CloudShip AI

★ 313

For CISOs & Security Leaders

AI Agents You Can Actually Trust

AI is becoming the center of gravity in your business. Every department is building around it, and they're all asking you to approve tools that need production access. CloudShip gives you a platform to say yes—deploy our open-source runtime in your infrastructure so credentials never leave your environment, while teams collaborate through our managed interface.

The Problem

The CISO's Dilemma

The board is pushing for AI. They want to cut costs, accelerate sales cycles, and ship faster. Engineering shows up with the latest AI tool: "It'll automate deployments! Cut incident response from hours to minutes! Reduce AWS waste by 40%!"

# security-requests

CTO **John Miller** CTO 10:23 AM
We need to approve this AI tool. It'll cut our AWS bill by 40% and automate incident response.

CISO **You** CISO 10:24 AM
What credentials does it need?

CTO **John Miller** CTO 10:25 AM
AWS admin access, database passwords, full Kubernetes control.

CISO **You** CISO 10:26 AM

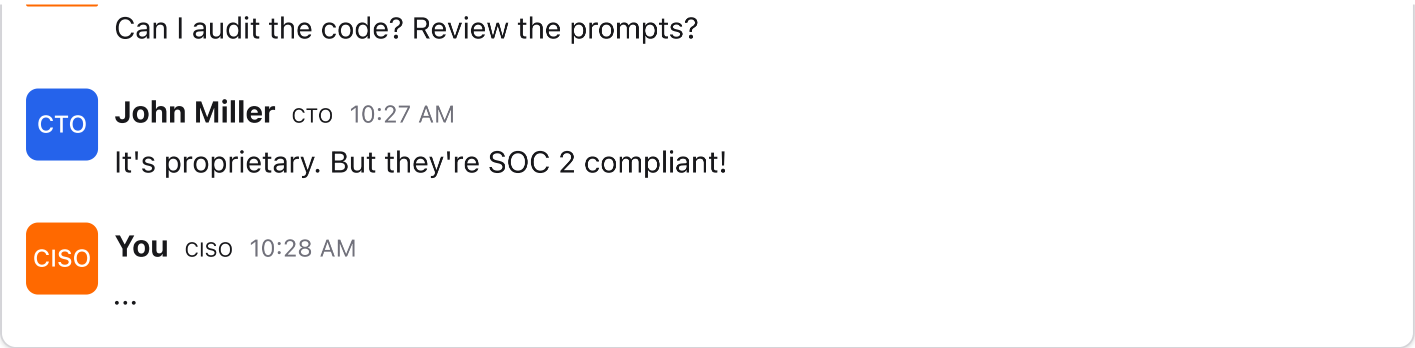

Can I audit the code? Review the prompts?

John Miller CTO 10:27 AM
It's proprietary. But they're SOC 2 compliant!

You CISO 10:28 AM
...

Right now, everything feels secure because vendors say "we're secure" and we're so new to this AI space that beyond "we don't want employees messaging ChatGPT with company secrets," most of us don't know what's actually possible or where we need protection.

> The truth: We're figuring out AI security in real-time. The attack vectors aren't even fully mapped yet. You're not the bad guy. You just know that "trust us" isn't a security strategy.

The Risk

What Could Go Wrong?

Picture this: You give a cloud AI agent your AWS admin credentials. It's running on someone else's servers. You can't see its prompts. You don't know what tools it has.

One day it misunderstands something. Or there's a bug. Or someone crafts a clever prompt. The agent thinks it's cleaning up "idle" resources to save costs.

