

# Aperture: AWS cost intelligence *workflow-shaped*, not dashboard-shaped.

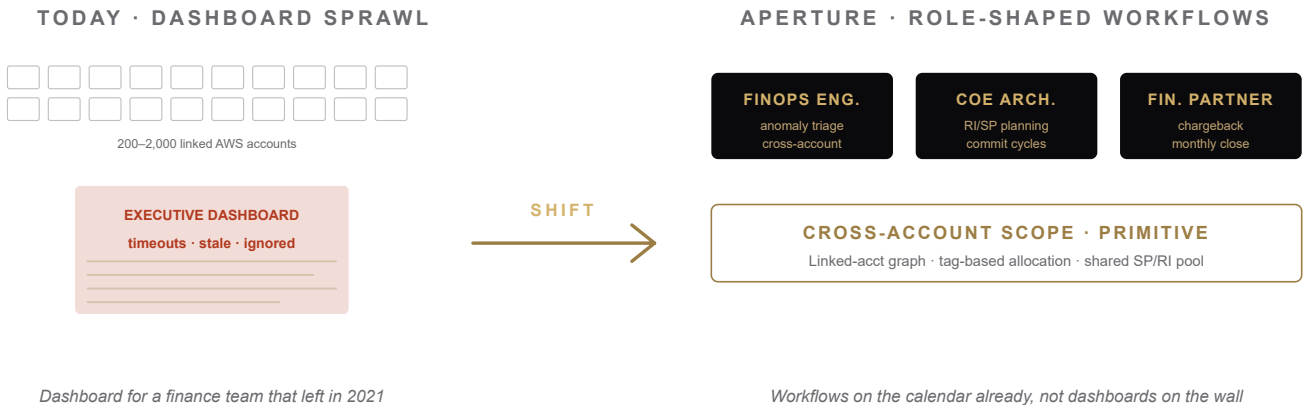
Fortune 500 enterprises waste 25–40% of a \$90M–\$400M annual AWS bill across 200–2,000 accounts. AWS-native tools collapse above ~100 accounts; vendor dashboards target a finance team that no longer exists. Aperture ships role-specific workflows, with cross-account scope as a primitive.

**Author** S. Ize-Iyamu    **Audience** AWS Enterprise + Cloud FinOps PMs    **Length** 3 pages    **Status** Concept  
**Targets** AWS · Azure · GCP

## The Problem

F500 enterprises spend \$90M–\$400M / yr on AWS across 200–2,000 linked accounts. Cost Explorer was built for one account, bolted to many; above ~100 accounts the UI times out and consolidation becomes a quarterly Python job. Existing FinOps vendors built executive dashboards for a finance team that no longer runs day-to-day cloud cost. The actual users are a triad — **FinOps engineer** (anomaly triage), **Cloud CoE architect** (commitment planning), **finance partner** (chargeback close) — and none of them want a dashboard.

FIGURE 1 · WORKFLOW SHIFT



Dashboard for a finance team that left in 2021

Workflows on the calendar already, not dashboards on the wall

Today: 200–2,000 accounts feed an executive dashboard nobody opens. Aperture ships three role-specific workflows on top of cross-account scope as a primitive — not as an enterprise-tier upsell.

## Why this matters now

Three forces converge: **multi-account architectures hit critical mass** (F500 averages went from ~12 accounts in 2019 to 200+ in 2025), **committed-spend waste is real money** (\$1B+ in unused Savings Plans across the F500 in 2024 per public 10-Q disclosures), and **FinOps Foundation seats grew 4x** in two years. The first vendor to ship workflows the buyer already runs wins distribution.

## Sizing the prize

Bottom-up: ~500 F500 × ~\$180M avg AWS × 30% recoverable = ~\$27B / yr recoverable spend. At a 5–8% platform fee on validated savings (industry-norm for Cloudability, Apptio, Vantage), TAM is ~\$1.4–2.2B / yr. Add Azure + GCP (~40% of AWS scale): blended TAM ~\$2–3B / yr. The single number we optimize for is validated \$ saved per customer, audited against a customer-attested baseline.

Directional sizing: Gartner forecasts + 10-Q AWS disclosures + 6 F500 FinOps interviews. Concept-brief ballpark.

F500 RECOVERABLE AWS WASTE  
**~\$27B / yr**  
500 × \$180M × 30%

PLATFORM TAM (AWS + CLOUDS)  
**~\$2–3B / yr**  
5–8% on validated savings

## Strategic insight

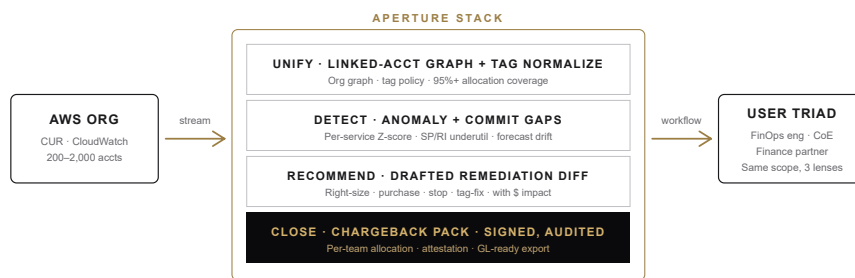
The category sells dashboards because it forgot who the user is. The CFO has Tableau. The user is the **FinOps engineer** who gets paged at 2am when EC2 jumps 40% on one linked account, currently needs four browser tabs and a Python notebook, and decides what gets fixed. Win the engineer's Monday ritual and the chargeback close walks itself.

