MR.PLANB FIELD GUIDE · AI PLATFORMS

OpenShift AI vs SUSE AI: Comparison Sheet

Ten categories side by side, who each platform fits, and what to weigh before choosing.

Updated **2026-10-29** · Full guide: www.mrplanb.com/comparisons/openshift-ai-vs-suse-ai

> Both are Kubernetes-native platforms for deploying and scaling AI and ML workloads. OpenShift AI is tied to the Red Hat ecosystem, with enterprise support and a steeper learning curve. SUSE AI is more open, runs on any Kubernetes distribution and is easier to start with. Both can run production AI/ML at scale.

10
categories compared

24/7
Red Hat enterprise support for OpenShift AI

Any
Kubernetes distribution SUSE AI can run on

2
platforms, both Kubernetes-native

Feature comparison

Category	OpenShift AI	SUSE AI
Licensing and cost	Subscription, enterprise pricing	Open source, flexible licensing
Ease of deployment	Complex, needs expertise	Straightforward, Rancher integration
Model management	Built-in registry and versioning	Flexible, through KubeFlow
GPU support	NVIDIA GPU Operator integration	Native GPU scheduling with K8s
MLOps integration	Tight OpenShift integration	Works with any Kubernetes
Enterprise support	24/7 Red Hat support	Community plus optional SUSE support
Learning curve	Steep, needs Red Hat knowledge	Gentle, standard K8s concepts
Multi-cloud	Strong	Cloud-agnostic, runs anywhere
Community	Active Red Hat community	Growing open-source community
Vendor lock-in	Red Hat ecosystem dependency	Minimal, standards-based

Best for: OpenShift AI

- [ ] Organizations already invested in the Red Hat ecosystem.
- [ ] Enterprises requiring 24/7 support and SLAs.
- [ ] Teams needing tight integration with OpenShift features.
- [ ] Regulated industries with compliance requirements.

Best for: SUSE AI

- [ ] Cloud-agnostic deployments.
- [ ] Organizations seeking cost-effective AI platforms.
- [ ] Teams familiar with standard Kubernetes.

- [ ] Homelabs and edge computing scenarios.

Trade-offs to weigh

OpenShift AI costs

Higher licensing costs, a steeper learning curve and potential vendor lock-in.

SUSE AI costs

Smaller enterprise footprint, community support as the main option and a less prescriptive architecture.

Foundation

OpenShift AI requires OpenShift, which adds complexity but brings built-in networking, security and monitoring.

Flexibility

SUSE AI runs on K3s, RKE2 and public cloud managed Kubernetes, with Rancher for multi-cluster management.

Model management and GPUs

- OpenShift AI: built-in model registry with versioning, lineage tracking and automated deployment pipelines.
- OpenShift AI: integrates with OpenShift Pipelines (Tekton) for ML CI/CD.
- SUSE AI: uses KubeFlow, a modular approach where teams pick their MLOps stack.
- GPUs: OpenShift AI uses the NVIDIA GPU Operator; SUSE AI uses native Kubernetes GPU scheduling and device plugins.
- Both support MPI operators for multi-GPU and distributed training.

Sources and notes: Red Hat OpenShift AI documentation and support matrix, as listed on the page · SUSE AI platform overview and KubeFlow guide, as listed on the page