Refine — Product Overview & Market Brief

Self-hosted multi-cloud cost intelligence for AWS, Azure, and Google Cloud. Your data stays in your accounts. Your team stays in control.

> *Blacktip Solutions · Product version v3.2.0 · Last updated May 2026 This single document consolidates the former Product Summary, Refine Overview, AWS/Cloud-Native Tools Comparison, and Competitive Analysis into one brief.*

Contents

- **Part I — Product Overview** (what Refine is, the self-hosted model, the FinOps engine)
- **Part II — Capabilities** (the full feature set across all three clouds)
- **Part III — Market Comparison** (native cloud tools, third-party FinOps platforms, decision guidance)
- **Part IV — Pricing & Licensing**
- **Part V — Architecture, Security & Deployment**
- **Part VI — FAQ**

Part I — Product Overview

What Refine Is

Refine is a self-hosted FinOps platform that turns multi-cloud waste into an executable checklist.

It runs entirely inside your environment — a laptop, EC2, an Azure VM, GCE, or an on-prem Linux host — connects read-only to your AWS accounts, Azure subscriptions, and GCP projects, and ranks every cost-optimization opportunity across all three clouds in one prioritized queue. Inventory data, recommendations, metrics, and licensing all stay inside the customer perimeter. There is no SaaS to send your billing data to, no per-seat licensing that punishes you for adding teammates, and no percentage-of-spend fee that grows with your cloud bill.

Refine is built for FinOps teams, platform engineers, and security-sensitive organizations that want centralized cost intelligence across cloud providers **without handing inventory data to a SaaS vendor.**

The industry baseline, per the FinOps Foundation's annual *State of FinOps* report, is that **20–30% of cloud spend is waste.** Refine surfaces that waste, ranks it by dollar impact, and hands you a one-click CLI command to execute each fix. Customers typically save an additional **12–25% on cloud spend** after deploying Refine — even when their cloud provider's own recommendations were already being acted on.

What "Self-Hosted" Means at Refine

This is the core differentiator, and it did not change with the move to multi-cloud:

- **Customer-owned data storage.** All collected inventory and metric summaries live in the customer's own object storage (S3, Azure Blob, or GCS). Refine reads summary files from those buckets — it does not maintain a copy outside the customer environment.

- **Customer-owned compute.** The Refine application runs as a Docker container in the customer's environment. Blacktip Solutions has no operational access.
- **Ed25519-signed offline license.** Every install validates an Ed25519-signed license file entirely offline. No license check-in, no phone-home, no telemetry by default.
- **No SaaS data egress.** Inventory, metrics, recommendations, and PII never leave the customer perimeter. The only outbound traffic is whatever the customer's network policy permits.
- **Instant disconnect.** Removing the IAM role / managed identity / service account immediately revokes Refine's ability to read the customer's environment.

The FinOps Engine — Why Refine's Recommendations Hold Up

What sets Refine apart isn't a single feature; it's that the FinOps Foundation's published best practices are wired into the recommendation engine, not bolted on as marketing copy:

- **p95 sizing over a 90-day lookback**, not mean utilization — so bursty workloads don't get under-provisioned into failure.
- **Multi-signal idle detection** — a resource is flagged idle only when CPU *and* memory *and* network are all low (≥ 2 confirming signals), not on CPU alone. A job that runs five minutes an hour looks idle by CPU but is critical.
- **Age-based exclusion** — resources younger than 14 days are excluded from teardown recommendations because their metrics are unstable.
- **Confidence scoring** — every recommendation carries a confidence (high / medium / low) with the reasons that drove it (metric depth, memory availability, resource age).
- **Structured "Why?" rationale** — click any recommendation to see the exact math: which metric, which lookback, observed value vs. threshold, target utilization, projected new-state utilization, and the confidence reasons. This is what passes a CFO review; most tools give you a single free-text reason and a savings number.

Part II — Capabilities

Cloud Coverage Matrix

This is what Refine collects and analyzes today. Coverage is per-cloud — each provider needs native SDK collectors, pricing data, and rightsizing heuristics.

