

AWS COST · SAMPLE DELIVERABLE

Week-Six Savings Measurement

Week 6 of 6: Measure · prepared for StrataForge Robotics, Inc.

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

October 2, 2026 · AWS Cost Analysis & Optimization: The Optimization (six weeks)

SAMPLE DELIVERABLE

This is a sample deliverable. StrataForge Robotics, 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 on October 13, 2026.

Contents

Result	2
The method	4
The query for one change	5
The daily cost	7
Measured against estimated, by kind	9
Where the measurement fell short	12
Changes that measured at or below zero	13
The 40 largest changes	15
Intelligent-Tiering: in place, not yet visible	18
One-time charges	18
Approved, then declined	19
What remains	20
Re-running this measurement	22

Result

StrataForge’s AWS bill fell by \$9.4 million a year, measured on your own Cost and Usage Report. That is 95.8% of the \$9.8 million the Analysis estimated for the same 10,504 changes, and 16.8% of the \$55.7 million annualized bill we started from. Counting the items approved at the readout and declined afterward, as the engagement defines it, the figure for the guarantee is **\$9,652,024: 275.8 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.

The gate also clears with the one-time charges taken out. The changes cost \$97,097 once; deducted from the counted figure, the result is \$9,554,927, or 273.0 times the fee. The gate is met on that basis too.


Measured savings are not the whole of what is in place. 161 S3 Intelligent-Tiering changes, estimated at \$1.3 million a year, are in place but cannot show on the bill yet, so they are not counted ([Intelligent-Tiering: in place, not yet visible](#)). With them, and every change at the Analysis’s estimate, \$11.1 million a year is in place, which is 20% of the annualized bill and 47% of the \$23.5 million the Analysis identified. The rest waits on production windows, owner reviews, and the queued Savings Plans purchase ([What remains](#)).

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, 10,504 changes, annualized	\$9,377,111
Approved at the readout, then declined by the owner, counted at the estimate	\$274,912
Counted for the gate	\$9,652,024
The gate: three times the fee of \$35,000	\$105,000
Multiple of the fee	275.8 times
Measured savings alone, as a multiple of the fee	267.9 times
One-time charges, listed in One-time charges	\$97,097
Counted for the gate, net of the one-time charges	\$9,554,927
Not counted: Intelligent-Tiering, in place but not yet visible	\$1,342,048
Not counted: deferred, still to be put in place	\$11,296,614
Not counted: declined by the owner	\$355,158

The second invoice

The Optimization's second \$17,500 is invoiced only if the annualized savings in place at the end of week 6 reach three times the fee. They reach 275.8 times. The approved-then-declined rule adds \$274,912 to the count, which is 2.8% of it. Without the rule the multiple would be 267.9 times, so the rule does not decide the outcome.



Measured annualized savings by kind of change, the ten largest.

The method

The method is the one the engagement defined in week one, applied the same way to every change. It is a measurement on your bill, not a model.

- 1. The cost.** Each change is measured on CUR 2.0 **net amortized cost**, from the CUR 2.0 data export in Parquet, queried with Athena; one query per change, the same for every change. Net amortized cost is the cost after the private pricing discount and with Savings Plans and reservations spread over the usage they cover, so a change that moves usage between covered and uncovered hours is measured at its real cost. The expression is below, and it is the same for every change.
- 2. The window.** For each changed resource, take its cost per day in the seven days before the change and the seven days after it. The day of the change is excluded from both. A change counts only when seven full days of data follow it.
- 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. The discount is already inside net cost, so nothing is multiplied by it afterward.
- 4. What is left out.** These line types are excluded from the cost of every change:
SavingsPlanRecurringFee, SavingsPlanNegation, RIFee, Fee, Tax, Credit, Refund. Recurring fees, negotiations, taxes, credits, and refunds are not usage a change moves. The commitment fees are read separately, in step 6.

