



AWS COST · SAMPLE DELIVERABLE

Cost Analysis and Savings Roadmap

Week 2 of 6: Find and rank · prepared for Copperwick Therapeutics, Inc.

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

March 5, 2026 · AWS Cost Analysis & Optimization: The Optimization (six weeks)

SAMPLE DELIVERABLE

This is a sample deliverable. Copperwick Therapeutics, Inc. is a fictional company: its people, programs, products, accounts, and figures are invented for this sample, and any resemblance to a real company, person, program, or medicine is coincidental. AWS prices are list prices from the AWS Price List API retrieved on October 13, 2026, and from the public AWS Price List bulk API on October 17, 2026; Spot prices are 30-day means for September 14 to October 14, 2026.

Contents

Executive summary	2
Method	6
How each change reaches production	9
Deliverable 2: the waste register	13
Deliverable 3: the rightsizing plan	19
Batch, Spot, and file systems	25
Storage and databases	29
Network and the data lake	33
Logging, governance, and Graviton	39
Deliverable 4: the commitment plan	43
Deliverable 6: the savings roadmap	47
What is not counted	55
Appendix A: the registers	57
Appendix B: assumptions	58
Appendix C: price provenance	58

Executive summary

We found **\$4,310,044 a year** of savings in Copperwick's AWS bill, net of the 6% private pricing discount, which is 23.5% of the \$18,315,196 a year the in-scope accounts cost on that basis (\$4,585,154 a year at list prices). They come from 36 findings, priced as 41 changes in the 48 in-scope accounts, and carried out in full they take the in-scope bill from **\$1,526,267 to \$1,167,095 a month**. These are estimates from Copperwick's own usage and AWS list prices, restated net of the discount because that is what the bill falls by. They become measured savings only in week six, change by change, on Copperwick's Cost and Usage Report.

Every change in this report is **proposed**. The report precedes tomorrow's sponsor readout, so nothing here has been approved, scheduled, deferred, or declined.


What clears the gate, in dollars. The Optimization's second invoice, \$17,500, is issued only if the savings in place at the end of week six, measured on Copperwick's Cost and Usage Report, annualized, and net of the discount, reach **\$105,000 a year**: 0.6% of the in-scope bill. The

roadmap proposes 27 changes for weeks 3 and 4, and the 19 of them rated High confidence are \$1,987,526 a year, a margin of \$1,882,526 over the gate. The five waste findings alone, W-1 through W-5, are \$644,816 a year, \$539,816 over it ([The guarantee](#)).

The headline numbers. Savings are estimates restated net of Copperwick's 6% private pricing discount; list-price figures are context.

Measure	Amount	What it is
In-scope bill at the February run rate	\$1,526,267 a month	\$18,315,196 a year, net of the 6% private pricing discount. At list prices, after the existing Savings Plan: \$1,623,688 a month, \$19,484,251 a year
Identified savings, net	\$4,310,044 a year	\$359,172 a month across all 41 changes; 23.5% of the bill. At list: \$4,585,154 a year
After the roadmap	\$1,167,095 a month	Net. Before the fall in Enterprise Support, which is not counted (What is not counted)
High confidence	\$2,328,959 a year	24 changes, 54% of the identified savings
Proposed for weeks 3 and 4	\$3,063,619 a year	27 changes, March 9 to 20, before the last day a change can be measured (The waves)
Waiting on a date	\$1,246,425 a year	14 changes that cannot land before April 2, 2026, the end of the engagement (Changes that wait for a date)
Touching validated or qualified systems	\$709,873 a year	15 of the 41 changes, 16.5% of the identified savings (How each change reaches production)
One-time charges	\$529 at list	4 changes, listed in One-time charges ; never netted into the annual savings

Three things shape how to read the number. **Nothing is counted twice:** the savings are an ordered stack, waste first and the Savings Plan last, and every finding is the difference between two prices computed the same way ([Method](#)). **Identified is not in place:** 14 of the changes (\$1,246,425 a year) wait on a freeze, a change request, a campaign, a release, a benchmark, or a test, and 15 touch validated or qualified systems, where every change goes through Copperwick's change control. Nothing here says any change will be approved; each owner decides their own. **S3-4 is identified and cannot be counted inside the six weeks:** Intelligent-Tiering moves objects only after 30 days without access, so its \$119,686 a year is not measurable until its first seven-day window, May 11 to May 17, 2026 (the result is ready May 19), and it is reported separately. A bigger number sits beside all of this and is in none of it: what the guardrails could keep out of Copperwick's planned growth is \$250,000 to \$410,000 a year, not measured, not counted toward the gate, and not a saving on the bill ([Cost avoidance on planned growth](#)).



From today's bill to the bill after the roadmap, in the fixed order, dollars a month, net of the private pricing discount.

What is in the number

Identified savings by category, in the order the report prices them. B-1, S-3, E-2, N-1, and S3-1 each split into two changes, so 36 findings are 41 changes.

Category	IDs	Changes	Per month, net	Per year, net	Per year, at list
Waste	W-1, W-2, W-3, W-4, W-5	5	\$53,734	\$644,816	\$685,973
Rightsizing	R-1, R-2, R-3, R-4, R-5	5	\$27,172	\$326,068	\$346,881
Schedules	S-1, S-2, S-3a, S-3b	4	\$31,805	\$381,657	\$406,017
Batch and Spot	B-1, B-1b, B-2, B-3	4	\$51,724	\$620,679	\$660,298
File systems	F-1, F-3	2	\$37,655	\$451,864	\$480,707
Storage	E-1, E-2a, E-2b, EF-1	4	\$17,616	\$211,379	\$224,871
Databases	D-1, D-2	2	\$10,847	\$130,169	\$138,478
Network	N-1a, N-1b	2	\$11,763	\$141,149	\$150,160
Data lake	S3-1, S3-1b, S3-2, S3-3, S3-4, S3-5	6	\$58,266	\$699,194	\$743,824
Logging and governance	L-1, L-2, L-3, L-4	4	\$18,651	\$223,807	\$238,092

Category	IDs	Changes	Per month, net	Per year, net	Per year, at list
Graviton	G-1	1	\$5,490	\$65,875	\$70,080
Commitments	C-1, C-2	2	\$34,449	\$413,387	\$439,773
Total		41	\$359,172	\$4,310,044	\$4,585,154

What happened this week

The clinical freeze began on Monday, March 2, 2026: Program Basalt's database lock is March 13 and topline is April 24, and the freeze stops every change to the clinical systems until then, with no exception for cost. On Tuesday, March 3, 2026, we watched the GxP change advisory board work through its agenda, the first of five Tuesdays, to learn the evidence it asks for. Hana Lindgren walked cryo-EM and molecular dynamics with us. Today, Thursday, March 5, 2026, the savings review puts every finding in front of its owner and Dr. Lorraine Osei; the sponsor readout with Adaeze Whitfield is tomorrow, Friday, March 6, 2026.

Why a capable team missed it. Nobody was careless. A research estate's bill follows experiments, and the habits that protect science cost money: the ML team keeps p4d instances running because capacity is hard to get back, a file system is provisioned for the peak that might come, a volume is kept because someone might open the project again, and anything in a validated system is left alone because every change there costs a ticket and a signature. The Analysis separates the changes that are safe from the ones that need a process, and prices both.

What we need from you

- 1. Tell us which findings an owner will not take on as written.** Every change has a named owner, and the owner decides; we rank by dollars, and owners know things we do not.
- 2. Dr. Lorraine Osei: confirm the change path of each of the 15 changes classed as GxP or frozen** ([The 15 changes that carry a GxP class](#)), and the board Tuesday for each change request: S-3a on Tuesday, March 17, S3-3 on Tuesday, March 31, N-1b on Tuesday, March 31, S-3b on Tuesday, April 7, B-1b on Tuesday, April 21, R-3 on Tuesday, April 28, R-4 on Tuesday, May 26, R-2 on Tuesday, June 2.
- 3. Confirm the change windows:** non-GxP changes in weeks 3 and 4, nothing in the clinical systems before topline on April 24, and nothing in the production manufacturing systems from March 16 to 27.
- 4. Look at the commitment plan.** Rafael Quintero owns the one purchase in this report, a one-year No Upfront Compute Savings Plan at \$50.00 an hour (C-1), and, with Mateo Varga, the question of a second, family-locked plan for the p4d fleet (C-2) ([Deliverable 4: the commitment plan](#)).

5. **Sign off the waste register before anything is deleted.** Besides the Savings Plan, the deletions are the only changes that cannot be undone, and each has a snapshot or a named way back ([Deliverable 2: the waste register](#)).

Method

The numbers come from one model of Copperwick's in-scope estate, rebuilt from the resource inventory and twelve months of Copperwick's Cost and Usage Report and priced hour by hour at AWS list prices. Every finding later in the report shows its own arithmetic. A finding's price is a quantity times a price from the committed tables, and the report names the quantity and the price, so each row recomputes.

Savings are net of the private pricing discount. Copperwick has a private pricing agreement with AWS: a flat 6% discount on eligible usage, Savings Plan fees included, for the term July 1, 2024 to June 30, 2027. The bill falls by the discounted amount, so the savings, the gate, and the share of the bill in this report are net of it, and list-price figures appear beside them, labeled "at list". The share of the bill is the same on either basis.

The fixed order

A resource can appear in more than one finding: an instance is rightsized and then scheduled. Each finding is priced on the estate as the findings before it leave it, so the later one is worth what remains and no dollar is counted twice. The order never changes:

The order every finding is priced in. Waste goes first because it needs no design and shrinks what everything after is sized on; rightsizing and schedules precede Spot, so hours come off the right size; Graviton is priced on the sizes rightsizing leaves; and the Savings Plan goes last, on the 24-hour floor that survives every change before it ([Deliverable 4: the commitment plan](#)).

Step	Category	Findings	Bill after the step, net a month
	Today's bill		\$1,526,267
1	Waste	W-1, W-2, W-3, W-4, W-5	\$1,472,533
2	Rightsizing	R-1, R-2, R-3, R-4, R-5	\$1,445,361
3	Schedules	S-1, S-2, S-3a, S-3b	\$1,413,556
4	Batch and Spot	B-1, B-1b, B-2, B-3	\$1,361,832
5	File systems	F-1, F-3	\$1,324,177
6	Storage	E-1, E-2a, E-2b, EF-1	\$1,306,561

Step	Category	Findings	Bill after the step, net a month
7	Databases	D-1, D-2	\$1,295,714
8	Network	N-1a, N-1b	\$1,283,951
9	Data lake	S3-1, S3-1b, S3-2, S3-3, S3-4, S3-5	\$1,225,685
10	Logging and governance	L-1, L-2, L-3, L-4	\$1,207,034
11	Graviton	G-1	\$1,201,544
12	Commitments	C-1, C-2	\$1,167,095

A saving is the difference between two prices

Each finding is a change applied to a copy of the estate, priced the same way the bill is, so a saving is always the difference between two prices computed the same way. The simplest case is [R-1: Data platform services: m6i.4xlarge to m6i.2xlarge](#), 36 hosts moving from an m6i.4xlarge to an m6i.2xlarge for 730 hours: $36 \times (\$0.768 - \$0.384) \text{ an hour} \times 730 \text{ hours} = \$10,091.52$ a month at list. Four rules keep the stack honest:

- **A later finding is priced on the earlier ones' result.** The Savings Plan (C-1) is sized on the floor that every earlier change leaves, including R-1's smaller instances and B-1's Spot hours, so it never pays for an hour an earlier finding removed.
- **Rates are list prices for a 730-hour month,** from the AWS Price List API, plus the prices that table lacks from the public AWS Price List bulk API: EFS Infrequent Access reads at \$0.01 a GB, Glacier Flexible Retrieval standard retrieval at \$0.01 a GB and bulk retrieval at \$0, the hpc7g.16xlarge instance at \$1.6832 an hour, and the AWS Enterprise Support tiers. Nothing in this report is an estimated price ([Appendix C: price provenance](#)).
- **Spot is the 30-day mean plus a retry allowance.** B-1 prices every Spot hour at the price table's 30-day mean, \$0.5815 an hour for c7i.8xlarge, plus a 12% allowance for retried work, and gives no credit for c7i's faster cores. Savings Plans do not apply to Spot, so C-1 is sized after this change ([B-1: Genomics Batch, research pipelines: Spot](#)).
- **S3 is priced at the top Standard tier.** S3 Standard is priced in volume tiers across the organization: 51,200 GB at \$0.023, the next 460,800 GB at \$0.022, and everything above at \$0.021 a GB-month. The in-scope accounts alone keep 6.5 PB in Standard, far above the first 500 TB, so every gigabyte that leaves comes off the top tier and saves \$0.021 a GB-month before the cost of wherever it goes.

The Savings Plan is sized last, on the floor of 24-hour non-GPU usage that survives every other change, so the plan never covers an hour another finding has already removed ([Deliverable 4: the commitment plan](#)). Copperwick's existing three-year plan (\$150 an hour, August 1, 2024 to July 31, 2027) covers \$295.77 an hour of On-Demand usage at 49.3% off and stays fully used after every change, so no change's saving is absorbed by an idle commitment.

Prices, rules, and ratings

Confidence High means the price is a known configuration or the measurement covers the case that matters; Medium means the saving depends on behavior the Analysis could not fully observe, such as incremental snapshot blocks, a throughput peak, or a benchmark. Of the 41 changes, 24 are High (\$2,328,959 a year net) and 17 Medium (\$1,981,085). **Change path** is the class each change carries under Copperwick's change control ([How each change reaches production](#)). Nothing changes until its owner approves it, and nothing in a validated system changes without its path.

How the savings will be counted at week six. The gate counts changes in place at the end of week six, measured change by change on Copperwick's Cost and Usage Report: the cost per day of each changed resource in the seven days before the change against the seven days after, the change day excluded from both, multiplied by 365. A change needs seven full days of data after it, and the report finalizes about 48 hours late, so the last day a change can land and be measured is March 22; the plan puts every change meant to count in weeks 3 and 4. Non-commitment changes are measured at public On-Demand cost, so neither the existing Savings Plan, C-1, nor the private pricing discount is counted twice; Spot (B-1) at the Spot price paid; volume-driven changes (B-1, N-1a) per genome processed, because genomics spend follows lab throughput; and C-1 by its Savings Plan ARN. One-time charges are listed beside the result and never annualized. Deferred and declined changes, and changes in place whose savings cannot yet be measured, count for nothing.

The private pricing agreement and the commitment

A CFO's next question is whether cutting spend risks a shortfall against the agreement's commitment of \$27 million a year in spend net of the discount. It does not. The whole organization runs at about \$3,099,688 a month at list, \$37,196,251 a year, which is **\$34,964,476 a year net of the discount**; the 164 out-of-scope accounts are about \$1,476,000 a month of that. The entire identified roadmap would take it to \$30,654,432, which stays \$3,654,432 above the commitment before any planned growth. The savings inform the 2027 renewal: Copperwick can commit to a lower annual figure, or the same figure at a deeper discount. That value is real and is not counted.

How each change reaches production


This is the part of the roadmap a quality leader reads first. Copperwick validates its GxP computerized systems under 21 CFR Part 11 and GAMP 5, using a risk-based computer software assurance approach for testing. The cloud platform underneath is qualified infrastructure: the landing zone, network, and shared services are built from code under an infrastructure qualification, so a change inside that qualification is controlled but does not revalidate the applications above it. At kickoff, Dr. Lorraine Osei, Senior Director, Quality Assurance, Computerized Systems, set four rules, and every finding in this report carries the class they give it.

