



Spectro Cloud for MSPs

MSPs, CSPs and sovereign providers: here's how to scale your AI infrastructure practice and drive operational leverage with the most advanced management platform on the market.

Are you ready for the new reality?

Your customers' infrastructure is changing faster than it has in a decade. They're moving off VMware, standing up AI factories, asking for sovereign and regulated-ready environments, and pushing workloads out to the edge.

Every one of those shifts is a managed-services opportunity, and every one of them lands on your team to deliver, support, and stand behind. How can you seize the opportunity profitably, at scale?

PaletteAI is Spectro Cloud's platform for managing full-stack Kubernetes, VMs, and AI workloads across cloud, data center, edge, and air-gapped environments.

For an MSP like you, it does two jobs at once: it's the foundation for innovating the services you sell, and the tooling that lets you run those services at a profit.

New demand, new markets to serve

With our powerful white-labeled PaletteAI platform, MSPs, CSPs, and sovereign cloud providers can scale out several distinct lines of business to serve enterprises and public sector organizations:



Managed Kubernetes

Standardized, multi-tenant clusters across any cloud, data center, or bare metal, on whatever distro and stack your customers prefer.



VM app modernization

A managed, low-risk way to reduce dependency on vSphere, wrapped in your SLA and your commercial model.



Edge services

Thousands of nodes provisioned and managed with zero-touch onboarding and zero-downtime upgrades, even over unreliable connections.



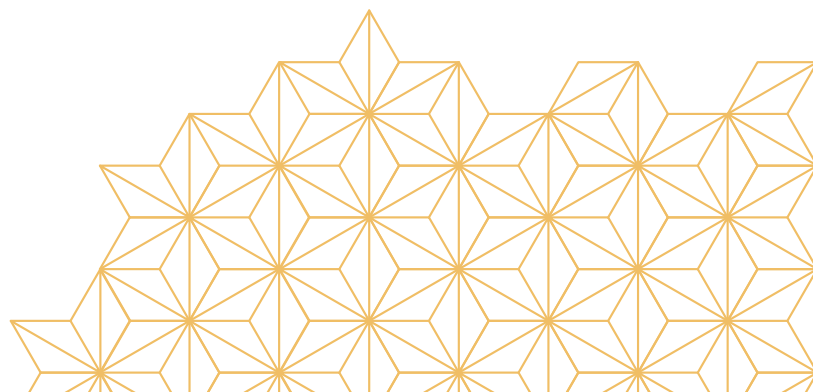
Managed AI and GPU services

GPU-as-a-service, Model-as-a-Service, and multi-tenant AI infrastructure delivered with PaletteAI, so your customers get enterprise AI without building the platform themselves.



Sovereign and regulated cloud

US- or EU-hosted, air-gap-capable, FIPS and FedRAMP-aligned environments for regulated industries that need more than the hyperscalers can offer.



The engine for better business

As an MSP, you know that the tooling you invest in can have a big impact not just on the services you can offer, but on the economics of your own business. That's where PaletteAI earns its place.



More margin from the same infrastructure

Automated provisioning, upgrades, and day-2 operations mean far fewer engineer-hours per cluster. Running PaletteAI at scale, T-Mobile's platform team took its engineer-to-cluster ratio from 10:1 to 100:1. That ratio is the difference between a thin managed-services margin and a healthy one.



Speed that shortens time to revenue.

Repeatable, profile-driven deployments let you stand up environments in hours rather than weeks. The sooner you can deliver, the sooner your investment starts paying you back. One federal technology provider [took a classified AI factory from loading dock to live in 30 days with PaletteAI.](#)



Consistency your customers can feel

Because every cluster is built from a standardized profile, upgrades are predictable and troubleshooting is faster. Consistent delivery keeps SLA performance up and earns renewals.



Less risk for you and your customers

PaletteAI provides strong security controls designed for multitenancy, with all the compliance badges to back it up, making your next customer infosec review that bit easier.



More from the people you can't easily hire

K8s and AI infrastructure talent is scarce and expensive. PaletteAI lets your specialists automate the repetitive work and spend their time on the engineering that sets your services apart.

Why MSPs choose us

We've built our company around partners, and today we have a rich ecosystem of tech alliances with the likes of NVIDIA, Everpure and Supermicro — so chances are we already work with the vendors in your stack.

PaletteAI is built in full conformance with Kubernetes standards and stays open across OS, distro, and tooling: we dislike lock-in as much as your customers do.

We're certified to ISO 27001, SOC 2 Type 2, UK Cyber Essentials Plus, and FIPS 140-3, with FedRAMP Low (Moderate in process) and EU-hosted and air-gapped options for regulated work. Our pricing is simple and transparent, so it scales with the value you create.

And most importantly, we're proven to deliver. We're already trusted by sovereign clouds and MSPs in Europe and the US, as well as some of the world's largest organizations like GE HealthCare, T-Mobile, Yum! Brands and the US Air Force.

Start where it makes sense for you

A platform decision doesn't get made in a single demo, and we wouldn't ask you to. Evaluate PaletteAI at your own pace: scope a first use case with our MSP team, run a proof of concept against your real environment, and lean on us for the architecture and commercial modeling along the way. When you're ready to build, our partner program is set up to support your go-to-market, not just your onboarding.

Talk to your Spectro Cloud partner representative or distributor, or start a conversation at spectrocloud.com/partners.

