



AWS COST · SAMPLE DELIVERABLE

Week-Six Savings Measurement

Week 6 of 6: Measure · prepared for Seatlark, Inc.

Prepared by Charles Sieg · charles.sieg@vantalect.com · vantalect.com

June 17, 2026 · AWS Cost Analysis & Optimization: The Optimization (six weeks)

SAMPLE DELIVERABLE

This is a sample deliverable. Seatlark, Inc. is a fictional company: every account, resource, finding, figure, and person in this report is simulated to show the form and depth of a Vantalect engagement. AWS prices are list prices from the AWS Price List API retrieved on October 13 and 16, 2026, and AWS Business Support+ pricing is from the AWS Support pricing page on October 16, 2026.

Contents

Result	2
The method	4
The query for one change	5
The daily cost	7
Measured against estimated	9
Every change	10
Where the measurement fell short	11
The gate	12
What is not counted, and when it comes	13
Robustness	15
One-time charges	15
Approved, then declined	16
What is deliberately not counted	16
Re-running this measurement	17

Result

The changes in place cut \$148,967 a year from Seatlark's AWS bill, measured on your own Cost and Usage Report. That is 96.1% of the \$155,076 the Analysis estimated for the same 20 changes, and 24.0% of the \$619,914 annual bill we started from. Counting the one change you approved at the readout and then declined, as the engagement's terms define it, the figure for the guarantee is **\$152,076: 4.35 times** the \$35,000 fee, against a gate of \$105,000. The gate is met, and the second invoice, \$17,500, is issued with this report.

Without that one change the multiple is 4.26 times, so the gate does not depend on it. The \$35,000 fee is everything Seatlark paid: the Analysis was the engagement's first two weeks, not a second fee. The measured savings repay it in 2.8 months, and the first year nets \$113,967 after it; the counted figure decides the gate and nothing else. The one-time charges of the changes, \$3,117.00, are listed beside the result in [One-time charges](#) and never annualized; with them deducted, the counted figure is \$148,959, still 4.26 times the fee.


The result against the gate. The full computation is in [The gate](#).

The result	Amount
Measured savings, 20 changes in place, annualized	\$148,967
Counted for the gate, with S3–3 at its estimate	\$152,076
The gate: three times the fee of \$35,000	\$105,000
Counted, as a multiple of the fee	4.35 times
Measured alone, as a multiple of the fee	4.26 times
Payback on the fee, from the measured savings	2.8 months
Margin over the gate	\$47,076

The second invoice

The Optimization’s second \$17,500 is invoiced only if the annualized savings in place at the end of week six reach three times the fee. They reach 4.35 times. The invoice is issued with this report. If you disagree with any figure here, the query in [The query for one change](#) reproduces it from your own data.

Estimated savings by category, and how much of each was measured



Each bar is the category’s estimate; the dark part is what the Cost and Usage Report measured. The detail is in [Measured against estimated](#).

The method

The terms of the engagement say how the second half of the fee is decided:

The Optimization is invoiced in full only if the annualized savings in place at the end of week six are at least three times the fee. Savings are measured change by change on the client's Cost and Usage Report: the daily cost of each changed resource or commitment in the week before the change against the week after it, annualized. A change the client approved at the week-two readout and later declines counts at the savings the roadmap estimated for it.

Each step below applies that sentence to your bill, the same way to every change.

