enterprise applications needs to be constantly updated, these costs can be incurred multiple times every month.

In short, data movement in hybrid cloud and multicloud environments is not just a technical task; it is a strategic challenge that impacts cost, performance, security, and compliance. Organizations need a way to move petabytes of data to be:

Solving the Top Four Challenges in AI Data Movement



Riverbed Data Express: SaaS-Based Simplicity

Sometimes, enterprises just want to move data. They want the easy button to move massive amounts of data across their multi-cloud infrastructure .

Riverbed Data Express Service enables enterprises to radically accelerate the movement of massive datasets now required to prepare and move enterprise data at scale—reaching data transfer speeds up to 10 times faster than current industry solutions —improving a customer’s time to value and lowering costs. With the new Riverbed Data Express Service, what once took months can now be completed in days—giving organizations the speed and security of data delivery now required to prosper in the multicloud and hybrid cloud era.



Fast: Move Petabytes Per Day—Literally

Riverbed Data Express Service is engineered for speed, offering transfer rates from 5TB to 40TB per hour—that’s up to 1PB per day. With support for concurrent transfers and continuous sync, you can scale to meet even the most demanding cloud migration workloads. Compared to alternative available solutions, Data Express Service delivers nearly 10x performance at peak.



Easy: SaaS Simplicity and Better Economics

Forget complex configurations or hidden costs. Data Express Service is a fully managed SaaS platform—no VMs to deploy or manage, no tuning required. The power of Riverbed SteelHead is under the cover, ready to move petabytes of data. You pay only for what you move, with a consumption-based pricing model that aligns with your cloud economics. For many data types, the service can reduce cloud egress costs up to 90%, reducing overall cost of data migration.

Setup takes minutes via a simple user interface and integration is seamless via REST APIs. It works across Oracle Cloud Infrastructure today, with planned services for multi-cloud (AWS, OCI, Azure, Google Cloud Platform) and data center connectivity.

Solving the Top Four Challenges in AI Data Movement



Secure: Built for Enterprise-Grade Protection

Security is non-negotiable. Data Express Service is multi-tenant by design, with RBAC, SSO/IAM and tenant isolation at both the control and data layers. Data in transit is protected with post-quantum cryptography (PQC) secure tunnels.

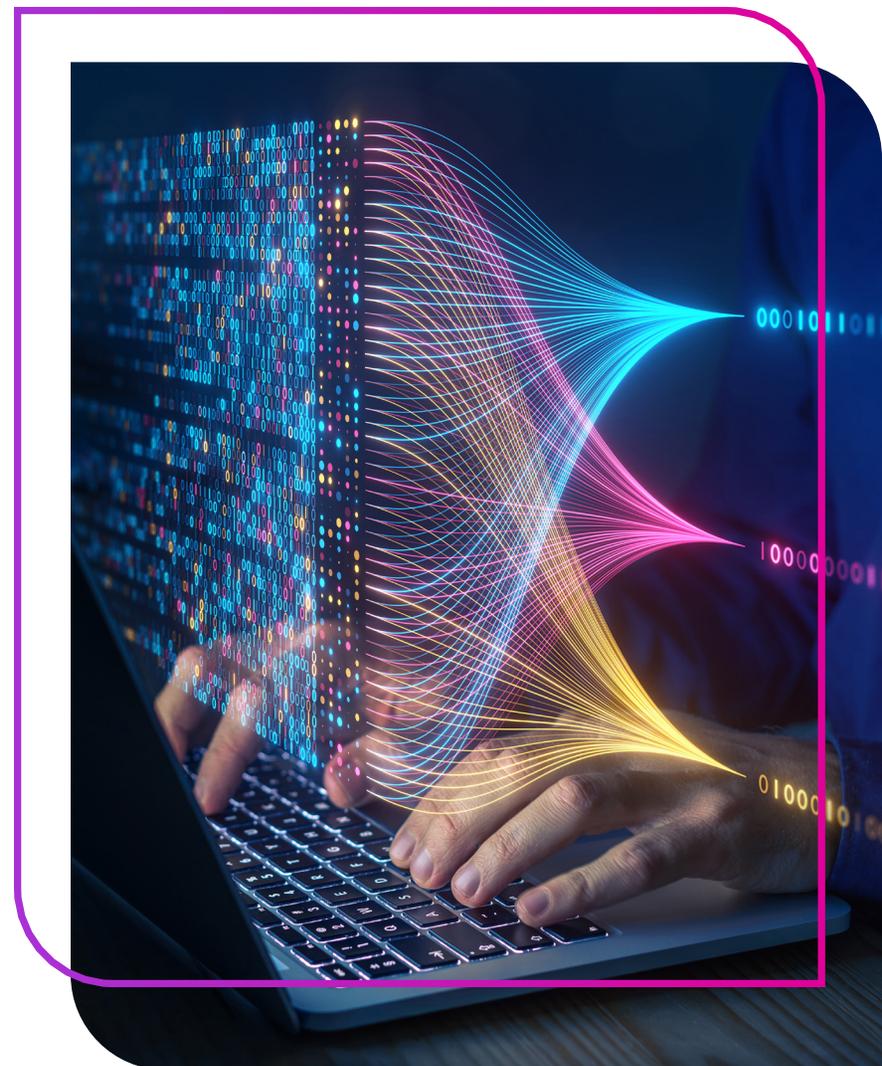
You can even deploy agents inside your own cloud or data center fabric to apply your existing security policies.