5. **The lag.** The CUR finalizes a day about 48 hours later, so the October 1 measurement used days through September 29. A change counts only when seven full days follow it, so with data through September 29, 2026 the last change that could be measured was made seven days before that. Week 5's changes ran only on September 21 and 22 for that reason.
6. **Commitments.** A commitment is measured by its fee, the usage it covers, and its utilization, by commitment ARN. 5 existing commitments were checked on the same days, and every one stayed fully used after the changes: ri-hana at 100%, ri-r5-8xl at 100%, sp-3yr-2024 at 100%, sp-1yr-2025 at 100%, ri-aurora at 100%. A rightsizing or Graviton move that left a reservation or a Savings Plan underused would have shown up there as a cost, not a saving.
7. **One-time charges** (lifecycle transition requests, final snapshots, and the like) are listed in [One-time charges](#), beside the run-rate, and are never annualized into the savings.
8. **Nothing is counted that is not in place.** Deferred and declined items count for nothing. The one exception is the approved-then-declined rule ([Approved, then declined](#)).

The resource a change is measured on depends on what the change did to it:

How each counted change was measured. Replaced fleets are measured by tag because their node and instance IDs change when a node group is replaced.

Measured by	Changes
resource ID	9,839
tags: node group or NodePool, and Team (node IDs change)	581
bucket: storage, requests, and retrieval	82
commitment ARN: fee, covered usage, and utilization	2
Total	10,504

A worked example

Change A-0161, CloudTrail data events scoped, applied September 22, 2026. The net daily cost of the resource was \$1,868.12 before and \$174.59 after. The daily saving is \$1,693.53. Times 365 is \$618,138, and the table of the largest changes shows \$618,138. The Analysis had estimated \$599,428.

The query for one change

A reader can run the measurement for any single change in Athena against the CUR 2.0 export that the week-5 showback uses. Replace the placeholders with the resource ID and the change date. The CASE expression is the cost expression from step 1, and the NOT IN list is the excluded line

types. Each window is summed and divided by seven, not averaged: a resource that was deleted 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 very 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 savings_plan_net_savings_plan_effective_cost
              WHEN 'DiscountedUsage'
                THEN reservation_net_effective_cost
              ELSE line_item_net_unblended_cost
            END) AS cost
  FROM cur2.strataforge_cur
  WHERE line_item_resource_id = '<resource-id>'
     AND line_item_line_item_type NOT IN (
        'SavingsPlanRecurringFee',
        'SavingsPlanNegation',
        'RIFee',
        'Fee',
        'Tax',
        'Credit',
        'Refund'
      )
     AND billing_period IN ('<yyyy-mm-before>', '<yyyy-mm-after>')
  GROUP BY 1
), windows AS (
  SELECT sum(CASE WHEN day BETWEEN date_add('day', -7, date '<change-date>')
                  AND date_add('day', -1, date '<change-date>') THEN cost ELSE 0
            END) / 7.0 AS daily_before,
         sum(CASE WHEN day BETWEEN date_add('day', 1, date '<change-date>')
                  AND date_add('day', 7, date '<change-date>') 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 replaced fleet, filter on the tags in place of the resource ID. User-defined tag keys appear in the resource_tags map with a user_ prefix.

SQL

```
WHERE element_at(resource_tags, 'user_<node-group-tag>') = '<node-group>'
      AND element_at(resource_tags, 'user_<team-tag>') = '<team>'
```

For a Savings Plan, read its utilization by ARN from the recurring-fee lines, which the cost measurement leaves out. The ratio is used commitment over total commitment to date, and it should stay at 100% when a change touches the usage it covers.

SQL

```
SELECT date(line_item_usage_start_date) AS day,  
       sum(savings_plan_used_commitment) / sum(savings_plan_total_commitment_to_date) AS  
utilization  
FROM cur2.strataforge_cur  
WHERE savings_plan_savings_plan_a_r_n = '<savings-plan-arn>'  
      AND line_item_line_item_type = 'SavingsPlanRecurringFee'  
      AND billing_period IN ('<yyyy-mm-before>', '<yyyy-mm-after>')  
GROUP BY 1  
ORDER BY 1;
```

For a reservation, read the same days from the reservation utilization report in Cost Explorer.

Two limits are worth stating. Cost, not usage, is what is measured, so a workload that grew or shrank between the two weeks moves the figure for reasons that are not the change; with 10,504 changes, such effects average out and for any single change they do not, which is why the table below shows the estimate beside the measurement. And the report is provisional until AWS finalizes the month's invoice: the data runs through September 29, 2026, and we will re-run the measurement against the finalized report at the first monthly review and tell you if it moves.