- 1. The cost.** Each change other than the Savings Plan is measured on **public On-Demand cost** (pricing_public_on_demand_cost on covered usage, the unblended cost elsewhere), from the CUR 2.0 data export in Parquet, queried with Athena; one query per change, the same for every change. C-1 started on June 3, 2026, and it is counted on its own, by its ARN (step 6). Amortized cost would price covered usage at the plan's rate, so a change whose after-window runs past June 3, 2026 (S-1, R-4, R-1, R-2) would show the plan's discount as its own saving, and C-1 would be counted twice. On public On-Demand cost those changes measure as they appear in this report. Seatlark has no private pricing discount, so public On-Demand cost is what uncovered usage bills. The expression is in [The query for one change](#).
- 2. The window.** For each changed resource, take its cost per day in the 7 days before the change and the 7 days after it. The day of the change is excluded from both, because a change lands partway through its day.
- 3. The difference and the annualization.** The daily saving is the daily cost before minus the daily cost after, and the annual saving is that difference times 365.
- 4. What is left out.** These line types are excluded from every change's cost: SavingsPlanRecurringFee, SavingsPlanNegation, RIFee, Fee, Tax, Credit, Refund. Recurring fees, negations, taxes, credits, and refunds are not usage a change moves.
- 5. The lag.** The Cost and Usage Report finalizes a day about 48 hours later. This measurement ran on June 17, 2026, so it uses data through June 15, 2026. A change needs 7 full days after it, so only changes made by June 8, 2026 could be measured. Every counted change was in place by June 3, 2026, with the first on May 27, 2026, which is why week 5 carries no savings: it fell inside the production freeze, from June 4 to June 14.
- 6. The Savings Plan.** A commitment is measured by its ARN: its fee, the usage it covered, and its utilization, not by a resource ID ([The Savings Plan, by its ARN](#)).
- 7. One-time charges** are listed beside the result and never annualized into it.

8. **Nothing is counted that is not in place.** Deferred and declined changes count for nothing. The one exception is the rule the terms state for a change approved at the readout and declined later ([Approved, then declined](#)).

How each counted change was selected on the Cost and Usage Report.

Measured by	Changes
resource ID	17
bucket: storage, requests, and retrieval	2
commitment ARN: fee, covered usage, and utilization	1
Total	20

A worked example

W-1, Unattached gp2 volumes from test cutovers (10 x 500 GB), was made on May 27, 2026. The cost per day of the 10 volumes was \$16.44 in the seven days before (May 20 to May 26) and \$0.00 in the seven days after (May 28 to June 3), because a deleted volume has no line. The daily saving is \$16.44; times 365 is \$6,000. The Analysis estimated \$6,000.

The query for one change

Anyone on your team can run the measurement for a single change in Athena against the CUR 2.0 export. The query below is W-1's, with the real volume IDs and dates; replace the IDs, the dates, and the billing periods for another change. The window is summed and divided by 7, not averaged: a deleted resource has no rows after the change, and a day with no rows is a day with no cost, so an average would return nothing for the changes that save the most.

SQL

```
WITH daily AS (
  SELECT date(line_item_usage_start_date) AS day,
         sum(CASE line_item_line_item_type
                WHEN 'SavingsPlanCoveredUsage' THEN pricing_public_on_demand_cost
                WHEN 'DiscountedUsage' THEN pricing_public_on_demand_cost
                ELSE line_item_unblended_cost
              END) AS cost
  FROM cur2.seatlark_cur
  WHERE line_item_resource_id IN (
    'vol-cutover-test-01',
    'vol-cutover-test-02',
    'vol-cutover-test-03',
    'vol-cutover-test-04',
    'vol-cutover-test-05',
    'vol-cutover-test-06',
```

```

        'vol-cutover-test-07',
        'vol-cutover-test-08',
        'vol-cutover-test-09',
        'vol-cutover-test-10'
    )
    AND line_item_line_item_type NOT IN (
        'SavingsPlanRecurringFee',
        'SavingsPlanNegation',
        'RIFee',
        'Fee',
        'Tax',
        'Credit',
        'Refund'
    )
    AND billing_period IN ('2026-05', '2026-06')
GROUP BY 1
), windows AS (
    SELECT sum(CASE WHEN day BETWEEN date '2026-05-20' AND date '2026-05-26' THEN cost ELSE 0
END) / 7.0 AS daily_before,
        sum(CASE WHEN day BETWEEN date '2026-05-28' AND date '2026-06-03' THEN cost ELSE 0
END) / 7.0 AS daily_after
    FROM daily
)
SELECT daily_before,
        daily_after,
        (daily_before - daily_after) * 365 AS annual_saving
FROM windows;

```

For a change that replaces resources, filter on tags in place of resource IDs; user-defined tag keys appear in the `resource_tags` map with a `user_prefix`, as in `element_at(resource_tags, 'user_Team') = 'marketplace-api'`.