Domain	AWS	Azure	GCP
Compute rightsizing (p95 + 90-day)	EC2	Virtual Machines	Compute Engine
Database rightsizing (CPU + connections + IOPS)	RDS	Azure SQL Database	Cloud SQL
Idle / orphan detection (multi-signal)	EC2, RDS, Lambda, Fargate	VMs, SQL, Functions, AKS	GCE, Cloud SQL, Cloud Functions
Block storage (over-provisioned + orphaned)	EBS volumes + snapshots	Managed Disks	Persistent Disks
Object storage (lifecycle + access-tier)	S3	Storage Accounts / Blob	Cloud Storage
Network waste (NAT, LB, public IPs)	NAT Gateways, ALB/NLB, EIPs	Public IPs, App Gateway	Cloud NAT, LB
Logs / observability (retention + ingestion)	CloudWatch Logs	Log Analytics	Cloud Logging
Caching (hit-rate + eviction)	ElastiCache	Cache for Redis	Memorystore
Serverless (memory tuning + invocation)	Lambda	Function Apps	Cloud Functions
Containers (CPU/mem + Spot eligibility)	Fargate	AKS node pools	GKE
Commitments	Savings Plans + RIs	Reserved Instances + Azure Savings Plan	Committed Use Discounts
Cost data	Cost & Usage Reports	Cost Management Export	BigQuery billing export
Cross-cloud advisor	✓	✓	✓

AWS (14 services): EC2, RDS, S3, EBS volumes, EBS snapshots, NAT Gateways, ELB (ALB/NLB), Lambda, CloudWatch Logs, Fargate, ElastiCache, Data Transfer, Batch, Lightsail. **Azure (7+ services):** Virtual Machines, Managed Disks, Storage Accounts, Azure SQL Database, Public IPs, AKS, Function Apps. **GCP (5+ services):** Compute Engine, Persistent Disks, Cloud Storage, Cloud SQL, Cloud Functions.

Every capability ships in **every** install — there is no per-cloud gating. Connect AWS, Azure, and GCP from day one without changing tools or upgrading a tier.

Cost Intelligence

Refine resolves each resource's monthly cost from two sources and blends them:

- **Actual** cost — AWS Cost & Usage Reports, Azure Cost Management Export → blob, GCP BigQuery billing export — when available.
- **Estimated** cost — built-in per-cloud pricing tables — when the actual isn't there yet (e.g., before an Azure export's first daily drop).

When both exist they're averaged; when only one exists, that one is used. This means the dashboard shows non-zero cost immediately after a sync instead of waiting 24 hours for the first export. Estimates are usage-based — a VM that ran 125 hrs/mo costs less than an always-on one.

The dashboard's **Cost by Service** breakdown reconciles to the headline total, including estimate-based spend, so cross-cloud numbers always add up.

Utilization Analysis

For VM-class resources and managed databases, Refine ingests, across **1-month / 3-month / 1-year** periods:

- CPU utilization (average and true peak)
- Memory utilization where the cloud's metric agent is enabled
- Network throughput
- IOPS and connection counts (for databases)
- Usage-pattern classification: steady / spiky / idle
- Workload-profile classification: CPU-bound / memory-bound / balanced / bursty / storage-optimized

Period metrics are collected by CloudWatch (AWS), Azure Monitor, and GCP Cloud Monitoring respectively, and surfaced behind the same 1mo / 3mo / 1yr toggle on every inventory page.

Rightsizing & Cleanup Recommendations

For each resource Refine produces one of: **Oversized** (downsize), **Undersized** (upsize), **Idle** (terminate / stop), **Generation Upgrade** (newer family at same or lower price), or **Optimal**. Each recommendation includes a target SKU, estimated monthly savings (provider public pricing for the resource's region), and a ready-to-run CLI command — `aws ec2 modify-instance-attribute`, `az vm resize`, or `gcloud compute instances set-machine-type` as appropriate.

The **Cost Cleanup** page generates per-provider cleanup scripts (`aws s3 rm`, `az storage blob delete`, `gsutil rm`, etc.) scoped by a multi-filter tag selection (AND across keys, OR within values) that works uniformly across all three clouds.

Commitment Instruments — AWS, Azure, and GCP

Refine analyzes each cloud's native commitment products, surfaces utilization and coverage, and recommends purchases or modifications. Recommendations are advisory — Refine does not execute purchases.