The four rules and the changes in each class. Every one of the 41 changes has exactly one class.

Rule	Class	What it needs	Changes	Per year, net
1	Non-GxP	The platform team's normal change process. Most of the savings are here. It also covers a change to a part of a qualified account that sits outside its qualification: sdp-lake-prod is qualified for the GxP raw-data buckets and the account controls that protect them only, so a change to its research buckets follows this rule.	26	\$3,600,171
2	GxP standard change	A pre-approved change type in the qualified infrastructure catalog: like-for-like, reversible, with no effect on a GxP function. It needs a ticket, an executed installation checklist, and notice to QA, not a new validation. Changing gp2 to gp3 on a validated server is one, because it is an online Elastic Volumes change with the same or better performance. A change that runs from the qualified log-archive and security-tooling accounts, such as log retention, Config recording, and CloudTrail event selectors, is one too: it goes with notice to QA and moves no date. (E-2a, L-2, L-3, L-4)	4	\$138,332
3	GxP change request	The GxP change advisory board, which meets on Tuesdays (March 3, 10, 17, 24, and 31; the first has passed). It needs a validation impact assessment, risk-based testing (unscripted where the risk is low, scripted operational qualification where a GxP function could be affected), QA approval, and a release window. A database resize, a storage-class change on GxP records, or a new network path is one. (S-3a, R-3, R-4, B-1b, D-2, N-1b, S3-3)	7	\$381,543

Rule	Class	What it needs	Changes	Per year, net
4	Freeze	A freeze stops every change in its scope, with no exception for cost (the one exception is a security change, such as removing a supplier's access): the clinical freeze (March 2 to April 24) and the GMP campaign freeze (March 16 to 27). Each covers named systems: the clinical freeze covers <code>clin-sce-prod</code> , <code>clin-sce-val</code> , <code>clin-cdr-prod</code> , and <code>clin-biomarker-prod</code> , and the GMP freeze covers <code>gmp-lims-prod</code> , <code>gmp-mes-prod</code> , <code>gmp-cds-prod</code> , and <code>gmp-historian-prod</code> . No development account, GMP val environment, or clinical operations account is frozen. The clinical freeze holds 4 changes, and the GMP freeze holds 0: E-2a is scheduled for Thursday, March 12, before the GMP freeze begins on Monday, March 16, and S-3a changes val environments, which the GMP freeze does not cover. (S-3b, R-2, R-5, E-2b)	4	\$189,998
Total			41	\$4,310,044

How a finding is routed: it is first classed by its account. A Non-GxP finding goes through the normal change process and can be in place in week 3 or 4. A GxP finding is checked against the freezes; if its system is frozen, it waits until the freeze ends. If not, a like-for-like change inside the infrastructure qualification is a standard change, done with a checklist and notice to QA. Anything else is a change request: an impact assessment, risk-based testing, the Tuesday board, QA approval, and execution in the system's release window, which for most of these systems falls after the six weeks.



The identified savings by change path, net a year. Non-GxP changes are 84% of the total.

What this does to the numbers

Of the 41 changes, **15 touch validated or qualified systems, worth \$709,873 a year net** (\$755,184 at list), 16.5% of the identified savings. Of the 15, 6 (S-3a, E-2a, D-2, L-2, L-3, L-4; \$306,635) are proposed for weeks 3 and 4, and 9 (\$403,238) wait on a freeze, a change request, a release, or a window. The validated accounts are about 8.7% of the in-scope bill, so a small share of the savings there is expected; the report says so plainly rather than stretching it. 6 more changes (W-5, EF-1, S3-1, S3-1b, S3-2, S3-4) are made in a qualified account but change only research resources, so they follow Rule 1 and are classed Non-GxP; the account's qualification covers the GxP raw-data buckets and the controls that protect them, and nothing else. The standard changes among the 15 (L-2, L-3, and L-4) are class 2 even though they act on research resources, because they run from the qualified log-archive and security-tooling accounts, and L-2 names its research log groups by prefix and account, never by the absence of a tag.

Records keep their retention. No finding deletes, shortens, or weakens a GxP record, an audit trail, or a backup of a validated system. GxP buckets keep Object Lock and versioning; their CloudTrail data events and their log groups keep their standard-operating-procedure retention; and snapshots in validated accounts are excluded from the waste work. Research data has retention obligations too: 10 years for raw data that backs invention records, and the lifecycle findings move it to colder storage, never out of retention.

The 15 changes that carry a GxP class

Every change classed GxP or frozen, largest first, with the system it touches and when it could go in.
Dollars are net of the private pricing discount.

ID	Change	System	Class	Proposed	Confidence	Per year, net
S-3a	GMP val environments: weekday schedule	LIMS, MES, and CDS val environments	GxP change request	week 4 (March 16 to 20)	High	\$87,540
S3-3	GxP raw data to Glacier Instant Retrieval	GxP raw-data buckets (Object Lock)	GxP change request	earliest April 14	Medium	\$84,451
D-2	Multi-AZ off: LIMS and MES val	LIMS and MES val databases	GxP change request	week 4 (March 16 to 20)	High	\$80,763
R-2	SCE grid: 8 to 6 nodes	SCE production grid	Frozen (clinical)	earliest June 22	Medium	\$77,785
L-2	Log retention: 90 days for research		GxP standard change	week 3 (March 9 to 13)	High	\$58,216
R-5	Clinical data repository: 8 to 6 nodes	Clinical data repository (Redshift)	Frozen (clinical)	earliest April 27	Medium	\$53,688
R-4	MES database: db.r5.4xlarge to db.r5.2xlarge	MES production database	GxP change request	earliest June 15	Medium	\$50,065
S-3b	SCE val environment: weekday schedule	SCE val environment	Frozen (clinical)	earliest April 27	High	\$45,819
B-1b	Biomarker pipelines: Spot	GCP-validated biomarker pipelines	GxP change request	earliest May 11	Medium	\$42,020
L-4	CloudTrail data events: writes only on research buckets		GxP standard change	week 3 (March 9 to 13)	High	\$35,194
L-3	AWS Config: daily recording in three accounts		GxP standard change	week 3 (March 9 to 13)	High	\$32,216

ID	Change	System	Class	Proposed	Confidence	Per year, net
R-3	LIMS database: db.r5.8xlarge to db.r6i.4xlarge	LIMS production database	GxP change request	earliest May 18	Medium	\$30,698
E-2a	gp2 to gp3: GMP systems, before their freeze	LIMS, MES, CDS, and historian production volumes, and the val environments	GxP standard change	week 3 (March 9 to 13)	High	\$12,706
E-2b	gp2 to gp3: clinical systems, in the freeze	SCE production and val volumes, and the biomarker pipelines	Frozen (clinical)	earliest April 27	High	\$12,706
N-1b	S3 gateway endpoints: validated VPCs	Network paths of the validated VPCs	GxP change request	earliest April 14	High	\$6,006
Total						\$709,873

Deliverable 2: the waste register

\$644,816 a year net (\$685,973 at list) of Copperwick's bill pays for resources nobody uses or for capacity that sits idle: 5 findings, W-1 through W-5, all in non-GxP accounts. Removing it needs no new design, only an owner's yes and a safe way back. The register is at the end of this section ([The register](#)); the CSV carries the same rows, docs/samples/copperwick/data/cost/waste-register.csv.

How every removal is done. The same five rules apply to every finding here. (1) The owner signs off first, and nothing is deleted on our word alone. (2) Tag the resource `cw:pending-removal=<date>` and wait a business day. (3) Keep a way back at the cost the finding states: a snapshot kept 90 days, an image of the instance, or a lifecycle rule that only an owner can remove. (4) Prove it in the bill a week later; that is what week 6 counts. (5) If the resource is in code, delete it from the repository in the same change, so the next deployment does not recreate it. Snapshots in validated accounts are excluded from this work.

The five findings

W-1: Unattached gp2 volumes

Owner: Priya Okonkwo. **Change path:** Non-GxP. **Proposed for:** week 3 (March 9 to 13).

Confidence: High. **Saves:** \$95,871 a year net of the private pricing discount (\$101,990 at list).

Status: proposed.

The finding. Unattached gp2 volumes in research and sandbox accounts (90 TB); 14 TB snapshotted first and kept 90 days. **What we saw.** No attachment in 90 days of CloudTrail; each volume goes to its owner for confirmation in the waste register.

The change. Snapshot the 14 TB owners asked for, then delete every volume. **Rollback:** Restore a kept volume from its snapshot.

The volumes sit in research and sandbox accounts, so the script runs once per account. A detached volume with no I/O in 90 days cannot be missed by a running host, and the owners asked to keep 14 TB as snapshots, which is why that cost is netted from the saving even though it ends after 90 days.

BASH

```
# Per account, as the remediation role. keep.txt holds the volumes whose owners asked for a
snapshot; delete.txt holds the rest.
# Only rows the owner signed off.
while read -r VOL; do
    SNAP=$(aws ec2 create-snapshot --volume-id "$VOL" --description "waste-register W-1" --query
    SnapshotId --output text \
        --tag-specifications 'ResourceType=snapshot,Tags=[{Key=cw:register-id,Value=W-1},
    {Key=cw:delete-after,Value=<yyyy-mm-dd>}]')
    aws ec2 wait snapshot-completed --snapshot-ids "$SNAP" && aws ec2 delete-volume --volume-id
    "$VOL"
done < keep.txt
while read -r VOL; do aws ec2 delete-volume --volume-id "$VOL"; done < delete.txt
```

How it was priced, at list. 90 TB × 1,024 × \$0.10 a GB-month, less 14 TB × 1,024 × \$0.05 for the snapshots owners asked to keep = \$8,499.20 a month, \$101,990.40 a year. Net of the 6% discount, **\$95,871 a year.**

W-2: SageMaker GPU notebooks: auto-stop

Owner: Mateo Varga. **Change path:** Non-GxP. **Proposed for:** week 3 (March 9 to 13). **Confidence:** High. **Saves:** \$182,152 a year net of the private pricing discount (\$193,778 at list). **Status:** proposed.

The finding. SageMaker GPU notebooks: auto-stop after 60 idle minutes (4 ml.g5.12xlarge billed about 22% of hours after). **What we saw.** Kernel activity shows each notebook busy about 22% of hours.

The change. A lifecycle configuration stops a notebook after 60 idle minutes; an exemption tag keeps a long training run alive. **Rollback:** Remove the lifecycle configuration.

The ML team keeps four ml.g5.12xlarge notebooks (\$7.09 an hour each) running around the clock for scientists who use them in bursts. A lifecycle configuration stops a notebook after 60 idle minutes, which leaves it billed about 22% of hours. A notebook carrying an exemption tag keeps running, so a long training run is never interrupted. A notebook must be stopped before a lifecycle configuration can be attached.

BASH

```
# autostop-on-start.sh runs on every start: it installs a five-minute cron job that stops the
instance after 3,600 idle seconds
# unless the instance carries the cw:no-autostop tag.
aws sagemaker create-notebook-instance-lifecycle-config --notebook-instance-lifecycle-config-
name cw-autostop-60 \
    --on-start Content="$(base64 < autostop-on-start.sh | tr -d '\n')"
aws sagemaker stop-notebook-instance --notebook-instance-name <notebook-name>
aws sagemaker wait notebook-instance-stopped --notebook-instance-name <notebook-name>
aws sagemaker update-notebook-instance --notebook-instance-name <notebook-name> --lifecycle-
config-name cw-autostop-60
aws sagemaker start-notebook-instance --notebook-instance-name <notebook-name>
```

How it was priced, at list. $4 \times \$7.09 \text{ an hour} \times 730 \text{ hours} \times (1 - 0.22) = \$16,148.18 \text{ a month,}$
\$193,778.21 a year. Net of the 6% discount, **\$182,152 a year.**

W-3: One idle p4d.24xlarge released

Owner: Mateo Varga. **Change path:** Non-GxP. **Proposed for:** week 3 (March 9 to 13). **Confidence:** High. **Saves:** \$180,808 a year net of the private pricing discount (\$192,349 at list). **Status:** proposed.

The finding. One of five p4d.24xlarge training instances, idle 84% of hours: released. **What we saw.** Ninety days of GPU utilization from the CloudWatch agent: one instance idle 84% of hours.

The change. Release one instance; keep four while p4d capacity is scarce. **Rollback:** Relaunch from the training AMI, if capacity allows.

Copperwick runs 5 p4d.24xlarge instances (\$21.958 an hour each) for sequence-design training because the capacity is hard to get back once released. Ninety days of GPU utilization from the CloudWatch agent show one of the five idle 84% of hours. We propose releasing one and keeping four while p4d capacity is scarce; Mateo Varga decides, and the image taken first is the way back if capacity allows.

BASH

```
# An image first, then release. The instance ID comes from the signed-off register row.
aws ec2 create-image --instance-id <instance-id> --name "cw-p4d-training-<yyyy-mm-dd>" --
description "waste-register W-3"
aws ec2 wait image-available --image-ids <image-id>
aws ec2 terminate-instances --instance-ids <instance-id>
```

How it was priced, at list. $1 \times \$21.958 \text{ an hour} \times 730 \text{ hours} = \$16,029.08 \text{ a month}, \$192,348.94 \text{ a year. Net of the 6\% discount, } \$180,808 \text{ a year.}$

W-4: Orphaned EBS snapshots

Owner: Priya Okonkwo. **Change path:** Non-GxP. **Proposed for:** week 3 (March 9 to 13).

Confidence: Medium. **Saves:** \$86,630 a year net of the private pricing discount (\$92,160 at list).

Status: proposed.

The finding. Orphaned EBS snapshots in research accounts (source volume and AMI gone, past retention): 150 TB billed. **What we saw.** Source volume and AMI deleted, past the retention policy; incremental blocks may be shared with kept snapshots, so the rating is medium.

The change. Delete after owner sign-off; snapshots in validated accounts are excluded.

Rollback: None after deletion: each is signed off in the waste register first.

A deleted snapshot is gone, so each deletion has a signed-off line in the register. **The honest limit:** EBS snapshots are incremental. Deleting one frees only the blocks no other snapshot references, so the estimate uses each snapshot's billed cost from the resource-level lines of Copperwick's Cost and Usage Report, not its volume's size, and a deleted snapshot's blocks can turn out to be shared with a survivor. That is why this finding is rated Medium. Snapshots in validated accounts are not in the set.

BASH

```
# From the signed-off list of snapshots whose source volume and image are gone and whose retention has passed.
while read -r SNAP; do aws ec2 delete-snapshot --snapshot-id "$SNAP"; done < approved-snapshots.txt
```

How it was priced, at list. $150 \text{ TB billed} \times 1,024 \times \$0.05 \text{ a GB-month, from each snapshot's billed cost in the CUR's resource lines} = \$7,680.00 \text{ a month}, \$92,160.00 \text{ a year. Net of the 6\% discount, } \$86,630 \text{ a year.}$

W-5: Noncurrent versions and incomplete uploads

Owner: Tomasz Reyes. **Change path:** Non-GxP. **Proposed for:** week 3 (March 9 to 13).

Confidence: High. **Saves:** \$99,355 a year net of the private pricing discount (\$105,696 at list).

Status: proposed.

