

OPSWERKS CASE STUDY

AI-READY INFRASTRUCTURE



Resilient AI Data Platforms at Petabyte Scale

How OpsWerks helps a Fortune 100 technology company keep the data platform behind its AI roadmap resilient.

Platform scale

Hundreds to a thousand Kubernetes clusters across major public and hybrid clouds

Petabyte-scale data

The platforms OpsWerks keeps in service process petabytes of data

Deadlines met

OpsWerks upgraded hundreds of Kubernetes clusters in a single month

Challenge

One Fortune 100 technology company's data platform for AI reached an inflection point: converge a legacy stack and a new stack, or maintain the legacy while the new one gets built. The platform's own developers and SREs were juggling user support for both stacks while doing the engineering work of the transition.

The scale left no slack. Staying in service meant operating hundreds to a thousand Kubernetes clusters across public and hybrid cloud, processing petabytes of data for the engineers who build on them.

Compliance added a deadline. Both stacks needed current Kubernetes versions, vulnerability fixes, and deprecation and migration work. Falling behind means paying extra for extended support, hitting compatibility issues, and in the worst case, taking service disruptions.

The scale and breadth of the platform demanded diverse skills, a significant number of engineers, and tight deadlines. The support bar was high: some of the platform's users are primary contributors to the major open source data projects it runs.

About the platform

An internal platform offering data engineering applications and cloud services to data engineers, scientists, and developers across the organization.

Technologies:

OpsWerks Ready for the Challenge

OpsWerks has owned 24/7 operations for the world's most demanding enterprise teams running mission-critical workloads since 2015. More than 250 engineers across the US and Philippines cover US, EMEA, and APAC around the clock.

This engagement demands breadth, because no data tool runs in isolation. OpsWerks brings data platform management, infrastructure management and optimization, and CI/CD for data and AI workflows in one team, with technical leaders who step in when complexity spikes.

OpsWerks Approach

OpsWerks runs operations for the platform's data engineering applications and the infrastructure hosting them. The customer's platform developers build the platform and own its roadmap; its SRE team owns the reliability targets; OpsWerks owns the operations, support, and upgrade work underneath.

From access setup and initial syncs to autonomous operations in a month.

Support where the users are OpsWerks answers platform users directly in Slack, correlates a user's issue to existing alerts, sizes the blast radius, and routes the incident to the owning team. The complexity varies by application; the team interfaces with the developers of each one.

Coverage matched to need The engagement began with 24/7 PagerDuty response and shifted to 24/5 direct user support in Slack. Teams that still need round-the-clock on-call keep it.

Upgrades as routine, not fire drills For more than two years OpsWerks has run the platform's release deployments, RDS upgrades, and managed Kubernetes cluster upgrades from version 1.24 through 1.32.

A faster upgrade pipeline OpsWerks streamlined the process for updating Kubernetes clusters, cutting the time to prepare for and execute each regular upgrade.

Deprecations raised early The team flags deprecation concerns ahead of deadlines, so planning and scheduling happen on time and maintenance lands as a better experience for the platform's users.

The team that was trained now does the training

The customer's developers and SREs trained only the initial OpsWerks onboarding team. OpsWerks now trains every new member itself, and it trains the customer's new SREs too.

Results

PLATFORM SCALE

Hundreds to a thousand
Kubernetes clusters across
major public and hybrid
clouds

PETABYTE-SCALE DATA

The platforms OpsWerks
keeps in service process
petabytes of data

DEADLINES MET

OpsWerks upgraded
hundreds of Kubernetes
clusters in a single month

RECURRING UPGRADES

More than two years of
managed Kubernetes
upgrades, version 1.24
through 1.32, with
deprecated releases
brought up to date

FAST TO AUTONOMOUS

A month from onboarding
to running operations under
light supervision

Impact

OpsWerks operates as the reliability and operations engine for the customer's data platform infrastructure at massive scale. The work spans infrastructure, orchestration, and data visualization support, because no tool on the platform runs in isolation.

Raising deprecations early kept planning and scheduling on time, made maintenance a better experience for the platform's users, and saved the customer a significant amount of time and resources.

The customer's developers and SREs kept their focus on converging two stacks into one instead of splitting time between user support and transition work.

The platform's users, the data engineers and scientists building on it, stay focused on models, products, and insights. The discipline has kept pace as the work shifted from DevOps to DataOps to MLOps.

This is autonomous execution and outcome ownership in practice: OpsWerks runs with minimal supervision, escalates high-impact concerns on time, and raises process improvements as patterns emerge.

Trained once. Autonomous since. At platform scale.

Facing similar challenges?

Talk with the team that delivered this work. We will map your environment, your constraints, and what OpsWerks would own.

[Schedule a discovery call](#)

PARTNERWITHUS@OPSWERKS.COM · [OPSWERKS.COM](https://opswerks.com)