Provider	Commitment Instruments Covered
AWS	Savings Plans (Compute, EC2 Instance, SageMaker) + Reserved Instances (EC2, RDS, ElastiCache)
Azure	Azure Savings Plan for Compute + Reserved Instances (VM, SQL Database, Storage)
GCP	Spend-based & Resource-based Committed Use Discounts (CUDs) + Sustained Use Discounts (informational)

For each cloud, Refine surfaces the **utilization rate** (% of purchased commitment actually consumed), the **coverage rate** (% of eligible spend covered), three **commitment-depth options** (conservative / balanced / aggressive) so finance can choose risk tolerance, and **downloadable purchase scripts** the customer reviews and runs.

> ***Now shipping:*** *Azure Reserved Instance analysis and GCP Committed Use Discount optimization are* ***live*** *— earlier briefs listed these as "in progress." Both compute their recommendations from your real VM spend, and the Savings Plans page has a per-cloud picker (AWS / Azure / GCP).*

Commitment guardrails. Refine tracks per-plan utilization and **blocks recommendations that would push an active commitment below 100% utilization** unless the operator explicitly accepts the trade-off. AWS will happily recommend a Savings Plan and separately recommend rightsizing the very instances that would underwater it; Refine projects post-change utilization and refuses to do that to you.

Customer-Tunable Optimizer Thresholds *(new in v3.x)*

Every threshold that drives a recommendation is now operator-tunable from **Savings Config → Optimizer Thresholds** — no code changes, no redeploy. Each resource class (Compute, Databases, Serverless, Block Storage, Object Storage, Caching, Containers) has:

- A **preset selector** — *Conservative*, *Default*, or *Aggressive*. Default matches Refine's built-in cutoffs, so changing nothing keeps today's behavior.
- **Advanced raw-knob overrides** (e.g., Compute's CPU-avg %, CPU-peak %, memory-avg %, network floor). Editing any knob flips the card to *Custom*.
- **Reset to preset** to fall back to the canned default.

Thresholds are **shared across clouds per resource class** — the Compute card governs EC2, Azure VMs, and GCE alike — and take effect on the next sync. Sliding Compute to *Aggressive* surfaces more idle/rightsizing wins; *Conservative* tightens them.

Resource Policies *(new in v3.x)*

Operators can set three independent, durable flags on any resource, inline on its inventory row, across every inventory type and all three clouds:

Flag	Effect
Permanent	Suppresses idle / stop / delete recommendations — the resource must stay up (e.g., contractual).
Size-locked	Suppresses downsize / rightsize / tier-change recommendations — it needs its exact size.
Exclude from commitments	Drops the resource from Savings Plan / RI / CUD sizing, so a short-lived resource doesn't inflate a commitment.

When a policy blocks a recommendation, the row still shows what it *would* have saved — **struck through, with a "blocked" marker** — but that amount is **excluded** from the page's Potential Savings total and idle/oversized count, so locking a resource never inflates your headline numbers. All policies are listed centrally on the **Exemptions & Resource Policies** page alongside dismissed recommendations.

Cross-Cloud Advisor

Refine compares workloads across providers (e.g., "this EC2 fleet would cost ~18% less on equivalent Azure VMs"). This ships **informational only** — no migration execution, no data-plane replication, no application-portability assessment.

Recommendation Lifecycle & Attributed Savings

Every recommendation moves through **Suggested → Accepted → Implemented → Verified**, with realized savings attributed to the user who marked it implemented and a follow-up verification against current cloud state. The audit trail records who exempted what, when, and why — across AWS, Azure, and GCP in one ledger.

Multi-User, Multi-Account, Role-Aware

- **ROOT / ADMIN / USER** roles with account-group scoping across all three clouds.
- **Account / subscription / project groups** restrict visibility per user.
- **LDAP / Active Directory / Azure AD / Entra ID** with JIT provisioning and group sync.
- **Azure Government** cloud supported for regulated tenants.
- Provider-aware account filter (AWS accounts, Azure subscriptions, GCP projects visually grouped).

Part III — Market Comparison

The Short Version

Your cloud gives you data. Refine gives you a plan.

Every major cloud now ships a respectable set of native cost tools — AWS Cost Explorer, Compute Optimizer, Trusted Advisor, Budgets; Azure Cost Management + Advisor; GCP Reports, Recommender, BigQuery billing export. They're all good at surfacing telemetry, flagging obvious waste, and projecting commitment purchases **inside their own walled garden.**

What none of them does — and what no single-cloud tool ever will — is unify recommendations across providers, rank by dollar impact in one prioritized list, project second-order effects on your commitments *before* you act, hand you executable commands, and verify each change end-to-end after it lands. Refine layers on top of the native tooling.

