

Sample report — synthetic demo estate, not a real client. Every number below is computed by the same scanner that would run against your account, with the same methodology and the same refusal to claim a saving it cannot verify. Only the estate is invented.

AWS Database Cost Audit

Generated 2026-07-26 - Lookback 30 days - Regions: us-east-1 - 8 accounts

USD 24,068

Estimated monthly savings

USD 288,813

Annualized

84

Actionable findings

69

Resources scanned

8

Accounts covered

Scan coverage: 41 RDS instances · 6 ElastiCache clusters · 5 Redshift clusters · 6 OpenSearch domains · 6 DynamoDB tables · 3 MemoryDB clusters · 2 DMS instances

WHERE THE SAVINGS ARE

RDS Right-Sizing Candidates		USD 3,368.95/mo
Redshift Node Reduction		USD 3,171.12/mo
Redshift Graviton (RG) Migration		USD 2,855.03/mo
RDS Reserved Instance Opportunities		USD 2,643.69/mo
Idle RDS Instances: Scheduling Candidates		USD 1,812.08/mo
Aurora I/O-Optimized Evaluation		USD 1,362.59/mo
Redshift Pause Schedules		USD 924.91/mo
MemoryDB Right-Sizing		USD 902.28/mo
Extended Support Surcharges (RDS, DocumentDB, OpenSearch)		USD 745.77/mo
OpenSearch Right-Sizing		USD 735.84/mo
RDS Storage: Provisioned IOPS to gp3		USD 726.40/mo
Redshift Reserved Nodes		USD 534.36/mo
DynamoDB Over-Provisioned Capacity		USD 503.92/mo
Idle RDS Read Replicas		USD 386.90/mo

RDS Snapshot Hygiene		USD 370.50/mo
OpenSearch Master Oversizing		USD 367.92/mo
DynamoDB Billing Mode		USD 255.86/mo
Leftover DMS Instances		USD 245.28/mo
Idle OpenSearch Domains		USD 243.82/mo
Valkey Engine Migration (ElastiCache, MemoryDB)		USD 194.18/mo
RDS Reserved Instance Waste		USD 188.50/mo
Idle ElastiCache Clusters		USD 157.68/mo
ElastiCache Right-Sizing		USD 153.30/mo
Neptune Stop Schedules		USD 148.19/mo
Storage: gp2 to gp3 (RDS, OpenSearch)		USD 133.47/mo
DynamoDB Standard-IA Candidates		USD 119.29/mo
Graviton Migration (RDS, DocumentDB, Neptune, OpenSearch)		USD 116.07/mo
DynamoDB Unused Indexes		USD 113.88/mo
RDS Multi-AZ on Non-Production		USD 113.15/mo
Idle MemoryDB Clusters		USD 112.57/mo
ElastiCache Reserved Nodes		USD 96.36/mo
OpenSearch Reserved Instances		USD 90.52/mo
Previous-Generation Upgrades (RDS, DocumentDB, Neptune)		USD 82.49/mo
DocumentDB Stop Schedules		USD 66.60/mo
Stopped RDS Instances Still Billing		USD 10.00/mo
DynamoDB PITR on Non-Production		USD 8.00/mo
Idle DynamoDB Tables		USD 6.25/mo
Right-Sizing		USD 9,459.19/mo
Idle Resources & Scheduling		USD 4,228.16/mo
Modernization		USD 4,106.69/mo
Commitments & Reservations		USD 3,553.43/mo
Consultative Reviews		USD 1,362.59/mo
Storage & Backups		USD 1,357.66/mo
682415903277 (acme-data-platform)		USD 8,316.99/mo
349072118463 (acme-prod-core)		USD 6,031.54/mo
517346802915 (acme-ml-research)		USD 3,249.59/mo
795013264871 (acme-dev-sandbox)		USD 1,778.82/mo
560839472016 (acme-legacy-lob)		USD 1,378.00/mo
418926075334 (acme-marketing)		USD 839.13/mo
904271635840 (acme-staging)		USD 522.79/mo
236580417962 (acme-qa)		USD 399.77/mo

Click an account to filter the whole report to it.

TOP 10 OPPORTUNITIES

	FINDING	ACCOUNT	RESOURCE	ENV	SAVING/MO
1	Redshift Node Reduction	682415903277 acme-data-platform	analytics-wh ra3.xlplus x 8 · us-east-1	production	USD 3,171.12
2	Redshift Graviton (RG) Migration	682415903277 acme-data-platform	etl-transform-wh ra3.4xlarge x 4 · us-east-1	production	USD 2,855.03
3	RDS Right-Sizing Candidates	517346802915 acme-ml-research	ml-feature-store db.r6i.8xlarge · us-east-1	production	USD 1,762.95
4	Aurora I/O-Optimized Evaluation	—	Aurora clusters (account level) all	—	USD 1,362.59
5	RDS Reserved Instance Opportunities	349072118463 acme-prod-core	core-orders-api, core-payments, core-analytics + 2 more r6g · us-east-1	—	USD 1,144.64*
6	Redshift Pause Schedules	795013264871 acme-dev-sandbox	bi-dev-wh ra3.xlplus x 2 · us-east-1	dev	USD 924.91
7	MemoryDB Right-Sizing	349072118463 acme-prod-core	session-store db.r6g.xlarge · us-east-1	production	USD 902.28
8	RDS Right-Sizing Candidates	682415903277 acme-data-platform	dp-warehouse db.r5.4xlarge · us-east-1	production	USD 832.20
9	RDS Right-Sizing Candidates	349072118463 acme-prod-core	core-analytics db.r6g.4xlarge · us-east-1	production	USD 773.80
10	OpenSearch Right-Sizing	349072118463 acme-prod-core	search-prod r6g.xlarge.search x 6 · us-east-1	production	USD 735.84

▼

RDS Reserved Instance Opportunities

7 findings

USD 2,643.69*/mo

ACCOUNT ↕	RESOURCE ↕	ENV ↕	RECOMMENDATION	SAVING/MO ↓	CONFIDENCE ↕	RISK
349072118463 acme-prod-core	core-orders-api, core-payments, core-analytics +2 more r6g · us-east-1	—	Buy size-flexible RIs covering 64.0 normalized units in family r6g (postgres, mixed Single-AZ and Multi-AZ). Verify 90 days of steady usage before a 1-year commitment. 80.0 NU running · 16.0 covered · gap 64.0 NU	USD 1,144.64*	high	low
517346802915 acme-ml-research	ml-feature-store r6i · us-east-1	—	Buy size-flexible RIs covering 32.0 normalized units in family r6i (postgres, Single-AZ). Verify 90 days of steady usage before a 1-year commitment. 32.0 NU running · 0.0 covered · gap 32.0 NU	USD 652.25*	high	low
682415903277 acme-data-platform	dp-warehouse, dp-etl-primary r5 · us-east-1	—	Buy size-flexible RIs covering 32.0 normalized units in family r5 (postgres, Single-AZ). Verify 90 days of steady usage before a 1-year commitment. 32.0 NU running · 0.0 covered · gap 32.0 NU	USD 616.12*	high	low