The finding. Noncurrent object versions and incomplete multipart uploads in research buckets expire after 30 days (0.4 PB). **What we saw.** S3 Storage Lens: 0.4 PB of noncurrent versions and incomplete uploads in research buckets.

The change. A lifecycle rule expires noncurrent versions and aborts incomplete uploads after 30 days; GxP buckets are not touched. **Rollback:** Remove the rule (expired versions do not come back).

Research buckets have no rule that expires noncurrent object versions or incomplete multipart uploads, and S3 Storage Lens shows 0.4 PB of them. The lifecycle rule below expires them after 30 days and leaves current objects alone. GxP buckets are not touched: they keep versioning and Object Lock. The command replaces a bucket's whole lifecycle configuration, so the rule is merged into the bucket's existing rules before it is applied. S3 is billed per bucket in the Cost and Usage Report, and W-5, S3-1, S3-2, and S3-4 change the same lake buckets within days of each other, so each is measured by prefix and storage class: S3 Inventory and S3 Storage Lens give the bytes in each class under each change's prefixes per day, times the published price, reconciled to the bucket's CUR total.

JSON

```
{
  "Rules": [
    {
      "ID": "cw-W-5-noncurrent-and-multipart",
      "Status": "Enabled",
      "Filter": { "Prefix": "" },
      "NoncurrentVersionExpiration": { "NoncurrentDays": 30 },
      "AbortIncompleteMultipartUpload": { "DaysAfterInitiation": 30 }
    }
  ]
}
```

BASH

```
aws s3api get-bucket-lifecycle-configuration --bucket <research-bucket> > current.json #
merge the rule above into it
aws s3api put-bucket-lifecycle-configuration --bucket <research-bucket> --lifecycle-
configuration file://merged.json
```

How it was priced, at list. 0.4 PB × 1,048,576 GB × \$0.021 a GB-month (the top Standard tier) = \$8,808.04 a month, \$105,696.46 a year. Net of the 6% discount, **\$99,355 a year**.

The register

The waste register: every finding with its account and owner. Dollars are net of the private pricing discount; the same rows at list are in the CSV.

ID	Finding	Account	Owner	Confidence	Per month, net	Per year, net
W-1	Unattached gp2 volumes	rc-genomics-prod	Priya Okonkwo	High	\$7,989	\$95,871
W-2	SageMaker GPU notebooks: auto-stop	rc-ml-prod	Mateo Varga	High	\$15,179	\$182,152
W-3	One idle p4d.24xlarge released	rc-ml-prod	Mateo Varga	High	\$15,067	\$180,808
W-4	Orphaned EBS snapshots	rc-genomics-prod	Priya Okonkwo	Medium	\$7,219	\$86,630
W-5	Noncurrent versions and incomplete uploads	sdp-lake-prod	Tomasz Reyes	High	\$8,280	\$99,355
Total					\$53,734	\$644,816

Deliverable 3: the rightsizing plan

Compute that is too big or too awake accounts for \$707,725 a year net of the identified savings: five rightsizing findings, R-1 through R-5 (\$326,068), and four schedule findings, S-1, S-2, S-3a, and S-3b (\$381,657). Rightsizing precedes schedules in the fixed order, so hours come off the right size. R-1, S-1, and S-2 are in non-GxP accounts, S-3a and S-3b are in val environments that host no GxP records, and R-2 through R-5 touch validated production systems; each carries the change path in [How each change reaches production](#).

What the measurements say

What each rightsizing and schedule change rests on, from 93 days of CloudWatch agent metrics, access-time scans, and the owners' own descriptions of the workloads.

ID	Change	Evidence
R-1	Data platform services: m6i.4xlarge to m6i.2xlarge	9% CPU and 31% memory at the 95th percentile over 93 days
S-1	Visualization workstations: idle-stop and a weekday schedule	The other 40 run overnight structure renders, so they stay on
S-2	Development accounts: weekday hours	No CPU activity outside working hours on weekdays, none on weekends
S-3a	GMP val environments: weekday schedule	The val environments host no GxP records; QA's condition is that none stops during an executed protocol
S-3b	SCE val environment: weekday schedule	The same evidence as S-3a; the SCE is in the clinical freeze
R-2	SCE grid: 8 to 6 nodes	The grid's queue never held more than six nodes' work outside the database-lock weeks
R-3	LIMS database: db.r5.8xlarge to db.r6i.4xlarge	14% CPU at the 95th percentile; FreeableMemory never below 61%; buffer-cache hit ratio above 99.6%
R-4	MES database: db.r5.4xlarge to db.r5.2xlarge	11% CPU at the 95th percentile; memory and wait statistics within the vendor's guidance at half the size
R-5	Clinical data repository: 8 to 6 nodes	Cluster CPU under 30% at the 95th percentile; managed storage is separate from the nodes

By group: the **data platform services** (R-1) run at 9% CPU and 31% memory at the 95th percentile over 93 days, so an m6i.2xlarge halves the cost of each and still leaves their memory comfortable. The **SCE grid** (R-2) is eight r6i.16xlarge nodes, and its queue never held more than six nodes' work outside the database-lock weeks, which is why the change waits for topline. **Development**

accounts (S-2) show no CPU activity outside working hours on weekdays and none on weekends. The **visualization workstations** (S-1) split: 20 of the 60 are idle most of the day and get idle auto-stop and a weekday schedule, while the other 40 stay on, because scientists run overnight structure renders on them. The **val environments** (S-3a, S-3b) mirror production but host no GxP records.

The nine findings

Each production resize is a rolling replacement, one host at a time, with the previous size one change away. The change goes in the launch template or the Auto Scaling group's definition so the next deployment keeps it; this command does the day-of roll for a group:

BASH

```
# For a rollback, point the launch template back at the old instance type and run it again.
aws autoscaling start-instance-refresh --auto-scaling-group-name <asg-name> \
  --preferences '{"MinHealthyPercentage": 90, "InstanceWarmup": 300}'
```

R-1: Data platform services: m6i.4xlarge to m6i.2xlarge

Owner: Tomasz Reyes. **Change path:** Non-GxP. **Proposed for:** week 3 (March 9 to 13).

Confidence: High. **Saves:** \$113,832 a year net of the private pricing discount (\$121,098 at list).

Status: proposed.

The finding. Data platform services: 36 m6i.4xlarge to m6i.2xlarge. **What we saw.** 9% CPU and 31% memory at the 95th percentile over 93 days.

The change. Change the instance type in the launch template and roll the Auto Scaling groups.

Rollback: Roll back the launch template version.

The 36 services are in an Auto Scaling group, so the roll keeps capacity throughout. Memory at 31% today would fill about 62% of the smaller size's RAM at the 95th percentile, so m6i.2xlarge still leaves room.

How it was priced, at list. $36 \times (\$0.768 - \$0.384) \text{ an hour} \times 730 \text{ hours} = \$10,091.52 \text{ a month, } \$121,098.24 \text{ a year. Net of the 6\% discount, } \mathbf{\$113,832 \text{ a year.}}$

R-2: SCE grid: 8 to 6 nodes

Owner: Ines Halvorsen. **Change path:** Frozen (clinical). **Proposed for:** earliest June 22. **Board vote:** Tuesday, June 2. **Confidence:** Medium. **Saves:** \$77,785 a year net of the private pricing discount (\$82,750 at list). **Status:** proposed.

The finding. SCE grid: 8 to 6 r6i.16xlarge nodes after topline. **What we saw.** The grid's queue never held more than six nodes' work outside the database-lock weeks.

The change. Remove two nodes after topline, by change request. It waits on the post-topline analyses, which peak right after topline: the scripted operational qualification regression runs first in the SCE val environment, then the change request (board vote June 2). **Rollback:** Add the nodes back.

Removing two nodes from the grid is a change to a validated production system, so it is a change request after topline. The grid's queue history is the evidence for the change request, and the rollback is to add the nodes back if a late analysis needs them.

How it was priced, at list. $2 \times \$4.7232 \text{ an hour (RHEL)} \times 730 \text{ hours} = \$6,895.87 \text{ a month, } \$82,750.46 \text{ a year. Net of the 6\% discount, } \$77,785 \text{ a year.}$

R-3: LIMS database: db.r5.8xlarge to db.r6i.4xlarge

Owner: Daniel Achterberg. **Change path:** GxP change request. **Proposed for:** earliest May 18. **Board vote:** Tuesday, April 28. **Confidence:** Medium. **Saves:** \$30,698 a year net of the private pricing discount (\$32,657 at list). **Status:** proposed.

The finding. LIMS database: db.r5.8xlarge to db.r6i.4xlarge, Multi-AZ (Oracle BYOL). **What we saw.** 14% CPU at the 95th percentile; FreeableMemory never below 61%; buffer-cache hit ratio above 99.6%.

The change. Downsize in the LIMS release, after an operational qualification regression and a month-end performance test. It waits on the LIMS release, with an operational qualification regression and a month-end performance test (board vote April 28). **Rollback:** Modify back to db.r5.8xlarge in a maintenance window.

The LIMS database runs Oracle on a bring-your-own-license basis. CPU alone would be thin evidence, so the evidence is memory: FreeableMemory never falls below 61% and the buffer-cache hit ratio stays above 99.6%, which is why db.r6i.4xlarge, with half the vCPUs and half the memory, is supported. It is a GxP change request: an operational qualification regression and a performance test at month-end load, executed in the LIMS release. The change also halves the Oracle processor licenses the instance needs. The report names that and does not count it, because licenses are not on the AWS bill ([What is not counted](#)).

BASH

```
# Run in the LIMS release window, Multi-AZ, after the regression passes. The previous class is the rollback.
aws rds modify-db-instance --db-instance-identifier <lims-db> --db-instance-class db.r6i.4xlarge --apply-immediately
aws rds wait db-instance-available --db-instance-identifier <lims-db>
```

How it was priced, at list. $(\$7.456 - \$3.728) \text{ an hour} \times 730 \text{ hours} = \$2,721.44 \text{ a month, } \$32,657.28 \text{ a year. Net of the 6\% discount, } \$30,698 \text{ a year.}$

R-4: MES database: db.r5.4xlarge to db.r5.2xlarge

Owner: Daniel Achterberg. **Change path:** GxP change request. **Proposed for:** earliest June 15.

Board vote: Tuesday, May 26. **Confidence:** Medium. **Saves:** \$50,065 a year net of the private pricing discount (\$53,261 at list). **Status:** proposed.

The finding. MES database: db.r5.4xlarge to db.r5.2xlarge, Multi-AZ (SQL Server SE). **What we saw.** 11% CPU at the 95th percentile; memory and wait statistics within the vendor's guidance at half the size.

The change. Downsize in the June maintenance window, after the GMP campaign and the annual periodic review. It waits on the MES maintenance window, after the GMP campaign and the annual periodic review (board vote May 26). **Rollback:** Modify back in a maintenance window.

The MES database runs SQL Server Standard with the license included, at 11% CPU at the 95th percentile. The resize goes in the maintenance window after the GMP campaign and the annual periodic review. No database in scope runs an engine version past standard support during the engagement, so none of these resizes meets an Extended Support surcharge. The engine versions and their support dates:

- **LIMS production and val (RDS):** Oracle Database 19c Enterprise Edition, bring your own license. Support under Copperwick's own Oracle contract; no RDS Extended Support charge applies to Oracle.
- **MES production and val (RDS):** SQL Server 2019 Standard Edition, license included. Microsoft extended support to January 8, 2030.
- **The six development databases (D-1):** RDS for PostgreSQL 16. Standard support to February 28, 2029.
- **Research and platform databases:** RDS for PostgreSQL 15 and 16, and Aurora PostgreSQL 16. Standard support to February 29, 2028 (15) and February 28, 2029 (16).

How it was priced, at list. (\$12.16 – \$6.08) an hour × 730 hours = \$4,438.40 a month, \$53,260.80 a year. Net of the 6% discount, **\$50,065 a year.**

R-5: Clinical data repository: 8 to 6 nodes

Owner: Ines Halvorsen. **Change path:** Frozen (clinical). **Proposed for:** earliest April 27.

Confidence: Medium. **Saves:** \$53,688 a year net of the private pricing discount (\$57,115 at list). **Status:** proposed.

The finding. Clinical data repository: 8 to 6 ra3.4xlarge nodes. **What we saw.** Cluster CPU under 30% at the 95th percentile; managed storage is separate from the nodes.

The change. Elastic resize to six nodes after topline. It waits on topline, which ends the clinical freeze, and then a change request. **Rollback:** Elastic resize back.

The clinical data repository is a Redshift cluster of RA3 nodes, whose managed storage is billed separately from the nodes, so removing two nodes does not touch the data. The resize is an elastic resize after topline; the clinical freeze rules it out until then, and Ines Halvorsen owns the question of whether to touch the repository at all.

How it was priced, at list. $2 \times \$3.260 \text{ an hour} \times 730 \text{ hours} = \$4,759.60 \text{ a month}, \$57,115.20 \text{ a year.}$ Net of the 6% discount, **\$53,688 a year.**

S-1: Visualization workstations: idle-stop and a weekday schedule

Owner: Mateo Varga. **Change path:** Non-GxP. **Proposed for:** week 3 (March 9 to 13). **Confidence:** Medium. **Saves:** \$117,574 a year net of the private pricing discount (\$125,078 at list). **Status:** proposed.

The finding. Visualization workstations: idle auto-stop and a weekday schedule on 20 of 60.

What we saw. The other 40 run overnight structure renders, so they stay on.

The change. Idle auto-stop and a weekday schedule on 20 workstations. **Rollback:** Remove the schedule tag.

Twenty workstations get an idle auto-stop and a weekday schedule; a scientist who needs one outside the schedule starts it from the console, and a workstation tagged for an overnight render is exempt.

How it was priced, at list. $20 \times \$1.212 \text{ an hour} \times (730 - 300) \text{ hours} = \$10,423.20 \text{ a month},$ \$125,078.40 a year. Net of the 6% discount, **\$117,574 a year.**

S-2: Development accounts: weekday hours

Owner: Priya Okonkwo. **Change path:** Non-GxP. **Proposed for:** week 3 (March 9 to 13). **Confidence:** High. **Saves:** \$130,724 a year net of the private pricing discount (\$139,068 at list). **Status:** proposed.

The finding. Development accounts: 12 hours on weekdays (260 of 730 hours). **What we saw.** No CPU activity outside working hours on weekdays, none on weekends.

The change. Instance Scheduler: 7 AM to 7 PM Eastern on weekdays. **Rollback:** Remove the schedule tag.

The schedule is run by AWS Instance Scheduler, and the tag on each instance is the whole interface: removing the tag removes the schedule. A host that must run overnight carries an exemption tag.

How it was priced, at list. $\$18,000 \text{ a month} \times (730 - 260) / 730 = \$11,589.04 \text{ a month}$, $\$139,068.49 \text{ a year}$. Net of the 6% discount, **\$130,724 a year**.

S-3a: GMP val environments: weekday schedule

Owner: Daniel Achterberg. **Change path:** GxP change request. **Proposed for:** week 4 (March 16 to 20). **Confidence:** High. **Saves:** \$87,540 a year net of the private pricing discount (\$93,128 at list). **Status:** proposed.

The finding. GMP val environments (LIMS, MES, CDS): weekday schedule except during executed protocols. **What we saw.** The val environments host no GxP records; QA's condition is that none stops during an executed protocol.

The change. The weekday schedule, with a hold tag the scheduler honors during an executed protocol. **Rollback:** Remove the schedule tag.