Native Tooling Per Cloud

AWS

In scope: Cost Explorer, Compute Optimizer, Trusted Advisor, Savings Plans recommendations, Budgets. **Strengths:** Cost Explorer is the gold-standard billing visualizer. Compute Optimizer's 14-day rightsizing is accurate on stable workloads. Trusted Advisor (Business/Enterprise) flags idle load balancers and underutilized EC2/RDS. **Gaps Refine fills:** cross-service ranking (one list spanning EC2 + RDS + Lambda + the long tail); Savings Plan guardrails; the long tail Trusted Advisor never computes (gp2→gp3, orphaned snapshots, idle NAT, abandoned ALBs, ElastiCache hit-ratio, CloudWatch Logs retention); executable scripts with rollback; the Suggested→Verified lifecycle with attributed savings. **Typical outcome:** an additional **15–25% on AWS spend** in the first 90 days, layered on top of native tools.

Azure

In scope: Cost Management + Billing, Azure Advisor (Cost subset), Reservations & Savings Plan recommenders. **Strengths:** the most polished native cost UI on any cloud; Advisor surfaces idle VMs, idle Public IPs, and App Service rightsizing. **Gaps Refine fills:** cross-resource correlation (linking the idle VM to its still-billed disk + Public IP + NAT); RI-vs-Savings-Plan interplay; the Azure long tail (orphaned managed disks, unattached Public IP SKUs that flipped to

billable, oversized Log Analytics ingestion, oversized SQL DTU/vCore tiers); a real exemption workflow; Azure Government coverage. **Typical outcome:** an additional **12–20% on Azure spend.**

GCP

In scope: Reports, Budgets, Recommender (Active Assist), Cloud Asset Inventory, BigQuery billing export. **Strengths:** BigQuery billing export is the most powerful raw-data foundation any hyperscaler offers; Recommender is the best-integrated native rightsizer of the three clouds. **Gaps Refine fills:** you still have to write the SQL (or build a Looker army) to get a ranked recommendation set; forward-looking commitment math; cross-project aggregation across a folder hierarchy; the GCP long tail (idle Cloud SQL replicas, GKE node-pool consolidation, idle Cloud NAT, unused BigQuery slot reservations, default-retention logging); a UI a platform engineer can hand to a finance partner. **Typical outcome:** an additional **15–22% on GCP spend.**

Cross-Cloud Differentiation

Even if you stitched together all three native consoles you still wouldn't have what Refine gives a multi-cloud customer — the native tools cannot, by design, talk to each other.

- **Unified Cost Cleanup ranking** — one ranked list of waste across AWS, Azure, and GCP, sorted by dollar impact. A $1,400/mo idle Azure App Gateway sits right next to $2,100/mo orphaned EBS snapshots and a $900/mo idle GCP Cloud NAT.
- **Cross-cloud commitment alerts** — if your AWS Compute SP is at 102% utilization while your Azure SP is at 67%, Refine fires the signal to rebalance. No native tool sees the other clouds.
- **One sign-on cost view** — LDAP / AD / Entra ID with account-group scoping; one console, role-aware, across all three providers.
- **Normalized recommendation language** — a Reserved Instance (AWS), a Reservation (Azure), and a Committed Use Discount (GCP) are normalized to one vocabulary a finance partner can read and approve.
- **Cross-cloud exemption & attribution** — exemption reason codes and realized savings carry across the whole portfolio in one ledger.

Third-Party Landscape

Autonomous-execution vendors (require write access)

- **Spot.io (NetApp)** — autonomous Spot management, Eco for commitments, Ocean for K8s. Share-of-savings (~15–25%) + platform fees.
- **Zesty** — real-time commitment trading on the secondary market + EBS disk autoscaling. Share-of-savings (~20–25%).
- **CAST AI** — autonomous Kubernetes optimization across clouds. Tiered + share-of-savings; ~$50K–$200K+ at enterprise scale.
- **ProsperOps** — autonomous Savings Plan / RI portfolio management. Share-of-savings (~25%), often $100K+/yr.

Where Refine differs: Refine is read-only by design and never writes to your account; these tools require production write access — a non-starter for regulated/gov customers. Refine is self-hosted with customer-owned data; they are SaaS-plus-agent. Refine is flat per-install pricing, not share-of-savings. The two models can coexist (Refine for rightsizing + long-tail waste; an executor for autopilot commitment rebalancing).