### THE UNLOCK

Treat **cross-account scope as a free primitive**, not an enterprise add-on. Ship three role workflows as the entire surface: anomalies open with a remediation diff already drafted; commitment plans simulate against rolling 90-day usage; chargeback closes on the last business day with a signed allocation pack. Customer-zero is Aperture's own AWS bill, self-attributed in the product.

## Architecture · Cross-account intelligence stack

FIGURE 2 · SYSTEM ARCHITECTURE



*CUR + CloudWatch streamed across the linked-account graph; tags are normalized to a customer-defined policy (95%+ coverage by month 6); detection runs per-service against rolling baselines; recommendations ship as drafted diffs (not "consider tagging") and the monthly close pack is signed and GL-ready.*

### WORKED EXAMPLE · F500 RETAILER, 340 ACCTS, \$210M / YR AWS

By month 4: anomaly triage closes **median 22 min vs 6+ hr** baseline (page → drafted-diff → owner approval). Commitment simulator finds **\$11.4M / yr** in SP rebalancing across 14 BUs. Allocation coverage 71% → 96.4%; chargeback close 9 days → 2. Audited savings: **\$26.8M / yr** (12.8% of bill).

## Sequenced GTM

PHASE	CUSTOMER WEDGE	FORCING-FUNCTION WORKLOAD	PROOF POINT
<b>Wedge</b> M0-6	F500 with 200+ accounts, \$100M+ AWS bill, named FinOps team	Quarterly commitment renewal + monthly chargeback close	> 8% audited bill reduction in 90 days · 95% allocation coverage
<b>Beachhead</b> M6-18	Mid-market (50-200 accts) on \$20M+ AWS, no FinOps team	Engineer-led optimization without finance org	30 paying customers · NRR > 120% · 10x ROI on platform fee
<b>Multi-cloud</b> M18+	Same buyers, with Azure + GCP added on the same scope primitive	Cross-cloud chargeback + multi-cloud commit reshuffle	~40% of revenue from non-AWS spend within Y2

## Tradeoffs we accept

- **CUR is 12-24h delayed; real-time runs on CloudWatch + delta-CUR.** Small accuracy gap vs month-end CUR is the cost; closing it requires an AWS partnership we don't yet have.
- **Tag-policy normalization is a 6-8 week onboarding lift per customer.** No synthetic substitution. Allocation accuracy is the moat; sloppy onboarding breaks the chargeback story permanently.
- **No first-party multi-cloud at GA.** v1 ships AWS-only — cross-cloud is Y2. Premature multi-cloud dilutes depth on the wedge.

## Metrics that matter

LAYER	METRIC	Y1 TARGET	WHY IT MATTERS
North-star	Annualized \$ saved / customer (audited)	> \$4M / customer	The number we get paid against
Quality	Anomaly false-positive rate	< 12%	Above this, the engineer mutes the page
Coverage	Tag-policy allocation coverage	> 95% by M6	Below this, chargeback stops being defensible
Speed	Median anomaly time-to-remediation	< 30 min	Page → drafted-diff approval
Liquidity	Paying enterprise logos	30+ by Y1	Logo density drives reference selling
Business	Net revenue retention	> 120%	Bill grows 18–25% / yr · fee compounds with it

## Risks & mitigations

### HIGH AWS ships a credible multi-account UI inside Cost Explorer.

**Mitigation:** AWS has a structural conflict — every saved dollar is revenue lost. Aperture's gainshare pricing is buyer-aligned; AWS's cannot be. Compete on workflow depth (simulator, drafted-diff) and multi-cloud, both of which AWS will not build.

### HIGH Tag-policy normalization fails on legacy tagging chaos.

**Mitigation:** dedicated 6–8 week onboarding pod; coverage is a reported number with a 95%-by-M6 SLA. Customers below SLA at M6 get fee credits + remediation. Drop customers we can't get above 80% by M3 — bad references kill the category.

### MED Cloudability / Apptio / Vantage incumbency in F500 procurement.

**Mitigation:** wedge on the FinOps engineer, not the CFO. Two-week trial with the customer's own CUR; the success criterion is the engineer keeps the tab open. Procurement follows the user.

### MED Pricing model: gainshare vs platform fee.

**Mitigation:** hybrid — flat platform floor (covers cost-to-serve) + gainshare above an attested baseline. The floor protects against zero-savings months on already-optimized accounts.

## 30 / 60 / 90, first quarter sprint plan

### 30 DAYS

#### Anomaly triage v0

- › CUR + CloudWatch ingestion (1 design-partner org, 200+ accts)
- › Cross-account anomaly detection · false-positive baseline
- › Drafted-diff remediation for top 5 anomaly classes

### 60 DAYS

#### Commitment simulator + tag normalize

- › SP/RI rebalance simulator on rolling 90-day usage
- › Tag-policy normalization · 95%-coverage SLA tracker
- › 5 design-partner orgs live · first audited savings cohort

### 90 DAYS

#### Chargeback close + GA

- › Signed, GL-ready monthly chargeback pack
- › Multi-tenant GA · self-serve onboarding for < 200 accts
- › 30 customers signed · first NRR cohort measurement

### DECISION ASKED

Authorize a 90-day build-and-prove sprint with a seven-person team (PM, four engineers, ML lead, FinOps domain expert) and a budget of ~\$3.6M. Success: 30 paying logos, > \$4M audited savings / customer median, false-positive < 12%, allocation > 95% by M6, NRR clearing 120%.