The daily cost

The check on the per-change method is the whole bill. A whole-bill comparison does not equal the sum of the measured changes, because the bill also moves with growth in usage and with Support, so the check has to account for both. The average daily net cost in the first seven days of the engagement (August 24, 2026 to August 30, 2026) was \$152,548. In the last seven days of data (September 23, 2026 to September 29, 2026) it was \$131,869. That is \$20,680 a day lower, or \$7.55 million a year.

Daily net cost, August 24 to September 29



Net cost per day across the engagement, through the last day the Cost and Usage Report had finalized. The step down follows the changes of weeks 3 to 5.

The fall in the bill should equal the measured savings of the changes in effect for the whole last week, plus the fall in Support, minus the growth in usage, minus any one-time charge billed in that week. Growth runs at the rate the baseline forecast, 3.6% a month. The parts, annualized:

The whole-bill check. The difference is the part of the fall that the changes, Support, and growth do not explain.

Part of the change in the bill, annualized	Amount
Measured savings of the changes in effect for all of the last week	\$9,377,112
Support, billed as a percentage of charges, falls with them: 3.0 cents for every dollar saved	\$281,313
Growth in usage over the same weeks, which raises the bill	-\$1,984,951
One-time charges billed in the last week	\$0
Expected fall in the bill	\$7,673,474
Fall in the bill, from the two averages	\$7,548,025
Difference: the fall from the averages minus the expected fall	-\$125,449

The difference is \$125,449 a year, 1.6% 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 \$424,096 a year, so the difference is well inside it. 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, by kind

The estimate and the measurement agree to within 4.2% overall. 10 of the 40 kinds of change measured more than 5% below their estimate, and 6 measured at or above it. 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.

Every kind of change that is counted, estimated and measured, annualized. Figures are net of the private pricing discount. S3 Intelligent-Tiering is not in the table because it is not yet visible.

Kind of change	Changes	Estimated	Measured	Measured vs estimated
S3 lifecycle and storage classes	82	\$1,227,646	\$1,169,178	95.2%
Rightsize and move to Graviton	318	\$1,178,678	\$1,110,848	94.2%
Instances off nights and weekends	307	\$1,013,492	\$915,170	90.3%
Idle instances	135	\$656,256	\$651,725	99.3%
CloudTrail data events scoped	1	\$599,428	\$618,138	103.1%
Unattached volumes	1,325	\$508,516	\$506,260	99.6%
Rightsize (same family)	305	\$557,628	\$486,476	87.2%
Orphaned snapshot chains	1,333	\$469,530	\$458,968	97.8%
Unused p5 capacity reservation	1	\$448,400	\$448,400	100.0%
Idle SageMaker notebooks and Studio apps	46	\$471,651	\$437,525	92.8%
ECR lifecycle policies	2	\$408,108	\$395,137	96.8%
Databases off nights and weekends	54	\$264,400	\$250,530	94.8%
Unused AMIs and their snapshots	2,502	\$238,410	\$233,339	97.9%

Kind of change	Changes	Estimated	Measured	Measured vs estimated
Unused Redshift cluster	1	\$156,419	\$154,930	99.0%
CloudWatch log retention	692	\$148,740	\$143,789	96.7%
EBS gp2 and io1 to gp3	2,362	\$138,492	\$138,505	100.0%
SageMaker Savings Plan	1	\$133,969	\$133,389	99.6%
DynamoDB to on-demand	38	\$143,768	\$124,168	86.4%
Move to Graviton, same size	274	\$122,177	\$118,559	97.0%
Kinesis shards sized to peak	32	\$105,858	\$105,910	100.0%
AWS Config recording scoped	2	\$105,462	\$105,030	99.6%
Aurora I/O-Optimized	8	\$93,397	\$88,268	94.5%
AWS Backup cold storage	2	\$84,988	\$83,888	98.7%
Lambda on Arm	192	\$83,997	\$80,543	95.9%
ElastiCache: Valkey on Graviton	6	\$76,388	\$76,672	100.4%
Idle non-production databases	13	\$66,569	\$66,252	99.5%
Debug logs to Infrequent Access	45	\$57,474	\$56,068	97.6%
Database Savings Plan	1	\$50,918	\$50,917	100.0%
Database engines out of Extended Support	38	\$48,167	\$48,009	99.7%
WorkSpaces billed by the hour	1	\$34,247	\$29,409	85.9%
Databases: rightsize and Graviton	46	\$22,673	\$22,454	99.0%
Kinesis streams with no traffic	2	\$14,078	\$14,032	99.7%
VPC endpoints instead of NAT	6	\$15,159	\$13,379	88.3%
Long-stopped instances	41	\$9,632	\$9,577	99.4%
Idle load balancers	38	\$7,136	\$7,100	99.5%
Unassociated Elastic IPs	157	\$6,395	\$6,371	99.6%
Kafka consumers read in their own AZ	2	\$6,964	\$5,993	86.1%