Visibility & FinOps-suite vendors (SaaS, read-only or minimal write)

- **Vantage** — modern multi-cloud visibility with strong integrations and anomaly detection. SaaS; your billing/tag data is pulled into their cloud.
- **Finout** — unified cost observability, MegaBill consolidated billing, K8s allocation. A showback/chargeback platform.
- **CloudZero** — unit-economics framing ("cost per customer/feature"); ~1.9% of spend.
- **Apptio Cloudability (IBM)** — the original enterprise FinOps platform; strongest on TBM reporting and chargeback at very large scale.
- **Flexera CMP** — broad cloud-management platform (cost + governance + ITAM); cost recs are often shallower than the native tools it aggregates.

Where Refine differs: all are SaaS — your data leaves your environment. Refine is self-hosted with no egress. Visibility tools are stronger on dashboards and anomaly detection; Refine is stronger on actionable, ranked, executable recommendations plus SP/RI/CUD guardrails. Many customers run both (e.g., Vantage for showback, Refine for action).

Five Structural Differences From the Entire Third-Party Category

1. **Self-hosted** — runs inside your AWS / Azure / GCP environment (or on-prem). Nearly every other commercial FinOps tool is SaaS.
2. **Customer-owned data** — no billing, tagging, inventory, or resource metadata ever leaves your environment. Blacktip has no network path to it.
3. **Read-only by design** — read-only on every cloud; we never request, accept, or use write credentials.
4. **Per-install Ed25519 license** — one flat annual fee per install, signed offline, bound to an installation ID. No phone-home, no telemetry. Cloud spend can 10× without changing your Refine cost.
5. **No SaaS data egress** — air-gapped and GovCloud-native delivery is bundled in the default product.

Capability Matrix at a Glance

Capability	AWS Native	Azure Native	GCP Native	Spot.io / Zesty / CAST	Vantage / Finout	Cloudability / Flexera	ProsperOps	Refine
Multi-cloud unified ranking	—	—	—	partial	✓	✓	—	✓
Cross-cloud commitment alerts	—	—	—	partial	partial	partial	partial	✓
Rightsizing across long-tail services	partial	partial	partial	compute-only	partial	partial	—	✓
Orphaned snapshots / disks / IPs	—	—	partial	—	partial	partial	—	✓
Idle gateways (NAT, ALB, AppGW, Cloud NAT)	partial	partial	partial	—	partial	partial	—	✓
Logs / observability retention right-sizing	—	partial	—	—	—	—	—	✓
Commitment guardrails (block bad recs)	—	—	—	partial	—	—	✓	✓
SP/RI/CUD utilization — all 3 clouds	partial	partial	partial	partial	—	partial	partial	✓
Customer-tunable optimizer thresholds	—	—	—	—	—	—	—	✓
Per-resource policy flags (permanent / size-lock)	—	—	—	—	—	—	—	✓
Recommendation lifecycle + attributed savings	—	—	—	partial	partial	partial	partial	✓
Exemption workflow with reason codes	—	—	—	—	partial	partial	—	✓
Executable bash / PowerShell / az / gcloud scripts	—	—	—	(auto-exec)	—	—	(auto-exec)	✓
Read-only by design	n/a	n/a	n/a	—	✓	✓	—	✓
Self-hosted (customer-owned data)	n/a	n/a	n/a	—	—	—	—	✓

Capability	AWS Native	Azure Native	GCP Native	Spot.io / Zesty / CAST	Vantage / Finout	Cloudability / Flexera	ProsperOps	Refine
Air-gap / GovCloud delivery	n/a	n/a	n/a	—	—	partial	—	✓
Flat per-install pricing	n/a	n/a	n/a	—	—	—	—	✓

