

Solving the Four Data Movement Challenges in the AI Era

Data at the Speed of AI with Riverbed Data Express

The Four Challenges of AI Data Movement

AI is hungry – and data is the feast. However, most of the data needed is not located in data centers with the specialized infrastructure required to run AI models. As a result, data needs to be moved to AI and not the other way around. This creates four specific challenges.

1. Time to Transfer: Perhaps the most critical challenge for AI data movement is the time it takes to transfer data to AI models. AI requires fast data access to keep models accurate, responsive, and valuable to enable business value and full utilization of compute resources. In AI, time to transfer isn't just a technical metric – it's a competitive advantage. Distance, inadequate network infrastructure and protocol overhead can slow down data movement, especially when data must be moved from remote geographic locations to optimized data centers running AI models.

Scale and complexity also creates additional time to transfer challenges.

2. Scale: Enterprise AI workloads routinely involve terabytes to petabytes of data which can take weeks to months to move. Moving this data consumes bandwidth, compute, and time – creating systemic bottlenecks for other mission critical applications.

Business Challenge

In the age of AI-driven transformation, data is no longer just an asset – it's the lifeblood of innovation.

AI models – especially large-scale ones – live in specialized GPU-rich environments. Your enterprise data doesn't. This disconnect creates delays, drives up costs, and undermines the agility needed to compete. For CIOs and business leaders, solving this challenge isn't a technical detail – it's a strategic imperative.

As enterprises scale AI across cloud, edge, and hybrid environments, a silent bottleneck threatens progress: data movement.

AI models require high-performance infrastructure – GPUs, TPUs, orchestration tools – housed in specialized data centers. Meanwhile, enterprise data is scattered across clouds, data centers, edge devices, and SaaS platforms.

This means data must move to the model – not the other way around because the data needed to fuel AI is rarely co-located with the compute power that processes it. And moving that data – securely, quickly, and cost-effectively – is now a business-critical function.

3. Complexity: Data is fragmented across environments, formats, and jurisdictions. Consolidating it introduces security, compliance, and governance risks. IT teams must navigate a maze of APIs, protocols, and privacy regulations – all while maintaining performance and accuracy.

The mere scale of the task adds operational complexity and places a heavy burden on resources, both human and infrastructure. This requires robust orchestration and governance and dedicated teams as the risk to the business of not doing it correctly is too great to do otherwise.

4. Cost: One of the most apparent data movement costs 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.

It can easily cost \$80,000 in egress charges to move one petabyte of data out of a cloud provider, although each of the four major hyperscalers have different rates depending on numerous factors including where the data resides, so actual costs do vary.

One hidden cost in AI deployments is the ‘Double Bubble’—when data migration between two cloud environments creates overlapping infrastructure expenses. While temporary, this duplication can run into millions, significantly inflating project budgets.

Opportunity costs include eroded competitive advantage—slowing value delivery of AI-driven insights.

Solving Data Movement Challenges with Riverbed Data Express

Riverbed Data Express Service is a high-speed, secure, and fully managed SaaS solution designed to accelerate the movement of massive datasets between public clouds, regions within the same cloud, and data centers. It enables enterprises to radically accelerate the movement of datasets to deploy AI models at scale—reaching data transfer speeds up to 10 times faster than current industry solutions—improving a 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 AI era.

The Riverbed Data Express Advantages

Riverbed’s Data Express Service is a SaaS-based solution designed to solve data movement problems at enterprise scale.

Speed That Matches AI Ambition

- 10x faster than traditional solutions
- Move up to 1PB per day
- Transfer rates of 5TB to 40TB per hour

Simplicity That Scales

- Fully managed SaaS – no VMs, no tuning
- Consumption-based pricing aligned with cloud economics
- Setup in minutes, integration via REST APIs

Better Economics That Makes Sense

- Reduce cloud egress costs up to 90%
- Lower operational costs

Faster AI Outcomes That Matter

- Protect first-mover advantages
- Easily refresh data for accurate results

Security That Meets Enterprise Standards

- Post-quantum cryptography (PQC) secure tunnels

Visibility That Drives Governance

- Real-time insights into throughput, job status, and trends
- Advanced analytics and alerting
- Unified portal for observability and control

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.