The condition on the schedule is QA's: no environment stops during an executed protocol. The scheduler enforces it with a hold tag that is set on an environment for the length of an executed protocol; the tag is the record. The GMP campaign freeze covers the production manufacturing systems; the val environments host no GxP records and sit outside it. Because they are validated systems, the schedule still goes to the Tuesday board as a change request, and Dr. Lorraine Osei confirms its class. The board meets on Tuesday, March 17, and the approved change runs in the next release window, Thursday, March 19; it does not go to production on the evening of approval.

How it was priced, at list. $\$13,175.04 \text{ a month for the GMP val app tiers} \times (730 - 300) / 730 = \$7,760.64 \text{ a month}$, $\$93,127.68 \text{ a year}$. Net of the 6% discount, **\$87,540 a year**.

S-3b: SCE val environment: weekday schedule

Owner: Ines Halvorsen. **Change path:** Frozen (clinical). **Proposed for:** earliest April 27. **Board vote:** Tuesday, April 7. **Confidence:** High. **Saves:** \$45,819 a year net of the private pricing discount (\$48,743 at list). **Status:** proposed.

The finding. SCE val environment: the same schedule. **What we saw.** The same evidence as S-3a; the SCE is in the clinical freeze.

The change. The same schedule after topline. It waits on topline, which ends the clinical freeze, and then a change request (board vote April 7). **Rollback:** Remove the schedule tag.

The same schedule and the same hold tag as S-3a, on the SCE val environment, which is inside the clinical freeze until topline.

How it was priced, at list. $4 \times \$2.3616 \text{ an hour} \times (730 - 300) \text{ hours} = \$4,061.95 \text{ a month}$, $\$48,743.42 \text{ a year}$. Net of the 6% discount, **\$45,819 a year**.

Batch, Spot, and file systems

Batch and file systems are \$1,072,543 a year net of the identified savings: four Batch and Spot findings, B-1, B-1b, B-2, and B-3 (\$620,679), and two file-system findings, F-1 and F-3 (\$451,864). Both are research workloads that were sized for the worst day and left there.

Spot, done honestly

Every genomics compute environment is On-Demand c6i.8xlarge, 102 on average. The alignment and variant-calling stages checkpoint and retry under Nextflow, so they tolerate interruption. Spot for c6i.8xlarge itself was only 33.6% below On-Demand in the price table’s 30-day window, so the new environments use the price-capacity-optimized strategy across c7i.8xlarge, whose 30-day mean is \$0.5815 an hour, 57.2% below c6i.8xlarge’s \$1.36. Both are 32-vCPU instances, so every genomics job fits on one. The estimate prices every Spot hour at the c7i.8xlarge mean plus a 12% allowance for retried work, so each Spot hour is worth \$0.7087 less than the On-Demand hour it replaces, and it gives no credit for c7i’s faster cores. Savings Plans do not apply to Spot, so the commitment (C-1) is sized after this change.

The instances considered for Spot and their prices, us-east-1. Spot prices are the price table’s 30-day means, from EC2 Spot price history ([Appendix C: price provenance](#)).

Instance	On-Demand an hour	Spot, 30-day mean	Below its On-Used for Demand price
c6i.8xlarge	\$1.360	\$0.9034	33.6%The current On-Demand genomics environments, and the second Spot pool
c7i.8xlarge	\$1.428	\$0.5815	59.3%The Spot environments (B-1, B-1b), priced at the mean plus retries
g5.12xlarge	\$5.672	\$4.5010	20.6%Considered for cryo-EM (B-3) and not used

The clinical biomarker pipelines are GCP-validated, and an instance-type change is a revalidation under their validation plan, so B-1b is the same move as B-1 on a later date. Spot was considered for the cryo-EM workers and rejected: g5.12xlarge Spot averaged only 20.6% off, not worth interruptions on multi-hour refinement jobs.

B-1: Genomics Batch, research pipelines: Spot

Owner: Priya Okonkwo. **Change path:** Non-GxP. **Proposed for:** week 4 (March 16 to 20).

Confidence: Medium. **Saves:** \$238,111 a year net of the private pricing discount (\$253,310 at list).

Status: proposed.

The finding. Genomics Batch, research pipelines: 40% of their work to Spot (c7i.8xlarge and c6i.8xlarge pools), with a 12% retry allowance. **What we saw.** Alignment and variant calling checkpoint and retry under Nextflow.

The change. Price-capacity-optimized Spot compute environments across c7i.8xlarge and c6i.8xlarge (32 vCPUs each, so every job fits one instance). **Rollback:** Point the job queue back at the On-Demand environment.

Moving 40% of the research genomics work, the equivalent of 40.8 of the 102 nodes, to new Spot compute environments leaves the On-Demand environment in place for the rest, so a Spot shortage degrades throughput and does not stop the pipeline. The job queue gets both environments, Spot first.

How it was priced, at list. $102 \times 0.40 \times 730 \text{ hours} \times (\$1.36 - \$0.58148 \times 1.12) = \$21,109.15$ a month, \$253,309.83 a year. Net of the 6% discount, **\$238,111 a year.**

B-1b: Biomarker pipelines: Spot

Owner: Ines Halvorsen. **Change path:** GxP change request. **Proposed for:** earliest May 11. **Board vote:** Tuesday, April 21. **Confidence:** Medium. **Saves:** \$42,020 a year net of the private pricing discount (\$44,702 at list). **Status:** proposed.

The finding. Clinical biomarker pipelines (GCP-validated), in their own Batch environment in clin-biomarker-prod: the same Spot move, 40% of their work. **What we saw.** The pipelines run in their own validated Batch environment, never on the research fleet; an instance-type change is a revalidation under their validation plan.

The change. Spot compute environments in clin-biomarker-prod, after the May revalidation. It waits on the biomarker pipelines' revalidation: an instance-type change is a revalidation under their validation plan (board vote April 21). **Rollback:** Point the job queue back.

The same move for 40% of the biomarker pipelines' work (7.2 of their 18 c6i.8xlarge nodes), in their own Batch environment in clin-biomarker-prod, never on the research fleet. It is a GxP change request because the instance type is part of the pipelines' validated configuration.

How it was priced, at list. $18 \times 0.40 \times 730 \text{ hours} \times (\$1.36 - \$0.58148 \times 1.12) = \$3,725.14$ a month, \$44,701.73 a year. Net of the 6% discount, **\$42,020 a year.**

B-2: Genomics Batch pinned to the Lustre zone

Owner: Priya Okonkwo. **Change path:** Non-GxP. **Proposed for:** week 4 (March 16 to 20).

Confidence: Medium. **Saves:** \$36,962 a year net of the private pricing discount (\$39,322 at list).

Status: proposed.

The finding. Genomics Batch pinned to the Lustre file system's Availability Zone (cross-AZ transfer down 80%). **What we saw.** 200 of the 300 TB a month of cross-AZ transfer is Batch reading Lustre.

The change. One subnet per compute environment, in the file system's zone. **Rollback:** Restore the subnets.

Pinning the genomics environments to the Lustre file system's Availability Zone would save cross-AZ transfer for 200 of the 300 TB a month, the Batch jobs reading Lustre. **The trade-off is real:** one zone means fewer Spot pools for B-1, so more interruptions. Priya Okonkwo owns the file system and the Spot environments, and the trade-off is theirs to weigh.

How it was priced, at list. $0.80 \times 200 \text{ TB} \times 1,024 \times \$0.01 \text{ a GB each way} \times 2 = \$3,276.80 \text{ a month}$, \$39,321.60 a year. Net of the 6% discount, **\$36,962 a year.**

B-3: Cryo-EM GPU workers: a Batch queue that scales to zero

Owner: Hana Lindgren. **Change path:** Non-GxP. **Proposed for:** earliest April 6. **Confidence:** Medium. **Saves:** \$303,586 a year net of the private pricing discount (\$322,964 at list). **Status:** proposed.

The finding. Cryo-EM GPU workers to an AWS Batch GPU queue that scales to zero (10 g5.12xlarge billed about 35% of hours). **What we saw.** 22% average GPU utilization; processing follows collection.

The change. A Batch GPU queue with a minimum of zero. It waits on the next cryo-EM collection campaign: moving the workers to a Batch queue is a workflow change for the cryo-EM team.

Rollback: Restore the standing workers.

Ten g5.12xlarge workers run around the clock at 22% average GPU utilization, because processing follows collection. An AWS Batch GPU queue with a minimum of zero scales the workers up when a job arrives and down when the queue empties, which bills about 35% of hours. The first job after an idle stretch waits for a worker to start, which is the cost the cryo-EM team trades.

How it was priced, at list. $10 \times \$5.672 \text{ an hour} \times 730 \text{ hours} \times 0.65 = \$26,913.64 \text{ a month}$, \$322,963.68 a year. Net of the 6% discount, **\$303,586 a year.**

Lustre provisioned for a peak that never came

The genomics file system is 64 TiB at 1,000 MB/s/TiB, 64 GB/s. CloudWatch shows 29 TiB used and a 90-day throughput peak of 14.8 GB/s. Lustre storage cannot shrink, so the change is a new persistent-2 file system of 48 TiB at 500 MB/s/TiB (24 GB/s, 62% headroom over the peak), with a data repository association on the same S3 prefix, cut over between pipeline batches. A further step to 250 MB/s/TiB would give 12 GB/s, below the measured peak, so the report does not recommend it.

The cryo-EM file system (60 TiB at 500 MB/s/TiB) stands idle between monthly collection campaigns. A scratch-2 file system (\$0.14 a GB-month) created from S3 for each 12-day campaign costs 16% of the persistent file system's monthly cost. That is a workflow change for the cryo-EM team, scheduled with B-3 for the next campaign.

F-1: Genomics Lustre: 64 TiB at 1,000 to 48 TiB at 500

Owner: Priya Okonkwo. **Change path:** Non-GxP. **Proposed for:** week 4 (March 16 to 20).

Confidence: High. **Saves:** \$255,040 a year net of the private pricing discount (\$271,319 at list).

Status: proposed.

The finding. Genomics Lustre: 64 TiB at 1,000 MB/s/TiB to 48 TiB at 500 MB/s/TiB (a new file system on the same S3 repository). **What we saw.** 29 TiB used; a 90-day throughput peak of 14.8 GB/s against 64 GB/s provisioned.

The change. A new persistent-2 file system of 48 TiB at 500 MB/s/TiB (24 GB/s), linked to the same S3 prefix, cut over between pipeline batches. **Rollback:** Create a larger file system from the same repository.

The new file system is created beside the old one, loaded from the S3 repository, and cut over between pipeline batches, so no job runs against a half-moved file system. The old file system is kept until the first full batch has run on the new one; the import requests are a one-time charge of \$40 ([One-time charges](#)).

How it was priced, at list. $64 \text{ TiB} \times 1,024 \times \$0.60 - 48 \text{ TiB} \times 1,024 \times \$0.34 \text{ a GB-month} = \$22,609.92 \text{ a month}, \$271,319.04 \text{ a year. Net of the 6\% discount, } \$255,040 \text{ a year.}$

F-3: Cryo-EM Lustre: scratch per campaign

Owner: Hana Lindgren. **Change path:** Non-GxP. **Proposed for:** earliest April 6. **Confidence:**

Medium. **Saves:** \$196,824 a year net of the private pricing discount (\$209,388 at list). **Status:** proposed.

The finding. Cryo-EM Lustre: a scratch file system per collection campaign instead of a standing persistent one. **What we saw.** The file system stands idle between monthly 12-day collection campaigns.

The change. A scratch-2 file system created from S3 for each campaign. It waits on the next cryo-EM collection campaign: a scratch file system per campaign is a workflow change for the cryo-EM team, scheduled with B-3. **Rollback:** Keep a persistent file system.

A scratch file system is not persistent, so results the team wants after a campaign must be exported back to S3, where the raw movies live, before the file system is deleted. The cryo-EM team owns that workflow change, and it is scheduled with B-3 so the campaign runs on both new patterns at once.

How it was priced, at list. $60 \text{ TiB} \times 1,024 \times \$0.34 - 60 \text{ TiB} \times 1,024 \times \$0.140 \times 12/30$ a GB-month = \$17,448.96 a month, \$209,387.52 a year. Net of the 6% discount, **\$196,824 a year.**

Storage and databases

Storage is \$211,379 a year net in four findings, E-1, E-2a, E-2b, and EF-1, and databases \$130,169 in two, D-1 and D-2.

Storage

Copperwick keeps 420 TB of gp2 in research and platform accounts and 110 TB in validated accounts, next to 260 TB already on gp3. gp3 costs \$0.08 a GB-month against gp2's \$0.10 and includes 3,000 IOPS and 125 MiB/s per volume. A volume that needs more is not stuck at the included amount: gp3 is provisioned up to 80,000 IOPS and 2,000 MiB/s per volume. A gp2 volume above 1 TiB has a baseline above 3,000 IOPS, so every volume is checked against its measured peak before it is converted, and any that needs more than the included amount is converted with the extra provisioned, which the estimate does not price. The conversion is online, and a volume can be modified once every six hours, so the rollback is a second modification.

BASH

```
# Per approved volume. Add --iops and --throughput where the pre-conversion check shows the
volume needs more than the included amount.
aws ec2 modify-volume --volume-id <volume-id> --volume-type gp3
aws ec2 describe-volumes-modifications --volume-ids <volume-id> --query
'VolumesModifications[0].ModificationState'
```

E-1: gp2 to gp3: research and platform

Owner: Priya Okonkwo. **Change path:** Non-GxP. **Proposed for:** week 3 (March 9 to 13).

Confidence: High. **Saves:** \$97,026 a year net of the private pricing discount (\$103,219 at list).

Status: proposed.

The finding. gp2 to gp3, research and platform accounts (420 TB). **What we saw.** Each volume's measured IOPS and throughput are checked against gp3's included 3,000 IOPS and 125 MiB/s before it is converted; the estimate prices the storage difference only.

The change. Elastic Volumes, online. **Rollback:** Modify back to gp2 after six hours.

Research and platform accounts, 420 TB. A non-GxP change through the platform team's normal process.

How it was priced, at list. $420 \text{ TB} \times 1,024 \times (\$0.10 - \$0.08) \text{ a GB-month} = \$8,601.60 \text{ a month}$, \$103,219.20 a year. Net of the 6% discount, **\$97,026 a year.**

E-2a: gp2 to gp3: GMP systems, before their freeze

Owner: Daniel Achterberg. **Change path:** GxP standard change. **Proposed for:** week 3 (March 9 to 13). **Confidence:** High. **Saves:** \$12,706 a year net of the private pricing discount (\$13,517 at list).

Status: proposed.

The finding. gp2 to gp3, the GMP systems before their freeze began (LIMS, MES, CDS, and historian production, and the val environments: 55 TB). **What we saw.** Like-for-like, reversible, with no effect on a GxP function: a standard change in the qualified infrastructure catalog, run on March 12, before the GMP freeze began on March 16; the same check as E-1: each volume's measured IOPS and throughput are checked against gp3's included 3,000 IOPS and 125 MiB/s before it is converted; the estimate prices the storage difference only.

The change. Elastic Volumes, online, with an executed installation checklist and notice to QA.

Rollback: Modify back.

In validated accounts this is a GxP standard change, the pre-approved type in the qualified infrastructure catalog: a ticket, an executed installation checklist, and notice to Dr. Lorraine Osei's team, with no new validation. It covers the GMP systems before their freeze begins: LIMS, MES, CDS, and historian production volumes and the val environments, run on Thursday, March 12, before the GMP campaign freeze begins on Monday, March 16.