Decision Matrix — Pick What Fits Your Shape

Customer profile	Native sufficient?	Add Refine?	Add third-party SaaS?
Single-cloud, <$25K/mo, ad-hoc cost work	Yes	Not yet	No
Single-cloud, $25K–$250K/mo, engineering-led	Basics only	Yes — pays for itself in the long tail	No
Multi-cloud (any two), $50K–$500K/mo	No — native is per-cloud	Yes — consolidates ranking + commitment math	Optional — Vantage/Finout for showback
Multi-cloud, $500K+/mo, finance-led with chargeback	No	Yes — engineering recommendations	Yes — Cloudability/Flexera for TBM chargeback
Regulated / Gov / Defense / Healthcare / Finance	Partial	Yes — the only self-hosted, read-only, air-gap option	No — SaaS usually disqualified by egress
Kubernetes >50% of spend, comfortable with autonomy	No	Yes — for everything outside K8s	Yes — CAST AI / Kubecost in-cluster
Commitment management is the dominant lever, OK with write access	No	Optional — surfaces the math without executing	Yes — Spot.io / Zesty / ProsperOps for autopilot
One tool, one annual fee, no SaaS egress, no share-of-savings	Partial	Yes — this is Refine's exact shape	No

When NOT to Pick Refine

We tell prospects this up front — it shortens the cycle.

1. **You need autonomous remediation.** Refine is read-only; it generates commands, your team runs them. If you want hands-off rebalancing on autopilot, look at Spot.io, Zesty, ProsperOps, or CAST AI — and accept that they require production write access.
2. **You need BU showback / chargeback.** Refine has account groups and tag-based filtering, not a full chargeback dashboard with custom allocation rules. For 50+ cost centers with custom logic, use Cloudability or CloudHealth.
3. **You're a single-cloud AWS shop with no expansion plans.** The native trio is free and may be sufficient; Refine's economics work best with multiple clouds, a meaningful commitment portfolio, or a sovereignty requirement.
4. **Kubernetes is more than half your bill.** Refine tracks Fargate / AKS / GKE node pools / Batch at the resource level, not pod-level cost-by-namespace. Pair with Kubecost or OpenCost.

5. **You need ML anomaly detection.** Refine ships 90-day baselines and threshold-based detection, not ML anomaly spotting. CloudZero, Anodot, and Datadog do that today.
6. **You need SaaS billing aggregation (Snowflake, Datadog).** Refine covers cloud infrastructure cost only; Vantage and Finout do SaaS billing.

Where Refine Is the Only Credible Answer

1. **Multi-cloud customers who need one ranked list and one commitment view.** No native tool crosses provider lines; most third-party tools do this only in SaaS.
2. **Customers who can't or won't send cost data to a SaaS vendor** — regulated, gov, defense, healthcare, finance.
3. **Air-gapped or GovCloud installs** — a signed offline bundle with no phone-home, no version-check API.
4. **Customers who don't want their FinOps bill to scale with their cloud bill** — flat per-install.
5. **Customers who want every recommendation traceable** — Suggested→Verified, attributed savings, exemption reasons, reviewer attribution, across all three clouds in one ledger.

Part IV — Pricing & Licensing

One product, all clouds. There is no per-cloud tier and no feature gating — every install can connect AWS accounts, Azure subscriptions, and GCP projects from day one. Pricing is a **flat per-monitored-account annual license** (an AWS account, an Azure subscription, or a GCP project). Refine does **not** charge a percentage of cloud spend, per seat, or per recommendation. A customer doubling their cloud bill does not double their Refine bill.

The license is an **Ed25519-signed JSON file** verified entirely offline against an embedded public key, bound to the install's `installation_id`. It carries a `sku` label for billing/telemetry only (it does **not** restrict which clouds you can connect) and an optional `allowed_cloud_accounts` allowlist of `{provider, account_id}` tuples. When an operator registers an account, Refine enforces one check: if the allowlist is non-empty, the account must be in it. A leaked license still can't be retargeted at arbitrary accounts.

Cost vs. percentage-of-spend competitors (illustrative)

Monthly cloud spend	Cloudability	CloudZero (~1.9%)	Vantage	Refine
$50K/mo	~$54K/yr	~$11K/yr	~$5K/yr	Flat per-account license
$250K/mo	~$97K/yr	~$57K/yr	~$25K/yr	Same flat license
$1M/mo	~$228K/yr	~$228K/yr	~$150K/yr	Same flat license

The economic argument sharpens at scale: a customer at $1M/mo pays the same per-account rate as one at $50K/mo. A customer running only AWS today can connect Azure or GCP tomorrow with no license change and no redeployment.

