

AWS COST · SAMPLE DELIVERABLE

Quick Wins Applied

Week 3 of 6: Quick wins · prepared for Copperwick
Therapeutics, Inc.

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

March 13, 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

Summary	2
What went in	4
Change by change	7
The waste register	11
E-2a, L-2, and L-3: GxP standard changes	12
S3-4: in place, not yet measurable	13
L-4: approved at the readout, reversed on Wednesday	14
The freezes and the database lock	14
Running totals	14
Honest limits, as of March 13	15
The exact mechanism	16
Waste: volumes, snapshots, and the idle GPU	16
SageMaker notebooks: auto-stop after 60 idle minutes	17
Schedules, rightsizing, and databases	18
gp2 to gp3 (E-1 and E-2a)	19
The data lake: lifecycle rules and the tar job	20
EFS, networking, logging, and Config	21
What moves to week 4	23

Summary


In week 3, from Monday, March 9, to Friday, March 13, Copperwick's engineers put **19 changes** in place. Eighteen of them count toward the guarantee at their estimates: **\$1,796,936 a year, net of the 6% private pricing discount, 9.8% of the in-scope bill** (\$1,911,633 at list). The other, S3-4, is in place but cannot be measured until its objects move tiers, so it counts \$0 and is reported on its own ([S3-4: in place, not yet measurable](#)). One more change, L-4, was approved at the readout and reversed by Rosalind Tran on Wednesday, March 11, before it was applied; under the terms it counts at its estimate, \$35,194 net ([L-4: approved at the readout, reversed on Wednesday](#)). **Counted so far: \$1,832,130 a year net, 10.0% of the bill and 42.5% of the roadmap.**

Every dollar in this report is an estimate. It is the Analysis's price for the change, annualized and net of the 6% discount, with the price at list beside it as context. Nothing is measured until week six, when each change is read off Copperwick's Cost and Usage Report, and the week-six figure will differ from this one. At the estimates, what is counted so far is \$1,727,130 over the \$105,000 gate, and the gate is decided by the measurement, not by this report.

Every change was approved by the owner of the workload at the readout, made by a Copperwick engineer through Copperwick's own change process, and drafted and reviewed by Vantalect. Fifteen were ordinary changes in non-GxP accounts; Three were GxP standard changes with notice to QA: E-2a, gp2 to gp3 on the GMP systems before their freeze, with an executed checklist, and L-2 and L-3, made from qualified accounts ([E-2a, L-2, and L-3: GxP standard changes](#)). No change touched a clinical system in the freeze, and Program Basalt's database lock is today ([The freezes and the database lock](#)).

Week 3 in numbers. Savings are estimates until the week-six measurement.

Week 3 at a glance	Result
Changes in place, March 10 to 13	19
Counted toward the gate	18
Savings counted, a year, net and estimated	\$1,796,936
The same, a month	\$149,745
The same, a year at list (context)	\$1,911,633
L-4, reversed March 11, counted at its estimate	\$35,194
Counted so far, net	\$1,832,130
S3-4, in place and not yet measurable (counts \$0)	\$119,686 estimated
Rated High confidence	14 of 18
One-time cost of making the changes, estimated, at list	\$149



The ten largest week-3 changes, estimated and annualized, net.

Decision for the sponsor

Adaeze Whitfield: no decision is needed to keep this week's savings. Hold Thursday, March 19, for the commitment decision with Rafael Quintero, and expect the GxP change advisory board to take S-3a on Tuesday, March 17, for execution in the release window on Thursday, March 19. Week 4 puts the largest changes in place ([What moves to week 4](#)).

What went in

Eight kinds of change are counted. The largest by value is waste, \$644,816 a year net across 5 changes; the data lake changes move data to colder storage without moving a record out of its retention.

The counted week-3 changes by kind. Net of the 6% discount; list is context.

