

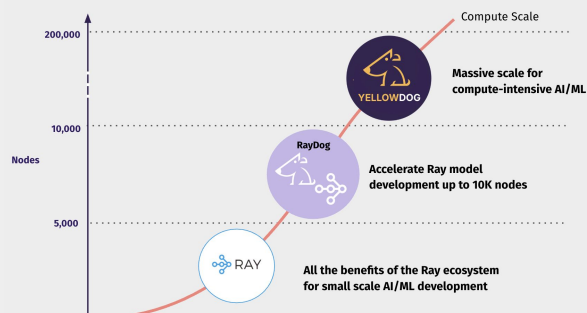


RayDog Delivers the Ease of Ray with the Power of YellowDog

Ray is an open source framework that makes it simpler to build distributed AI and ML applications. RayDog takes it further, delivering instant orchestration for compute, cost efficiency, and seamless scalability—removing infrastructure barriers so teams can innovate faster.

Powered by YellowDog's dynamic workload manager, RayDog is purpose-built for financial services companies, AI innovators and research teams, enabling them to leverage the simplicity of Ray while accelerating development with full access to the highest levels of compute—driving faster experimentation and innovation.

Powering Ray Clusters at Massive Scale



RAYDOG ADVANTAGE

Infrastructure Management - Fully managed compute infrastructure

Scalability - Seamless scaling from local datasets to massive clusters up to 10K nodes

Autoscaling - Intelligent autoscaling with dynamic upscaling and automated downscaling

Multi-Cloud Support - Supports multi-cloud provisioning across different regions

Cost Optimisation - Bids for compute across clouds; automated reclamation of idle resources

Security & Compliance - Integrated security and governance with policy-based scheduling

Developer Experience - Developer-friendly Ray integration that is "Ray-ready out of the box"

Observability & Monitoring - Detailed telemetry on Ray cluster performance and cost

BUSINESS OUTCOMES

Frees teams from infrastructure overhead, allowing data scientists and research engineers to focus on innovation

Accelerates time-to-market by enabling workloads to grow instantly to enterprise scale without rewriting code.

Ensures compute resources match demand in real-time, reducing costs and improving workload efficiency

Provides resiliency and flexibility, allowing workloads to run where it's optimal for cost, performance, or compliance

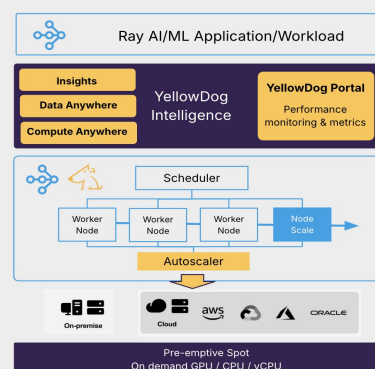
Reduces cloud spend on ML and AI workloads, maximizing ROI and enabling cost-effective scaling

Ensures regulatory compliance and data protection, critical for finance, healthcare, and other regulated industries

Speeds innovation by letting developers use familiar Ray APIs and ML frameworks without worrying about underlying infrastructure

Enables data-driven decisions, optimises workload distribution, and proves ROI to stakeholders

RayDog extends the open-source Ray framework with enterprise-grade orchestration. Under the hood, YellowDog's scheduler and autoscaler sit beneath Ray, managing low-level provisioning, scaling and workload distribution, across clouds and regions. From the Ray user's view, it behaves like Ray—while YellowDog handles the complexity of large-scale, cost-efficient, and stable compute up to 10,000 nodes.



Accelerate Innovation with RayDog

Leading organisations focused on strategic innovation using AI/ML are accelerating workload development with RayDog, powered by YellowDog's high-scale, orchestrated compute platform. Learn more [here](#).