ACCOUNT ↓	RESOURCE ↓	ENV ↓	RECOMMENDATION	SAVING/MO ↓	CONFIDENCE ↓	RISK
517346802915 acme-ml-research	ml-aurora-dev r6g · us-east-1	—	Buy size-flexible RIs covering 4.0 normalized units in family r6g (aurora-postgresql, Single-AZ). Verify 90 days of steady usage before a 1-year commitment. 4.0 NU running · 0.0 covered · gap 4.0 NU	USD 76.65*	high	low
418926075334 acme-marketing	mkt-cdp, mkt-attribution r6g · us-east-1	—	Buy size-flexible RIs covering 4.0 normalized units in family r6g (mysql, Single-AZ). Verify 90 days of steady usage before a 1-year commitment. 12.0 NU running · 8.0 covered · gap 4.0 NU	USD 71.54*	high	low
560839472016 acme-legacy-lob	lob-finance m5 · us-east-1	—	Buy size-flexible RIs covering 4.0 normalized units in family m5 (mysql, Single-AZ). Verify 90 days of steady usage before a 1-year commitment. 4.0 NU running · 0.0 covered · gap 4.0 NU	USD 45.99*	high	low
418926075334 acme-marketing	mkt-tracker t3 · us-east-1	—	Buy size-flexible RIs covering 4.0 normalized units in family t3 (postgres, Single-AZ). Verify 90 days of steady usage before a 1-year commitment. 4.0 NU running · 0.0 covered · gap 4.0 NU	USD 36.50*	high	low

* Reserved Instance purchases are a commitment decision that belongs to you: savings assume the workload keeps running for the full term, and only you have insight into future usage plans (migrations, re-platforming, decommissions). The scanner sees history, not roadmaps. Mechanics: groups running size-flexible instances by account, region, family, engine, and AZ config; subtracts active RI coverage in normalized units. Savings compare live public On-Demand rates against matching RI offering rates (upfront amortized) - the rate math is exact. Confidence reflects workload longevity: high when every instance in the group has run 6+ months, medium when younger. Instances too young to prove steady demand are excluded entirely, as are instances this report recommends scheduling, deleting, moving to another family, or converting from Multi-AZ; instances flagged for a same-family downsize are sized at their target class, never today's

larger one. Also weigh the Database Savings Plans alternative (see that section): RIs discount deeper on stable shapes, a Savings Plan follows the estate through engine, family, and region changes.

▼ ElastiCache Reserved Nodes

1 finding

USD 96.36*/mo

ACCOUNT ↓	RESOURCE ↓	ENV ↓	RECOMMENDATION	SAVING/MO ↓	CONFIDENCE ↓
349072118463 acme-prod-core	api-cache cache.r7g.large · us-east-1	production	Reserve 2x cache.r7g.large (redis). ElastiCache RIs match exact node types - re- check after any modernization moves land. 2 running · 0 covered · gap 2 x cache.r7g.large	USD 96.36*	medium

* Commitment decision belongs to the client (see Reserved Instance note). ElastiCache reserved nodes match exact node types with no size flexibility; clusters already flagged for deletion, right-sizing, or migration are excluded from sizing. Savings compare live On-Demand against the lowest matching offering, upfront amortized.

▼ Redshift Reserved Nodes

1 finding

USD 534.36*/mo

ACCOUNT ↓	RESOURCE ↓	ENV ↓	RECOMMENDATION	SAVING/MO ↓	CONFIDENCE ↓
682415903277 acme-data-platform	reporting-wh ra3.xlplus x 4 · us-east-1	production	Reserve 2x ra3.xlplus nodes. Size the commitment after any pending resize or RA3 migration lands. 4 running · 2 covered · gap 2 x ra3.xlplus	USD 534.36*	medium

* Commitment decision belongs to the client (see Reserved Instance note). Exact node-type matching; clusters pending pause, resize, or RA3 migration are excluded. Savings compare live On-Demand against the lowest matching offering, upfront amortized.

▼ OpenSearch Reserved Instances

1 finding

USD 90.52*/mo

ACCOUNT ↕	RESOURCE ↕	ENV ↕	RECOMMENDATION	SAVING/MO ↓	CONFIDENCE ↕
418926075334 acme- marketing	api-search r6g.large.search x 2 · us-east-1	production	Reserve 2x r6g.large.search data nodes. Size after pending right-sizing or Graviton moves land. 2 running · 0 covered · gap 2 x r6g.large.search	USD 90.52*	medium

* Commitment decision belongs to the client (see Reserved Instance note). Exact instance-type matching; actioned domains excluded. Savings compare live On-Demand against the lowest matching offering.

Database Savings Plans Review 1 finding

ACCOUNT ↕	RESOURCE ↕	ENV ↕	RECOMMENDATION	SAVING/MO ↓	CONFIDENCE ↕	RISK ↕
—	Database Savings Plans all	—	USD 3,880.00 of On-Demand spend across 2 DSP-eligible database service(s) with no Savings Plan coverage. Evaluate a 1-year, no-upfront Database Savings Plan (up to 20% off provisioned, 35% off serverless) alongside the per-service reservation findings in this report - commit to whichever fits your roadmap. 3,880 On-Demand across 2 eligible service(s) · no SP coverage	TBD*	medium	low

* Commitment decision belongs to the client (see Reserved Instance note). Database Savings Plans (1-year, no upfront) apply automatically across engines, instance families, sizes, and regions - so they survive the modernization moves this report recommends. Discounts are published as up-to ranges (20% provisioned / 35% serverless), not per-resource rates the scanner can verify, so no dollar amount is claimed. RIs and reserved nodes discount deeper on stable shapes; do not stack both over the same usage.

RDS Reserved Instance Waste 1 finding

USD 188.50/mo

ACCOUNT ↓	RESOURCE ↓	ENV ↓	RECOMMENDATION	SAVING/MO ↓	CONFIDENCE ↓	RISK ↓
—	RDS Reserved Instances all	—	RI utilization is 87.0% (640 purchased hours unused over the lookback). Re-match reservations to running instances or exchange/resell where the term allows. 87.0% utilization · 640 unused hours	USD 188.50	medium	low

Purchased reservation hours going unused - usually family drift after migrations or scheduling of RI-covered instances. Re-match, exchange, or resell where the term allows.