Kind of change	IDs	Changes	Net, a year	At list, a year	Share of week 3
Waste	W-1, W-2, W-4, W-5, W-3	5	\$644,816	\$685,973	35.9%
Schedules	S-2, S-1	2	\$248,298	\$264,146	13.8%
Data lake	S3-2, S3-1	2	\$218,155	\$232,080	12.1%
Storage	E-1, EF-1, E-2a	3	\$198,673	\$211,354	11.1%
Logging and governance	L-1, L-2, L-3	3	\$188,613	\$200,652	10.5%
Network	N-1a	1	\$135,143	\$143,770	7.5%
Rightsizing	R-1	1	\$113,832	\$121,098	6.3%
Databases	D-1	1	\$49,406	\$52,560	2.8%
Total		18	\$1,796,936	\$1,911,633	100.0%

Fourteen of the eighteen are rated High confidence: the resource is plainly unused, or the change is a configuration with a known price. Four are Medium (W-4, S-1, EF-1, S3-2): W-4, because EBS snapshots are incremental and deleting one frees only the blocks no other snapshot references; S-1, because the other 40 workstations run overnight renders and the scientists decide; EF-1, because the Infrequent Access read charge is not in the estimate; and S3-2, because Cryo-EM reprocessing decides how much comes back. Two one-time costs came with the week, at list: **\$149** for the two jobs that have run (\$141 on S3-1 (the tar job's compute and requests) and \$8 on S3-2 (lifecycle transition requests)), and \$340 on S3-4 (lifecycle transition requests to Intelligent-Tiering). One-time costs are listed again in week six beside the result and never netted into the annual savings.

When each change went in. Monday, March 9, was the first day of the tar job for the oldest sequencer runs (S3-1), which finished on Thursday, March 12. S3-4 is in place on Wednesday and counts \$0.

Day	ChangesIn place	Counted, net a year
Tuesday, March 10	8W-1, W-2, S-2, E-1, D-1, N-1a, L-1, L-2	\$846,719
Wednesday, March 11	6W-4, W-5, S-1, EF-1, L-3, S3-4	\$424,716
Thursday, March 12	4W-3, R-1, E-2a, S3-2	\$430,830
Friday, March 13	1S3-1	\$94,671
Total	19	\$1,796,936

Change by change

Each row names the owner who approved the change, the account, the date, what it touched, the evidence that made it safe, and the way back. The route is Copperwick's normal change process for every row except E-2a. [The exact mechanism](#) shows how each one was made.

Every counted week-3 change, in the order it went in. S3-4 and L-4 are in their own sections.

Change	Resource scope	Evidence	Way back	Net, a year	At list
W-1 Unattached gp2 volumes (90 TB). Priya Okonkwo; rc-genomics-prod; normal change process; Tuesday, March 10.	90 TB of unattached gp2 volumes in the research and sandbox accounts; 14 TB snapshot first at the owners' request	No attachment in 90 days of CloudTrail; each volume goes to its owner for confirmation in the waste register	Restore a kept volume from its snapshot	\$95,871	\$101,990
W-2 SageMaker GPU notebook auto-stop. Mateo Varga; rc-ml-prod; normal change process; Tuesday, March 10.	4 ml.g5.12xlarge SageMaker notebooks, stopped after 60 idle minutes (billed about 22% of hours after)	Kernel activity shows each notebook busy about 22% of hours	Remove the lifecycle configuration	\$182,152	\$193,778
S-2 Development accounts on a weekday schedule. Priya Okonkwo; rc-genomics-dev; normal change process; Tuesday, March 10.	Every development account, including the clin- and gmp-development accounts (in no freeze): 12 hours on weekdays (260 of 730 hours a month)	No CPU activity outside working hours on weekdays, none on weekends	Remove the schedule tag	\$130,724	\$139,068
E-1 gp2 to gp3, research and platform accounts. Priya Okonkwo; rc-genomics-prod; normal change process; Tuesday, March 10.	420 TB of gp2 volumes in the research and platform accounts, to gp3	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	Modify back to gp2 after six hours	\$97,026	\$103,219

Change	Resource scope	Evidence	Way back	Net, a year	At list
D-1 Multi-AZ off on six development databases. Priya Okonkwo; rc-genomics-dev; normal change process; Tuesday, March 10.	6 development PostgreSQL databases, Multi-AZ off	Development databases need no standby	Modify back	\$49,406	\$52,560
N-1a S3 gateway endpoints in research VPCs. Grace Mbeki; rc-genomics-prod; normal change process; Tuesday, March 10.	S3 gateway endpoints in the research VPCs; 260 TB a month leaves the NAT gateways	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	Remove the route	\$135,143	\$143,770
L-1 Batch job logs to Infrequent Access log groups. Priya Okonkwo; rc-genomics-prod; normal change process; Tuesday, March 10.	Research Batch job logs (34 TB a month) to new log groups in the Infrequent Access class, 30-day retention; the biomarker pipelines keep SOP retention	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	Point the job definitions back	\$98,181	\$104,448
L-2 Research log retention of 90 days. Grace Mbeki; cw-log-archive; GxP standard change; Tuesday, March 10.	Research log groups named on a positive list (prefix and account): 90-day retention (220 of 300 TB); every other log group, GxP audit logs included, keeps its retention	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	None after expiry (the logs are gone)	\$58,216	\$61,932
W-4 Orphaned EBS snapshots (150 TB billed). Priya Okonkwo; rc-genomics-prod; normal change process; Wednesday, March 11.	150 TB billed of orphaned EBS snapshots in research accounts (source volume and AMI gone, past retention)	Source volume and AMI deleted, past the retention policy; incremental blocks may be shared with kept snapshots, so the rating is medium	None after deletion: each is signed off in the waste register first	\$86,630	\$92,160