The Savings Plan, by its ARN

C-1 is measured on the plan, not on a resource. The first query reads the usage the plan covered and what that usage would have cost On-Demand, and the second its utilization, from the recurring-fee lines that the cost measurement leaves out. The plan costs \$5.00 an hour, \$120.00 a day. It covered \$6.82 an hour of On-Demand usage, \$163.70 a day, so it saved \$43.70 a day, \$15,950 a year, at 100% utilization. It was bought on June 3, 2026, and its windows are May 27 to June 2 and June 4 to June 10.


SQL

```
SELECT date(line_item_usage_start_date) AS day,  
       sum(CASE WHEN line_item_line_item_type = 'SavingsPlanCoveredUsage' THEN  
pricing_public_on_demand_cost ELSE 0 END) AS covered_on_demand,  
       sum(CASE WHEN line_item_line_item_type = 'SavingsPlanCoveredUsage' THEN  
savings_plan_savings_plan_effective_cost ELSE 0 END) AS covered_effective,  
       sum(CASE WHEN line_item_line_item_type = 'SavingsPlanRecurringFee' THEN  
savings_plan_used_commitment ELSE 0 END)  
       / nullif(sum(CASE WHEN line_item_line_item_type = 'SavingsPlanRecurringFee' THEN  
savings_plan_total_commitment_to_date ELSE 0 END), 0) AS utilization  
FROM cur2.seatlark_cur  
WHERE savings_plan_savings_plan_a_r_n = '<savings-plan-arn>'  
      AND billing_period IN ('2026-05', '2026-06')  
GROUP BY 1  
ORDER BY 1;
```

Two limits are worth stating. The measurement is of cost, not usage, so a workload that grew or shrank between the two weeks moves a single change's figure for reasons that are not the change; across 20 changes those effects average out, and for any one change they do not, which is why the next sections show the estimate beside the measurement. And the report is provisional until AWS finalizes the month's invoice: we re-run the measurement at the first monthly review and say so if it moves.

The daily cost

The check on the per-change method is the whole bill. A whole-bill comparison cannot equal the sum of the changes, because the bill also moves with growth, so the check has to account for it. Business Support+ is billed once a month and has no daily line in the Cost and Usage Report, so it is in neither average below and stays out of the check. The average daily cost in the first seven days of the engagement (May 11 to May 17) was \$1,583. In the last seven days of data (June 9 to June 15) it was \$1,245. That is \$338 a day lower, or \$123,531 a year, even though wave 6 added \$2,600 a month from May 30, 2026.



Cost per day before Business Support+, which bills monthly, May 11 to June 15, against the same line without the engagement's changes (the May run rate plus wave 6). Each change's measured saving starts the day after it; one-time charges show on the days they billed.

The fall in the bill should equal the measured savings of the changes in effect for the whole last week, minus the growth from wave 6, minus any one-time charge billed in that week.

The whole-bill check, on usage before Business Support+. The difference is the part of the fall that the changes and wave 6 do not explain. Business Support+, billed monthly, falls by about 7% of each dollar saved, \$10,428 a year, on the monthly invoice; it is not in this check and not counted.

Part of the change in the bill, annualized	Amount
Measured savings of the changes in effect for all of the last week	\$148,967
Wave 6, which raises the bill	-\$31,200
One-time charges billed in the last week	\$0
Expected fall in the bill	\$117,767
Fall in the bill, from the two averages	\$123,531
Difference: the fall from the averages minus the expected fall	\$5,764

The difference is \$5,764 a year, 4.9% of the expected fall. Two seven-day averages of a bill that swings day to day carry noise of their own: from the spread of the daily figures in these two weeks, one standard error on the fall is about \$4,969 a year, so the difference is inside two standard

errors. The check does not prove any single change. It shows that no large saving is missing from the bill and that none is counted twice, and that is all we claim for it. We do not count the whole-bill fall as savings; the register counts only the per-change measurements.