▼ Expiring RDS Reservations 1 finding

ACCOUNT ↓	RESOURCE ↓	ENV ↓	RECOMMENDATION	SAVING/MO ↓	CONFIDENCE ↓	RISK
418926075334	2x db.r6g.large (mysql) db.r6g.large · us-east-1	—	Reservation expires in 44 days (2026-09-09). Decide now: renew, resize the commitment, or let it lapse deliberately. ends 2026-09-09 (44d)	TBD	high	med

Reservations ending inside the warning window. A lapsed RI silently reverts covered instances to On-Demand rates; renewal analysis should start now, not at expiry.

● Idle Resources & Scheduling 25 findings

USD 4,228/mo

▼ Idle RDS Instances: Scheduling Candidates 14 findings USD 1,812.08/mo

ACCOUNT ↓	RESOURCE ↓	ENV ↓	RECOMMENDATION	SAVING/MO ↓	CONFIDENCE ↓
517346802915 acme-ml-research	ml-experiments db.r6i.2xlarge · us-east-1	dev	Stop 21:00-07:00 America/New_York plus weekends via AWS Instance Scheduler. 30/30 idle windows · 21:00-07:00 America/New_York	USD 515.26	high
236580417962 acme-qa	qa-perf db.r5.xlarge · us-east-1	qa	Stop 21:00-07:00 America/New_York plus weekends via AWS Instance Scheduler. 30/30 idle windows · 21:00-07:00 America/New_York	USD 242.72	high
904271635840 acme-staging	stg-batch db.m5.xlarge · us-east-1	staging	Stop 21:00-07:00 America/New_York plus weekends via AWS Instance Scheduler. 30/30 idle windows · 21:00-07:00 America/New_York	USD 145.63	high
560839472016 acme-legacy-lob	lob-reports db.m5.xlarge · us-east-1	production	Stop 21:00-07:00 America/New_York plus weekends via AWS Instance Scheduler. ⚠ No non-prod environment tag; confirm with owner before scheduling 30/30 idle windows · 21:00-07:00 America/New_York	USD 145.63	low
904271635840 acme-staging	stg-web db.m6g.xlarge · us-east-1	staging	Stop 21:00-07:00 America/New_York plus weekends via AWS Instance Scheduler. 30/30 idle windows · 21:00-07:00 America/New_York	USD 132.01	high

ACCOUNT ↓	RESOURCE ↓	ENV ↓	RECOMMENDATION	SAVING/MO ↓	CONFIDENCE ↓
795013264871 acme-dev-sandbox	dev-shared db.t4g.xlarge · us-east-1	dev	Stop 21:00-07:00 America/New_York plus weekends via AWS Instance Scheduler. 30/30 idle windows · 21:00-07:00 America/New_York	USD 129.45	high
517346802915 acme-ml-research	ml-notebooks db.m6g.xlarge · us-east-1	dev	Stop 21:00-07:00 America/New_York (weekdays and weekend nights; weekend daytime stays up) via AWS Instance Scheduler. ⚠ Weekend daytime traffic observed; saving covers nightly windows only, not full weekends 30/30 idle windows · 21:00-07:00 America/New_York · nights only (weekends busy)	USD 94.29	high
236580417962 acme-qa	qa-main db.m5.large · us-east-1	qa	Stop 21:00-07:00 America/New_York plus weekends via AWS Instance Scheduler. 30/30 idle windows · 21:00-07:00 America/New_York	USD 72.82	high
795013264871 acme-dev-sandbox	dev-legacy db.m5.large · us-east-1	—	Stop 21:00-07:00 America/New_York plus weekends via AWS Instance Scheduler. ⚠ No non-prod environment tag; confirm with owner before scheduling 30/30 idle windows · 21:00-07:00 America/New_York	USD 72.82	low

ACCOUNT ↓	RESOURCE ↓	ENV ↓	RECOMMENDATION	SAVING/MO ↓	CONFIDENCE ↓
904271635840 acme-staging	stg-orders db.m6g.large · us-east-1	staging	Stop 21:00-07:00 America/New_York plus weekends via AWS Instance Scheduler. 30/30 idle windows · 21:00-07:00 America/New_York	USD 66.00	high
904271635840 acme-staging	stg-payments db.m6g.large · us-east-1	staging	Stop 21:00-07:00 America/New_York plus weekends via AWS Instance Scheduler. 30/30 idle windows · 21:00-07:00 America/New_York	USD 66.00	high
236580417962 acme-qa	qa-automation db.t4g.large · us-east-1	qa	Stop 21:00-07:00 America/New_York plus weekends via AWS Instance Scheduler. 30/30 idle windows · 21:00-07:00 America/New_York	USD 64.73	high
795013264871 acme-dev-sandbox	dev-alpha db.t4g.medium · us-east-1	dev	Stop 21:00-07:00 America/New_York plus weekends via AWS Instance Scheduler. 30/30 idle windows · 21:00-07:00 America/New_York	USD 32.36	high
795013264871 acme-dev-sandbox	dev-beta db.t4g.medium · us-east-1	dev	Stop 21:00-07:00 America/New_York plus weekends via AWS Instance Scheduler. 30/30 idle windows · 21:00-07:00 America/New_York	USD 32.36	high

An instance qualifies when it shows zero connections during the stop window on 95%+ of observed windows over the full lookback. Each row states its window in the region's business timezone (e.g. Asia/Kolkata for ap-south-1) - never assume UTC. Savings cover the nightly windows plus weekends, and weekend daytime is checked separately: any weekend traffic drops the estimate to the nightly windows alone rather than assuming the weekend is free. Stopped compute hours only; storage and backups keep billing. RDS auto-starts instances stopped 7+ days; AWS Instance Scheduler handles that rotation.

DocumentDB Stop Schedules 1 finding

USD 66.60/mo

ACCOUNT ↕	RESOURCE ↕	ENV ↕	RECOMMENDATION	SAVING/MO ↓	CONFIDENCE ↕	RISK ↕
795013264871 acme-dev-sandbox	docs-dev db.t3.medium · us-east-1	dev	All 2 cluster member(s) idle on 30/30 observed windows. Stop the cluster 21:00-07:00 America/New_York plus weekends. 30/30 idle windows · 21:00-07:00 America/New_York	USD 66.60	high	low

DocumentDB clusters stop like RDS instances (auto-restart after 7 days; storage keeps billing). Window evaluated per member in the region's business timezone shown on each row.

Neptune Stop Schedules 1 finding

USD 148.19/mo

ACCOUNT ↕	RESOURCE ↕	ENV ↕	RECOMMENDATION	SAVING/MO ↓	CONFIDENCE ↕	RISK ↕
517346802915 acme-ml-research	graph-dev db.r5.large · us-east-1	dev	All 1 cluster member(s) idle on 30/30 observed windows. Stop the cluster 21:00-07:00 America/New_York plus weekends. 30/30 idle windows · 21:00-07:00 America/New_York	USD 148.19	high	low

Neptune clusters support stop/start with the same 7-day auto-restart. Window evaluated per member in the region's business timezone shown on each row.