Change	Resource scope	Evidence	Way back	Net, a year	At list
W-5 Noncurrent versions and incomplete uploads expire (0.4 PB). Tomasz Reyes; sdp-lake-prod; normal change process; Wednesday, March 11.	0.4 PB of noncurrent object versions and incomplete multipart uploads in research buckets, expiring after 30 days	S3 Storage Lens: 0.4 PB of noncurrent versions and incomplete uploads in research buckets	Remove the rule (expired versions do not come back)	\$99,355	\$105,696
S-1 Visualization workstation auto-stop and schedule. Mateo Varga; rc-workstations; normal change process; Wednesday, March 11.	20 of the 60 visualization workstations (g5.2xlarge): idle auto-stop and a weekday schedule; the other 40 stay on	The other 40 run overnight structure renders, so they stay on	Remove the schedule tag	\$117,574	\$125,078
EF-1 EFS lifecycle to Infrequent Access. Tomasz Reyes; sdp-lake-prod; normal change process; Wednesday, March 11.	40 TB of EFS: files untouched for 30 days to Infrequent Access (70% of the data)	Access-time scans: 70% of files untouched for 30 days	Set the policy to none	\$88,941	\$94,618
L-3 Config daily recording in three accounts. Grace Mbeki; cw-security-tooling; GxP standard change; Wednesday, March 11.	AWS Config: daily recording for EC2 instances, network interfaces, and volumes in rc-genomics-prod, rc-genomics-dev, and rc-eks-prod	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	Set the recording frequency back to continuous	\$32,216	\$34,272
W-3 One idle p4d.24xlarge released. Mateo Varga; rc-ml-prod; normal change process; Thursday, March 12.	1 of the 5 p4d.24xlarge training instances, idle 84% of hours; the other 4 stay	Ninety days of GPU utilization from the CloudWatch agent: one instance idle 84% of hours	Relaunch from the training AMI, if capacity allows	\$180,808	\$192,349

Change	Resource scope	Evidence	Way back	Net, a year	At list
R-1 Data platform services to m6i.2xlarge. Tomasz Reyes; sdp-catalog-prod; normal change process; Thursday, March 12.	36 data platform service instances, m6i.4xlarge to m6i.2xlarge	9% CPU and 31% memory at the 95th percentile over 93 days	Roll back the launch template version	\$113,832	\$121,098
E-2a gp2 to gp3, GMP systems before their freeze. Daniel Achterberg; gmp-cds-prod; GxP standard change; Thursday, March 12.	55 TB of gp2 volumes on the GMP systems (LIMS, MES, CDS, and historian production, and the val environments), to gp3, before their freeze began	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	Modify back	\$12,706	\$13,517
S3-2 Cryo-EM raw movies to Glacier Flexible Retrieval. Hana Lindgren; sdp-lake-prod; normal change process; Thursday, March 12.	Cryo-EM raw movies older than 180 days (0.6 PB) to Glacier Flexible Retrieval	Movies are one to four gigabytes each; reprocessing restores in bulk	Bulk restore (free per GB) or standard (\$0.01 a GB)	\$123,484	\$131,366
S3-1 Oldest sequencer runs, tarred, to Deep Archive. Tomasz Reyes; sdp-lake-prod; normal change process; Friday, March 13.	Raw sequencer run folders older than 18 months (0.4 PB), tarred per run with a checksum manifest, to Glacier Deep Archive	Ten-year retention, and a 12-hour restore is acceptable for research raw data	Restore the tars (12 hours, standard retrieval)	\$94,671	\$100,714
Total	18 changes			\$1,796,936	\$1,911,633

The waste register

Waste is the one group where a change can destroy something. The rule for the register is that every deletion has a signed line from the owner of the resource, and a copy is kept wherever someone could want the thing again. Snapshots in validated accounts are excluded from the waste work altogether. The five waste changes are \$644,816 a year net.