Kind of change	Changes	Estimated	Measured	Measured vs estimated
MSK brokers on Graviton	2	\$5,866	\$5,845	99.6%
Public IPv4 addresses removed	82	\$3,340	\$3,323	99.5%
Move to the current generation	9	\$3,027	\$3,038	100.4%
Total	10,504	\$9,787,440	\$9,377,111	95.8%

Where the measurement fell short

The largest differences in dollars, with our read of each. These are readings, not findings: we will check each against a second week of data at the first monthly review.

The eight kinds of change with the largest shortfall against their estimate.

Kind of change	Estimated	Measured	Shortfall	Measured vsOur read estimated
Instances off nights and weekends	\$1,013,492	\$915,170	\$98,322	90.3%4 changes measured at or below zero, and 2 more measured far under the estimate (Changes that measured at or below zero). Incident 1: the owning team suspended the schedule for a release test and restored it on October 2.
Rightsize (same family)	\$557,628	\$486,476	\$71,152	87.2%9 changes measured at or below zero, and 2 more measured far under the estimate (Changes that measured at or below zero). Incident 2: the service's next deploy rolled its launch template back to the old size; fixed in the template on October 1.
Rightsize and move to Graviton	\$1,178,678	\$1,110,848	\$67,830	94.2%Within the normal spread between an estimate and a measurement.
S3 lifecycle and storage classes	\$1,227,646	\$1,169,178	\$58,468	95.2%Savings build as objects age into the cheaper class, and objects still read after they move are billed retrieval, which the access logs priced in the estimate; the seven-day window catches only the objects that had moved by then.
Idle SageMaker notebooks and Studio apps	\$471,651	\$437,525	\$34,126	92.8%1 change measured at or below zero (Changes that measured at or below zero). Incident 4: its owner restarted the notebook and left it running; idle shutdown was turned back on October 2.
DynamoDB to on-demand	\$143,768	\$124,168	\$19,600	86.4%2 changes measured at or below zero (Changes that measured at or below zero). Incident 3: a backfill ran against the table that week, so on-demand cost more than the capacity it replaced.
Databases off nights and weekends	\$264,400	\$250,530	\$13,870	94.8%Within the normal spread between an estimate and a measurement.

Kind of change	Estimated	Measured	Shortfall	Measured vsOur read estimated
ECR lifecycle policies	\$408,108	\$395,137	\$12,971	96.8%Within the normal spread between an estimate and a measurement.

Changes that measured at or below zero

20 of the 10,504 counted changes measured at or below zero, and 6 more measured positive but well under their estimates. We list every one with its reason, because a measurement that only reports its wins is not a measurement. The 26 changes trace to 5 incidents, not 26 separate failures; each incident is described once, and each change in the table below points to its incident. Where a fix was made, the incident gives its date. These rows are in the totals above at their measured values, not their estimates.

The incidents behind the changes that missed their estimates. Each one is a single event.

Incident	Changes	Of them, at or below zero	Estimated	Measured	What happened
1	6	4	\$26,076	-\$1,947	The owning team suspended the schedule for a release test and restored it on October 2.
2	11	9	\$21,420	-\$1,618	The service's next deploy rolled its launch template back to the old size; fixed in the template on October 1.
3	2	2	\$9,445	-\$1,719	A backfill ran against the table that week, so on-demand cost more than the capacity it replaced.
4	1	1	\$8,263	-\$658	Its owner restarted the notebook and left it running; idle shutdown was turned back on October 2.
5	6	4	\$1,253	-\$93A	A new caller doubled the function's invocations that week, which hid the per-invocation saving.