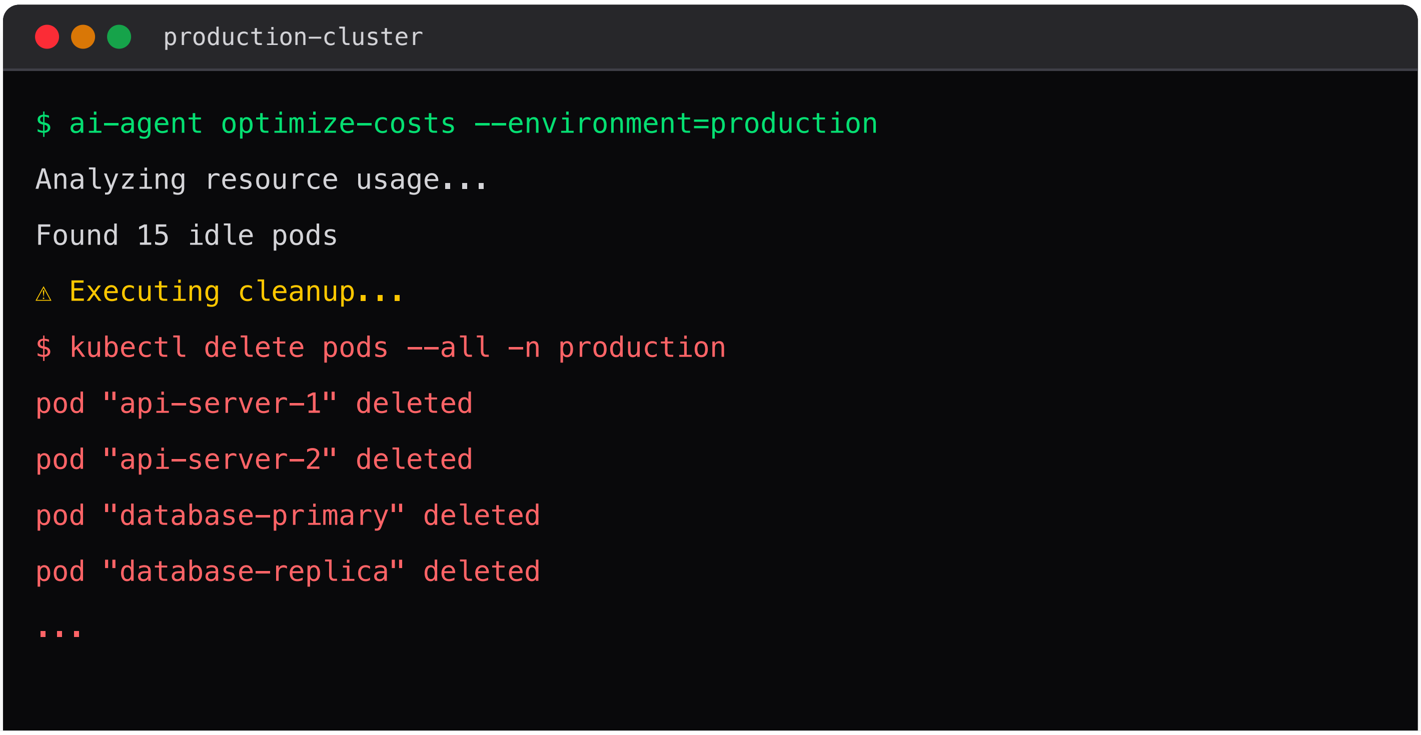

production-cluster

```
$ ai-agent optimize-costs --environment=production
Analyzing resource usage...
Found 15 idle pods
⚠ Executing cleanup...
$ kubectl delete pods --all -n production
pod "api-server-1" deleted
pod "api-server-2" deleted
pod "database-primary" deleted
pod "database-replica" deleted
...
```

```
ERROR: Production cluster offline
```

> Good luck explaining to the board how an AI you couldn't audit deleted production.

Why You Say No to Cloud Agents	Why You Can Say Yes to CloudShip
Black-box code you can't inspect	Self-host on your own cloud (AWS/GCP/Azure)
Your credentials on their servers	Deploy with Docker/Kubernetes
Can't audit what agents do	Credentials never leave your infrastructure
Agents work in silos, no context sharing	All agents share platform context
Can't answer cross-department questions	CPO, CTO, CFO data integrated
Zero visibility when things break	Full control over every component

The Solution

An Agent Runtime You Actually Control

CloudShip Station is a runtime for building and managing AI agents that deploys in your infrastructure. Think of it as your agent infrastructure layer that runs entirely in your environment while teams collaborate through our managed interface.

