

■ ZOPDEV RESEARCH • JUNE 2026

The State of Cloud 2026

An Emerging Markets view of global cloud, run as a mission.
Waypoints, field notes, and a checklist you leave with.

The trail ahead, in one glance.

Every finding in the research maps to something you can do. So the trail is organized by capability, not by chapter. Walk it in order — each waypoint builds on the last.

01

Know your trajectory

Optimizer, Drifter or Scaler — every cloud bill is on one of three paths. Find yours.

≈ 6 MIN • 1 SELF-TEST

02

Find the waste

87% of waste is forgotten, not chosen. Ten fixes recover 80% of it.

≈ 7 MIN • 10 CHECKBOXES

03

Close the action gap

Teams detect waste 30× faster than they fix it. Put a name on the list.

≈ 5 MIN • 4 CHECKBOXES

04

Choose your footprint

Regions, metadata, resilience — make them decisions, not accidents.

≈ 6 MIN • 3 DIG-DEEPERS

Know which trajectory you're on — and pick it on purpose.

The "average enterprise cloud bill" is a meaningless number. In the same quarter of the same cohort, one customer halved spend while another quadrupled it. The story is in the spread.

OBJECTIVE

Place your cloud bill on one of three trajectories — and decide whether it's the one you actually chose.

- ✓ Classify your org as an Optimizer, Drifter or Scaler using real run-rate data
- ✓ Explain why 5–10% quarterly drift is the default state of enterprise cloud
- ✓ Answer the one question every CTO should be able to answer

Every enterprise in the cohort — 13 firms, 25 production accounts, 32 regions — fell into one of three trajectories. Not three maturity levels. Three directions of travel:

~8% OF ENTERPRISES • DECLINING

The Optimizer

\$1,565/day → \$786/day • 69 days

Active remediation, FinOps embedded in engineering.

Recommendations get triaged, owners get assigned, savings get tracked. No migration. No rebuild. Pure hygiene.

~50% OF ENTERPRISES • FLAT-RISING

The Drifter

\$5,787/wk → \$6,361/wk • 8 weeks

Detection in place, action absent. Cost grows quietly, 5–10% a quarter, from accumulated unfixed waste. No new launch. Just drift. The default state of enterprise cloud.

~42% OF ENTERPRISES • 2–4× GROWTH

The Scaler

\$1,429/wk → \$5,918/wk • 6 weeks

Growth-stage spend. Every new team gets its own workspace, region, NAT gateway, snapshot policy. The growth is legitimate. The waste embedded in it is not.

FIELD NOTE №1

One Optimizer cut daily spend from **\$1,565 to \$786** in 69 days.

Minus 50%, in nine weeks — with no migration and no rebuild. Their secret wasn't better detection; they had the same tools as everyone else. Someone simply owned the list, every day, and worked through it.

OBSERVED RUN-RATE • COHORT TELEMETRY • MAR-MAY 2026

Most enterprises are Drifters, and drift compounds. At 5–10% a quarter, a flat workload’s bill grows 20–40% a year with zero new product shipped. That’s the quiet budget overrun the flashy AI line items get blamed for.



DIG DEEPER • THE BIMODAL ECONOMY IN NUMBERS

Optimizers shed waste fast. Scalars accumulate it fast. Drifters stay stuck in between — which is why averages lie.

ARCHETYPE	SHARE	OBSERVED MOVE	WINDOW
Optimizer	~8%	-50% daily spend	69 days
Drifter	~50%	+10% weekly run-rate	8 weeks
Scaler	~42%	4.1× weekly run-rate	6 weeks

All ten of the report’s pre-registered hypotheses were confirmed by the telemetry — including “customers detect waste 10× faster than they act.” The real ratio turned out to be ~30×. The conventional wisdom about cloud is broadly correct; what’s missing is the will to act on it.

Trajectory audit — do these four things this week

- ☐ **Pull 12 weeks of daily spend** for your top accounts — the raw run-rate, not the invoice summary.

EFFORT • 30 MIN

- ☐ **Plot the slope and name it** — declining, flat-rising, or 2–4× growth. Write the word down.

EFFORT • 15 MIN

- ☐ **Ask the CTO question out loud** in your next leadership sync: which trajectory are we on, and did we choose it?

EFFORT • 1 MEETING

- ☐ **Assign a trajectory owner** — one person who reports the slope monthly, the way someone already reports uptime.

EFFORT • 1 DECISION



— TRAIL STOP • SELF-TEST

Which archetype are you?

? TAP THE CARD THAT SOUNDS MOST LIKE YOUR LAST TWO QUARTERS

OPTION A

"Our bill went down — and we can name who made it go down."

Recommendations get owners. Savings get tracked in the same standup as features.

OPTION B

"Flat-ish... creeping up a few percent a quarter. Nothing launched."