Redshift Pause Schedules 1 finding

USD 924.91/mo

ACCOUNT ↓	RESOURCE ↓	ENV ↓	RECOMMENDATION	SAVING/MO ↓	CONFIDENCE ↓	RISK ↓
795013264871 acme-dev-sandbox	bi-dev-wh ra3.xlplus x 2 · us-east-1	dev	Zero connections on 30/30 observed windows. Use Redshift's native pause schedule (pause 21:00, resume 07:00 America/New_York, paused weekends). 30/30 idle windows · 21:00-07:00 America/New_York · 2 nodes	USD 924.91	high	low

Redshift's native pause/resume schedule stops per-hour compute billing during the stated stop window and weekends; managed storage continues. Each row states its window and timezone. Resume takes minutes - built for exactly these non-production warehouses.

▼ Idle RDS Read Replicas

1 finding

USD 386.90/mo

ACCOUNT ↓	RESOURCE ↓	ENV ↓	RECOMMENDATION	SAVING/MO ↓	CONFIDENCE ↓
349072118463 acme-prod-core	core-orders-replica db.r6g.xlarge · us-east-1	production	Read replica took zero connections across 720 observed hours. Delete it; recreate from the source in minutes if ever needed. 0 connections over 720 hours	USD 386.90	high

Read replicas with zero connections across the whole window serve no traffic. Deleting one saves its full compute and storage; recreating from the source later takes minutes.

▼ Idle ElastiCache Clusters

1 finding

USD 157.68/mo

ACCOUNT ↕	RESOURCE ↕	ENV ↕	RECOMMENDATION	SAVING/MO ↓	CONFIDENCE ↕	RISK
795013264871 acme-dev-sandbox	dev-cache cache.r5.large · us-east-1	dev	Zero connections across 720 observed hours on 1 node(s). ElastiCache cannot be stopped - delete the cluster (snapshot first, Redis restores in minutes). 0 connections over 720 hours · 1 node(s)	USD 157.68	high	low

Caches with zero connections across the whole window bill 24/7 and cannot be stopped. Redis/Valkey snapshot to S3 and restore in minutes; memcached holds no durable data. Delete after confirming ownership.

Idle OpenSearch Domains

1 finding

USD 243.82/mo

ACCOUNT ↕	RESOURCE ↕	ENV ↕	RECOMMENDATION	SAVING/MO ↓	CONFIDENCE ↕	RISK
795013264871 acme-dev-sandbox	logs-dev r5.large.search x 2 · us-east-1	dev	Zero search and indexing traffic across the whole window. Take a manual snapshot to S3, then delete the domain. 0 search · 0 indexing all window	USD 243.82	high	low

Domains with zero search and indexing traffic all window bill 24/7 and cannot be stopped. Snapshot to S3, delete, restore in under an hour if ever needed.

Idle MemoryDB Clusters

1 finding

USD 112.57/mo

ACCOUNT ↕	RESOURCE ↕	ENV ↕	RECOMMENDATION	SAVING/MO ↓	CONFIDENCE ↕	RIS
795013264871 acme-dev-sandbox	dev-tokens db.t4g.medium · us-east-1	dev	Zero connections across 720 observed hours on 2 node(s). Take a final snapshot and delete the cluster. 0 connections over 720 hours · 2 node(s)	USD 112.57	high	lo

MemoryDB is durable Redis: a final snapshot preserves the full dataset, and idle clusters bill every node 24/7. Delete after confirming ownership.

Leftover DMS Instances

1 finding

USD 245.28/mo

ACCOUNT ↓	RESOURCE ↓	ENV ↓	RECOMMENDATION	SAVING/MO ↓	CONFIDENCE ↓	RISK
560839472016 acme-legacy-lob	legacy-migration dms.c5.xlarge · us-east-1	—	Replication instance has 1 task(s), none active. Export the task configurations and delete the instance; recreate in minutes if another migration comes. 1 task(s), none active	USD 245.28	high	low

DMS replication instances bill 24/7 whether tasks run or not - the classic leftover from a finished one-time migration. Export task configurations, delete the instance, recreate in minutes when the next migration comes.

Idle DynamoDB Tables

1 finding

USD 6.25/mo

ACCOUNT ↓	RESOURCE ↓	ENV ↓	RECOMMENDATION	SAVING/MO ↓	CONFIDENCE ↓	RISK
795013264871 acme-dev-sandbox	temp-migration PAY_PER_REQUEST · us-east-1	—	No reads or writes across the whole window. Export to S3 (or take an on-demand backup), then delete the table. 25.0 GB	USD 6.25	medium	low

No reads or writes across the whole window. Export to S3 or take an on-demand backup, then delete. Confirm no seasonal consumer first.

DynamoDB Unused Indexes

1 finding

USD 113.88/mo

ACCOUNT ↕	RESOURCE ↕	ENV ↕	RECOMMENDATION	SAVING/MO ↓	CONFIDENCE ↕
682415903277 acme-data-platform	user-profiles / email-index PROVISIONED · us-east-1	production	GSI 'email-index' served zero reads all window while replicating every table write. Confirm no consumer, then delete the index. 0 reads all window	USD 113.88	high

A GSI that serves no reads still replicates every table write and bills its own capacity and storage. Deleting is safe; recreating later is a background backfill.

▼ Stopped RDS Instances Still Billing 1 finding USD 10.00/mo

ACCOUNT ↕	RESOURCE ↕	ENV ↕	RECOMMENDATION	SAVING/MO ↓	CONFIDENCE ↕	RISK ↕
682415903277 acme-data-platform	dp-old-migration db.m5.large · us-east-1	—	Stopped instance still bills 500 GB of storage every month. If nobody claims it, take a final snapshot and delete. 500 GB still billing	USD 10.00	low	medi

Stopped instances keep billing storage, and RDS restarts them every 7 days. Long-stopped usually means forgotten: snapshot and delete after confirming ownership.

