

ACTIVATE for research, education, and industry

ACTIVATE is a control plane that connects researchers and engineers to computing resources on premises, in the cloud, and at GPU providers under one identity. Parallel Works runs it as a service, or the customer runs it in its own cloud account or data center.



Commercial and academic

Cloud, on premises, and GPU providers

Argonne spin-out, 2015

HOW ACTIVATE WORKS

One login to compute, data, workflows, and AI



WORKFLOWS

One definition runs on any connected resource

- Form:** published workflows launch from a generated input form.
- Definition:** steps, inputs, and resources in one YAML file.
- Graph:** the run graph shows each step and its dependencies.
- Marketplace:** a catalog of workflows to run, fork, and adapt.

WAYS TO DEPLOY

Same code base, hardened containers, and base images in every deployment

Shared platform

An organization on the Parallel Works platform.

Operated by Parallel Works

Dedicated cloud

A deployment in the customer's own cloud account.

AWS, Azure, Google Cloud, Oracle

On premises

A deployment on the customer's hardware or Kubernetes cluster.

Customer data center

One GPU reservation, shared by many teams

Parallel Works runs a dedicated GPU reservation as one cluster with per-team accounts and quotas.

Researchers and engineers

Browser, CLI and API, and AI assistants under one login

ACTIVATE

Group sign-on, Slurm accounts and QOS, Kubernetes namespaces, usage by group

Software environment

NVIDIA drivers, CUDA, NCCL tuned to the fabric, Apptainer and Docker

GPU reservation

GPU nodes, InfiniBand fabric, storage, and the data center

Parallel Works


User-experience SLA, 99.5 percent monthly, and first-line support

GPU provider

Infrastructure SLA, passed through

Workloads move between providers under the same login, group mapping, Slurm accounts, and storage paths, with no re-onboarding. In production, dedicated H100 clusters have opened to users within a week of node handover.

COST AND ALLOCATIONS



- Groups:** filter spend by user, group, or project.
- Daily spend:** compute, storage, and AI usage by day.
- Allocations:** budgets with thresholds that notify, freeze, or stop.
- Reports:** saved and printed views for showback and chargeback.

UNIVERSITIES AND RESEARCH CENTERS

- Campus single sign-on through SAML or OIDC, with groups mapped to Slurm accounts and storage.
- Burst from campus clusters to the cloud under the same accounts.
- Notebooks, remote desktops, and VS Code sessions for courses and labs.
- Department and project chargeback from one usage record.

LICENSING

Teams up to about 100 users license per user on the shared platform. Larger programs take a program license per deployment, with unlimited users and connected resources and a Parallel Works mission engineering team. Academic institutions pay 50 percent of commercial list.

Prices at parallelworks.com/pricing

SALES POC

Phyllis Rhodes, Director of Sales
prhodes@parallelworks.com

TECHNICAL POC

Matthew Shaxted, CEO
shaxted@parallelworks.com

SEE IT RUN

Schedule a demo
parallelworks.com/contact/demo