The dashboards exist. The recommendations exist. Nobody's job is to act on them.

OPTION C

"It's growing 2–4× — new regions, new teams, new accounts every week."

Growth is real and welcome. But every new service ships with its own NAT gateway.

YOUR READ · ~8% OF THE COHORT

You're an Optimizer.

Trajectory: declining · rare air

1. Protect the habit — the moment the list loses its owner, you become a Drifter within two quarters.
2. Graduate from hygiene to architecture: footprint, commitment strategy, and the GPU line (Waypoint 04 and beyond).
3. Set a GPU utilization SLO now, before AI spend crosses 15% of the bill in 2027 — Optimizers who move early stay Optimizers.

YOUR READ · ~50% OF THE COHORT

You're a Drifter — the default state.

Trajectory: flat-rising · 5-10% quarterly drift

1. Your problem is not detection — skip straight to Waypoint 03 and put a name on the remediation list.
2. Run the top-10 fix list at Waypoint 02: ten items recover ~80% of your waste, most in minutes each.
3. Book a 30-minute weekly remediation slot. The Optimizer's only real difference was that someone showed up to it.

YOUR READ · ~42% OF THE COHORT

You're a Scaler.

Trajectory: 2-4× growth · waste compounding inside it

1. Don't slow the growth — install guardrails at provisioning time: default schedules, snapshot lifecycle, sunset tags on every experiment.
2. Watch fragmentation (Waypoint 04): each new workspace spinning up its own NAT gateway and scratch volumes is tomorrow's 16.9%.
3. Audit quarterly. Scalers who wait discover the waste after it has three more zeros.

Find the waste nobody chose.

Vendor narratives push big strategic moves — Reserved Instances, Savings Plans, Spot. Those account for just 6% of flagged items. The real waste is mundane: things nobody remembers creating.

OBJECTIVE

Locate the 87% of waste that's forgotten, not strategic — and recover ~80% of it with ten fixes.

- ✓ Recite the top-10 anti-pattern list from memory (your platform team should too)
- ✓ Explain why orphans dominate every cloud estate
- ✓ Defend the 16.9% number — and why the true figure is likely 30–40%

Across 25,225 live resources, 16.9¢ of every cloud dollar was provably recoverable waste, sitting in plain sight — before any architectural change, renegotiation, or repatriation. And 87.4% of the flagged items were orphans: resources with no owner, no purpose, and a running meter.

ORPHAN		87.4%
DISCOUNT		6.3%
RIGHTSIZING		3.7%
IDLE		1.3%
SCHEDULE		0.9%

Why do orphans win? Three reasons, all human. Provisioning is frictionless while de-provisioning is scary. Accounts are organized by

team or environment — almost never by lifecycle, so nobody owns the cleanup. And snapshots compound silently: every dev snapshots before a risky change, and almost nobody deletes afterwards.

FIELD NOTE №2

One customer was carrying 3,807 EBS snapshots.

Snapshot storage can exceed live volumes by 14×. Orphan snapshots — source volume deleted, still billed forever — showed up in 4 of 6 AWS-heavy enterprises. The biggest cloud waste problem isn't buying the wrong thing. It's forgetting to delete the right thing.

DETECTION RULE RC-021 · COHORT TELEMETRY



DIG DEEPER · HOW 16.9% WAS DERIVED (AND WHY IT'S THE FLOOR)

STEP	FIGURE	DETAIL
Annualized cohort spend	\$2,222,930	\$420,554 measured over a 69-day window, annualized
Proven recoverable savings	\$375,217	Only high-confidence recommendations counted
Waste-to-spend ratio	16.9%	$\$375,217 \div \$2,222,930$

Excluded on purpose: architectural rebuilds, repatriation, commitment renegotiation, storage-tier moves without access evidence, and application-layer efficiency. Include those and the true recoverable figure is likely 30–40% of spend. For a \$10M/yr cloud bill, 16.9% alone is \$1.69M on the table; across a ~\$680B global market, north of \$100B a year. In rupee terms, a single mid-sized enterprise can be sitting on ₹50 crore — sharper still, be-

cause the bill is USD-denominated and the revenue deferring it is not.

— TRAIL STOP • THE MAIN CHECKLIST

Fix these ten things. Recover 80% of the waste.

None of these are strategic decisions to revisit. They're hygiene to install — most take about five minutes each. Your progress saves automatically; print this page and it becomes a blank worksheet.