Insightful: Real-Time Visibility and Governance

From day one, you get deep visibility into job status, throughput, file counts, and history—all through a unified portal. Over time, Riverbed is adding outcome-oriented insights like time saved, top data movers, and trend analysis.

Advanced analytics and alerting integrate with your workflows, giving you the observability and governance you need to stay in control.



Use Cases That Matter to CIOs

Riverbed Data Express Service is purpose-built for the modern CIO agenda:

- Cloud migration
- Disaster recovery and backup
- Data replication for high availability
- Real-time analytics
- Application portability
- Compliance and data residency
- Dev/test environments
- AI/ML model training
- Edge data processing
- Multicloud optimization

Regardless of your industry, whether you're in financial services, healthcare, media, manufacturing, or some other segment, Data Express Service adapts to your enterprise needs. You can subscribe via Riverbed sales or partners, with options for pay-as-you-go or enterprise subscriptions. Training and onboarding are available through Riverbed Academy.

The Riverbed Data Express Service is currently in limited availability, with general availability coming soon. If you're a CIO looking to future-proof your data infrastructure, now is the time to engage.

Data Movement As a Strategic Priority

Data movement no longer has to be cumbersome, backend task—it can be a strategic enabler. Riverbed Data Express Service gives you the speed, simplicity, security, and insights needed to lead in an AI data-driven world.

Ready to accelerate your AI journey? [Contact us](#) for a Riverbed Data Express demo.

About Riverbed

Riverbed, the leader in AI observability, helps organizations optimize their users' experiences by leveraging AI automation for the prevention, identification, and resolution of IT issues. With over 20 years of experience in data collection and AI and machine learning, Riverbed's open and AI-powered observability platform and solutions optimize digital experiences and greatly improve IT efficiency. Riverbed also offers industry-leading Acceleration solutions that provide fast, agile, secure acceleration of any app, over any network, to users anywhere. Together with our thousands of market-leading customers globally – including 95% of the FORTUNE 100 – we are empowering next-generation digital experiences. Learn more at riverbed.com

© 2026 Riverbed Technology LLC. All rights reserved. Riverbed and any Riverbed product or service name or logo used herein are trademarks of Riverbed. All other trademarks used herein belong to their respective owners. The trademarks and logos displayed herein may not be used without the prior written consent of Riverbed or their respective owners. CS-547_BreakingBarriers_DataMovement_Ebook_022326