Every change that measured at or below zero, or far under its estimate, with the incident that explains it. The measured figure is what counts; the totals above include these rows as measured.

ID	Kind of change	Resource	Estimated	Measured	Incident
W-6040	Idle SageMaker notebooks and Studio apps	studio-studio-app-59a2c	\$8,263	-\$658	4
R-0019	Rightsize (same family)	i-0e9b3d9093da769a3	\$6,925	-\$253	2
R-0022	Rightsize (same family)	i-028fccce91de9e0123	\$2,770	-\$280	2
R-0036	Rightsize (same family)	i-015d26311dad912f6	\$2,770	\$21	2
R-1219	Rightsize (same family)	i-0ab1b5ec6e2422e81	\$782	\$19	2
R-1331	Rightsize (same family)	i-0c16b187a3567e38b	\$1,313	-\$248	2
R-1341	Rightsize (same family)	i-09273d4e58aa22acc	\$1,356	-\$109	2
R-1354	Rightsize (same family)	i-09ec88ec95643c95a	\$1,313	-\$155	2
R-1363	Rightsize (same family)	i-0e587ff57fe681991	\$782	-\$121	2
R-1377	Rightsize (same family)	i-04cc0d0d5689be310	\$1,313	-\$245	2
R-1392	Rightsize (same family)	i-062d2769388b457d9	\$1,313	-\$199	2
R-1615	Rightsize (same family)	i-0ff1551f1528f4d2b	\$782	-\$48	2
R-1994	Lambda on Arm	fleet-api-f65faa	\$106	-\$5	5
R-2132	Lambda on Arm	fleet-api-92eb42	\$69	-\$16	5
R-2145	Lambda on Arm	fleet-api-259c71	\$91	\$1	5
R-2184	Lambda on Arm	perception-api-eb85f2	\$95	\$1	5
R-2193	Lambda on Arm	perception-api-e2050f	\$790	-\$65	5

ID	Kind of change	Resource	Estimated	Measured	Incident
R-2194	Lambda on Arm	perception-api-9a95d1	\$103	-\$9	5
S-0115	Instances off nights and weekends	i-039479eebfd4a1759	\$27	-\$2	1
S-0197	Instances off nights and weekends	i-0a0169a3a44f924a6	\$27	-\$3	1
S-0211	Instances off nights and weekends	i-0ac6e9abfcfad9c93	\$2,051	\$65	1
S-0270	Instances off nights and weekends	i-0faf4adc97af919c0	\$8,518	\$7	1
S-0278	Instances off nights and weekends	i-0b22026df20a6b1c8	\$5,277	-\$201	1
S-0340	Instances off nights and weekends	i-03f0f6079a2a3a0eb	\$10,176	-\$1,812	1
T-3430	DynamoDB to on-demand	fleet-api-9b178	\$6,557	-\$1,615	3
T-3440	DynamoDB to on-demand	fleet-api-1278a	\$2,888	-\$104	3

The 40 largest changes

The 40 largest changes by measured savings account for \$3.2 million a year, 34% of the total. The tail is long: 10,504 changes in all, most of them small, which is why the table is a sample of the evidence and the register is the whole of it.

The 40 largest changes, with the net daily cost before and after and the annualized savings.

#ID	Change	Account	Applied	Net daily before	Net daily after	Estimated	Measured	vs estimate
1A-0161	CloudTrail data events scoped	sf-management	Sep 22	\$1,868.12	\$174.59	\$599,428	\$618,138	103%
2W-6059	Unused p5 capacity reservation	crucible-training-prod	Sep 8	\$2,528.14	\$1,299.65	\$448,400	\$448,400	100%
3T-3444	ECR lifecycle policies	sf-cicd	Sep 10	\$1,291.66	\$393.18	\$339,005	\$327,944	97%
4W-6060	Unused Redshift cluster	corp-data-warehouse	Sep 11	\$424.47	\$0.00	\$156,419	\$154,930	99%