- ☐ **Delete orphan EBS snapshots** — source volume gone, charged forever.
100% OF AWS USERS • \$80K+/YR AT STAKE
- ☐ **Remove unattached EBS volumes** — a single volume leaks up to \$2,500/yr.
100% OF AWS USERS • \$25K+/YR
- ☐ **Downgrade Premium SSD in non-prod Databricks (Azure)** — the predictable anti-pattern.
100% OF DATABRICKS-ON-AZURE • \$66K+/YR
- ☐ **Cover long-running EC2 with Savings Plans** — on-demand for 30+ days is a choice.
~90% OF PROD FLEETS • \$50K+/YR
- ☐ **Move Single-AZ production RDS to Multi-AZ** — not cost, risk. A 5-minute outage vs a multi-hour one.
EVERY AWS ENTERPRISE • RESILIENCE RISK
- ☐ **Schedule dev/test workloads** — running 24/7 by default wastes up to 67% on dev VMs.
100% OF DEV ENVIRONMENTS • \$40K+/YR
- ☐ **Apply Azure Hybrid Benefit to Windows VMs** — licenses you already own.
COMMON IN REGULATED INDUSTRIES • 40% OFF LICENSE
- ☐ **Move stateless / batch off on-demand** — spot and flexible pricing exist for exactly this.
COMMON IN CLOUD-NATIVE SHOPS • \$30K+/YR

- ☐ **Migrate x86 to Graviton/ARM where it just works —**
most fleets qualify.

80% OF COMPUTE FLEETS • \$25K+/YR

- ☐ **Kill EBS charges on stopped instances — stopped ≠**
free.

FOUND IN EVERY ACCOUNT • \$5K+/YR

IF YOUR PLATFORM TEAM CAN'T RECITE THIS LIST FROM
MEMORY, YOU HAVE A 16.9% PROBLEM.

Close the gap between seeing and doing.

Detection has been solved. Execution has not. Cloud teams identify waste roughly 30× faster than they act on it — and the reasons are human, not technical.

OBJECTIVE

Turn a wall of recommendations into a worked list with a named owner — the only thing Optimizers do differently.

- ✓ Diagnose the three human blockers: owner ambiguity, risk asymmetry, no apply-path
- ✓ Stand up a weekly remediation loop that survives busy quarters
- ✓ Track applied-vs-detected as a first-class metric

Here's the funnel the telemetry exposed, across the whole cohort:

8,000+

waste signals detected

6,300+

auto-resolved by churn and turnover —
luck, not intent

~250

explicitly reviewed by a
human

<24

formally applied
via a remediation
workflow

<1%

OF DETECTED
WASTE
FORMALLY
REMIEDIATED

~80%

AUTO-RESOLVES
THROUGH
CHURN, NOT
INTENT

30×

FASTER TO
DETECT THAN
TO ACT

Why the gap exists — all human, none technical. **Owner ambiguity:** the account belongs to platform, the workload to the app team, the cost to finance; nobody owns the recommendation. **Risk asymmetry:** deleting a snapshot has a small chance of breaking something visible and personal; keeping it has a 100% chance of costing money, diffusely. **No apply-path:** most FinOps platforms show recommendations; almost none execute them safely. The platform stops at the dashboard — the work happens, or doesn't, in the cloud console.

The Optimizer's secret wasn't better tooling. **Someone owned the list.**

Same detection tools as everyone else. The difference: one person triaged the recommendations every day and worked through them. That alone produced -50% in 69 days. The market for FinOps that detects is mature. The market for FinOps that acts is wide open.

CHAPTER 5 · THE DETECTION-TO-ACTION GAP

Install the action loop — four moves

- ☐ **Name one owner for the recommendation queue** — a person, not a committee. Put it in their goals.
FIXES · OWNER AMBIGUITY
- ☐ **Book a standing 30-minute weekly remediation slot** — the same discipline as patch Tuesday.
FIXES · DRIFT BY NEGLECT
- ☐ **Define blast-radius tiers** — pre-approve zero-risk deletions (orphan snapshots, unattached volumes) so fear of prod stops blocking them.
FIXES · RISK ASYMMETRY
- ☐ **Report applied-vs-detected monthly** — if it's under 10%, you have a Drifter's funnel no matter what the dashboard says.
FIXES · NO FEEDBACK LOOP

Choose your footprint like it's a decision. Because it is.

Multi-cloud, region spread, metadata sprawl, resilience posture — in the cohort, roughly half of these “strategies” turned out to be accidents nobody went back to revisit.

OBJECTIVE

Audit where you actually run — providers, regions, metadata — and make each part deliberate.

- ✓ Separate the multi-cloud narrative from the 1-in-13 reality
- ✓ Spot the metadata iceberg growing under a flat workload
- ✓ Retire resilience theater and the observability tax

65.5%

OF COHORT
SPEND ON AWS
(53% OF
ACCOUNTS)

1 in 13

GENUINELY
MULTI-CLOUD
IN PRODUCTION

14×

SNAPSHOT
STORAGE VS
LIVE VOLUMES

Start with the provider picture: AWS takes two-thirds of every cloud dollar; Azure punches above its weight on fewer accounts; GCP is the long tail concentrated in product and SaaS. And “multi-cloud strategy” is mostly talk — exactly one enterprise in thirteen ran production across all three hyperscalers. Six of thirteen are single-cloud AWS, four single-cloud Azure, two single-cloud GCP. The rest of the conversation is failover plans nobody tested.