Measured against estimated

The measurement and the Analysis's estimates agree to within 3.9% overall. We expected the measurements to fall a little under the estimates: an estimate is made from the 30 days before the Analysis, and a measurement is made from the week after a change, when some usage has moved. Of the 20 changes, 8 measured at their estimates, 9 measured between 94% and 99% of them, and 3 measured below 92%, which the next section explains.

8 changes measured at their estimate to the dollar (D-1, E-1, S3-2, W-1, W-3, E-2a, N-2, C-1), and that is what they should do: each one deletes a whole resource, changes the price of a quantity that did not move inside its windows, or, for C-1, is the arithmetic of the plan's fee against the usage it covered at full utilization. Every change whose quantity moves with traffic measured under its estimate.

The two rightsized production tiers, tested again on June 10, 2026. Report 5 said the API tier's June 5 peak on m5.xlarge reached 61%, over the Analysis's 60% ceiling, and promised the next on-sale's result. On June 10 the API tier peaked at 57%, under the ceiling, and the 75% alarm added after June 5 never fired. The partner APIs ran their busiest feed on m5.xlarge at 41%, against the 42% the Analysis projected. Neither result changes a measurement; both are why R-1 and R-2 stay in place.

Every counted change, by category, estimated and measured, annualized. Graviton has no row because both of its changes are deferred.

Category	Changes	Estimated	Measured	Measured vs estimated
Rightsizing (R-4, R-1, R-2)	3	\$29,644	\$28,445	96.0%
Waste (W-1, W-2, W-3, W-4)	4	\$25,733	\$24,433	95.0%
Schedules (S-1)	1	\$24,183	\$22,007	91.0%
Logging (L-1, L-3, L-2, L-4)	4	\$18,112	\$17,233	95.2%
Commitments (C-1)	1	\$15,950	\$15,950	100.0%
S3 lifecycle (S3-2, S3-1)	2	\$13,136	\$12,796	97.4%
Network (N-1, N-2)	2	\$10,342	\$10,126	97.9%
Databases (D-1)	1	\$9,552	\$9,552	100.0%
Storage (E-1, E-2a)	2	\$8,424	\$8,424	100.0%
Total	20	\$155,076	\$148,967	96.1%

Every change

The 20 changes in place, in the order they were made. Every window is the seven full days of finalized data after the change and the seven before it, the change day excluded. Basis: resource ID, except the Savings Plan (its ARN), and the two S3 lifecycle changes (the bucket).

ID	Change	Owner	Applied	Estimated	Measured	vs est.	Daily saving
D-1	Multi-AZ off on non-production databases	Beth Morrow	May 27	\$9,552	\$9,552	100%	\$26.17
E-1	gp2 to gp3 on attached volumes	Beth Morrow	May 27	\$4,224	\$4,224	100%	\$11.57
L-1	Log-group retention	Beth Morrow	May 27	\$7,004	\$6,934	99%	\$19.00
L-3	Flow logs to S3	Beth Morrow	May 27	\$2,460	\$2,386	97%	\$6.54
N-1	S3 gateway endpoints	Beth Morrow	May 27	\$7,188	\$6,973	97%	\$19.10
S3-2	Migration backups to Deep Archive	Owen Pradhan	May 27	\$7,462	\$7,462	100%	\$20.44
W-1	Unattached gp2 volumes deleted	Beth Morrow	May 27	\$6,000	\$6,000	100%	\$16.44
W-2	MGN replication servers retired	Beth Morrow	May 27	\$7,927	\$7,768	98%	\$21.28
W-3	Idle load balancers deleted	Aaron Feld	May 27	\$394	\$394	100%	\$1.08

