



Parallel Works ACTIVATE

ACTIVATE AI

Operationalizing AI Resource Adoption at Scale

What is Parallel Works' ACTIVATE AI?

Parallel Works, provider of the ACTIVATE control plane for hybrid multi-cloud computing resources, includes ACTIVATE AI, to help organizations operationalize artificial intelligence (AI) and machine learning (ML) workloads.

Operationalizing AI presents an array of challenges. Many enterprises manage fragmented infrastructure across Kubernetes, cloud and legacy systems; face high costs and underutilization of GPU resources; and lack a consistent path from research to scalable, production-ready AI. ACTIVATE AI addresses these hurdles by unifying infrastructure under one control plane, enabling GPU fractionalization with cost tracking across Kubernetes, virtual machines and cloud-native environments, and simplifying the move from experimentation to deployment across hybrid and distributed environments.

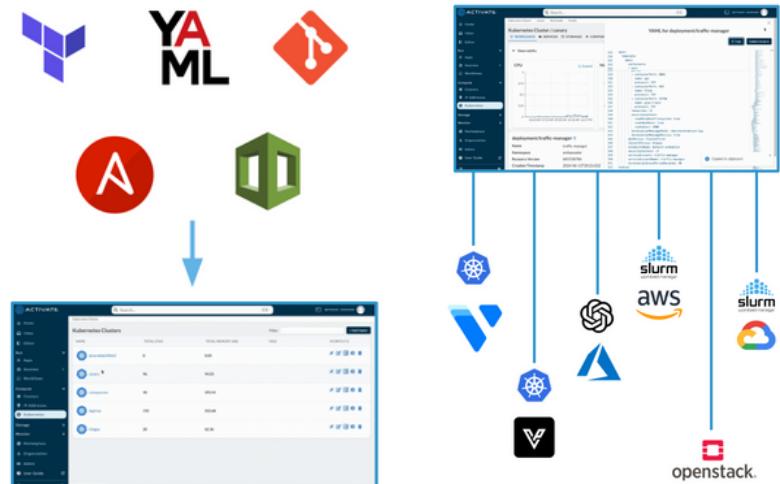
Technical Features

Kubernetes Support

Kubernetes is becoming essential for running modern enterprise, AI, and ML workloads, but it often requires specialized knowledge to use. ACTIVATE's new Kubernetes provider changes that. With just a few clicks, organizations can integrate their Kubernetes clusters into the same intuitive environment they already use to run remote desktop, HPC, or virtualized workloads.

AI Resource Integrations

ACTIVATE enables seamless provisioning and federated user access to AWS SageMaker, Azure Machine Learning, and Open-AI Compatible chat apps.



Flexible configurations

Unified resources

Neocloud Support

We've partnered with a diverse set of GPU providers to deliver high-performance, cost-effective, and globally available infrastructure for AI workloads. These partnerships enable vendor-neutral, flexible deployment options across public cloud, on-premises systems, and hybrid environments. Parallel Works does not resell neocloud compute cycles but instead provides reference architectures that let you bring your accounts.



parallelworks.com



info@parallelworks.com



Chicago, IL



Parallel Works ACTIVATE

ACTIVATE AI

Operationalizing AI Resource Adoption at Scale

Benefits of Kubernetes with ACTIVATE

AUTOMATED USER ACCESS	Plug in existing Kubernetes clusters, on-premises or in the cloud, and gain fine-grained control over resource allocation (GPU, CPU, RAM) across projects and teams. ACTIVATE identifies and gives permissions directly integrated into RoleBindings for Kubernetes user access. ACTIVATE Groups can be mapped to the Namespace-level isolation to ensure secure, team-specific environments.
DYNAMIC RESOURCE QUOTAS	Define compute, memory, GPU, and storage limits at the Kubernetes namespace-level. Then dynamically apply quota adjustments in real-time for distributed project teams. Fair-share, auto-scaling, and group-level policies will be available soon.
CHARGEBACK, SHOWBACK & BUDGET ENFORCEMENT	Assign usage-based pricing and track internal resource consumption across Kubernetes clusters for accurate budgeting, optimization and accountability.
SEAMLESS FRACTIONALIZATION OFGPUS	Manage and allocate GPU resources across multiple users on both MIG-enabled and non-MIG GPUs, with dynamic partitioning powered by Juice Labs integration.
MULTI-INFRASTRUCTURE ORCHESTRATIONS & PERSISTENT SESSIONS	Seamlessly run and migrate workloads across Kubernetes, batch schedulers, OpenStack, VMware and virtual desktops, making ACTIVATE AI a true hybrid orchestrator.



Access Parallel Works ACTIVATE AI today!

Contact us at info@parallelworks.com to learn more or schedule a [live demo](#).



parallelworks.com



info@parallelworks.com



Chicago, IL