■ Right-Sizing 11 findings USD 9,459/mo

▼ RDS Right-Sizing Candidates 3 findings USD 3,368.95/mo

ACCOUNT ↕	RESOURCE ↕	ENV ↕	RECOMMENDATION	SAVING/MO ↓	CONFIDENCE ↕
517346802915 acme-ml-research	ml-feature-store db.r6i.8xlarge · us-east-1	production	Downsize db.r6i.8xlarge to db.r6i.4xlarge in a maintenance window (same family keeps size-flexible RI coverage). p95 CPU 12% · peak 25% · 60d data	USD 1,762.95	high

ACCOUNT ↕	RESOURCE ↕	ENV ↕	RECOMMENDATION	SAVING/MO ↓	CONFIDENCE ↕
682415903277 acme-data-platform	dp-warehouse db.r5.4xlarge · us-east-1	production	Downsize db.r5.4xlarge to db.r5.2xlarge in a maintenance window (same family keeps size-flexible RI coverage). p95 CPU 13% · peak 24% · 60d data	USD 832.20	high
349072118463 acme-prod-core	core-analytics db.r6g.4xlarge · us-east-1	production	Downsize db.r6g.4xlarge to db.r6g.2xlarge in a maintenance window (same family keeps size-flexible RI coverage). p95 CPU 11% · peak 22% · 60d data	USD 773.80	high

Decision uses p95 of hourly CPU maxima over the lookback with a spike guard, plus verified memory headroom where instance RAM is known — never plain averages. Targets stay one size down in the same family, preserving size-flexible RI coverage. Execute in a maintenance window; expect a brief failover/restart.

▼ ElastiCache Right-Sizing

1 finding

USD 153.30/mo

ACCOUNT ↕	RESOURCE ↕	ENV ↕	RECOMMENDATION	SAVING/MO ↓	CONFIDENCE ↕
349072118463 acme-prod-core	feature-cache cache.r6g.xlarge · us-east-1	production	Downsize cache.r6g.xlarge to cache.r6g.large across 1 node(s). Redis scales down online; memcached needs a replacement cluster. p95 CPU 14% · memory 38% · target cache.r6g.large	USD 153.30	high

Decision uses p95 CPU (EngineCPUUtilization where the engine reports it) with a spike guard, plus memory usage where available. Redis scales down online; memcached needs a replacement cluster. Savings use live public node rates.

▼ Redshift Node Reduction

1 finding

USD 3,171.12/mo

ACCOUNT ↕	RESOURCE ↕	ENV ↕	RECOMMENDATION	SAVING/MO ↓	CONFIDENCE ↕
682415903277 acme-data- platform	analytics- wh ra3.xlplus x 8 · us-east-1	production	Elastic-resize from 8 to 4 nodes: p95 CPU 18% and disk 32% leave headroom even after halving. Resize completes online in minutes. p95 CPU 18% · disk 32% · 8 to 4 nodes	USD 3,171.12	high

Elastic resize sheds nodes online in minutes. Recommended only when p95 CPU is low AND disk usage leaves room for each remaining node's share to double after halving.

▼ OpenSearch Right-Sizing 1 finding USD 735.84/mo

ACCOUNT ↕	RESOURCE ↕	ENV ↕	RECOMMENDATION	SAVING/MO ↓	CONFIDENCE ↕
349072118463 acme-prod- core	search-prod r6g.xlarge.search x 6 · us-east-1	production	Downsize r6g.xlarge.search to r6g.large.search across 6 data node(s). Blue/green deployment handles the change without downtime. p95 CPU 14% · JVM 60% · target r6g.large.search	USD 735.84	high

p95 CPU with a spike guard plus verified JVM memory pressure headroom - heap-bound domains never shrink. Blue/green deployment applies the change without downtime.

▼ OpenSearch Master Oversizing 1 finding USD 367.92/mo

ACCOUNT ↕	RESOURCE ↕	ENV ↕	RECOMMENDATION	SAVING/MO ↓	CONFIDENCE ↕
349072118463 acme-prod-core	search-prod r6g.xlarge.search x 6 · us-east-1	production	3x r6g.xlarge.search dedicated masters front only 6 data node(s). r6g.large.search masters are AWS's guidance at this scale. 3x r6g.xlarge.search · target r6g.large.search	USD 367.92	medium

Dedicated masters coordinate cluster state; they serve no queries. AWS guidance sizes them by data-node count, and small domains need small masters.

▼ MemoryDB Right-Sizing

1 finding

USD 902.28/mo

ACCOUNT ↕	RESOURCE ↕	ENV ↕	RECOMMENDATION	SAVING/MO ↓	CONFIDENCE ↕
349072118463 acme-prod-core	session-store db.r6g.xlarge · us-east-1	production	Downsize db.r6g.xlarge to db.r6g.large across 4 node(s); MemoryDB scales node type online per shard. p95 CPU 15% · target db.r6g.large	USD 902.28	high

p95 engine CPU with a spike guard plus memory usage where reported. MemoryDB scales node type online per shard.

▼ DynamoDB Over-Provisioned Capacity

1 finding

USD 503.92/mo

ACCOUNT ↕	RESOURCE ↕	ENV ↕	RECOMMENDATION	SAVING/MO ↓	CONFIDENCE ↕
349072118463 acme-prod-core	orders-events PROVISIONED · us-east-1	production	Lower provisioned capacity to ~390 RCU / 260 WCU (p95 consumption + 130% headroom) and attach auto scaling. 17% utilized · 2000/1000 RCU/WCU → 390/260	USD 503.92	high

p95 of hourly consumed capacity against the provisioned floor. Targets are p95 plus headroom with auto scaling attached to absorb anything above.

▼ DynamoDB Billing Mode

1 finding

USD 255.86/mo

ACCOUNT ↓	RESOURCE ↓	ENV ↓	RECOMMENDATION	SAVING/MO ↓	CONFIDENCE ↓
418926075334 acme- marketing	clickstream PROVISIONED · us-east-1	production	Switch to on-demand: p95 consumption is 2% of provisioned capacity. On-demand bills actual requests only - right answer for this traffic shape. to on-demand · 2% utilized	USD 255.86	high

Provisioned tables running near-empty pay for capacity nobody uses; on-demand tables with steady heavy traffic pay a premium per request. Each row states the direction. One-click change, reversible once per 24h.

▼ RDS CPU Spikes: Fix Before Right-Sizing

1 finding

potential USD 416.10/mo

ACCOUNT ↓	RESOURCE ↓	ENV ↓	RECOMMENDATION	POTENTIAL/MO ↓	CONFIDENCE ↓
682415903277 acme-data- platform	dp-etl- primary db.r5.2xlarge · us-east-1	production	p95 CPU is only 14% but the window contains a 93% spike. Find and fix the spike source (batch job, missing index, backup overlap) before downsizing to db.r5.xlarge. Potential saving after the fix: ~USD 416.10/month. p95 CPU 14% · peak 93% · 60d data	USD 416.10	high

These instances look oversized on p95 CPU, but at least one spike above the threshold occurred in the window. Downsizing now would amplify the spike on smaller hardware. Fix the cause first; each then becomes a normal right-sizing candidate at the stated potential saving.