ID	Change	Owner	Applied	Estimated	Measured	vs est.	Daily saving
E-2a	Staging search io1 to gp3	Imani Brooks	May 28	\$4,200	\$4,200	100%	\$11.51
L-2	Debug logging off	Aaron Feld	May 28	\$3,600	\$3,168	88%	\$8.68
N-2	Non-production NAT gateways consolidated	Beth Morrow	May 28	\$3,154	\$3,154	100%	\$8.64
S-1	Non-production schedules	Beth Morrow	May 28	\$24,183	\$22,007	91%	\$60.29
W-4	Snapshot and AMI sprawl trimmed	Beth Morrow	May 28	\$11,412	\$10,271	90%	\$28.14
L-4	Agent metrics trimmed	Beth Morrow	Jun 1	\$5,047	\$4,745	94%	\$13.00
R-4	Non-production rightsizing	Beth Morrow	Jun 1	\$9,461	\$9,272	98%	\$25.40
S3-1	Ticket scans to Glacier Instant Retrieval	Kara Lindqvist	Jun 1	\$5,674	\$5,334	94%	\$14.61
R-1	Marketplace API tier to m5.xlarge	Aaron Feld	Jun 2	\$10,092	\$9,587	95%	\$26.27
R-2	Partner APIs to m5.xlarge	Diego Saldana	Jun 2	\$10,092	\$9,587	95%	\$26.27
C-1	Compute Savings Plan	Lydia Brandt	Jun 3	\$15,950	\$15,950	100%	\$43.70
Total				\$155,076	\$148,967	96.1%	

Where the measurement fell short

3 changes measured below 92% of their estimates. Each has a reason, and we list every one, because a measurement that reports only its wins is not a measurement. These rows are in the totals at their measured values, not their estimates.

The three changes that measured under 92% of their estimates.

ID	Change	Estimated	Measured	Short by	vs estimate
L-2	Debug logging off	\$3,600	\$3,168	\$432	88%
W-4	Snapshot and AMI sprawl trimmed	\$11,412	\$10,271	\$1,141	90%
S-1	Non-production schedules	\$24,183	\$22,007	\$2,176	91%

L-2, Debug logging off (88%). Two services read their log level only at deploy, and their deploys were queued behind the release train, so they logged at debug through the seven-day after-window. The release train held those deploys, and the production freeze from June 4 to June 14 began only as the window closed, so it is not the cause. The two services pick the change up with their next release, and the part that has not happened is not counted.

W-4, Snapshot and AMI sprawl trimmed (90%). EBS snapshots are incremental, so deleting one frees only the blocks no other snapshot references (the estimate used billed cost per snapshot). The saving is therefore smaller than the sum of the deleted snapshots' billed cost, and this is the reason the finding was rated medium confidence.

S-1, Non-production schedules (91%). The schedule works; what the plan missed is the perf-test fleet. Three on-sale rehearsals ran the perf-test fleet about 32 hours in the seven-day after-window, against the 22 a week in the plan (a pace of about 140 hours a month, not 96). Rehearsals are the purpose of that account, so we did not tighten the schedule: the monthly review plans the fleet at 140 hours.

The other 9 changes below their estimates. L-1, L-3, N-1, W-2, L-4, R-4, S3-1, R-1, R-2 measured between 94% and 99% of their estimates, \$2,360 a year in all. We cannot name a specific cause for any of them, and we will not invent one: this is the ordinary spread between an estimate from the 30 days before the Analysis and a measurement from the week after the change. The monthly review re-measures them with a second week of data.

The gate

The gate computation. The second invoice is issued because the counted figure is above the gate, with or without the one-time charges.