How it was priced, at list. $55 \text{ TB} \times 1,024 \times (\$0.10 - \$0.08) \text{ a GB-month} = \$1,126.40 \text{ a month}$, \$13,516.80 a year. Net of the 6% discount, **\$12,706 a year.**

E-2b: gp2 to gp3: clinical systems, in the freeze

Owner: Ines Halvorsen. **Change path:** Frozen (clinical). **Proposed for:** earliest April 27.

Confidence: High. **Saves:** \$12,706 a year net of the private pricing discount (\$13,517 at list).

Status: proposed.

The finding. gp2 to gp3, the clinical systems in the freeze (SCE production and val, and the biomarker pipelines: 55 TB). **What we saw.** The same like-for-like conversion as E-2a, a GxP standard change with no board vote, on systems inside the clinical freeze, where nothing changes until topline; the same check as E-1: each volume's measured IOPS and throughput are checked against gp3's included 3,000 IOPS and 125 MiB/s before it is converted; the estimate prices the storage difference only.

The change. Elastic Volumes, online, as a standard change on April 27, after topline. It waits on topline, which ends the clinical freeze, and then a GxP standard change with notice to QA and no board vote. **Rollback:** Modify back.

The same standard change on the clinical systems in the freeze: the SCE production and val volumes and the biomarker pipelines' volumes, owned by Ines Halvorsen and run from `clin-sce-prod`. A freeze stops every change in its scope, so these volumes wait for topline. After it, E-2b is a GxP standard change (a ticket, an executed installation checklist, and notice to QA) on April 27, with no board vote. It is not a change request.

How it was priced, at list. $55 \text{ TB} \times 1,024 \times (\$0.10 - \$0.08) \text{ a GB-month} = \$1,126.40 \text{ a month}$, \$13,516.80 a year. Net of the 6% discount, **\$12,706 a year.**

EF-1: EFS lifecycle to Infrequent Access

Owner: Tomasz Reyes. **Change path:** Non-GxP. **Proposed for:** week 3 (March 9 to 13).

Confidence: Medium. **Saves:** \$88,941 a year net of the private pricing discount (\$94,618 at list).

Status: proposed.

The finding. EFS lifecycle: files untouched for 30 days to Infrequent Access (70% of 40 TB). **What we saw.** Access-time scans: 70% of files untouched for 30 days.

The change. A lifecycle policy to Infrequent Access after 30 days, back to Standard on first access. **Rollback:** Set the policy to none.

EFS Infrequent Access costs \$0.025 a GB-month against Standard's \$0.30, and charges \$0.01 a GB each time an Infrequent Access file is read. That read charge is not in the estimate, because how much of the cold data will be read again is not knowable in advance; the measurement at week six includes the actual charge. The policy moves a file back to Standard on its first access, so a project that wakes up is not billed the read charge for long.

BASH

```
aws efs put-lifecycle-configuration --file-system-id <file-system-id> \  
  --lifecycle-policies TransitionToIA=AFTER_30_DAYS  
TransitionToPrimaryStorageClass=AFTER_1_ACCESS
```

How it was priced, at list. $0.70 \times 40 \text{ TB} \times 1,024 \times (\$0.30 - \$0.025)$ a GB-month; the \$0.01 a GB read charge is not estimated = \$7,884.80 a month, \$94,617.60 a year. Net of the 6% discount, **\$88,941 a year.**

Databases

D-1: Multi-AZ off: development PostgreSQL

Owner: Priya Okonkwo. **Change path:** Non-GxP. **Proposed for:** week 3 (March 9 to 13).

Confidence: High. **Saves:** \$49,406 a year net of the private pricing discount (\$52,560 at list).

Status: proposed.

The finding. Multi-AZ off on six development PostgreSQL databases. **What we saw.**

Development databases need no standby.

The change. Modify to Single-AZ. **Rollback:** Modify back.

Development databases need no standby. The change is one modification per database, and a standby can be restored the same way.

BASH

```
aws rds modify-db-instance --db-instance-identifier <database> --no-multi-az --apply-  
immediately
```

How it was priced, at list. $6 \times (\$2.000 - \$1.000)$ an hour $\times 730$ hours = \$4,380.00 a month, \$52,560.00 a year. Net of the 6% discount, **\$49,406 a year.**

D-2: Multi-AZ off: LIMS and MES val

Owner: Daniel Achterberg. **Change path:** GxP change request. **Proposed for:** week 4 (March 16 to 20). **Confidence:** High. **Saves:** \$80,763 a year net of the private pricing discount (\$85,918 at list). **Status:** proposed.

The finding. Multi-AZ off on the LIMS and MES val databases. **What we saw.** The val databases mirror production.

The change. Modify to Single-AZ. **Rollback:** Modify back.

The val databases mirror production, and a val environment that differs from production is a question for Dr. Lorraine Osei's board, not a technical one: turning Multi-AZ off makes val a different configuration from the production system it validates against. It is a GxP change request and the saving is the standby, half of the val databases' cost.

How it was priced, at list. Half the LIMS and MES val database cost (the standby) = \$7,159.84 a month, \$85,918.08 a year. Net of the 6% discount, **\$80,763 a year**.

Network and the data lake

Network is \$141,149 a year net in one finding priced as two changes, N-1a and N-1b. The data lake is \$699,194 in five findings priced as six changes, S3-1 through S3-5, and it is the category where retention obligations decide what is allowed.

Network

Batch jobs read the data lake through NAT gateways at \$0.045 a GB, because the research VPCs have no S3 gateway endpoint: 36 NAT gateways process 300 TB a month. VPC flow logs show 260 TB of it is S3 in the same Region. A gateway endpoint is free and carries only same-Region traffic, so reads from the us-west-2 replica bucket still use NAT, and the estimate counts the same-Region traffic only.

BASH

```
# One endpoint per VPC, added to the route tables the private subnets use. Rollback is deleting the endpoint.
aws ec2 create-vpc-endpoint --vpc-id <vpc-id> --vpc-endpoint-type Gateway \
  --service-name com.amazonaws.us-east-1.s3 --route-table-ids <route-table-id> <route-table-id>
```

N-1a: S3 gateway endpoints: research VPCs

Owner: Grace Mbeki. **Change path:** Non-GxP. **Proposed for:** week 3 (March 9 to 13). **Confidence:** High. **Saves:** \$135,143 a year net of the private pricing discount (\$143,770 at list). **Status:** proposed.

The finding. S3 gateway endpoints in research VPCs (260 TB a month leaves their NAT gateways).

What we saw. VPC flow logs: 260 TB a month of the research VPCs' NAT traffic is S3 in the same Region (cross-Region traffic is left out; a gateway endpoint does not carry it). Each research VPC runs its own NAT gateways (36 in all); only the validated VPCs use the shared NAT in cw-network.

The change. A gateway endpoint in each research VPC's route tables. **Rollback:** Remove the route.

The research VPCs are not GxP systems, so this goes through the platform team's normal change process, from rc-genomics-prod. Each research VPC runs its own NAT gateways (36 in all), so the endpoint changes a route in those VPCs and touches neither the shared NAT in cw-network nor the Transit Gateway. The endpoint leaves the buckets, their policies, and every job's code alone.

How it was priced, at list. $260 \text{ TB} \times 1,024 \times \$0.045 \text{ a GB} = \$11,980.80 \text{ a month}, \$143,769.60 \text{ a year. Net of the 6\% discount, } \$135,143 \text{ a year.}$

N-1b: S3 gateway endpoints: validated VPCs

Owner: Grace Mbeki. **Change path:** GxP change request. **Proposed for:** earliest April 14. **Board vote:** Tuesday, March 31. **Confidence:** High. **Saves:** \$6,006 a year net of the private pricing discount (\$6,390 at list). **Status:** proposed.

The finding. S3 gateway endpoints in validated VPCs (8 TB a month). **What we saw.** The validated VPCs reach S3 through the shared NAT in cw-network, over the Transit Gateway, so both charges fall; a new network path for a qualified system.

The change. The same endpoint, by change request. It waits on GxP change control: an endpoint is a new network path for a qualified system, so it is a change request (board vote March 31).

Rollback: Remove the route.

In validated VPCs the endpoint changes the network path of a qualified system, so it is a change request however small the dollars are, voted by the board on Tuesday, March 31. The traffic is 8 TB a month. The validated VPCs are the ones that use the shared NAT in cw-network, so both the NAT and the Transit Gateway charges fall.

How it was priced, at list. $8 \text{ TB} \times 1,024 \times (\$0.045 \text{ NAT} + \$0.02 \text{ Transit Gateway}) \text{ a GB} = \$532.48 \text{ a month}, \$6,389.76 \text{ a year. Net of the 6\% discount, } \$6,006 \text{ a year.}$

The data lake: lifecycle with retention obligations

The in-scope accounts keep 6.5 PB in S3 Standard. The findings move data by what it is and what must be done with it, never by age alone:

Where each class of data goes and what keeps it safe. Savings are priced at the top Standard tier, \$0.021 a GB-month, less the price of the class the data goes to.

Findings	Data	PB	Goes to	Why	Retention obligation
S3-1, S3-1b	Raw sequencer run folders older than 18 months	0.70	Glacier Deep Archive, one tar per run	A 12-hour restore is acceptable for research raw data	Ten-year retention kept
S3-2	Cryo-EM raw movies older than 180 days	0.60	Glacier Flexible Retrieval	Reprocessing restores in bulk, which is free per GB	Ten-year retention kept
S3-3	GxP raw data older than two years (QC instrument data and GLP study data)	0.42	Glacier Instant Retrieval, Object Lock kept	Millisecond reads meet the record-retrieval SOP; a 12-hour restore does not	Object Lock and retention kept
S3-4	Analysis results	1.60	Intelligent-Tiering, automatic tiers only	Access logs show 55% of objects untouched for 30 days, and 20% of all objects, inside that 55%, untouched for 90: 35% bill at the Infrequent tier and 20% at Archive Instant	None; objects stay readable at once
S3-5	Research replicas in us-west-2	0.80	Deep Archive; regenerable intermediates are expired and stop replicating	It is a copy; the GxP replicas sit in the Standard-IA line and are not touched	GxP replicas unchanged

Three details keep these from being wrong. **Small objects:** Deep Archive adds 40 KB of overhead per object (8 KB billed at the Standard rate and 32 KB at the Deep Archive rate) and charges \$0.05 per 1,000 transitions, and raw run folders hold thousands of small files each, so a Batch job tars each run with a checksum manifest, verifies it, and only then expires the originals, with the data owner's sign-off. **Minimum durations:** Glacier Flexible Retrieval bills a 90-day minimum, which never applies to data held for 10 years. **Restore costs:** Flexible Retrieval restores at \$0.01 a GB on the standard tier and nothing on the bulk tier, and a Deep Archive restore takes about 12 hours. GxP records are different: Deep Archive is not an option for them, because a 12-hour restore does not meet the record-retrieval SOP.

S3-1: Raw run folders, tarred, to Deep Archive

Owner: Tomasz Reyes. **Change path:** Non-GxP. **Proposed for:** week 3 (March 9 to 13).

Confidence: High. **Saves:** \$94,671 a year net of the private pricing discount (\$100,714 at list).

Status: proposed.

The finding. Raw sequencer run folders older than 18 months, tarred per run, to Glacier Deep Archive (0.4 PB). **What we saw.** Ten-year retention, and a 12-hour restore is acceptable for research raw data.

The change. A Batch job tars each run with a checksum manifest, verifies it, and transitions it; after the owner's sign-off it deletes the originals by version ID, so no noncurrent copy stays billed in the versioned bucket. **Rollback:** Restore the tars (12 hours, standard retrieval).

BASH

```
# Per run folder, in the tar job. The originals expire only after the verification passes and
the data owner signs off.
tar -cf "$RUN.tar" "$RUN/" && sha256sum "$RUN.tar" > "$RUN.tar.sha256"
aws s3 cp "$RUN.tar" "s3://<lake-bucket>/archive/$RUN.tar" --storage-class DEEP_ARCHIVE --
checksum-algorithm SHA256
aws s3 cp "$RUN.tar.sha256" "s3://<lake-bucket>/archive/$RUN.tar.sha256"
```

The tar job's compute and requests are a one-time charge of \$141 ([One-time charges](#)). After the data owner signs off, the job deletes each original by its version ID, so no noncurrent copy of an original stays billed under W-5's 30-day rule.

How it was priced, at list. 0.4 PB × 1,048,576 GB × (\$0.021 – \$0.00099) a GB-month = \$8,392.80 a month, \$100,713.63 a year. Net of the 6% discount, **\$94,671 a year**.

S3-1b: Remaining raw run folders, second pass

Owner: Tomasz Reyes. **Change path:** Non-GxP. **Proposed for:** earliest April 20. **Confidence:** High. **Saves:** \$71,003 a year net of the private pricing discount (\$75,535 at list). **Status:** proposed.

The finding. The remaining raw run folders (0.3 PB) after the tar job's second pass. **What we saw.** The same as S3-1: ten-year retention, and a 12-hour restore is acceptable for research raw data.

The change. The tar job's second pass. It waits on the tar job's second pass, which is a question of the job's capacity. **Rollback:** Restore the tars (12 hours, standard retrieval).

The second half of the raw run folders, priced separately because the tar job's capacity, not a decision, sets when it can run: the first pass covers the folders in S3-1 and the second pass the rest.

How it was priced, at list. 0.3 PB × 1,048,576 GB × (\$0.021 – \$0.00099) a GB-month = \$6,294.60 a month, \$75,535.22 a year. Net of the 6% discount, **\$71,003 a year**.

S3-2: Cryo-EM raw movies to Glacier Flexible Retrieval

Owner: Hana Lindgren. **Change path:** Non-GxP. **Proposed for:** week 3 (March 9 to 13).

Confidence: Medium. **Saves:** \$123,484 a year net of the private pricing discount (\$131,366 at list).

Status: proposed.

The finding. Cryo-EM raw movies older than 180 days to Glacier Flexible Retrieval (0.6 PB). **What we saw.** Movies are one to four gigabytes each; reprocessing restores in bulk.

The change. A lifecycle transition at 180 days. **Rollback:** Bulk restore (free per GB) or standard (\$0.01 a GB).

Cryo-EM movies are one to four gigabytes each, so the per-object overhead is immaterial, and reprocessing restores a whole campaign in bulk. The lifecycle rule transitions objects at 180 days; the transition requests are a one-time charge of \$8.

JSON

```
{ "Rules": [ { "ID": "cw-S3-2-cryoem-raw", "Status": "Enabled", "Filter": { "Prefix": "cryoem/
raw/" },
                "Transitions": [ { "Days": 180, "StorageClass": "GLACIER" } ] } ] }
```

How it was priced, at list. 0.6 PB × 1,048,576 GB × (\$0.021 – \$0.0036) a GB-month = \$10,947.13 a month, \$131,365.60 a year. Net of the 6% discount, **\$123,484 a year.**

S3-3: GxP raw data to Glacier Instant Retrieval

Owner: Daniel Achterberg. **Change path:** GxP change request. **Proposed for:** earliest April 14.

Board vote: Tuesday, March 31. **Confidence:** Medium. **Saves:** \$84,451 a year net of the private pricing discount (\$89,842 at list). **Status:** proposed.