#ID	Change	Account	Applied	Net daily before	Net daily after	Estimated	Measured	vs estimate
5C-0003	SageMaker Savings Plan	sf-management	Sep 18	\$365.45	\$0.00	\$133,969	\$133,389	100%
6T-3447	ECR lifecycle policies	crucible-models	Sep 10	\$262.73	\$78.64	\$69,103	\$67,193	97%
7T-2369	S3 lifecycle and storage classes	sf-log-archive	Sep 10	\$228.60	\$44.51	\$72,549	\$67,192	93%
8S-0410	Instances off nights and weekends	fleet-perf	Sep 8	\$232.21	\$50.34	\$73,435	\$66,381	90%
9T-2469	S3 lifecycle and storage classes	fleet-eu-prod	Sep 10	\$202.69	\$23.45	\$67,303	\$65,420	97%
10T-3445	AWS Backup cold storage	sf-backup	Sep 10	\$324.18	\$148.67	\$65,090	\$64,060	98%
11A-0172	AWS Config recording scoped	sf-cicd	Sep 22	\$234.90	\$65.36	\$60,934	\$61,883	102%
12T-2461	S3 lifecycle and storage classes	fleet-eu-prod	Sep 10	\$201.25	\$36.59	\$65,199	\$60,104	92%
13T-2438	S3 lifecycle and storage classes	fleet-data-prod	Sep 11	\$241.18	\$99.78	\$54,722	\$51,612	94%
14C-0001	Database Savings Plan	sf-management	Sep 18	\$139.70	\$0.20	\$50,918	\$50,917	100%
15S-0407	Instances off nights and weekends	fleet-perf	Sep 9	\$174.88	\$37.32	\$55,051	\$50,210	91%
16T-2418	S3 lifecycle and storage classes	fleet-data-prod	Sep 10	\$212.82	\$79.17	\$49,323	\$48,785	99%
17T-2414	S3 lifecycle and storage classes	fleet-data-prod	Sep 8	\$238.05	\$104.47	\$53,974	\$48,757	90%
18T-2518	S3 lifecycle and storage classes	crucible-training-prod	Sep 8	\$153.16	\$20.77	\$47,010	\$48,319	103%
19S-0409	Instances off nights and weekends	fleet-perf	Sep 11	\$154.43	\$25.69	\$48,934	\$46,989	96%
20W-0157	Idle instances	sandbox-ml	Sep 8	\$124.09	\$0.00	\$46,307	\$45,293	98%
21A-0173	AWS Config recording scoped	crucible-sim-prod	Sep 22	\$172.75	\$54.54	\$44,528	\$43,147	97%
22W-6029	Idle SageMaker notebooks and Studio apps	crucible-research	Sep 8	\$142.22	\$24.19	\$42,205	\$43,080	102%
23T-2367	S3 lifecycle and storage classes	sf-log-archive	Sep 11	\$141.24	\$27.28	\$45,657	\$41,595	91%
24T-2426	S3 lifecycle and storage classes	fleet-data-prod	Sep 9	\$179.77	\$76.29	\$39,437	\$37,772	96%