Then look under the compute layer. The visible workload is the tip; the metadata is the iceberg — 22 CloudWatch alarms per EC2 instance, 3 IAM roles per workload, 2.3 snapshots per volume per year. Metadata has no lifecycle, so the bill grows linearly with it even when workload is flat.

FIELD NOTE №4

One production account ran **6,922 CloudWatch alarms** — the vast majority unmonitored.

No SNS destination, no escalation, no on-call rotation listening. Noise generators paying their own bill. Treat your monitoring stack like your production stack: life-cycle policies, deprecation windows, ownership tags.

CHAPTER 6 · THE OBSERVABILITY TAX



DIG DEEPER · FOUR FOOTPRINTS, TWO ACCIDENTS

FOOTPRINT PATTERN	EXAMPLE (ANONYMIZED)	REGIO NS	\$ / REGION / MO
Global sprawl, one tenant	Global sporting goods · Azure	32	~\$60
Concentrated production	Global CPG · Azure	3	~\$5,900
Single-region monolith	Consumer internet · AWS	2	~\$12,500
Spread without intent	E-commerce · AWS	19	~\$560

Two of these reflect deliberate architecture. Two reflect accident. Thirty-two regions at ~\$60 each is not a strategy — it’s entropy with a billing account.



DIG DEEPER • CONSOLIDATORS VS FRAGMENTERS (AZURE)

Consolidators run 3–5 large subscriptions (~\$15–20K/month each) — governance is tractable. **Fragmenters** run 20–35 small ones (\$500–1,500/month each) — governance scales linearly with team count and cleanup becomes hopeless. One customer ran 33 distinct subscriptions in a single tenant; another ran 3 for comparable spend. The most visible symptom: Databricks workspace proliferation, each workspace spinning up its own NAT gateway and Premium SSD scratch volumes.



DIG DEEPER • RESILIENCE THEATER

Every AWS enterprise in the cohort had at least one Single-AZ production database — 36+ flagged in one customer alone. The fix is trivial, the cost a small uptick in storage, the benefit the difference between a 5-minute outage and a multi-hour one when an AZ fails. The reason is never “we made a deliberate choice.” It’s always “we ran the migration script in 2021 and never went back.”

One more thing before the debrief — the emerging-markets lens sharpens all of this. Cloud bills are USD-denominated; local revenue usually isn’t, so a 10% currency swing is a 10% cost increase with zero new resources. That’s a treasury risk, not a tech line item — and it’s why EM CFOs increasingly avoid 3-year RI commitments, why sovereign cloud priced in local currency commands a premium, and why enterprises that skipped the legacy-migration era are, on average, more architecturally coherent than their developed-market peers.

You walked the trail. Here's what you're carrying out.

Four waypoints, one field kit: a named trajectory, a ten-item fix list, an action loop with an owner, and a footprint you chose on purpose.



WAYPOINT 01 COMPLETE

Know your trajectory

Optimizer ~8%, Drifter ~50%, Scaler ~42%.
Most enterprises drift 5–10% a quarter without intent — the slope is now a number someone owns.

WAYPOINT 02 COMPLETE

Find the waste

16.9% provably recoverable (likely 30–40% truly); 87% of it is orphans nobody chose. Ten fixes recover ~80% of the total.

WAYPOINT 03 COMPLETE

Close the action gap

Detection outruns action 30×; under 1% of signals get formally fixed. The blockers are human — so the fix is an owner, a cadence, and pre-approved blast radii.

WAYPOINT 04 COMPLETE

Choose your footprint

Genuine multi-cloud is 1 in 13. Metadata is the iceberg, fragmentation is governance debt, and Single-AZ prod databases are theater — all now on your audit list.

THE QUESTION EVERY CTO SHOULD ASK

”Which trajectory are we on — and is it the one we chose?”

If you can answer that in one sentence, with data, you got what this manual came to give you. If you can't yet, the trail is right there above you — and the checklist saves your place. One more thing to watch: the GPU shift replays every anti-pattern in this manual with three more zeros. A forgotten CPU weekend costs \$37; a forgotten 8×H100 weekend costs \$4,704. The teams that install these habits in 2026 will look like Optimizers in 2027. There is no third option.

■ YOUR NEXT MOVE

Don't file this manual. Deploy it.

ZopDev builds the apply-path — the part of FinOps the whole industry skipped. Bring your trajectory, your checklist score, and your footprint audit; we'll help you work the list the way the Optimizer did.

THE STATE OF CLOUD 2026 • AN EMERGING MARKETS
VIEW • **PUBLISHED JUNE 2026** • CHALLENGE IT:
CONTACT@ZOP.DEV