The finding. GxP raw data older than two years to Glacier Instant Retrieval, Object Lock kept (0.42 PB). **What we saw.** Millisecond reads meet the record-retrieval SOP; a 12-hour Deep Archive restore does not.

The change. A lifecycle transition by change request, with a documented retrieval test. It waits on GxP change control: a change request with a documented retrieval test before any object moves (board vote March 31). **Rollback:** Copy back to Standard.

QC instrument data and GLP study data keep Object Lock and versioning; a lifecycle transition changes the storage class and does not touch the retention. Glacier Instant Retrieval keeps millisecond reads for an inspector's request. It is a GxP change request with a documented retrieval test before any object moves.

JSON

```
{ "Rules": [ { "ID": "cw-S3-3-gxp-raw", "Status": "Enabled", "Filter": { "Prefix": "gxp/
raw/" },
    "Transitions": [ { "Days": 730, "StorageClass": "GLACIER_IR" } ] } ] }
```

How it was priced, at list. $0.42 \text{ PB} \times 1,048,576 \text{ GB} \times (\$0.021 - \$0.004) \text{ a GB-month} = \$7,486.83 \text{ a month}, \$89,841.99 \text{ a year. Net of the 6\% discount, } \mathbf{\$84,451 \text{ a year.}}$

S3-4: Analysis results to Intelligent-Tiering

Owner: Tomasz Reyes. **Change path:** Non-GxP. **Proposed for:** week 3 (March 9 to 13).

Confidence: Medium. **Saves:** \$119,686 a year net of the private pricing discount (\$127,326 at list).

Status: proposed.

The finding. Analysis results to S3 Intelligent-Tiering (1.6 PB). **What we saw.** Access logs show 55% of the objects untouched for 30 days; 20% of all objects, inside that 55%, untouched for 90. So 35% bill at the Infrequent Access tier and 20% at the Archive Instant Access tier.

The change. A lifecycle transition to Intelligent-Tiering (the automatic tiers only). **Rollback:** Transition back is not needed: Intelligent-Tiering returns objects to its frequent tier on access.

Intelligent-Tiering's monitoring fee on about 34 million objects is \$85 a month, and it is netted from the saving. **It cannot count inside the six weeks:** objects move to the infrequent tier only after 30 days without access, and to the archive instant tier after 90, so a transition made in week 3 produces nothing measurable before week 6. The change is identified and proposed, and its savings are reported separately, first measurable on May 11, 2026. That first window measures the Infrequent tier only: the first objects reach the Archive Instant tier on June 9, 2026.

JSON

```
{ "Rules": [ { "ID": "cw-S3-4-results", "Status": "Enabled", "Filter": { "Prefix":
"results/" },
    "Transitions": [ { "Days": 0, "StorageClass": "INTELLIGENT_TIERING" } ] } ] }
```

How it was priced, at list. $1.6 \text{ PB} \times 1,048,576 \text{ GB} \times (0.35 \times (\$0.021 - \$0.0125) + 0.20 \times (\$0.021 - \$0.004))$, less \$85 a month of monitoring on about 34 million objects = \$10,610.48 a month, \$127,325.70 a year. Net of the 6% discount, **\$119,686 a year.**

S3-5: Research replicas to Deep Archive

Owner: Tomasz Reyes. **Change path:** Non-GxP. **Proposed for:** earliest April 21. **Confidence:** High.

Saves: \$205,899 a year net of the private pricing discount (\$219,041 at list). **Status:** proposed.

The finding. Research replicas in us-west-2: raw data and results to Deep Archive, regenerable intermediates expired and no longer replicated (the GxP replicas, in the Standard-IA line, unchanged). **What we saw.** Regenerable intermediates (0.25 PB) can be recomputed from raw data. The 0.8 PB is the research replicas only; the GxP replicas sit in Standard-IA under Object Lock and are counted in that line.

The change. The replication rule's destination class becomes Deep Archive, existing replicas are moved by a lifecycle rule, and a lifecycle rule in the replica bucket expires the existing intermediates, which stop replicating. It waits on the disaster-recovery test: the change alters the research data's recovery time in the disaster-recovery plan. **Rollback:** Change the rules back (expired intermediates are recomputed from raw data).

The us-west-2 replicas are research data kept at Standard. The replication rule's destination storage class becomes Deep Archive, a lifecycle rule on the replica bucket moves what is already there, and regenerable intermediates (0.25 PB) are expired on the replica bucket and stop replicating. The GxP replicas are not touched: they sit in the Standard-IA line, under Object Lock, and this change leaves that line as it is. The change alters the recovery time of research data in the disaster-recovery plan, so it waits for the disaster-recovery test.

JSON

```
"Destination": { "Bucket": "arn:aws:s3:::<replica-bucket>", "StorageClass": "DEEP_ARCHIVE" }
```

How it was priced, at list. 0.8 PB at the us-west-2 Standard tiers, less 0.55 PB × 1,048,576 GB × \$0.00099, plus 35% of 90 TB × 1,024 × \$0.02 of replication transfer = \$18,253.45 a month, \$219,041.37 a year. Net of the 6% discount, **\$205,899 a year.**

Logging, governance, and Graviton

Logging and governance is \$223,807 a year net in four findings, L-1 through L-4, and Graviton \$65,875 in one, G-1.

Logging and governance

Research Batch job logs arrive at 34 TB a month. The biomarker pipelines log in `clin-biomarker-prod` and keep their SOP retention, so L-1 does not touch them. A log group's class is set when it is created, so L-1 points the job definitions at new log groups in the Infrequent Access class (\$0.25 a GB instead of \$0.50), which give up metric filters, subscription filters, and Live Tail, none of which job logs use; the old groups expire. Log groups hold 300 TB with no expiry at \$0.03 a GB-month. AWS Config records every short-lived Batch and Karpenter instance continuously at \$0.003 an item, and CloudTrail records 4.0 billion S3 data events a month at \$0.10 per 100,000.

Governance data is different from job data. GxP audit logs keep their standard-operating-procedure retention, and GxP buckets keep their CloudTrail data events, read and write, for the Part 11 audit trail. No finding in this category touches either.

L-1: Batch job logs to the Infrequent Access class

Owner: Priya Okonkwo. **Change path:** Non-GxP. **Proposed for:** week 3 (March 9 to 13).

Confidence: High. **Saves:** \$98,181 a year net of the private pricing discount (\$104,448 at list).

Status: proposed.

The finding. Research Batch job logs to new log groups in the Infrequent Access class, 30-day retention (the biomarker pipelines' logs are not touched). **What we saw.** The 34 TB a month is research Batch job logs, which use no metric filters, subscription filters, or Live Tail; the biomarker pipelines log in clin-biomarker-prod and keep their SOP retention.

The change. Job definitions point at new log groups in the Infrequent Access class. **Rollback:** Point the job definitions back.

BASH

```
aws logs create-log-group --log-group-name <new-job-log-group> --log-group-class
INFREQUENT_ACCESS
aws logs put-retention-policy --log-group-name <new-job-log-group> --retention-in-days 30
# Then point the job definition's awslogs-group at the new group; the old group gets a short
retention and expires.
```

How it was priced, at list. 34 TB × 1,024 × (\$0.50 – \$0.25) a GB = \$8,704.00 a month, \$104,448.00 a year. Net of the 6% discount, **\$98,181 a year.**

L-2: Log retention: 90 days for research

Owner: Grace Mbeki. **Change path:** GxP standard change. **Proposed for:** week 3 (March 9 to 13).

Confidence: High. **Saves:** \$58,216 a year net of the private pricing discount (\$61,932 at list).

Status: proposed.

The finding. Log retention: 90 days for a named list of research log groups (220 of 300 TB stored); GxP audit logs keep their SOP retention. **What we saw.** Research log groups had no retention set; after the change they keep 90 days (18 TB at 6 TB a month), and the new Batch job-log groups keep 30 days (34 TB). Run from the log archive's tooling in a qualified account, so a GxP standard change with notice to QA.

The change. A 90-day retention on a positive list of research log groups, named by prefix and account; a group not on the list is never touched. **Rollback:** None after expiry (the logs are gone).

Research log groups had no retention set. The change applies 90 days to research log groups only: 300 TB are stored in all, and the 220 TB in research groups is what the retention affects. GxP audit logs keep their standard-operating-procedure retention and are not in the set. The set is a

positive list: research log groups named by prefix and account, so a group that is not on the list is never touched, whether or not it carries a tag. The last 90 days of those groups stay stored; only the older logs are removed. L-2 runs from the qualified log-archive account, so it is a GxP standard change with notice to QA, and no date moves. Expired logs are gone, which is why the rollback is "none after expiry".

BASH

```
aws logs put-retention-policy --log-group-name <research-log-group> --retention-in-days 90
```

How it was priced, at list. (220 TB – 18 TB of the other research groups' last 90 days – 34 TB of Batch job logs kept 30 days by L-1) $\times 1,024 \times \$0.03$ a GB-month = \$5,160.96 a month, \$61,931.52 a year. Net of the 6% discount, **\$58,216 a year.**

L-3: AWS Config: daily recording in three accounts

Owner: Grace Mbeki. **Change path:** GxP standard change. **Proposed for:** week 3 (March 9 to 13).

Confidence: High. **Saves:** \$32,216 a year net of the private pricing discount (\$34,272 at list).

Status: proposed.

The finding. AWS Config: daily recording for EC2 instances, network interfaces, and volumes in the three high-churn research accounts (rc-genomics-prod, rc-genomics-dev, and rc-eks-prod). **What we saw.** Every short-lived Batch and Karpenter instance, with its network interface and volume, is a recorded item: the 1.6 million are those three types. The recorder settings are deployed from the security tooling account, which is qualified, so a GxP standard change with notice to QA.

The change. Daily recording for EC2 instances, network interfaces, and volumes in those three accounts (a recording-mode override). **Rollback:** Set the recording frequency back to continuous.

Config records an item every time a short-lived instance changes, which for Batch and Karpenter means constantly. The change applies to three accounts, rc-genomics-prod, rc-genomics-dev, and rc-eks-prod, and to three resource types in them: EC2 instances, network interfaces, and volumes. Daily recording costs \$0.012 per item per day and covers about 1,800 resources in each of the three accounts. Every other resource type, and every other account, keeps continuous recording. The recorder settings are deployed from the qualified security-tooling account, so L-3 is a GxP standard change with notice to QA.

How it was priced, at list. 1.6 million items $\times \$0.003$, less 3 accounts $\times 1,800$ resources $\times 30$ days $\times \$0.012$ = \$2,856.00 a month, \$34,272.00 a year. Net of the 6% discount, **\$32,216 a year.**

L-4: CloudTrail data events: writes only on research buckets

Owner: Grace Mbeki. **Change path:** GxP standard change. **Proposed for:** week 3 (March 9 to 13).

Confidence: High. **Saves:** \$35,194 a year net of the private pricing discount (\$37,440 at list).

Status: proposed.

The finding. CloudTrail S3 data events: write events only on research buckets; read and write kept on GxP buckets. **What we saw.** 78% of data events are reads on research buckets.

The change. Advanced event selectors for write events only on research buckets. **Rollback:** Restore the selectors.

CloudTrail S3 data events on research buckets record every read; logging writes only removes 78% of the events. GxP buckets keep reads and writes for the Part 11 audit trail. Rosalind Tran, Chief Information Security Officer, owns CloudTrail and detection and is consulted before this change is applied, because a read event can feed a detection.

BASH

```
aws cloudtrail put-event-selectors --trail-name <trail> --advanced-event-selectors file://research-write-only.json
```

How it was priced, at list. $0.78 \times 4.0 \text{ billion events} \div 100,000 \times \$0.10 = \$3,120.00$ a month, \$37,440.00 a year. Net of the 6% discount, **\$35,194 a year.**

Graviton

G-1: Molecular dynamics to Graviton

Owner: Hana Lindgren. **Change path:** Non-GxP. **Proposed for:** earliest April 20. **Confidence:** Medium. **Saves:** \$65,875 a year net of the private pricing discount (\$70,080 at list). **Status:** proposed.

The finding. Molecular dynamics to c7g.16xlarge (Graviton) after the benchmark. **What we saw.** The code has an arm64 build; throughput per vCPU on Copperwick's systems is unproven.

The change. Benchmark, then two c7g.16xlarge per c6i.32xlarge. It waits on the molecular dynamics benchmark: throughput per vCPU on Copperwick's systems is unproven. **Rollback:** Point the cluster back at c6i.

The molecular dynamics code has an arm64 build. A c7g.16xlarge (\$2.32 an hour, 64 vCPUs) costs 14.7% less per vCPU than a c6i.32xlarge (\$5.44 an hour, 128 vCPUs) at list, and two of them replace one. Throughput per vCPU on Copperwick's systems is unproven, because a Graviton vCPU is a full core and an Intel vCPU is a hyperthread, so the change waits for the benchmark and is rated Medium.

hpc7g.16xlarge is cheaper per hour than c7g and goes into the benchmark too. The HPC-optimized hpc7g.16xlarge is \$1.6832 an hour (AWS Price List bulk API), 27.4% below c7g.16xlarge's \$2.32, with the same 64 vCPUs. The benchmark covers both. G-1 is priced at c7g because it is the instance the estimate can stand on today; if hpc7g wins the benchmark, two of them replace one c6i.32xlarge at \$3.3664 an hour, which at list would save \$15,137 a month across the cluster instead of \$5,840. None of that difference is in the estimate or the roadmap. hpc6a.48xlarge, hpc7a.48xlarge, and hpc7a.96xlarge are not offered in us-east-1, so they are not options.

How it was priced, at list. $10 \times (\$5.44 - 2 \times \$2.32)$ an hour $\times 730$ hours = \$5,840.00 a month, \$70,080.00 a year. Net of the 6% discount, **\$65,875 a year.**

Deliverable 4: the commitment plan

Copperwick already holds a three-year No Upfront Compute Savings Plan of \$150 an hour, running from August 1, 2024 to July 31, 2027. At 49.3% off it covers \$295.77 an hour of On-Demand usage, saves \$106,413 a month at list, and is used 100% every month. Savings Plans apply first to the usage with the highest discount, so it sits on the non-GPU floor; the GPU fleet is all On-Demand. This section proposes one purchase on top of it, **C-1, a one-year No Upfront Compute Savings Plan at \$50.00 an hour**, worth **\$148,682 a year net** (\$158,172 at list), and raises a second, **C-2**, for a decision. Rafael Quintero, Vice President, Financial Planning and Analysis, owns the commitment decision, and Copperwick's finance team buys in cw-cicd, a member account inside the sharing group described below; Vantalect does not purchase commitments.

What cannot be undone. A Savings Plan of \$100 an hour or less can be returned within seven days of purchase, in the same calendar month; after that it cannot be canceled. Both plans proposed here are under that line, so either can be returned in its first week; each is also sized below the floor that survives every change, which is why the purchase comes last.

How the size was found

The plan is sized last, on the estate after every change in the roadmap, so it never covers hours another finding has already removed.

Who the plan applies to. A Savings Plan floats to any eligible usage in the accounts it is shared with. Since January 12, 2026, Copperwick uses restricted group sharing: sharing is restricted to the 48 in-scope accounts (a cost category), so the other 164 accounts in the organization receive none of the plan. The payer account cannot belong to a sharing group, so C-1 and C-2 would be bought in cw-cicd, a member account inside the group, as the existing plan was bought there in

2024. The floor below is therefore sized on usage in the 48 in-scope accounts, and C-1 is measured on covered usage in those 48. The floor of 24-hour non-GPU On-Demand usage after the changes proposed for weeks 3 and 4:

The floor, by workload, in dollars an hour at list On-Demand prices. The 9 rows sum to the total.