Part V — Architecture, Security & Deployment

Architecture

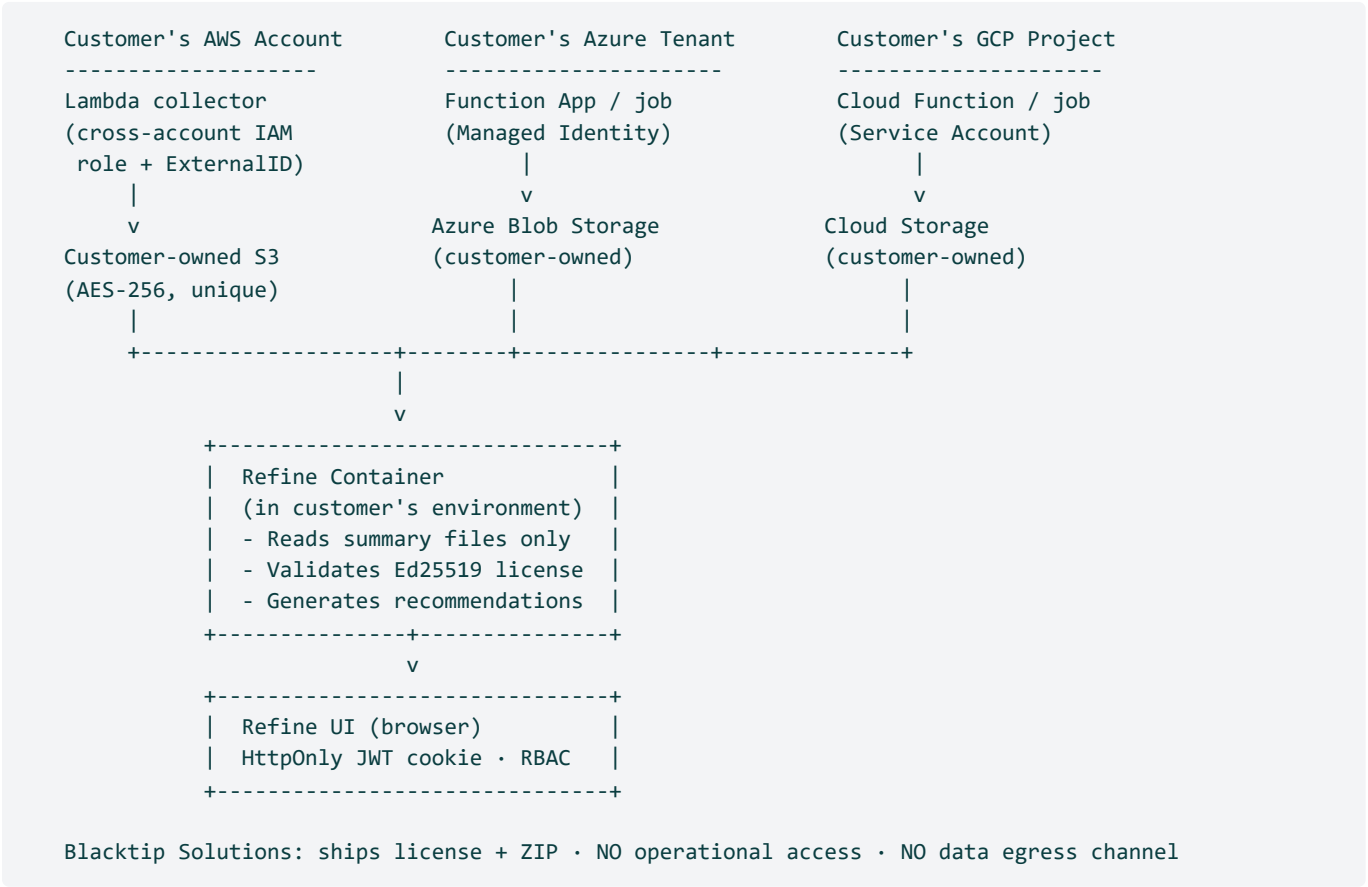

Collectors run in the customer's own cloud accounts using least-privilege identities. Refine reads pre-summarized JSON from customer-owned storage rather than querying cloud APIs directly — keeping the container's footprint minimal and air-gap-friendly. There is no socket from Refine back to Blacktip.

