CASE STUDY

Bridgewater Associates

YellowDog Enables High-Scale Global Compute To Accelerate Time to Results

One of the world's largest hedge funds partners with YellowDog to operate at a scale of compute that fundamentally changes the economics and velocity of quantitative research.

AT A GLANCE

35,000+	7×	99.7%	70%
Nodes in a typical production run	Increase in research velocity	Spot utilisation across multiple AWS regions	Cost reduction vs on-demand pricing

Executive Summary

Bridgewater Associates is one of the world's largest hedge funds. As its quant teams built larger ML models, they needed to scale their infrastructure to support them.

Working closely with YellowDog gives individual research teams "on tap" access to parallel compute across multiple AWS regions, drawing from dynamically sourced, cost-efficient Spot capacity at scale.

High-scale workload execution is embedded directly into Bridgewater's research workflows, giving each team the autonomy to build and run complex models without requiring infrastructure teams to pre-provision resources or manually intervene. Production clusters now regularly surge beyond 35,000 nodes so results are achieved faster.

- **Increased speed of research** — by running massively in parallel across tens of thousands of nodes, YellowDog delivers more than 7× faster research velocity.
- **Unlocking new possibilities** — increased compute scale helps drive the creation of new models and ideas.
- **Cost-efficient scale** — consistently high Spot utilisation across multiple AWS regions delivers over 70% cost savings versus on-demand pricing.

> ***"What differentiates YellowDog from any other provider is the efficiency and speed with which it achieves scale. It automatically orchestrates compute so our research teams can run extremely large workloads without the complexity and cost of building that capability ourselves."***
>
> **— Artem Kazakov, Lead Architect at AIA Labs - Bridgewater Associates**

YellowDog Enables High-Scale Global Compute

Quantitative research is, at its core, about maximizing compute efficiency to run more experiments, test more hypotheses, and act on signals faster than the competition. The ability to access and orchestrate compute at scale - within acceptable cost envelopes – is a decisive factor in research velocity.

Core Infrastructure Requirements for Quantitative Finance

- **Research velocity** — the ability to move from hypothesis to result with no infrastructure friction.
- **Massive parallel scale** — tens of thousands of nodes for large batch or high task volume workloads.
- **Cost efficiency** — operating almost exclusively on Spot means cost is never a barrier to experimentation.
- **Elastic cloud infrastructure** — ensure cost-efficient high-scale compute is available exactly when needed.

Bridgewater selected YellowDog because true performance at scale demands intelligent, global workload management. YellowDog helped increase Bridgewater's compute ceiling while maintaining close to 100% Spot utilisation, expanding the potential for new ideas more quickly and cost-efficiently.

The Solution

Intelligent Global Orchestration at Production Scale

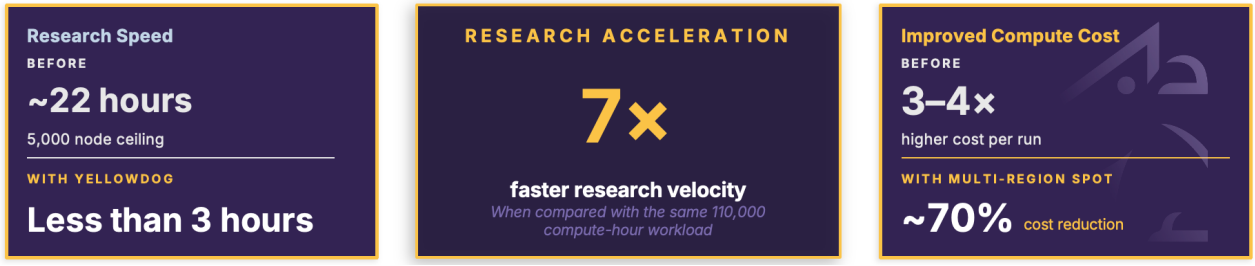

The production evidence is unambiguous. Bridgewater's cloud infrastructure regularly scales to more than 35,000 nodes across multiple AWS regions, executing large-scale workloads with the most compute-intensive portions completing in less than 3 hours, representing 7× faster research velocity.

Multi-Region Spot Strategy

By dynamically distributing cloud worker nodes across multiple AWS regions, YellowDog enables Bridgewater to "follow the moon," drawing from regions where Spot capacity is most available and cost-efficient at any given time. This geographic agility ensures that supply constraints in a single region do not force a fallback to on-demand pricing. The result: Bridgewater achieves 99.7% Spot utilisation with an estimated 70% cost savings compared with on-demand pricing.

What Performance at Scale Looks Like in Production

YELLOWDOG CAPABILITY	WHAT IT MEANS FOR RESEARCH TEAMS	OBSERVED IN PRODUCTION
Rapid Platform Integration	Minimal infrastructure overhead to drive greater quant research agility.	*Live in production clusters — integrated directly into existing research workflows.*
Massive Parallelisation	Schedule and execute extremely large workloads seamlessly within time-critical windows.	*Production runs regularly reaching 35,000+ nodes across the cluster.*
Instance Type Diversity	Orchestrate global Spot capacity across multiple instance types using a single logical compute fabric.	*29–32 different instance types leveraged on Spot across a single production run. 70% Arm Graviton.*
Multi-Region Execution	Reduce exposure to regional capacity constraints and absorb supply variability automatically.	*Multiple AWS regions; 99.7% Spot utilisation sustained across the production cluster.*
Dynamic Spot Management	Auto-recover from interruptions without researcher involvement.	*Spot continuity assured with auto-recovery across regions in less than 4 minutes.*

Orchestration & Infrastructure That Accelerates Alpha

YellowDog enables Bridgewater to operate a compute engine that would be unachievable using conventional infrastructure. By dynamically drawing from multi-region Spot capacity across dozens of instance types, Bridgewater's quant teams execute workloads that deliver results faster.

The numbers are clear: production runs regularly reaching 35,000+ nodes; validated 7× faster research velocity; consistently high Spot utilisation across multiple AWS regions; and an estimated 70% savings versus on-demand pricing.

> ***"YellowDog enables us to continuously push compute scale further. Using global cloud Spot availability has enabled us to use compute efficiently - which is key to sustaining our research velocity and competitive performance."***
>
> **— Artem Kazakov, Lead Architect at AIA Labs - Bridgewater Associates**