The 24/7 non-GPU On-Demand floor	Per hour
EKS platform On-Demand base (30% of its spend)	\$128.219
Clinical operations apps (80% always on)	\$92.055
Genomics research pipelines On-Demand minimum (60% On-Demand share after B-1, times a 60% minimum)	\$49.939
Clinical biomarker pipelines On-Demand minimum (a 60% minimum)	\$14.688
Validated production (Linux-equivalent compute)	\$41.472
Shared services (90% always on)	\$32.055
Molecular dynamics minimum (4 nodes)	\$21.760
Lab sandboxes minimum (40%)	\$18.630
Data platform services, after R-1	\$13.824
Floor after the changes proposed for weeks 3 and 4	\$412.642

The roadmap's later changes remove more of the floor, and the plan is sized on what remains:

From the floor to the plan. C-1 covers 68.2% of what is left uncovered after every change, so it stays fully used with 31.8% of that to spare.

Step	Per hour
Floor after the changes proposed for weeks 3 and 4	\$412.642
Less R-2, SCE grid: 8 to 6 nodes	-\$8.064
Less B-1b, Biomarker pipelines: Spot	-\$5.875
Less G-1, Molecular dynamics to Graviton	-\$3.200
Floor after every change on the roadmap	\$395.503
Less what the existing plan covers	-\$295.771
Left uncovered after every change	\$99.732
C-1 at \$50.00 an hour covers, at the one-year m6i rate (26.5% off)	\$68.056

C-1 pays for \$68.06 an hour of On-Demand usage: $(\$68.05615 - \$50.00) \times 730 \text{ hours} = \$13,180.99$ a month at list. It is a **Compute** plan, not an EC2 Instance plan, so moving molecular dynamics to Graviton (G-1) and the biomarker pipelines to Spot (B-1b) strand nothing.

Commitment sizes against the \$99.73 an hour left uncovered after every change. The bold row is C-1. A size above the line pays for hours that are not there.

Commitment an hour	Covers (On-Demand an hour)	Share of what is left uncovered	Saves a month, at list	Plan used after every change	Unused a month then
\$30.00	\$40.83	41%	\$7,909	100%	\$0
\$40.00	\$54.44	55%	\$10,545	100%	\$0
\$50.00	\$68.06	68%	\$13,181	100%	\$0
\$60.00	\$81.67	82%	\$15,817	100%	\$0
\$70.00	\$95.28	96%	\$18,453	100%	\$0
\$80.00	\$108.89	109%	\$14,404	92%	\$4,912

At \$50.00 an hour the plan stays fully used even if every change on the roadmap lands, with 32% to spare. The margin is deliberate: the floor can fall further than this roadmap assumes, through autoscaling or a project that ends, and a plan that stays fully used when it does is worth more than the extra dollars of a larger one.

C-1: Compute Savings Plan, one year

Owner: Rafael Quintero. **Change path:** Non-GxP. **Proposed for:** week 4 (March 16 to 20).

Confidence: High. **Saves:** \$148,682 a year net of the private pricing discount (\$158,172 at list).

Status: proposed.

The finding. Compute Savings Plan, one year, No Upfront, \$50 an hour, on the non-GPU floor.

What we saw. The floor left uncovered after every change on the roadmap is \$99.73 an hour; \$50 an hour covers 68% of it.

The change. A one-year No Upfront Compute Savings Plan. **Rollback:** Return it within seven days of purchase, in the same calendar month (a plan of \$100 an hour or less); after that it cannot be canceled.

Risk. Final once the seven days pass, and low only in the sense that it is sized below even the floor that remains after every change. Confidence is High: the saving does not depend on anything the other findings might fail to deliver. The way to buy it:

BASH

```
# 1. Find the offering for the type, term, and payment option. One year is 31536000 seconds.
aws savingsplans describe-savings-plans-offerings --region us-east-1 \
  --plan-types Compute --payment-options "No Upfront" --durations 31536000 --currencies USD \
  --query 'searchResults[0].{id:offeringId,type:planType,term:durationSeconds,payment:paymentOption}'

# 2. Purchase in cw-cicd, a member account of the sharing group, once the owner has approved it.
aws savingsplans create-savings-plan --region us-east-1 \
  --savings-plan-offering-id <offering-id> --commitment 50.00 --client-token cw-compute-sp-1yr \
  --tags '{"cw:approved-by": "<approver>", "cw:change-ticket": "<ticket-id>"}'

# 3. Check Cost Explorer's own recommendation before buying, and the plan's utilization after.
aws ce get-savings-plans-purchase-recommendation --region us-east-1 \
  --savings-plans-type COMPUTE_SP --term-in-years ONE_YEAR --payment-option NO_UPFRONT --
  lookback-period-in-days SIXTY_DAYS
```

At week six the plan is judged by its ARN: its fee, the usage it covered, and its utilization.

How it was priced, at list. \$50 an hour covers \$68.05615 an hour of On-Demand usage at the one-year m6i rate (26.5% off): $(\$68.05615 - \$50.00) \times 730 \text{ hours} = \$13,180.99$ a month, \$158,171.84 a year. Net of the 6% discount, **\$148,682 a year.**

A second plan, for a decision

C-2: EC2 Instance Savings Plan, p4d, one year

Owner: Rafael Quintero. **Change path:** Non-GxP. **Proposed for:** week 4 (March 16 to 20).

Confidence: Medium. **Saves:** \$264,705 a year net of the private pricing discount (\$281,601 at list).

Status: proposed.

The finding. EC2 Instance Savings Plan for the four remaining p4d.24xlarge, one year, No Upfront.

What we saw. The four p4d run around the clock.

The change. A one-year EC2 Instance Savings Plan for the p4d family. **Rollback:** Return it within seven days of purchase, in the same calendar month (a plan of \$100 an hour or less); after that it cannot be canceled.

After W-3 releases one instance, 4 p4d.24xlarge remain, at \$21.958 an hour each, \$87.83 an hour On-Demand. A one-year EC2 Instance Savings Plan for the p4d family would cost \$55.68 an hour, 36.6% off. **The risk is the lock:** an EC2 Instance plan is bound to the instance family in the Region, so if the ML roadmap moves training to a newer accelerator during the year, the plan strands. That is a question about the roadmap, not about the arithmetic, and it belongs to Mateo Varga and Rafael Quintero. The GPU fleet stays On-Demand unless they decide otherwise, and the report treats that as a decision, not an oversight.

How it was priced, at list. $(4 \times \$21.957642 - 4 \times \$13.9211) \text{ an hour} \times 730 \text{ hours} = \$23,466.72$ a month, \$281,600.66 a year. Net of the 6% discount, **\$264,705 a year**.

What waits

Three-year plans wait for the 2027 private pricing renewal, when the commitment and the plan can be sized together; a three-year Compute plan would save 49.3% against 26.5% for the same coverage. Database Savings Plans do not apply today: A Database Savings Plan covers Generation 7 and newer database instances only. Every in-scope database runs db.r5 or db.r6i, and R-3 and R-4 keep those generations, so none is eligible; a plan becomes worth sizing only if a database moves to a Generation 7 class. SageMaker Savings Plans are not worth it after W-2: the notebooks would bill about a fifth of the hours.

Deliverable 6: the savings roadmap

The roadmap is every change in this report, ranked by annual savings, with its owner, its change path, its confidence, and the week it could go in. **41 changes are worth \$4,310,044 a year net** (\$4,585,154 at list). The ten largest are \$2,110,950 of it (49%), and the five largest \$1,267,341. Every change has a named owner, and the owner decides. Every status in this report is "proposed".

The waves

The waves are what the Optimization's weeks 3 and 4 would carry out, clear of the freezes, and what Copperwick's own teams do afterward.

The waves and what each carries. Week 5 (the week of March 23) carries guardrails and no new savings, and week 6 is measurement and handover.

Wave	Dates	What lands	Changes	Per year, net
Week 3	March 9 to 13, before the database lock on March 13	Waste, idle GPUs, schedules, gp3 on the GMP systems before their freeze, S3 gateway endpoints, log classes and retention, the data lake, Config, and the CloudTrail selectors	20	\$1,951,816
Week 4	March 16 to 20, inside the GMP campaign freeze of March 16 to 27	Spot for genomics, the Lustre cutover, the val schedule and val database change through the Tuesday board, and the Savings Plan last	7	\$1,111,803
Later	After April 2	Changes that wait on a freeze, a change request, a campaign, a release, a benchmark, or a test	14	\$1,246,425
Total			41	\$4,310,044

Changes that wait for a date

\$1,246,425 a year net, 14 changes, cannot land inside the six weeks. Each waits on something outside the Optimization's control, and the dates are the earliest the calendar allows:

The changes that cannot land inside the six weeks, by earliest date. The dates follow from the freeze calendar, the change advisory board, and the owners' own calendars.

Earliest	ID	Change	Waits on	Per year, net
April 6	B-3	Cryo-EM GPU workers: a Batch queue that scales to zero	The next cryo-EM collection campaign: moving the workers to a Batch queue is a workflow change for the cryo-EM team	\$303,586
April 6	F-3	Cryo-EM Lustre: scratch per campaign	The next cryo-EM collection campaign: a scratch file system per campaign is a workflow change for the cryo-EM team, scheduled with B-3	\$196,824
April 14	S3-3	GxP raw data to Glacier Instant Retrieval	GxP change control: a change request with a documented retrieval test before any object moves (board vote March 31)	\$84,451
April 14	N-1b	S3 gateway endpoints: validated VPCs	GxP change control: an endpoint is a new network path for a qualified system, so it is a change request (board vote March 31)	\$6,006

Earliest	ID	Change	Waits on	Per year, net
April 20	S3-1b	Remaining raw run folders, second pass	The tar job's second pass, which is a question of the job's capacity	\$71,003
April 20	G-1	Molecular dynamics to Graviton	The molecular dynamics benchmark: throughput per vCPU on Copperwick's systems is unproven	\$65,875
April 21	S3-5	Research replicas to Deep Archive	The disaster-recovery test: the change alters the research data's recovery time in the disaster-recovery plan	\$205,899
April 27	R-5	Clinical data repository: 8 to 6 nodes	Topline, which ends the clinical freeze, and then a change request	\$53,688
April 27	S-3b	SCE val environment: weekday schedule	Topline, which ends the clinical freeze, and then a change request (board vote April 7)	\$45,819
April 27	E-2b	gp2 to gp3: clinical systems, in the freeze	Topline, which ends the clinical freeze, and then a GxP standard change with notice to QA and no board vote	\$12,706
May 11	B-1b	Biomarker pipelines: Spot	The biomarker pipelines' revalidation: an instance-type change is a revalidation under their validation plan (board vote April 21)	\$42,020
May 18	R-3	LIMS database: db.r5.8xlarge to db.r6i.4xlarge	The LIMS release, with an operational qualification regression and a month-end performance test (board vote April 28)	\$30,698
June 15	R-4	MES database: db.r5.4xlarge to db.r5.2xlarge	The MES maintenance window, after the GMP campaign and the annual periodic review (board vote May 26)	\$50,065
June 22	R-2	SCE grid: 8 to 6 nodes	The post-topline analyses, which peak right after topline: the scripted operational qualification regression runs first in the SCE val environment, then the change request (board vote June 2)	\$77,785
Total				\$1,246,425

The 41 changes, ranked

Every change, ranked by annual savings, net of the private pricing discount. The same rows at list are in the CSV.

Rank	ID	Change	Owner	Change path	Proposed	Confidence	Per year, net	Running share
1	B-3	Cryo-EM GPU workers: a Batch queue that scales to zero	Hana Lindgren	Non-GxP	earliest April 6	Medium	\$303,586	7%
2	C-2	EC2 Instance Savings Plan, p4d, one year	Rafael Quintero	Non-GxP	week 4 (March 16 to 20)	Medium	\$264,705	13%
3	F-1	Genomics Lustre: 64 TiB at 1,000 to 48 TiB at 500	Priya Okonkwo	Non-GxP	week 4 (March 16 to 20)	High	\$255,040	19%
4	B-1	Genomics Batch, research pipelines: Spot	Priya Okonkwo	Non-GxP	week 4 (March 16 to 20)	Medium	\$238,111	25%
5	S3-5	Research replicas to Deep Archive	Tomasz Reyes	Non-GxP	earliest April 21	High	\$205,899	29%
6	F-3	Cryo-EM Lustre: scratch per campaign	Hana Lindgren	Non-GxP	earliest April 6	Medium	\$196,824	34%
7	W-2	SageMaker GPU notebooks: auto-stop	Mateo Varga	Non-GxP	week 3 (March 9 to 13)	High	\$182,152	38%
8	W-3	One idle p4d.24xlarge released	Mateo Varga	Non-GxP	week 3 (March 9 to 13)	High	\$180,808	42%
9	C-1	Compute Savings Plan, one year	Rafael Quintero	Non-GxP	week 4 (March 16 to 20)	High	\$148,682	46%
10	N-1a	S3 gateway endpoints: research VPCs	Grace Mbeki	Non-GxP	week 3 (March 9 to 13)	High	\$135,143	49%
11	S-2	Development accounts: weekday hours	Priya Okonkwo	Non-GxP	week 3 (March 9 to 13)	High	\$130,724	52%
12	S3-2	Cryo-EM raw movies to Glacier Flexible Retrieval	Hana Lindgren	Non-GxP	week 3 (March 9 to 13)	Medium	\$123,484	55%

Rank	ID	Change	Owner	Change path	Proposed	Confidence	Per year, net	Running share
13	S3-4	Analysis results to Intelligent-Tiering	Tomasz Reyes	Non-GxP	week 3 (March 9 to 13)	Medium	\$119,686	58%
14	S-1	Visualization workstations: idle-stop and a weekday schedule	Mateo Varga	Non-GxP	week 3 (March 9 to 13)	Medium	\$117,574	60%
15	R-1	Data platform services: m6i.4xlarge to m6i.2xlarge	Tomasz Reyes	Non-GxP	week 3 (March 9 to 13)	High	\$113,832	63%
16	W-5	Noncurrent versions and incomplete uploads	Tomasz Reyes	Non-GxP	week 3 (March 9 to 13)	High	\$99,355	65%
17	L-1	Batch job logs to the Infrequent Access class	Priya Okonkwo	Non-GxP	week 3 (March 9 to 13)	High	\$98,181	68%
18	E-1	gp2 to gp3: research and platform	Priya Okonkwo	Non-GxP	week 3 (March 9 to 13)	High	\$97,026	70%
19	W-1	Unattached gp2 volumes	Priya Okonkwo	Non-GxP	week 3 (March 9 to 13)	High	\$95,871	72%
20	S3-1	Raw run folders, tarred, to Deep Archive	Tomasz Reyes	Non-GxP	week 3 (March 9 to 13)	High	\$94,671	74%
21	EF-1	EFS lifecycle to Infrequent Access	Tomasz Reyes	Non-GxP	week 3 (March 9 to 13)	Medium	\$88,941	76%
22	S-3a	GMP val environments: weekday schedule	Daniel Achterberg	GxP change request	week 4 (March 16 to 20)	High	\$87,540	78%
23	W-4	Orphaned EBS snapshots	Priya Okonkwo	Non-GxP	week 3 (March 9 to 13)	Medium	\$86,630	80%
24	S3-3	GxP raw data to Glacier Instant Retrieval	Daniel Achterberg	GxP change request	earliest April 14	Medium	\$84,451	82%
25	D-2	Multi-AZ off: LIMS and MES val	Daniel Achterberg	GxP change request	week 4 (March 16 to 20)	High	\$80,763	84%