Graviton Migration (RDS, DocumentDB, Neptune, OpenSearch)

7 findings

USD 116.07/mo

ACCOUNT ↕	RESOURCE ↕	ENV ↕	RECOMMENDATION	SAVING/MO ↓	CONFIDENCE ↕
418926075334 acme-marketing	sess-cache-prod cache.r5.xlarge · us-east-1	production	Move cache.r5.xlarge to cache.r7g.xlarge across 2 node(s). Redis/Memcached run identically on Graviton; modify with a rolling replacement. ⚠ Also evaluate the Valkey engine (~20% below Redis OSS) in the same maintenance window - one migration event, two savings target cache.r7g.xlarge	USD 52.56	high
560839472016 acme-legacy-lob	legacy-mc cache.m4.large · us-east-1	production	Move cache.m4.large to cache.m6g.large across 3 node(s). Redis/Memcached run identically on Graviton; modify with a rolling replacement. target cache.m6g.large	USD 19.71	high
560839472016 acme-legacy-lob	legacy-es m5.large.search x 3 · us-east-1	production	Move m5.large.search to m6g.large.search across 3 node(s). Blue/green deployment, no client changes. target m6g.large.search	USD 17.52	high

ACCOUNT ↕	RESOURCE ↕	ENV ↕	RECOMMENDATION	SAVING/MO ↓	CONFIDENCE ↕
682415903277 acme-data-platform	docs-prod-1 db.r5.large · us-east-1	production	Move db.r5.large to db.r6g.large (Graviton). In-place modify with a brief failover. target db.r6g.large	USD 13.14	high
682415903277 acme-data-platform	docs-prod-2 db.r5.large · us-east-1	production	Move db.r5.large to db.r6g.large (Graviton). In-place modify with a brief failover. target db.r6g.large	USD 13.14	high
682415903277 acme-data-platform	dp-metadata db.r5.xlarge · us-east-1	production	Move db.r5.xlarge to db.r8g.xlarge (Graviton). Same size, lower rate, typically better throughput; modify in a maintenance window. target db.r8g.xlarge	TBD	high
560839472016 acme-legacy-lob	lob-crm db.m5.2xlarge · us-east-1	production	Move db.m5.2xlarge to db.m8g.2xlarge (Graviton). Same size, lower rate, typically better throughput; modify in a maintenance window. target db.m8g.2xlarge	TBD	high

Same-size move from x86 to AWS Graviton, typically 10-20% cheaper with equal or better throughput. Managed engines need no application changes; execute as a normal class modification in a maintenance window. Instances already flagged for right-sizing are excluded here - do one class move, not two.

▼ Previous-Generation Upgrades (RDS, DocumentDB, Neptune)

2 findings

USD 82.49/mo

ACCOUNT ↕	RESOURCE ↕	ENV ↕	RECOMMENDATION	SAVING/MO ↓	CONFIDENCE ↕
560839472016 acme-legacy-lob	lob-archive db.m4.xlarge · us-east-1	production	Upgrade previous-generation db.m4.xlarge to db.m6i.xlarge. Newer generation is cheaper per unit of performance and keeps current-gen availability. target db.m6i.xlarge	USD 42.34	high
418926075334 acme-marketing	campaign-db db.r6g.xlarge · us-east-1	production	Refresh db.r6g.xlarge to db.r8g.xlarge (newer Graviton generation). Same size at a lower live rate with better per-vCPU performance; in-place modify. target db.r8g.xlarge	USD 40.15	high

Previous-generation instance classes cost more per unit of performance than current ones. Listed instances run engines without a Graviton path; the upgrade stays on x86.

▼ Valkey Engine Migration (ElastiCache, MemoryDB)

2 findings

USD 194.18/mo

ACCOUNT ↕	RESOURCE ↕	ENV ↕	RECOMMENDATION	SAVING/MO ↓	CONFIDENCE ↕
349072118463 acme-prod-core	profile-store db.r6g.large · us-east-1	production	Switch the engine from Redis to Valkey on the same db.r6g.large nodes (2 node(s)). MemoryDB upgrades the engine in place; durability and clients are unaffected. redis → valkey · 2 node(s), same node type	USD 134.32	high

ACCOUNT ↕	RESOURCE ↕	ENV ↕	RECOMMENDATION	SAVING/MO ↓	CONFIDENCE ↕
349072118463 acme-prod-core	leaderboard-cache cache.r6g.large · us-east-1	production	Switch the engine from Redis OSS to Valkey on the same cache.r6g.large nodes (2 node(s)). ElastiCache upgrades Redis 7.x to Valkey in place with no data loss; clients keep working unchanged. redis → valkey · 2 node(s), same node type	USD 59.86	high

Valkey is the Linux Foundation fork of Redis OSS 7.2 - same API, same tools - and AWS lists it ~20% below Redis OSS on ElastiCache nodes (~30% on MemoryDB, 33% on serverless). The engine upgrades in place with no data loss. Savings compare live rates for both engines on the exact node type; clusters already flagged for deletion, right-sizing, or a class move are excluded - one change at a time.

Extended Support Surcharges (RDS, DocumentDB, OpenSearch)

4 findings

USD 745.77/mo

ACCOUNT ↕	RESOURCE ↕	ENV ↕	RECOMMENDATION	SAVING/MO ↓	CONFIDENCE ↕
418926075334 acme-marketing	mkt-tracker db.t3.large · us-east-1	production	postgres 11.22 is past standard support and bills the RDS Extended Support surcharge at the doubled year-3 rate. Upgrade the engine major version to stop the charge. engine version 11.22	USD 292.00	high

ACCOUNT ↕	RESOURCE ↕	ENV ↕	RECOMMENDATION	SAVING/MO ↓	CONFIDENCE ↕
560839472016 acme-legacy-lob	lob-finance db.m5.large · us-east-1	production	mysql 5.7.44 is past standard support and bills the RDS Extended Support surcharge at the doubled year-3 rate. Upgrade the engine major version to stop the charge. engine version 5.7.44	USD 292.00	high
560839472016 acme-legacy-lob	legacy-docs-1 db.r5.large · us-east-1	production	DocumentDB 3.6 bills Extended Support at a 80% premium on the instance rate (doubles in year 3). Upgrade the cluster to 5.0+ to remove the surcharge. engine version 3.6.0	USD 161.77	high
560839472016 acme-legacy-lob	legacy-es m5.large.search x 3 · us-east-1	production	elasticsearch 7.10 bills the OpenSearch extended-support uplift on every node-hour. Upgrade to a current OpenSearch version to stop the charge. engine version 7.10	TBD	high

Engine versions past standard support bill the RDS Extended Support surcharge per vCPU-hour, doubling in year 3 - versions already in year 3 are estimated at the doubled rate. DocumentDB 3.6 instead bills a percentage premium on the instance rate. Because the surcharge accrues only while an instance runs, instances this report also recommends scheduling off are sized on their remaining running hours. Upgrading the engine removes the charge and restores full security patching.