The guarantee's computation	Amount
Measured savings, 20 changes in place, annualized	\$148,967
Approved at the readout, then declined, counted at the estimate (S3-3)	\$3,109
Counted for the gate	\$152,076
The gate: three times the fee of \$35,000	\$105,000
Multiple of the fee	4.35 times
Measured savings alone, as a multiple of the fee	4.26 times
One-time charges, listed in One-time charges	\$3,117.00
Counted, net of the one-time charges	\$148,959 (4.26 times)

The guarantee's computation	Amount
Not counted: approved, deferred (D-3, E-2b, D-2, N-3, G-1, G-2)	\$52,808
Not counted: declined at the readout (N-4, R-3, S-2)	\$35,846
Identified by the Analysis	\$246,839

The identified savings reconcile to the dollar. Every dollar the Analysis identified is in exactly one of four places:

In place plus deferred plus declined plus approved-then-declined equals identified; the difference is \$0.

Where the identified savings are	Changes	Annualized
In place, at the Analysis's estimates	20: D-1, E-1, L-1, L-3, N-1, S3-2, W-1, W-2, W-3, E-2a, L-2, N-2, S-1, W-4, L-4, R-4, S3-1, R-1, R-2, C-1	\$155,076
Approved, deferred past the engagement	6: E-2b, D-2, D-3, N-3, G-1, G-2	\$52,808
Declined at the readout	3: R-3, S-2, N-4	\$35,846
Approved at the readout, then declined	1: S3-3	\$3,109
Identified by the Analysis		\$246,839

What is not counted, and when it comes

Deferred and declined changes count for nothing toward the gate. They are the handover report's roadmap, each with an owner and a date.

Approved at the readout and deferred. Nothing here is counted until it is in place and measured.

ID	Change	Owner	AnnualizedTarget	Waits on
D-3	Catalog database to db.r6g.2xlarge	Imani Brooks	\$19,290Jul 9	a load test that replays an on-sale, then the July maintenance window
E-2b	Production search io1 to gp3	Imani Brooks	\$10,800Jun 22	the production freeze (June 4 to 14) and the biggest on-sale of the summer: Imani Brooks will not touch the search cluster in the week before it
D-2	Partner database to db.r6g.xlarge	Diego Saldana	\$9,294Jun 27	a load test against a restored copy at the new class, then ten business days' notice of a maintenance window, which Seatlark Connect's partner contracts require
N-3	Cache the colo order-status calls	Aaron Feld	\$7,369Aug 14	an application change (the order-status cache), approved for the Q3 sprint
G-1	Marketplace to Graviton	Aaron Feld	\$4,541Jul 23	the July container replatform, which moves the marketplace to arm64 images
G-2	Partner APIs to Graviton	Diego Saldana	\$1,514Jul 10	an arm64 build of the partner APIs, validated in staging
Total	Deferred		\$52,808	

Declined at the readout. These count for nothing.

ID	Change	Owner	AnnualizedReason, as the owner gave it
N-4	Second Direct Connect port to 1 Gbps	Owen Pradhan	\$17,082The orders database cutover in July needs failover at full rate. The ports are revisited at colo exit.
R-3	Ingestion workers to m5.2xlarge	Imani Brooks	\$13,455Compute Optimizer never saw an on-sale. Show me a load test first. Rescheduled as a perf-test exercise after the freeze, on the remaining roadmap.
S-2	Staging schedule	Aaron Feld	\$5,308Staging runs the overnight on-sale rehearsals and the partner acceptance tests in other time zones; the schedule would interrupt both.
Total	Declined		\$35,846

Robustness

The gate does not rest on the soft numbers. The 17 high-confidence changes in place total \$134,390 a year at their estimates, and they measured at \$130,194. Measured at only 90% of their estimates they would be \$120,951, which clears the gate by \$15,951 on those changes alone, without the medium-confidence changes (L-2, W-4, S3-1) and without S3-3.

The cushion is wide. The counted figure is \$47,076 above the gate, so the measured savings would have to be \$101,891 instead of \$148,967, a miss of 31.6%, before the gate failed. The week-two stop rule exists for the case where the roadmap is too small; here it was 2.4 times the gate.