Rank	ID	Change	Owner	Change path	Proposed	Confidence	Per year, net	Running share
26	R-2	SCE grid: 8 to 6 nodes	Ines Halvorsen	Frozen (clinical)	earliest June 22	Medium	\$77,785	86%
27	S3-1b	Remaining raw run folders, second pass	Tomasz Reyes	Non-GxP	earliest April 20	High	\$71,003	88%
28	G-1	Molecular dynamics to Graviton	Hana Lindgren	Non-GxP	earliest April 20	Medium	\$65,875	89%
29	L-2	Log retention: 90 days for research	Grace Mbeki	GxP standard change	week 3 (March 9 to 13)	High	\$58,216	91%
30	R-5	Clinical data repository: 8 to 6 nodes	Ines Halvorsen	Frozen (clinical)	earliest April 27	Medium	\$53,688	92%
31	R-4	MES database: db.r5.4xlarge to db.r5.2xlarge	Daniel Achterberg	GxP change request	earliest June 15	Medium	\$50,065	93%
32	D-1	Multi-AZ off: development PostgreSQL	Priya Okonkwo	Non-GxP	week 3 (March 9 to 13)	High	\$49,406	94%
33	S-3b	SCE val environment: weekday schedule	Ines Halvorsen	Frozen (clinical)	earliest April 27	High	\$45,819	95%
34	B-1b	Biomarker pipelines: Spot	Ines Halvorsen	GxP change request	earliest May 11	Medium	\$42,020	96%
35	B-2	Genomics Batch pinned to the Lustre zone	Priya Okonkwo	Non-GxP	week 4 (March 16 to 20)	Medium	\$36,962	97%
36	L-4	CloudTrail data events: writes only on research buckets	Grace Mbeki	GxP standard change	week 3 (March 9 to 13)	High	\$35,194	98%
37	L-3	AWS Config: daily recording in three accounts	Grace Mbeki	GxP standard change	week 3 (March 9 to 13)	High	\$32,216	99%
38	R-3	LIMS database: db.r5.8xlarge to db.r6i.4xlarge	Daniel Achterberg	GxP change request	earliest May 18	Medium	\$30,698	99%
39	E-2a	gp2 to gp3: GMP systems, before their freeze	Daniel Achterberg	GxP standard change	week 3 (March 9 to 13)	High	\$12,706	100%
40	E-2b	gp2 to gp3: clinical systems, in the freeze	Ines Halvorsen	Frozen (clinical)	earliest April 27	High	\$12,706	100%

Rank	ID	Change	Owner	Change path	Proposed	Confidence	Per year, net	Running share
41	N-1b	S3 gateway endpoints: validated VPCs	Grace Mbeki	GxP change request	earliest April 14	High	\$6,006	100%
Total							\$4,310,044	100%

Owners, and what it costs to get there

Every change has a named owner. Priya Okonkwo (Director, Research Computing) owns 9 changes worth \$1,087,951 a year; Tomasz Reyes (Director, Scientific Data Platform) owns 7 changes worth \$793,387 a year; Hana Lindgren (Associate Director, Structural Biology Computing) owns 4 changes worth \$689,769 a year; Mateo Varga (Director, Computational Design) owns 3 changes worth \$480,534 a year; Rafael Quintero (Vice President, Financial Planning and Analysis) owns 2 changes worth \$413,387 a year; Daniel Achterberg (Director, Manufacturing and Quality Systems) owns 6 changes worth \$346,223 a year; Grace Mbeki (Senior Director, Cloud Platform) owns 5 changes worth \$266,775 a year; Ines Halvorsen (Director, Clinical Data Science Systems) owns 5 changes worth \$232,018 a year. Four changes carry a one-time charge, listed below; every other change costs an engineer’s time and nothing else.

One-time charges

One-time charges, never annualized and never netted into the annual savings. The private pricing discount applies to them too. S3-4’s charge is the one with no savings measurable inside the six weeks.

ID	What	At list
S3-1	the tar job’s compute and requests	\$141.00
S3-2	lifecycle transition requests	\$8.00
F-1	the repository import requests	\$40.00
S3-4	lifecycle transition requests to Intelligent-Tiering	\$340.00
Total		\$529.00

The order that matters

Waste goes first, because it needs no design and shrinks what everything after is sized on. Rightsizing and schedules come before Spot and the Savings Plan, so the production resizes land first and the plan is bought last, in week 4. The freezes are walls: nothing in the clinical systems changes before topline on April 24, and the production manufacturing systems are left alone from March 16 to 27. A change needs seven full days of data after it to be measured, and the report finalizes about 48 hours late, so the last day a change can land and be measured is March 22; every change meant to count is therefore proposed for weeks 3 and 4. The engagement closes on April 2, 2026.

The guarantee

Where this stands against the guarantee

The Optimization's second invoice, \$17,500, is issued only if the savings in place at the end of week six, measured change by change on Copperwick's Cost and Usage Report, annualized, and net of the private pricing discount, reach the gate of **\$105,000 a year**, 0.6% of the in-scope bill. Copperwick pays \$35,000 in total for the six weeks, half at signature and half only past the gate. This Analysis, weeks one and two of the Optimization, is inside that fee; bought alone, it would cost \$25,000.

The roadmap identifies \$4,310,044 a year net. The 19 High-confidence changes proposed for weeks 3 and 4 are \$1,987,526 a year, a margin of \$1,882,526 over the gate. That figure is an estimate and not a measurement. What counts at week six is what is in place and visible on the bill: deferred and declined changes, and changes in place whose savings cannot yet be measured, count for nothing, and a change approved at the readout and later declined counts at its estimate. The cost avoidance on Copperwick's planned growth (Cost avoidance on planned growth) is not measured, not counted toward the gate, and not a saving on the bill, and it is in no figure here.

What is not counted

The savings in this report are the changes a team makes to the bill. Five things stay out of the gate arithmetic, and the report names each one.

The Oracle processor licenses R-3 frees. Moving the LIMS database from db.r5.8xlarge to db.r6i.4xlarge halves the vCPUs, and Oracle licenses are counted by processor, so the change halves the licenses the instance needs. The licenses are not on the AWS bill, so they are not in R-3's saving or in any total here. They are real money to Copperwick, and a licensing review should price them.

The fall in AWS Enterprise Support that follows any saving. Support is billed to the whole organization as a percentage of all usage, in tiers, and it is not part of the in-scope bill. Copperwick's usage is far above the top tier, so each dollar saved lowers the fee by 3% of a dollar.

The Enterprise Support tiers, from the AWS Price List bulk API ([Appendix C: price provenance](#)).

Monthly AWS usage	Support rate
\$0 to \$150,000	10%
\$150,000 to \$500,000	7%
\$500,000 to \$1,000,000	5%
Above \$1,000,000	3%

At the organization's usage of about \$3,099,688 a month at list, support is about \$127,491 a month. Carrying out the whole roadmap would lower it by about **\$129,301 a year**. That is the AWS fee following the bill down, not a change anyone makes, so it is counted nowhere, including in the gate.

S3–4 until its tiers move. The \$119,686 a year net in the identified total is real, and Intelligent-Tiering produces none of it for 30 days after the transition, so it cannot be measured before May 11, 2026. It is reported separately and is \$0 toward the gate.

Growth during the engagement. Copperwick's bill has risen all year, most recently with a third weekly sequencing run in January. Savings are measured change by change against the resources they touch, and the changes that follow lab throughput (B-1 and N-1a) are measured per genome processed, so a busy lab week does not hide a saving or a slow one invent it.

Cost avoidance on planned growth

Copperwick's 2026 plan adds about \$260,000 a month to the in-scope bill by March 2027: a long-read sequencing core in June, a second cryo-EM microscope in August, Feldspar's Phase 2 biomarker work, and ML growth. Week five builds the patterns that found this engagement's savings into the defaults for that growth, as nine guardrails with one owner each:

The nine guardrails the Optimization builds in week five, each with one owner.

ID	Guardrail	Owner
GR-1	The VPC template with S3 gateway endpoints	Grace Mbeki
GR-2	An SCP that refuses new gp2 volumes	Grace Mbeki
GR-3	Spot-by-default Batch compute environments; On-Demand needs an exception tag	Priya Okonkwo
GR-4	Log class and retention in the account baseline, with SOP retention on every GxP-tagged log group in any account	Grace Mbeki
GR-5	The data-class lifecycle standard, co-signed by Dr. Lorraine Osei for its GxP class	Tomasz Reyes
GR-6	A throughput and capacity review before any new Lustre file system	Priya Okonkwo
GR-7	Idle-GPU alarms on every GPU instance	Mateo Varga
GR-8	Notebook auto-stop and workstation idle-stop as account defaults	Mateo Varga
GR-9	Budgets and anomaly monitors per account family, routed to each owner	Rafael Quintero

Together they are expected to hold back 8% to 13% of the planned growth: \$20,800 to \$33,800 a month, about **\$250,000 to \$410,000 a year** at the March 2027 run rate. That range is not measured, not counted toward the gate, and not a saving on the bill: there is no "before" on Copperwick's Cost and Usage Report for a workload that does not exist yet. GxP data is never auto-archived by a default; its lifecycle rule is the change-controlled one.

Appendix A: the registers

Four CSV files carry the full detail behind this report. They sit under docs/samples/copperwick/data/cost/ and are the registers as they stand at the end of the engagement: the Analysis columns are this report's, and the status columns are maintained by the later reports.

waste-register.csv (5 rows, Deliverable 2), rightsizing-plan.csv (5 rows, the rightsizing findings of Deliverable 3), gxp-change-register.csv (15 rows, [How each change reaches production](#)), and savings-roadmap.csv (41 rows, every change, Deliverable 6). The columns are the finding's id, category, title, account, owner, change path, savings at list and net, one-time cost, and confidence, then the status columns the later reports maintain.

Appendix B: assumptions

- **Estimates at list prices, restated net.** Every saving is priced at AWS list prices at the February run rate of \$1,623,688 a month, then restated net of the 6% private pricing discount (\$1,526,267 a month). Week 6 replaces each estimate with a measurement.
- **Support follows the bill** by about \$129,301 a year at full roadmap, and is not counted ([What is not counted](#)).
- **The honest limits** are in the findings: incremental snapshots (W-4), the volumes that need more than gp3's included IOPS (E-1), the EFS read charge (EF-1), a Lustre file system that cannot shrink (F-1), the Spot retry allowance (B-1), the Intelligent-Tiering delay (S3-4), the benchmark that gates Graviton (G-1), and the Oracle licenses that are not on the bill (R-3).
- **Change paths** follow Copperwick's four rules, set at kickoff ([How each change reaches production](#)); Dr. Lorraine Osei confirms each class.
- **The Savings Plans (C-1, C-2)** are proposals until finance approves them, and **the cost avoidance** ([Cost avoidance on planned growth](#)) is not measured, not counted toward the gate, and not a saving on the bill.

Appendix C: price provenance

Every price in the model and in this report is a us-east-1 AWS list price from the AWS Price List API, retrieved October 13, 2026, except the prices that table lacks, which come from the public AWS Price List bulk API retrieved October 17, 2026: EFS Infrequent Access reads, Glacier Flexible Retrieval retrievals, hpc7g.16xlarge, and the Enterprise Support tiers. Spot prices are 30-day means for September 14 to October 14, 2026. Savings Plans effective rates come from the AWS Savings Plans API, retrieved October 13, 2026. For this sample, every price was retrieved in October 2026, after the engagement's dates; a live report carries the prices of its own analysis week. hpc6a.48xlarge, hpc7a.48xlarge, and hpc7a.96xlarge are not offered in us-east-1.

The prices behind the model, in the order the arithmetic uses them.

Price	Value	Unit	Source
EC2 c6i.8xlarge / c6i.32xlarge / m6i.4xlarge / m6i.2xlarge	\$1.360 / \$5.440 / \$0.768 / \$0.384	an hour	AWS Price List API, 2026–10–13
EC2 g5.12xlarge / g5.2xlarge / p4d.24xlarge	\$5.672 / \$1.212 / \$21.958	an hour	AWS Price List API, 2026–10–13
EC2 Graviton c7g.16xlarge; hpc7g.16xlarge	\$2.3200 / \$1.6832	an hour	AWS Price List API, 2026–10–13; AWS Price List bulk API, 2026–10–17
Spot, 30-day mean: c7i.8xlarge / c6i.8xlarge	\$0.5815 / \$0.9034	an hour	EC2 Spot price history, September 14 to October 14, 2026
SageMaker ml.g5.12xlarge notebook	\$7.09	an hour	AWS Price List API, 2026–10–13
EBS gp2 / gp3 storage; snapshot	\$0.10 / \$0.08 / \$0.05	a GB-month	AWS Price List API, 2026–10–13
FSx for Lustre persistent–2 at 1,000 / 500 MB/s/TiB; scratch–2	\$0.60 / \$0.34 / \$0.14	a GB-month	AWS Price List API, 2026–10–13
EFS Standard / Infrequent Access; Infrequent Access read	\$0.300 / \$0.025 / \$0.010	a GB-month / a GB-month / a GB	AWS Price List API, 2026–10–13; AWS Price List bulk API, 2026–10–17
S3 Standard tiers (organization)	\$0.023 / \$0.022 / \$0.021	a GB-month	AWS Price List API, 2026–10–13
S3 Intelligent–Tiering infrequent / archive instant; Glacier Instant / Flexible / Deep Archive	\$0.01250 / \$0.00400 / \$0.00400 / \$0.00360 / \$0.00099	a GB-month	AWS Price List API, 2026–10–13
Glacier Flexible Retrieval standard / bulk retrieval	\$0.01 / \$0.00	a GB	AWS Price List bulk API, 2026–10–17
NAT gateway data; cross–AZ each way; cross–Region	\$0.045 / \$0.010 / \$0.020	a GB	AWS Price List API, 2026–10–13
CloudWatch Logs ingest standard / Infrequent Access; storage	\$0.50 / \$0.25 / \$0.03	a GB / a GB / a GB-month	AWS Price List API, 2026–10–13

Price	Value	Unit	Source
AWS Config item, continuous / periodic; CloudTrail data events	\$0.003 / \$0.012 / \$0.100	an item / an item / 100,000 events	AWS Price List API, 2026-10-13
Redshift ra3.4xlarge	\$3.26	an hour	AWS Price List API, 2026-10-13

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

AWS, Amazon Web Services, AWS Graviton, and the AWS service names are trademarks of Amazon.com, Inc. or its affiliates. Oracle and Java are registered trademarks of Oracle Corporation and its affiliates. Microsoft, Windows, and SQL Server are trademarks of the Microsoft group of companies. Red Hat and Red Hat Enterprise Linux are trademarks of Red Hat, Inc. PostgreSQL is a registered trademark of the PostgreSQL Community Association of Canada. Kubernetes is a registered trademark of The Linux Foundation. Nextflow is a trademark of Sequera Labs. Vantalect is not affiliated with or endorsed by any of them. Copperwick Therapeutics, Inc. is fictional.