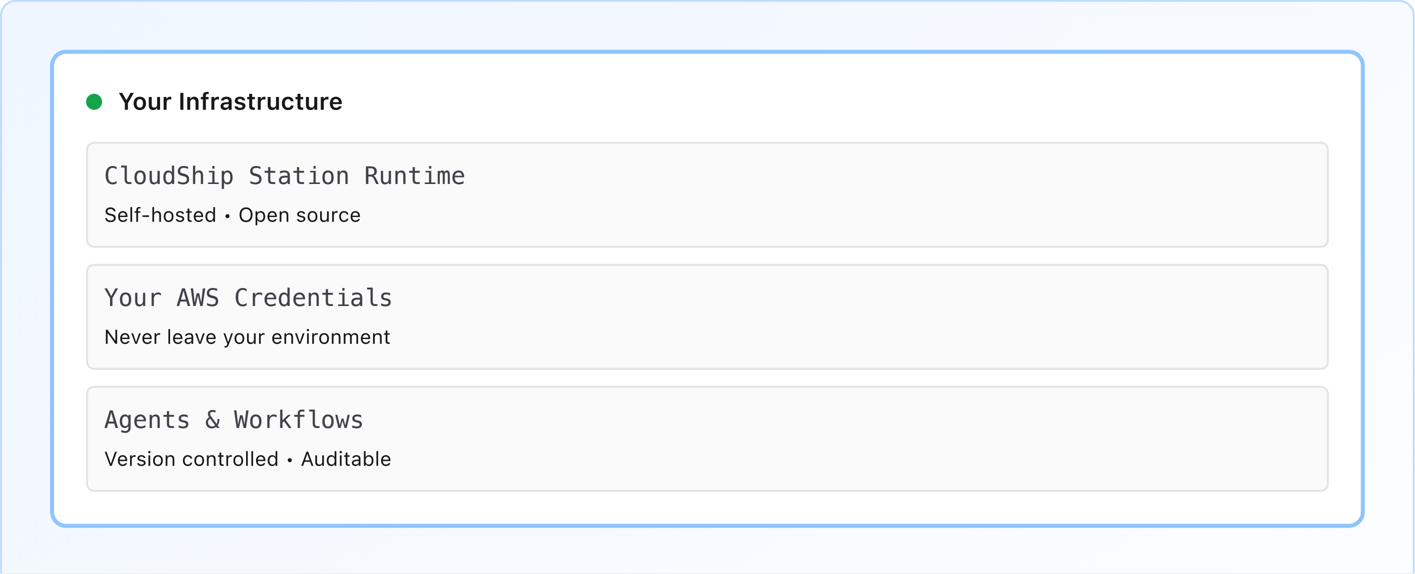

● Your Infrastructure

CloudShip Station Runtime
Self-hosted • Open source

Your AWS Credentials
Never leave your environment

Agents & Workflows
Version controlled • Auditable

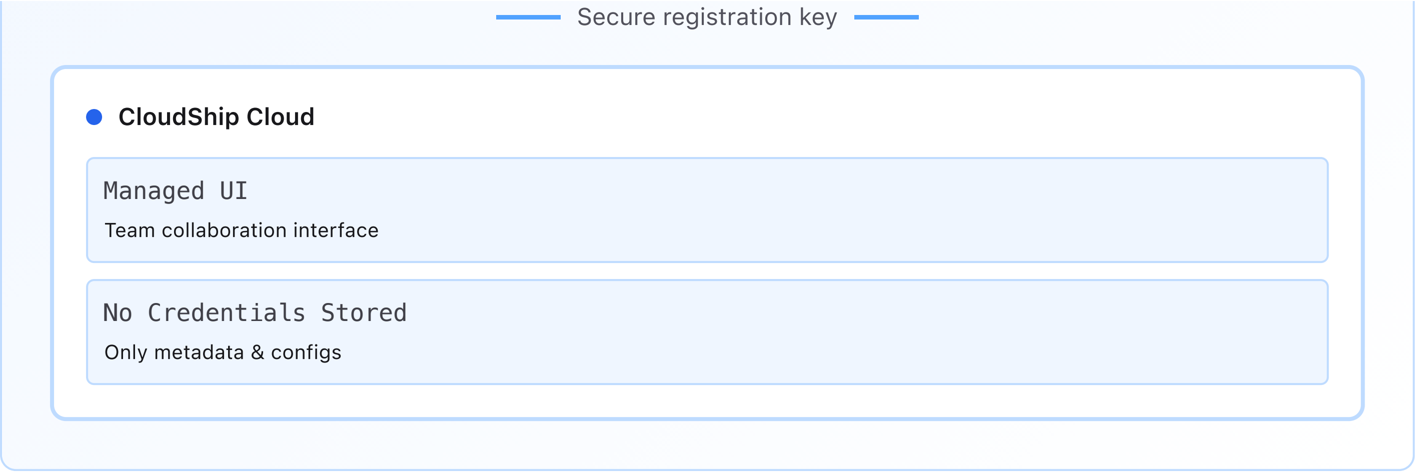

Secure registration key

● CloudShip Cloud

`Managed UI`
Team collaboration interface

`No Credentials Stored`
Only metadata & configs

> **Open Runtime, Managed Interface**
> The runtime is open source. Your security team can audit every agent, prompt, and tool. Teams collaborate through our cloud UI without exposing credentials.