Security & Compliance

- Inventory data stored in customer-owned object storage with provider-native encryption-at-rest (S3 AES-256 / Blob SSE / Google-managed GCS).
- Least-privilege access: cross-account IAM role + ExternalID (AWS), Managed Identity + RBAC scope (Azure), Service Account + IAM policy (GCP). Each grants only object-storage read + collector-invoke. Delete it and access is revoked instantly.
- Ed25519 offline license validation; RBAC (ROOT / ADMIN / USER); rate-limited auth; HttpOnly JWT cookies; HSTS / CSP / X-Frame-Options.
- Non-root container; bcrypt password hashing; parameterized SQL; tenant data isolation validated by the security-audit process; UUID input validation.

Deployment Options

Refine deploys on any Linux host that can run Docker. The host machine never holds cloud-native identity — it calls each cloud's public APIs over HTTPS using stored credentials, which is why any host cloud works with any monitored cloud.

Target	Provisioning template shipped
AWS EC2	CloudFormation (`refine-host.yaml`)
Azure VM	Bicep (`refine-host.bicep`)
GCP Compute Engine	Terraform (`refine-host.tf`)
On-prem / air-gapped	Docker Compose, no cloud provisioning

Onboarding per cloud: AWS — one-click CloudFormation per account (creates IAM role + Lambda collector + S3 bucket). Azure — one-line `curl | bash` in Cloud Shell (creates Entra ID app registration, Reader role, storage account + container, daily Cost Management Export; auto-detects Commercial vs US Government). GCP — a Terraform module (service account + viewer roles + BigQuery export + JSON key). First recommendations appear within minutes of the first sync.

Three delivery tracks

Track	Description	Best for
AWS Marketplace (Commercial CFN)	One-click subscribe → CloudFormation product launches in customer's AWS account. EC2 pulls Refine from AWS Public ECR (`public.ecr.aws/blacktip/refine`); Blacktip uploads only the ~500-byte Ed25519 license file to the customer's CFN-provisioned S3 bucket; a license-watcher daemon installs it within 60s. BYOL pricing day-1; Annual Contract pricing planned.	Commercial enterprises procuring through AWS Marketplace / EDP drawdown
Self-Hosted (Government)	Complete ZIP with bundled Docker image; zero internet at any point; Ed25519 offline license; all host templates. Distributed direct-sales today and via the GovCloud Marketplace listing (Phase 2).	Federal, DoD, air-gapped, classified
Self-Hosted (Direct-Sales Commercial)	Docker image pulled from registry; optional opt-in savings telemetry; same host templates; Watchtower-based "Update Now".	Enterprises, FinOps consultancies, security-sensitive orgs preferring direct procurement

Part VI — Frequently Asked Questions

Can we start with one cloud and add others later? Yes. Every install supports all three clouds with no license change — connect an Azure subscription or GCP project whenever you're ready. Just register it and sync.

Does Refine support GovCloud and Azure Government? Yes. AWS GovCloud is a supported region set; Azure Government tenants are supported as standard Azure subscriptions. The Government delivery track is air-gap-compatible.

Does Refine make changes to my cloud environment? No — strictly read-only. Every recommendation is a human-readable CLI command your administrators review and execute.

How are savings calculated? From each cloud provider's public pricing for your region and the recommended target SKU, blended with actual billing-export cost where available.

How often is data updated? Collectors run nightly per provider (2 AM UTC default); manual sync is available on-demand from the dashboard.

What's the difference between Refine and Spot.io / Zesty / CAST AI? Refine analyzes and recommends; those tools execute remediations and require production write access. Refine is read-only by design — typical buyers want the human-in-the-loop model because automated-remediation roles fail their security reviews.

Does Refine do showback or chargeback? No. Refine supports tag-based filtering across business-unit dimensions but does not produce chargeback invoices. Pair it with CloudHealth, Apptio Cloudability, or native cost-allocation tools.

Is my data secure? All inventory data stays in your own cloud accounts; per-cloud collectors do the collection. Blacktip never sees your data. Remove the role/identity/service account and access is revoked instantly.

Contact

Ready for a multi-cloud cost-intelligence demo? We'll walk through your actual spend with Refine loaded against a sandbox or single account on whichever cloud(s) you run — specific dollar recommendations across providers, not a generic demo.

- Email: **AMelcher@blacktip-ops.com** · **support@blacktip-ops.com**
- Website: **www.blacktip-ops.com**

Blacktip Refine — See Your Costs. Optimize Your Clouds.

© 2026 Blacktip Solutions · Indian Land, SC · Corporate@blacktip-ops.com
WOSB · VOSB · CAGE 14QN0 · UEI SU55FSWCWK98