The waste register for week 3. Each line was signed in the change ticket on the day the work ran, by the person named.

ID	What was done	Size	Signed by	Date	Kept	If it is wanted again	Net, a year
W-1	Snapshot the 14 TB owners asked for, then delete every volume	90 TB	Priya Okonkwo	Tuesday, March 10	14 TB of snapshots, for 90 days (until June 8)	Restore a volume from its snapshot	\$95,871
W-2	Stop each notebook after 60 idle minutes; an exemption tag keeps a long training run alive	4 notebooks	Mateo Varga	Tuesday, March 10	Nothing deleted; the notebooks and their volumes stay	Remove the lifecycle configuration	\$182,152
W-3	Release one instance; keep four while p4d capacity is scarce	1 of 5 instances	Mateo Varga	Thursday, March 12	The training AMI	Relaunch from the training AMI, if capacity allows	\$180,808
W-4	Delete after the owner's sign-off; snapshots in validated accounts are excluded	150 TB billed	Priya Okonkwo	Wednesday, March 11	Nothing: source volume and AMI are gone and retention has passed	None after deletion; the signed line is the record	\$86,630
W-5	Expire noncurrent versions and abort incomplete uploads after 30 days; GxP buckets are not touched	0.4 PB	Tomasz Reyes	Wednesday, March 11	Current versions, and every GxP bucket	Remove the rule (expired versions do not come back)	\$99,355
Total							\$644,816

Two safeguards carry most of the weight. **Snapshots first (W-1)**. Owners asked to keep 14 TB of the unattached volumes, so each of those was snapshotted and the snapshot kept 90 days, until June 8; its cost, 14 TB at \$0.05 a GB-month, is netted from W-1 even though it ends after three

months. **A verified tar manifest first (S3-1)**. Raw sequencer run folders hold thousands of small files each, and Deep Archive adds 40 KB of overhead per object (8 KB billed at Standard and 32 KB at Deep Archive) and charges \$0.05 per 1,000 transitions. A Batch job therefore tars each run with a checksum manifest, extracts it to a second location and compares every file against the manifest, and only then transitions the tar. The original files expire only after the data owner, Tomasz Reyes, signs the verified run; the originals' ten-year retention carries over to the tar, and a 12-hour restore is acceptable for research raw data.

Some things cannot be undone

Deleted snapshots (W-4) are not recoverable, which is why none was deleted that had a volume or image behind it: each orphaned snapshot's source volume and AML were already gone and its retention had passed, and the owner signed the line. Noncurrent object versions that expire (W-5) do not come back. A log event older than its new retention period (L-2) is gone the moment the retention applies. Each has a signed line, nothing was removed from a resource an owner had not approved, and no GxP record, audit trail, or backup of a validated system was touched.

E-2a, L-2, and L-3: GxP standard changes

gp2 to gp3 on a validated server is a pre-approved change type in the qualified infrastructure catalog: like-for-like, reversible, and with no effect on a GxP function. It needs a ticket, an executed installation checklist, and notice to QA, not a new validation. E-2a covers 55 TB of gp2 volumes on the GMP systems (LIMS, MES, CDS, and historian production, and the val environments), to gp3, before their freeze began, and it was made on Thursday, March 12, by Daniel Achterberg's team: \$12,706 a year net (\$13,517 at list), with the notice to Dr. Lorraine Osei's group.

The installation checklist for E-2a. The same change on the clinical systems in the freeze waits for topline.

Checklist step	What the executed checklist records
Scope	The volume IDs and accounts, matched to the validated inventory; run before the GMP freeze began, and none on a clinical system
Performance	gp3 includes 3,000 IOPS and 125 MiB/s, at or above every volume's gp2 baseline, so performance is the same or better
Execution	Elastic Volumes, online, one volume at a time; the old type, size, and IOPS written to a ledger first
Verification	Each volume reports gp3 and completed; the system owner confirms the application's health check
Way back	Modify the volume back to gp2; a volume can be modified once every six hours
Notice	Notice to QA filed with the ticket; no new validation, because the change is like-for-like and reversible

E-2a ran on the GMP systems on Thursday, March 12, before the GMP freeze began. Its twin, E-2b, is the same change on the clinical systems in the freeze (SCE production and val, and the biomarker pipelines; owner Ines Halvorsen): \$12,706 a year net (\$13,517 at list). It is a GxP standard change with no board vote, made on Monday, April 27, after topline ([What moves to week 4](#)).