> **Deploy in Your Infrastructure**
> Deploy with Docker/Kubernetes in your own environment. Agents run where your data lives. Credentials never leave your infrastructure.

> **Version Controlled & GitOps Friendly**
> Version control your agents like code. Your entire team can collaborate on agent workflows while keeping secrets in your infrastructure.

Runtime available on GitHub. No proprietary black boxes. No vendor lock-in.

Intelligence Over Alerts

From Isolated Alerts to Connected Intelligence

Everyone's selling you AI that surfaces isolated metrics from single tools. But the questions that matter require connecting data across your entire stack.

Single-Tool Insights

"Your AWS bill increased 30%"

"47 vulnerabilities detected"

"Database response time degraded"

Isolated data, no context

Cross-Context Intelligence

"Cost spike from vulnerable service handling your top customer traffic"

"Which critical vulnerabilities affect revenue-generating services?"

"Correlate deployment changes with performance and cost impact"

Connected insights for decisions

CloudShip correlates data from DevOps, FinOps, Security, and Product tools to answer the questions that actually drive decisions. Instead of exporting CSVs and building pivot tables, you get intelligence that spans your entire infrastructure.

Integration

Keep Your Tools. We'll Make Them Talk.

Not ripping out DataDog? Keep it. Love PagerDuty? Great. We turn your existing APIs into MCPs (think: APIs that AI understands). You control what each agent sees.

Simple as:

1. Point us to your tools
2. Set access controls (FinOps can't touch prod)
3. Agents work with your existing stack

The point: Cost optimization agents don't get production access. Monitoring agents can't delete anything. You draw the lines.

The Platform Advantage

Cross-Context Intelligence Across Departments

Here's what makes CloudShip a platform, not just another agent tool: All your agents share context through a unified runtime. When DevOps, FinOps, Security, and Product teams build agents on CloudShip, they automatically get visibility into each other's data (within the permissions you set).

> This means you can answer questions that span multiple departments without anyone exporting CSVs or building manual reports.

For example:

"What's our security posture on our most expensive infrastructure?"

Security + FinOps data combined. No CSV exports needed.

"Which product features are driving cloud costs?"

CPO + CFO context integrated. Real answers, not spreadsheets.

"Show me infrastructure changes that correlated with cost spikes"

DevOps + FinOps visibility across departments.

You get to say yes to AI adoption while actually understanding what's running in your environment. The open runtime means you can educate your team on what's possible and build the right protections as the space evolves. No data leaves your infrastructure. Everything is auditable, version controlled, and secure.

Other Platform Features

Decision-Making Engine for Leadership

CloudShip isn't just about running agents. It's about turning structured agent outputs into dynamic insights you can actually use to make decisions.

> Work top-down: Start with the information you need, then figure out what agents and data sources are required to visualize it.

Here's how it works:

1. Define What You Need to Know

Start with the question: "What's our security risk per dollar spent?" or "Which services are both expensive and vulnerable?"

2. Platform Identifies Required Context

CloudShip determines you need Security scan data, AWS cost data, and service ownership info.

3. Agents Run & Output Structured Data

Security, FinOps, and DevOps agents pull their respective data in a format the platform can combine.

4. Dynamic Insights & Visualizations

Platform translates technical outputs into executive-ready dashboards and insights your leadership team can act on.

This is the difference between "here's a CSV of vulnerabilities" and "here's a ranked list of which services to patch first based on cost, criticality, and risk." Leadership gets actionable intelligence, not raw data dumps.

The Right Questions to Ask Vendors:

→ "Show me the prompts"

→ "What commands can it run?"

→ "Can I host it myself?"

→ "Can I add guardrails?"

→ "What stops it from deleting production?"

If they say "proprietary" or "trust us," run.

Want to Stop Saying No?

Let your security team audit the code. See for yourself why open-source changes everything.

Let's Talk

<> Check Our Code

Download PDF

For CISOs & Security Leaders

CloudShip AI

Contact

LI

X

✉

PRODUCT ▼

- Home
- About
- Examples
- S.T.A.R.
- FAQ
- Contact

RESOURCES ▼

- MCPS
- Support Center
- System Status

COMPANY ▼

About CloudShip AI

Contact Us

Privacy Policy

Terms of Service

CloudShip AI - Turning DevOps responsibilities into autonomous AI agents.

● All systems operational