#ID	Change	Account	Applied	Net daily before	Net daily after	Estimated	Measured	vs estimate
25T-2370	S3 lifecycle and storage classes	sf-log-archive	Sep 8	\$117.57	\$16.19	\$37,718	\$37,004	98%
26T-2382	S3 lifecycle and storage classes	sf-cicd	Sep 10	\$112.03	\$11.44	\$38,046	\$36,717	97%
27W-6030	Idle SageMaker notebooks and Studio apps	crucible-research	Sep 9	\$135.67	\$35.27	\$39,855	\$36,646	92%
28W-6042	Idle SageMaker notebooks and Studio apps	crucible-research	Sep 9	\$130.62	\$34.26	\$38,225	\$35,170	92%
29T-2430	S3 lifecycle and storage classes	fleet-data-prod	Sep 11	\$177.94	\$84.67	\$37,264	\$34,044	91%
30W-6033	Idle SageMaker notebooks and Studio apps	crucible-research	Sep 11	\$124.88	\$34.34	\$36,650	\$33,047	90%
31T-2368	S3 lifecycle and storage classes	sf-log-archive	Sep 11	\$155.63	\$69.73	\$31,872	\$31,353	98%
32T-2387	S3 lifecycle and storage classes	sf-cicd	Sep 11	\$90.30	\$7.18	\$30,929	\$30,339	98%
33A-0171	WorkSpaces billed by the hour	corp-apps	Sep 11	\$339.21	\$258.64	\$34,247	\$29,409	86%
34S-0412	Instances off nights and weekends	crucible-dev	Sep 9	\$130.82	\$52.36	\$30,817	\$28,640	93%
35W-6027	Idle SageMaker notebooks and Studio apps	crucible-research	Sep 9	\$116.23	\$38.45	\$32,512	\$28,389	87%
36T-2372	S3 lifecycle and storage classes	sf-log-archive	Sep 11	\$118.63	\$43.38	\$29,141	\$27,466	94%
37S-0402	Instances off nights and weekends	fleet-dev	Sep 9	\$116.26	\$41.61	\$27,393	\$27,246	99%
38T-2371	S3 lifecycle and storage classes	sf-log-archive	Sep 11	\$114.74	\$40.66	\$28,025	\$27,041	96%
39T-2530	S3 lifecycle and storage classes	crucible-data	Sep 9	\$113.85	\$46.91	\$25,308	\$24,432	97%
40T-2516	S3 lifecycle and storage classes	crucible-training-prod	Sep 10	\$77.48	\$14.95	\$24,120	\$22,825	95%

Intelligent-Tiering: in place, not yet visible

161 S3 buckets moved to S3 Intelligent-Tiering in week 3. The Analysis estimated \$1.3 million a year for them, and none of it is counted here. Intelligent-Tiering moves an object only after 30 days without access in the new class (90 for Archive Instant Access), so the first saving shows about October 10 and the full saving by December 9. Until then the bill shows only the one-time charge of putting the objects in the class, and a measurement taken now would show a loss that is not real.

The plan is the same query on a bucket basis (storage, requests, and retrieval) from October 10, 2026, when the first saving should show, with the full saving expected from December 9, 2026. We will measure it at the first monthly review after that date and add it to the register at the measured figure, whatever it is. Until it is measured, the guarantee does not rely on it: the counted figure above clears the gate without it.

One-time charges

Some changes cost money once. The charges below are not netted into the savings above; they are reported on their own, as the engagement defines, and the gate was checked both ways.

One-time charges of the changes put in place. They were all incurred in week 3.

Kind of change	One-timeWhat it is charge
S3 Intelligent-Tiering	\$28,433Lifecycle transition requests into S3 Intelligent-Tiering
S3 lifecycle and storage classes	\$56,431Lifecycle transition requests into Glacier Instant Retrieval and Deep Archive
Idle instances	\$1,727AMIs of instances before they were terminated
Unattached volumes	\$9,714Archive-tier snapshots kept for 90 days before the volumes were deleted
Idle non-production databases	\$597The 90-day final snapshot kept for each deleted database
Long-stopped instances	\$196AMIs of instances before they were terminated
Total	\$97,097

The total, \$97,097, is 1.0% of one year’s measured savings. At the measured rate it is repaid in about 4 days of savings. Deducted from the counted figure, it leaves \$9,554,927, which is 273.0 times the fee.

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 StrataForge is not penalized for an owner’s reasonable second thought.

It applies to 10 items this engagement, all of one kind, the S3 lifecycle and storage classes change, for buckets in crucible-sim-prod. Hiro Tanaka’s team approved them at the readout and later chose to keep a year of results for a regression study. We did not push, and we report the savings separately from the measured figure, at the estimate, because there is no measurement of a change that was never made.

The items approved at the readout and declined afterward, counted at their estimates.