L-2 and L-3 are standard changes too. They are made from the log archive and security tooling accounts, which are qualified infrastructure, so each had a ticket and notice to QA; no date moved. L-4, which the CISO reversed before it was applied, is the same class.

S3-4: in place, not yet measurable

S3-4 moved the analysis results (1.6 PB) to S3 Intelligent-Tiering on Wednesday, March 11, with Tomasz Reyes's approval. It is estimated at \$119,686 a year net (\$127,326 at list), and **it counts \$0 toward the gate**. Intelligent-Tiering moves an object to a cheaper tier only after 30 days without access, so nothing on the bill can change before week six, and the monitoring fee on about 34 million objects (\$0.0025 per 1,000 objects a month) starts at once. The estimate is already net of that fee. The change is reported separately, never added to the counted total, and first measured over Monday, May 11, to May 17, once enough objects have moved tiers for a seven-day comparison; the result is ready May 19. The one-time cost of the transition requests, \$340, is listed above and is not annualized.

L-4: approved at the readout, reversed on Wednesday

L-4 would have limited CloudTrail S3 data events on research buckets to write events, saving \$35,194 a year net (\$37,440 at list, \$2,933 a month net). Grace Mbeki owned the change and approved it at the readout. On Wednesday, March 11, before it was applied, **Rosalind Tran, the Chief Information Security Officer, reversed it**. The read events on research buckets feed the insider-risk detection for sequence designs, and those detections matter more to Copperwick than \$2,933 a month. Read and write events stay on for all research and GxP buckets, as before, and nothing about the CloudTrail configuration changed.

The terms cover this case: "a change the client approved at the week-two readout and later declines counts at the savings the roadmap estimated for it." So L-4 **counts at its estimate, \$35,194 a year net**, in the running total below. It is a saving that will never appear on the bill, and we say so plainly. The gate does not depend on it: without L-4, the counted changes are \$1,691,936 over the gate at their estimates. Copperwick does not intend to revisit it, and the detection it protects stays as it was.

The freezes and the database lock

Program Basalt's database lock is Friday, March 13. The clinical freeze has been in force since March 2 and runs to topline on April 24; the GMP campaign freezes the manufacturing systems from Monday, March 16, to March 27. A freeze covers systems: the clinical freeze covers SCE production and val, the clinical data repository, and the biomarker pipelines, and the GMP freeze covers LIMS, MES, CDS, and historian production. Every week-3 change sits outside both: none touched a clinical system, E-2a ran on the GMP systems on Thursday, March 12, before their freeze began, and S-2 ran on Tuesday, March 10, on every development account, including the clin- and gmp- development accounts, which are in no freeze. The thirteen approved changes that wait (S-3b, R-2, R-3, R-4, B-1b, B-3, F-3, E-2b, N-1b, S3-1b, S3-3, S3-5, and G-1) are the ones inside a freeze, behind a GxP change request, or waiting for an event; none moved this week.

Running totals

The roadmap, as of March 13. Savings are net of the 6% discount, and the counted figure is the plan until week six measures it.

The 41 changes by where they stand on March 13. The "Counted so far" line is a subtotal of the two above it and is left out of the whole.

Where the roadmap stands	Changes	Net, a year	At list, a year
In place and counted, March 10 to 13	18	\$1,796,936	\$1,911,633
L-4, approved at the readout and reversed March 11: counts at its estimate	1	\$35,194	\$37,440
Counted so far	19	\$1,832,130	\$1,949,073
Planned for week 4	4	\$729,373	\$775,929
S3-4: in place, not yet measurable (counts \$0; first measured in May)	1	\$119,686	\$127,326
Approved and waiting on a freeze, a change request, or an event	13	\$1,192,737	\$1,268,870
Declined at the readout	4	\$436,118	\$463,956
The whole roadmap	41	\$4,310,044	\$4,585,154

If week 4 lands as planned, 23 changes are counted: \$2,561,503 a year net, 14.0% of the bill (\$2,725,002 at list). That is the most the gate can count, because the thirteen that wait and the four declined at the readout (\$436,118) count for nothing.