▼ RDS Multi-AZ on Non-Production

1 finding

USD 113.15/mo

ACCOUNT ↓	RESOURCE ↓	ENV ↓	RECOMMENDATION	SAVING/MO ↓	CONFIDENCE ↓	RI
904271635840 acme-staging	stg-integration db.m6g.large · us-east-1	staging	Convert Multi-AZ to Single-AZ: 'staging' environments rarely justify a hot standby. Halves the compute rate.	USD 113.15	high	I



Multi-AZ bills the standby at the full instance rate. Non-production environments that tolerate minutes of downtime rarely justify it; converting to Single-AZ halves compute cost.

▼ Redshift Graviton (RG) Migration 1 finding USD 2,855.03/mo

ACCOUNT ↓	RESOURCE ↓	ENV ↓	RECOMMENDATION	SAVING/MO ↓	CONFIDENCE ↓
682415903277 acme-data-platform	etl-transform-wh ra3.4xlarge x 4 · us-east-1	production	Migrate 4x ra3.4xlarge to 3x rg.4xlarge (AWS Graviton) per AWS's RA3-to-RG mapping. Elastic resize completes the move online in minutes; validate via snapshot-restore first if preferred. ⚠ 2,400 GB of manual snapshots currently bill at full size on RA3; on RG only unique blocks bill 4 nodes → 3x rg.4xlarge · 2,400 GB manual snapshots	USD 2,855.03	medium



Graviton RG nodes follow AWS's published RA3 mapping: rg.xlarge replaces ra3.xlplus one-for-one, and three rg.4xlarge replace four ra3.4xlarge (each RG node carries 1.33x the vCPU and memory). Savings compare live node rates; findings appear only in regions where a live RG rate exists. Three further RG advantages are counted as zero: Spectrum's per-TB scan charge disappears, warehouse workloads benchmark up to 2.2x faster, and manual snapshots bill on unique data blocks only (incremental snapshot billing, June 2026) where RA3 keeps paying full size per snapshot - the deduplication ratio cannot be known before migrating, so it is evidence, not a claim. Elastic resize performs the move online; clusters already flagged for pause or resize are excluded.

▼ Redshift RA3 Migration 1 finding

ACCOUNT ↕	RESOURCE ↕	ENV ↕	RECOMMENDATION	SAVING/MO ↓	CONFIDENCE ↕
560839472016 acme-legacy-lob	legacy-dw dc2.large x 4 · us-east-1	production	dc2.large is a previous-generation node. Evaluate RA3 with managed storage: compute and storage scale independently, which usually beats dc2.large economics and unblocks data sharing and cross-AZ features.	TBD	medium

Previous-generation node types couple compute to storage. RA3 with managed storage scales them independently and usually wins economically; sizing depends on the storage profile, so no automated saving is claimed.

▼ OpenSearch Optimized Instances 1 finding

ACCOUNT ↕	RESOURCE ↕	ENV ↕	RECOMMENDATION	SAVING/MO ↓	CONFIDENCE ↕
682415903277 acme-data-platform	ingest-logs r7g.xlarge.search x 3 · us-east-1	production	Indexing-heavy domain on r7g.xlarge.search: evaluate OpenSearch Optimized OR2/OM2 instances. S3-backed replicas remove replica EBS volumes entirely, and OR2 indexes up to 70% faster than r7g at equal size. peak indexing 1,400/s vs search 300/s	TBD	low

OpenSearch Optimized (OR2/OM2) instances use S3-backed replicas - replica EBS volumes disappear - and index up to 70% faster than r7g at equal size. Flagged domains are indexing-heavy on general-purpose families; sizing depends on the workload mix, so the finding is consultative with no claimed saving.

Storage: gp2 to gp3 (RDS, OpenSearch)

4 findings

USD 133.47/mo

ACCOUNT ↕	RESOURCE ↕	ENV ↕	RECOMMENDATION	SAVING/MO ↓	CONFIDENCE ↕
349072118463 acme-prod-core	core-orders-api db.r6g.2xlarge · us-east-1	production	Convert 2,000 GB gp2 volume to gp3: lower per-GB rate and a free 3,000-IOPS baseline. Online change, no downtime. ⚠ gp2 baseline above 3,000 IOPS at 2,000 GB; verify IOPS needs before converting 2,000 GB gp2	USD 46.00	high
560839472016 acme-legacy-lob	legacy-es m5.large.search x 3 · us-east-1	production	Convert 1,536 GB of gp2 EBS (across 3 data nodes) to gp3. Online change per node. 1,536 GB gp2	USD 41.47	high
682415903277 acme-data-platform	dp-metadata db.r5.xlarge · us-east-1	production	Convert 1,500 GB gp2 volume to gp3: lower per-GB rate and a free 3,000-IOPS baseline. Online change, no downtime. ⚠ gp2 baseline above 3,000 IOPS at 1,500 GB; verify IOPS needs before converting 1,500 GB gp2	USD 34.50	high
236580417962 acme-qa	qa-perf db.r5.xlarge · us-east-1	qa	Convert 500 GB gp2 volume to gp3: lower per-GB rate and a free 3,000-IOPS baseline. Online change, no downtime. 500 GB gp2	USD 11.50	high

gp3 costs ~20% less per GB than gp2 and includes a 3,000-IOPS baseline regardless of size. Conversion is online with no downtime. Savings use live per-GB rates for both volume types.

▼ RDS Storage: Provisioned IOPS to gp3

1 finding

USD 726.40/mo

ACCOUNT ↓	RESOURCE ↓	ENV ↓	RECOMMENDATION	SAVING/MO ↓	CONFIDENCE ↓
349072118463 acme-prod-core	core-ledger db.r6g.xlarge · us-east-1	production	Move io1 (8,000 provisioned IOPS, 800 GB) to gp3 with the same IOPS. Identical performance ceiling at this level, far cheaper IOPS. 8,000 provisioned IOPS · peak observed 2,900	USD 726.40	high

io1/io2 volumes provisioned at 16,000 IOPS or less fit inside gp3's range, where IOPS above the free baseline cost a fraction of provisioned-IOPS rates. Peak observed IOPS are shown as evidence that the ceiling is safe.

▼ RDS Snapshot Hygiene

2 findings

USD 370.50/mo

ACCOUNT ↓	RESOURCE ↓	ENV ↓	RECOMMENDATION	SAVING/MO ↓	CONFIDENCE ↓	RISK ↓
560839472016	3 manual snapshots us-east-1	—	Review 3 manual snapshots older than 90 days or orphaned (source DB deleted). Keep what compliance requires, delete the rest. 2,600 GB old · 2,600 GB orphaned	USD 247.00	medium	medium
682415903277	2 manual snapshots us-east-1	—	Review 2 manual snapshots older than 90 days or orphaned (source DB deleted). Keep what compliance requires, delete the rest. 1,300 GB old · 1,300 GB orphaned	USD 123.50	medium	medium

Manual snapshots never expire on their own. Old and orphaned snapshots (source instance deleted) accumulate storage cost silently. Savings are approximate - snapshots bill on used, not allocated, bytes. Deletion is irreversible; check

retention policy first.

