

CLIENT PROFILE

MeridianWealth

A Series C wealth management SaaS serving independent RIAs — portfolio reporting, client billing, rebalancing, and a client portal used by advisors at 200+ firms. Strong API foundation, compliance-conscious engineering team, and a CTO who had been fielding questions about AI agent readiness for six months without a clear answer.

9 days

Zero to production CLI

\$0

Engineers diverted

SOC 2

Compliant from day one

VERTICAL

Wealth Management
SaaS

COMPANY SIZE

~140 employees ·
Series C

ENGINEERING
TEAM

22 engineers

COMPLIANCE REQUIREMENTS

SOC 2 Type II · SEC Rule 17a-4 ·
FINRA

ENGAGEMENT TYPE

Growth Plan · \$3.5K/mo
retainer

⚡ THE SITUATION

MeridianWealth's advisors were increasingly hearing from their own clients: "Can I just ask my AI assistant to pull my portfolio?" The company's API existed — but it was built for internal consumption. Auth was session-based, outputs were unstructured HTML fragments mixed with JSON, and there was no rate-guarding or audit trail. An AI agent hitting it would either fail silently, expose sensitive financial data unpredictably, or exhaust rate limits without any visibility.

"We knew we needed to be agent-ready. We just didn't want to pull two engineers off our rebalancing engine to build plumbing. That felt like the wrong trade."

— CTO, MeridianWealth (composite)

✗ BEFORE AGENTARMOR

⚠ Agents failed on first call

Session-based auth rejected AI agent requests. No API key path, no OAuth flow designed for machine clients.

⚠ No audit trail for AI access

SEC Rule 17a-4 requires records of data access. An agent calling the API left zero logs — a compliance gap they couldn't ship to production with.

⚠ Unstructured outputs

Portfolio data returned as mixed HTML/JSON. Agents couldn't parse it reliably — and hallucinated figures when they tried.

⚠ No rate protection

A single agent in a loop could exhaust API capacity — with no circuit-breaker, no alerting, and no visibility until something broke.

★ THE BUSINESS IMPACT

"We went from 'we're figuring out our AI strategy' to demoing a live agent pulling portfolio data in a prospect meeting — in under two weeks. Two enterprise prospects cited agent-readiness as a differentiator. Compliance signed off without a single revision request."

— VP Product, MeridianWealth (composite)

→ HOW IT HAPPENED

1

DAY 1 — AGENT READINESS AUDIT

We mapped every gap before writing a line of code

Async API review identified 4 blockers: auth architecture, output format, PII exposure, and missing audit trail. Written build plan delivered same day.

2

DAYS 2-5 — POC BUILD

Live agent demo against their actual API

MCP-powered CLI covering portfolio queries, rebalancing status, and billing. Structured JSON, OAuth2 machine-client auth. Live demo in second meeting.

3

DAYS 6-9 — PRODUCTION HARDENING

Compliance-ready, not just functional

Immutable audit logging (SEC 17a-4), PII scrubbing, rate guards with Slack alerts, and a monitoring dashboard. Compliance team signed off on first review.

✓

DAY 9 — LIVE

Announced publicly. Compliance signed off.

CLI announced publicly. AgentArmor retains ongoing maintenance — API changes, schema updates, 8-hr SLA. Zero engineers diverted.

>

WHAT AN AGENT CAN NOW DO

```
meridian-cli · agent session
$ meridian portfolio get --client clt_4821 --as-of today
{
  "client_id": "clt_4821",
  "total_value": 1284503.22,
  "ytd_return": 0.0842,
  "drift_alert": "equities +3.2% over target",
  "audit_log_id": "log_9f3c12a"
}
# Typed. Structured. Logged. Compliance-ready.
```