ID	Bucket	Owner	EstimatedReason annual savings
T-2506	sf-episode-results-sim-results-7019-us-east-1-000	Hiro Tanaka	\$37,711approved at the readout; the Crucible team later chose to keep a year of results for a regression study
T-2507	sf-episode-results-sim-results-7019-us-east-1-001	Hiro Tanaka	\$59,685approved at the readout; the Crucible team later chose to keep a year of results for a regression study
T-2508	sf-episode-results-sim-results-7019-us-east-1-002	Hiro Tanaka	\$10,108approved at the readout; the Crucible team later chose to keep a year of results for a regression study
T-2509	sf-episode-results-sim-results-7019-us-east-1-003	Hiro Tanaka	\$17,079approved at the readout; the Crucible team later chose to keep a year of results for a regression study
T-2510	sf-episode-results-sim-results-7019-us-east-1-004	Hiro Tanaka	\$6,412approved at the readout; the Crucible team later chose to keep a year of results for a regression study

ID	Bucket	Owner	EstimatedReason annual savings
T-2511	sf-episode-results-sim-results-7019-us-east-1-005	Hiro Tanaka	\$3,300approved at the readout; the Crucible team later chose to keep a year of results for a regression study
T-2512	sf-episode-results-sim-results-7019-us-east-1-006	Hiro Tanaka	\$2,209approved at the readout; the Crucible team later chose to keep a year of results for a regression study
T-2513	sf-episode-results-sim-results-7019-us-east-1-007	Hiro Tanaka	\$46,691approved at the readout; the Crucible team later chose to keep a year of results for a regression study
T-2514	sf-episode-results-sim-results-7019-us-east-1-008	Hiro Tanaka	\$10,031approved at the readout; the Crucible team later chose to keep a year of results for a regression study
T-2515	sf-episode-results-sim-results-7019-us-east-1-009	Hiro Tanaka	\$81,685approved at the readout; the Crucible team later chose to keep a year of results for a regression study
Total			\$274,912

If Hiro Tanaka's team later reverses the decision and the rules go in, those buckets are measured the same way as every other change, and we replace the estimate with the measurement in the record.

What remains

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

Identified equals in place plus remaining plus declined plus approved–then–declined. The approved–then–declined figure is the amount counted for the gate on top of the measured savings.

Where the identified savings are	Annualized
Identified by the Analysis	\$23,520,211
In place, at the Analysis’s estimates: the counted changes and Intelligent–Tiering, which is not yet visible	\$11,129,490
Remaining, still to be put in place (includes the unapplied part of lines applied in part)	\$11,760,651
Declined by the owner	\$355,158
Approved at the readout, then declined	\$274,912

Of the remaining \$11.8 million a year, \$11.3 million is in lines not yet put in place at all, and the rest is the part of partially applied lines that is still waiting, such as the production node groups behind the one Graviton pilot. Neither is in the count. The handover report turns the remaining items into a roadmap with an owner and a date for each. The largest groups of lines:

The six largest groups of items still to put in place, estimated and annualized.

Still to put in place	Changes	Annualized
Commitments resized at renewal	1	\$2,100,001
Simulation on Spot through AWS Batch	334	\$1,422,523
Rightsize and move to Graviton	290	\$1,018,811
Robot telemetry to S3 through Firehose	96	\$818,009
Per–robot metrics to site metrics	4	\$814,735
S3 lifecycle and storage classes	53	\$778,099
All lines not yet put in place		\$11,296,614

The largest of them is the queued three-year Compute Savings Plan purchase, which starts November 1, 2026, after the measurement window. When it starts it will be measured the same way as the rest, in the first monthly review after.

Re-running this measurement

Everything above can be re-run by StrataForge's own team. The inputs are the CUR 2.0 export and the register of changes (which resource or tag, which date). The method is 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 (the export finalizes about 48 hours late, so that is nine days), run the query, and add the row to the register. The monthly cost review in the handover report keeps the register and re-runs the whole set each quarter, so the claim "the bill fell by this much" can always be checked by someone who was not in the room.

© 2026 Vantalect, LLC. Sample deliverable; StrataForge Robotics is fictional.

AWS, Amazon Web Services, AWS Graviton, and the AWS service names are trademarks of Amazon.com, Inc. or its affiliates. Kubernetes is a registered trademark of The Linux Foundation. GitHub is a trademark of GitHub, Inc. NVIDIA and Isaac Sim are trademarks of NVIDIA Corporation. SAP and SAP S/4HANA are trademarks of SAP SE. Vantalect is not affiliated with or endorsed by any of them. StrataForge Robotics is fictional.