▼ DynamoDB Standard-IA Candidates

1 finding

USD 119.29/mo

ACCOUNT ↓	RESOURCE ↓	ENV ↓	RECOMMENDATION	SAVING/MO ↓	CONFIDENCE ↓
560839472016 acme-legacy- lob	audit- archive PROVISIONED · us-east-1	production	Storage (800 GB) dominates this table's cost. Switch to the Standard-IA table class: ~60% cheaper storage, net of the higher request prices. 800.0 GB	USD 119.29	medium

Standard-IA cuts storage ~60% but raises request and capacity prices ~25%. Flagged only where storage cost exceeds throughput cost, and the saving shown is net: the storage discount minus the request premium on current throughput.

▼ DynamoDB PITR on Non-Production

1 finding

USD 8.00/mo

ACCOUNT ↓	RESOURCE ↓	ENV ↓	RECOMMENDATION	SAVING/MO ↓	CONFIDENCE ↓	RISK
236580417962 acme-qa	qa-fixtures PROVISIONED · us-east-1	qa	Point-in-time recovery is enabled on a 'qa' table. Non-prod data that can be regenerated rarely justifies continuous backup billing. 40.0 GB	USD 8.00	high	low

Point-in-time recovery bills per GB-month continuously. Non-production data that can be regenerated rarely justifies it.

■ Consultative Reviews 7 findings

USD 1,363/mo

▼ Aurora I/O-Optimized Evaluation

1 finding

USD 1,362.59/mo

ACCOUNT ↕	RESOURCE ↕	ENV ↕	RECOMMENDATION	SAVING/MO ↓	CONFIDENCE ↕	RISK ↕
—	Aurora clusters (account level) all	—	Aurora I/O is 78% of Aurora spend (USD 1,800.00/mo). Evaluate switching clusters to I/O-Optimized: instances +30%, storage +125%, I/O free. I/O is 78% of Aurora spend	USD 1,362.59	low	low

Aurora I/O-Optimized trades +30% instance and +125% storage cost for free I/O. Worth it when I/O runs above ~25% of cluster spend. Account-level estimate; evaluate and switch per cluster.

▼ RDS Burstable Credit Overruns 1 finding

ACCOUNT ↕	RESOURCE ↕	ENV ↕	RECOMMENDATION	SAVING/MO ↓	CONFIDENCE ↕
418926075334 acme-marketing	mkt-tracker db.t3.large · us-east-1	production	Instance was charged 5,200 surplus CPU credits over the lookback: it runs beyond the T-family baseline permanently. Move to a fixed-performance class (m-family) or the surcharge continues indefinitely. 5,200 surplus credits charged	TBD	medium

T-family instances charged surplus CPU credits run beyond baseline permanently - a hidden hourly surcharge. Fixed-performance classes usually cost less at sustained load.

▼ Aurora Serverless Candidates 1 finding

ACCOUNT ↕	RESOURCE ↕	ENV ↕	RECOMMENDATION	SAVING/MO ↓	CONFIDENCE ↕	RISK ↕
517346802915 acme-ml-research	ml-aurora-dev db.r6g.large · us-east-1	dev	Non-prod Aurora instance with p95 CPU 12%: a candidate for Aurora Serverless v2, which scales to near zero between test runs instead of billing a fixed class. p95 CPU 12%	TBD	low	medi

Spiky, mostly-idle non-production Aurora workloads fit Serverless v2, which scales capacity to near zero between runs. Model ACU consumption before converting.

▼ OpenSearch Serverless Candidates 1 finding

ACCOUNT ↕	RESOURCE ↕	ENV ↕	RECOMMENDATION	SAVING/MO ↓	CONFIDENCE ↕
795013264871 acme-dev-sandbox	search-sandbox r6g.large.search x 1 · us-east-1	sandbox	Mostly-idle non-prod domain (p95 CPU 8%): evaluate OpenSearch Serverless NextGen, which scales OCU's to zero after 10 minutes of inactivity and bills per usage with no minimum. p95 CPU 8% · 60d data	TBD	low

OpenSearch Serverless NextGen scales indexing and search OCU's to zero after 10 minutes of inactivity with no minimum. Flagged domains are mostly-idle non-production. Model OCU consumption before migrating - steady workloads stay cheaper on provisioned domains.

▼ RDS Commercial License Review 1 finding

ACCOUNT ↕	RESOURCE ↕	ENV ↕	RECOMMENDATION	SAVING/MO ↓	CONFIDENCE ↕
560839472016 acme-legacy-lob	lob-erp db.r5.2xlarge · us-east-1	production	oracle-ee on License Included: review whether Standard edition (SE2) covers the feature set, or whether existing enterprise agreements make BYOL cheaper. oracle-ee · license- included	TBD	low

Enterprise-edition license premiums often exceed infrastructure cost. Feature-usage review with the DBA team decides whether Standard edition or BYOL applies. Consultative - no automated saving claimed.

Other Database Services Review 2 findings

ACCOUNT ↕	RESOURCE ↕	ENV ↕	RECOMMENDATION	SAVING/MO ↓	CONFIDENCE ↕	RISK ↕
—	Amazon Timestream all	—	USD 840.00 over the lookback on a service the scanner does not deep-dive yet. Review memory-store retention (the expensive tier) versus magnetic-store retention per table; most estates keep far more in memory than queries touch. 840.00 in lookback	TBD	low	low

ACCOUNT ↕	RESOURCE ↕	ENV ↕	RECOMMENDATION	SAVING/MO ↓	CONFIDENCE ↕	RISK ↕
—	Amazon Quantum Ledger Database all	—	USD 310.00 on QLDB, which AWS has deprecated (end-of-support announced; no new features or new accounts). Plan the migration now - AWS's documented path is Aurora PostgreSQL with an audit table pattern. ⚠ Service deprecated by AWS - this is a migration deadline, not an optimization 310.00 in lookback	TBD	high	high

Services with real spend but no deep-dive analyzer yet: the figure marks the opportunity, the checklist starts the review. Deprecated services (QLDB) are migration deadlines, not optimizations.

Findings are estimates based on the stated lookback window and public pricing; private rates, credits, and tax are not reflected. Savings marked * are commitment decisions that depend on future usage plans only you can know. Execute changes in maintenance windows. Read-only access only; no infrastructure was modified.