



Self-Hosted AI Platform Deployment Support

Support for AI Vendors Deploying Into Self-Hosted Kubernetes Environments

THE CHALLENGE

Self-hosted deployment often becomes part of the buying motion when enterprise customers want AI platforms running in their own Kubernetes environments. Even with Helm charts, installation guidance, and customer documentation in place, deployments can still be hard to scale because every environment brings its own prerequisites, security requirements, cloud configuration, and operational constraints.

This shows up as:

- Long install calls and repeated troubleshooting sessions.
- Too much time spent by product, engineering, and customer-facing teams on infrastructure work adjacent to the software.
- Long cycles between initial purchase and production-ready deployment.
- Missing or incomplete customer prerequisites.
- Deployment challenges in air-gapped or highly restricted environments.
- Protracted, environment-specific security, networking, and database reviews.

WHAT FAIRWINDS PROVIDES

Fairwinds helps AI software vendors reduce the operational friction around self-hosted customer deployments. Depending on what your team needs, Fairwinds can:

Audit and improve the install process.

Support your internal team during active deployments.

Work directly with your customers during installation.

Across these engagements, Fairwinds helps teams:

- Tighten prerequisites before deployment begins.
- Document blockers and escalation points.
- Improve runbooks and deployment guidance.
- Work through environment-specific issues that slow installs down.



SERVICE TIERS

SILVER

Deployment Auditing and Runbook Development

Best for teams that need to make self-hosted installs more repeatable before expanding customer-facing deployment support.

- Proactive documentation refinement and install validation.
- Project-based, fixed-scope engagement.
- Audited installation guide, patched manifests, and best practices report.
- No live install support.
- No direct customer interaction.
- Milestone-based delivery schedule.

GOLD

Rapid-Response Deployment Support

Best for teams that already own the customer relationship during installation but need Kubernetes-specific escalation support when environment issues appear.

- Reactive, on-demand support for infrastructure questions and deployment blockers.
- Ongoing deployment support during active customer installs.
- Dedicated communications channel, verified solutions, and architectural answers.
- Live install support through your team.
- No direct customer interaction.
- One business-hour response for urgent deployment blockages.

PLATINUM

White-Glove Installation Support

Best for teams that want Fairwinds directly involved in customer deployments for higher-touch or more complex accounts.

- Direct white-glove installation support for self-hosted environments.
- Installation-based engagement per customer deployment.
- Dedicated communications channel with each customer to support installation.
- Live, customer-facing deployment support.
- Direct customer interaction.
- One business-hour response for urgent deployment blockages.

HOW THE TIERS DIFFER

CAPABILITY	SILVER	GOLD	PLATINUM
Primary Objective	Audit and improve the install process.	Support your team during active installs.	Work directly with your customer during active installs.
Who Fairwinds Works With	Your internal team.	Your internal team.	Your team and your customer.
Live Install Support	No.	Through your team.	Customer-facing.
Direct Customer Interaction	No.	No.	Yes.

CUSTOMER OUTCOMES



Streamlined onboarding so you can start capturing revenue immediately.



More consistent support during active self-hosted deployments.



Better documentation, runbooks, and operational handoffs over time.



Clearer pre-install requirements and deployment readiness.



Faster resolution of infrastructure blockers in customer environments.



Less time spent by internal teams translating between customer infrastructure teams and product engineering.

WHY FAIRWINDS

Fairwinds is strongest where a software vendor already has a product, deployment assets, and customer demand for self-hosted installs, but the delivery motion is manual, variable, or difficult to scale. AI platforms deployed in self-hosted Kubernetes environments need a partner with the Kubernetes expertise needed to handle messy real-world infrastructure, help customers work through blockers, and reduce the communication burden on internal teams.

Fairwinds supports customer environments with a wide range of security controls, prerequisites, and infrastructure patterns, including cases that don't fit neatly into product documentation. That support gives teams a clearer way to move from ad hoc install work to a more structured model with documented requirements, clearer expectations, and better operational handoffs.

NEXT STEPS

To discuss which service tier fits your deployment model and what a pilot could look like, [contact Fairwinds](#).