



Miria vs Starfish:  
Enterprise-Grade Data Movement vs.  
Lightweight Orchestration



- Enterprise software platform
- PostgreSQL-based metadata engine
- Designed for sustained, large-scale operations
- Full ACL and metadata preservation
- Built for billion-file environments



- Namespace automation layer
- Policy-based data placement
- Lightweight control plane
- Often deployed in multiple clusters to scale
- Primarily focused on file tiering workflows

### When it Matters Miria Wins

*When AI, compliance, and metadata integrity matter, and failure is not an option, Miria's architecture wins:*

• Multi-petabyte migrations

• Billion-file environments

• Small-file HPC workloads

• Regulated enterprise environments

• Complex metadata preservation

• Cold archive workflows (incl. tape)

### Practical Difference at Scale

*Miria provides architectural headroom without requiring cluster gymnastics.*

• Metadata scale pushes lightweight control planes to limit

• Database bottlenecks introduce operational risk

• Multi-cluster deployments increase complexity and risk

• Archive and cold workflows extend beyond simple object tiering

### The Bottom Line

Miria: Enterprise-scale data movement infrastructure.

Starfish: Namespace automation and policy management.



Let's Make It AI-Ready. [Book a Demo.](#)