gruve

SOLUTION BRIEF

FinOps cloud cost agent

End-to-end AP automation with command center for finance operations.

01

Business problem

Cloud costs are rising faster than almost any other operating expense, yet most teams still can’t see what they are actually spending across AWS, Azure, and Google Cloud. Fragmented tagging, inconsistent cost attribution, and manual reporting makes it hard to know where money is going or why. Engineering receives conflicting recommendations; Finance can’t trust forecasts, and FinOps struggles to turn insights into real savings. Kubernetes adds even more noise, with spend often impossible to tie back to services or owners. The result: budget overruns, delayed projects, and millions of missed savings across multi-cloud environments.

02

Why now

By 2025, a majority of large enterprises have adopted FinOps, yet multi-cloud visibility remains a top challenge, creating urgency for standardization and automation. The FinOps Open Cost and Usage Specification (FOCUS) gained broad cloud provider support in 2024–2025, enabling consistent cross-cloud reporting. AI agents announced at recent FinOps events to signal a market shift toward proactive, automated cost governance. Organizations implementing FinOps practices typically realize meaningful first-year cost reductions, strengthening the CFO mandate to act now.

03

Solution overview

Gruve's FinOps AI Agent delivers unified cost intelligence and autonomous optimization across AWS, Azure, and Google Cloud. The agent ingests provider billing and usage data, normalizes it using FOCUS, and applies explainable AI to detect anomalies, forecast spend, and recommend actions for rightsizing, reservations, and waste elimination. It orchestrates human-in-the-loop workflows via Jira, Slack, and ITSM tools, turning insights into tracked savings with governance controls. Finance gains accurate forecasting, Engineering receives actionable, risk-scored changes, and FinOps operationalizes continuous optimization with measurable outcomes.

Gruve solution items

- **Unified Cost Lake (FOCUS-Native):** A single source of truth for AWS, Azure, GCP, and Kubernetes to spend consistent allocation, chargeback, and team ownership.
- **Real-Time Anomaly Detection:** AI identifies spending spikes within minutes, explains root causes, and quantifies financial impact.
- **Continuous Optimization and Planning:** Rightsizing, idle cleanup, reservation modeling, and spot coverage, supported by risk scoring and policy-based automation.
- **Engineering Workflow Integration:** Recommendations delivered as Jira stories and Slack alerts with clear impact, priority, and rollback guidance—driving real engineering adoption .
- **Forecasting and Reporting:** Rolling forecasts, variance explanations, KPIs, and CFO-ready dashboards for predictable budgets and financial confidence.

Solution benefits

- 15–30% cost reduction in 90 days through rightsizing, idle cleanup, and smarter commitments.
- 100x faster anomaly detection—from days to minutes.
- Reliable budgets and forecasts, powered by a FOCUS-standardized, cross-cloud cost lake.
- 70% higher engineering adoption with Jira/Slack workflows and explainable recommendations.
- Consistent multi-cloud visibility across AWS, Azure, GCP, and Kubernetes from a single source of truth.

04

Solution benefits

- 15–30% cost reduction in 90 days through rightsizing, idle cleanup, and smarter commitments.
- 100x faster anomaly detection—from days to minutes.
- Reliable budgets and forecasts, powered by a FOCUS-standardized, cross-cloud cost lake.
- 70% higher engineering adoption with Jira/Slack workflows and explainable recommendations.
- Consistent multi-cloud visibility across AWS, Azure, GCP, and Kubernetes from a single source of truth.

05

Competitors & differentiation

Competitors

- **CloudZero:** Granular cost intelligence and unit economics across clouds. Differentiation: Gruve adds agentic automation with embedded workflows to close the loop from insight to action.
- **Apptio Cloudability:** Mature allocation and forecasting for enterprises. Differentiation: Gruve emphasizes explainable AI recommendations and faster time-to-value with prebuilt Jira/Slack playbooks.
- **Finout:** Strong Kubernetes and multi-cloud allocation. Differentiation: Gruve unifies allocation with autonomous commitment planning and anomaly triage in one agent.
- **Kubecost:** Open-source Kubernetes cost tracking. Differentiation: Gruve extends beyond K8s to full multi-cloud optimization with governance and approvals.
- **CloudBolt (Augmented FinOps):** AI-enabled lifecycle management. Differentiation: Gruve is FOCUS-native and tuned for human-in-the-loop change management across engineering tea

Gruve differentiation

Prebuilt AP agents for 3-way match and spend reclassification, fast time-to-value, and a governance-first command center that unifies exceptions, evidence, and SLAs across multi-ERP estates. Human-in-the-loop and full lineage provide audit-ready confidence

06

Case study

Workflow before & after

Before:
- Cloud spend spread across AWS, Azure, GCP, and Kubernetes with inconsistent allocation and no single source of truth.
- Anomalies discovered days or weeks late, leading to budget overruns and unplanned spend.
- Engineering receives conflicting recommendations with low trust and low adoption.
- Commitment planning is manual and risky, causing under-coverage or wasted spend.
- Finance, Engineering, and FinOps operate with different data and different tools.

After:
- All cloud and Kubernetes spend is normalized into a single FOCUS (FinOps Open Cost and Usage Specification)-standardized cost lake.
- AI detects spend spikes within minutes and quantifies financial impact.
- Rightsizing and reservation actions come with explainability, confidence scoring, and business impact
- Jira stories are created automatically with ownership and playbooks
- Finance receives accurate, scenario-based forecasts aligned to business goals.

At-a-glance results

- 15–30% cost reduction in 90 days
- Anomalies detected 100x faster (minutes → days)
- Reliable, predictable budgets powered by a unified focus-standardized cost lake

15-30%

reduction in cost in 90 days

About the client

North America–based SaaS company operating across AWS and Azure with Kubernetes workloads and rapidly scaling data services.

Challenges

- Inconsistent tagging and allocation
- Delayed anomaly detection
- Underutilized commitments
- Limited engineering adoption of cost actions

Solutions

- Focus-standardized cost lake
- Agentic anomaly detection
- Automated commitment modeling
- K8s cost allocation
- Jira/Slack workflow automation
- Executive dashboards.

Results & benefits

Sustained cost savings, improved coverage and utilization, reduced variance to budget, and higher engineering participation through explainable, low-friction workflows.

About Gruve

Gruve partners with leading enterprises to transform data into measurable business impact. Our team brings deep expertise in enterprise data architecture, AI and analytics strategy, cloud modernization, and organizational change. We combine technical rigor with business acumen, ensuring recommendations are both architecturally sound and executable within your organizational constraints. With proven success across financial services, healthcare, manufacturing, and technology sectors, Gruve delivers data and AI solutions that drive growth, efficiency, and competitive advantage.

Contact Gruve to operationalize FinOps with AI-driven, FOCUS-native multi-cloud cost management.

Email: sales@gruve.ai
Web: www.gruve.ai