One-time charges

Some changes cost money once. The charges below are not netted into the savings above, and they are never annualized; the gate was checked both ways.

One-time charges of the changes in place. The payback is the charge against that change’s measured saving.

ID	Change	One-time charge	What it is	Repaid in
S3-1	Ticket scans to Glacier Instant Retrieval	\$3,102.00	Lifecycle transition requests into S3 Glacier Instant Retrieval, about 155 million objects at \$0.02 per 1,000	7.0 months
S3-2	Migration backups to Deep Archive	\$15.00	One-time charge of the change	Within a few days
Total		\$3,117.00		

The total is 2.1% of one year’s measured savings, repaid in about 8 days of savings at the measured rate. S3-1 carries almost all of it, and its lifecycle rule was priced with the charge in view: Intelligent-Tiering was considered and rejected because its monitoring fee would have eaten most of the saving.

Approved, then declined

The rule

A change you approved at the week-two readout and later decline counts at the savings the roadmap estimated for it. The rule exists so that the guarantee cannot be defeated by a change being approved and then not made, and so that Seatlark is not penalized for an owner's reasonable second thought.

It applies to one change this engagement.

The change approved at the readout and declined afterward, counted at its estimate. It is the only reason the counted figure exceeds the measured one.

ID	Change	Owner	EstimatedReason annual savings
S3-3	Event renditions deleted	Diego Saldana	\$3,109White-label partners hotlink rendition URLs on archived event pages, and deleting them would break partner pages. Reversed in week five.

We did not push. It is reported separately from the measured figure, at the estimate, because there is no measurement of a change that was never made. If Diego Saldana's team later goes ahead, it is measured like every other change and the measurement replaces the estimate.

What is deliberately not counted

- **The fall in Business Support+.** Support is a percentage of usage, so any saving lowers it: at 7% on the next dollar, about \$869 a month, \$10,428 a year, on these savings. It is real and it is not counted.
- **Wave 6's growth.** Wave 6 added \$2,600 a month (\$31,200 a year) from May 30, 2026. The measurement is per change, so growth is neither a saving nor a loss against any of them.
- **Reductions the migration would have made anyway.** Work in your own migration plan that would have lowered the bill without us is not in these figures; the changes are the ones in the register.

- **Cost avoidance on the remaining waves.** The guardrails and the wave runbook (the week-5 report) are worth \$100,000 to \$170,000 a year at the end-state run rate on a model. That figure is not measured, not counted toward the gate, and not a saving on the bill. It is in no total above.

Re-running this measurement

Everything here can be re-run by Seatlark's own team. The inputs are the CUR 2.0 export, the register of changes (the resource, tag, or ARN, and the date), and the cost expression, the excluded line types, and the queries in [The method](#). To measure a change made after this report, wait for seven full days of finalized data, which with the 48-hour lag is nine days, run the query, and add the row. The monthly cost review in the handover report keeps the register and re-runs the whole set each quarter, so the claim that the bill fell by this much can always be checked by someone who was not in the room.

© 2026 Vantalect, LLC. Sample deliverable; Seatlark, Inc. is fictional.

AWS, Amazon Web Services, AWS Graviton, and the AWS service names are trademarks of Amazon.com, Inc. or its affiliates. Ansible is a trademark of Red Hat, Inc. Elasticsearch is a trademark of Elasticsearch B.V. PostgreSQL is a registered trademark of the PostgreSQL Community Association of Canada. MySQL is a trademark of Oracle Corporation and its affiliates. Redis is a trademark of Redis Ltd. Microsoft and SQL Server are trademarks of the Microsoft group of companies. VMware and vCenter are trademarks of Broadcom Inc. NetApp is a trademark of NetApp, Inc. Java is a trademark of Oracle Corporation and its affiliates. Vantalect is not affiliated with or endorsed by any of them. Seatlark, Inc. is fictional.