Honest limits, as of March 13

- **Estimates, not measurements.** Every figure above is the Analysis's price, annualized. Week six reads each change off Copperwick's Cost and Usage Report on Wednesday, April 1, on data through March 30, using public On-Demand cost so the private pricing discount and the Savings Plans are never counted twice.
- **Dollars follow the analysis order.** Waste is priced first, then rightsizing and schedules, storage, databases, network, the data lake, and logging, so no dollar is counted twice. The S3 changes use the top Standard tier (\$0.021 a GB-month), because the organization stores far more than 500 TB in Standard.
- **W-4 can miss.** Incremental snapshots mean the saving depends on which blocks are shared with the snapshots kept. It is rated Medium for that reason, and the estimate uses each snapshot's billed cost, not its volume size.

- **EF-1 leaves out one charge.** Reading a file back from Infrequent Access costs \$0.01 a GB, which the estimate does not include; the measurement includes the actual charge.
- **N-1a counts same-Region traffic only.** A gateway endpoint carries S3 traffic in the same Region, so reads of the us-west-2 replica bucket still use the NAT gateways.
- **The Cost and Usage Report finalizes about 48 hours late.** The seven days after a change made on March 10 are final on March 19, and after a change on March 13, on March 22.

The exact mechanism

Everything below is the code we used, reduced to what you need to run it again. Nothing changes a resource until it runs against a list its owner approved. Infrastructure that lives in code goes through the landing-zone repository's pipeline as AWS CloudFormation under a normal change ticket; everything else is a runbook that Copperwick's operators run from the shared services account (cw-shared-services) with a role assumed in the target account. Vantalect holds no standing write access: we write the pull request or the script and its dry-run output, and a Copperwick engineer merges or runs it.

Waste: volumes, snapshots, and the idle GPU

The input is the approved rows of the waste register as a CSV:

id,kind,resource,profile,ticket,keep. The profile assumes the remediation role in the resource's account. The script runs on the Linux runner (bash, GNU date, jq, AWS CLI v2) and writes one ledger line per action, which stays with the change ticket. Each branch first checks that the resource is still in the state the register recorded, so something attached or started since the Analysis is skipped, not removed. A snapshot that does not complete keeps its volume.

BASH

```
#!/usr/bin/env bash
# waste-remove.sh <approved.csv> [--execute]      (dry run unless --execute)
set -uo pipefail
PLAN="$1"; MODE="dry-run"; [[ "${2:-}" == "--execute" ]] && MODE="execute"
LEDGER="ledger-$(date -u +%Y%m%dT%H%M%S).tsv"
log() { printf '%s\t%s\n' "$(date -u +%FT%TZ)" "$*" | tee -a "$LEDGER" >&2; }

while IFS=, read -r id kind resource profile ticket keep; do
    [[ "$id" == "id" ]] && continue
    export AWS_PROFILE="$profile" AWS_DEFAULT_REGION=us-east-1
    case "$kind" in
        unattached-volume)                                     # W-1
            state=$(aws ec2 describe-volumes --volume-ids "$resource" --query 'Volumes[0].State' --
output text)
            [[ "$state" == "available" ]] || { log "SKIP $id $resource is $state"; continue; }
            [[ "$MODE" == "execute" ]] || { log "DRY $id snapshot=$keep, then delete $resource";
continue; }
    }
```

```

snap=""
if [[ "$keep" == "yes" ]]; then
    snap=$(aws ec2 create-snapshot --volume-id "$resource" --description "$ticket $id" \
        --tag-specifications "ResourceType=snapshot,Tags=[{Key=Retain-Until,Value=$(date -u \
-d '+90 days' +%F)}]" \
        --query SnapshotId --output text) || { log "FAIL $id snapshot"; continue; }
    aws ec2 wait snapshot-completed --snapshot-ids "$snap" || { log "FAIL $id $snap \
incomplete; volume kept"; continue; }
    fi
    aws ec2 delete-volume --volume-id "$resource" && log "OK $id deleted $resource ${snap: \
+kept $snap}" ;;
    orphaned-snapshot) # W-4
        vol=$(aws ec2 describe-snapshots --snapshot-ids "$resource" --query \
'Snapshots[0].VolumeId' --output text)
        aws ec2 describe-volumes --volume-ids "$vol" >/dev/null 2>&1 && { log "SKIP $id \
$resource: volume $vol still exists"; continue; }
        amis=$(aws ec2 describe-images --owners self --filters "Name=block-device- \
mapping.snapshot-id,Values=$resource" --query 'length(Images)' --output text)
        [[ "$amis" == "0" ]] || { log "SKIP $id $resource: $amis AMI(s) still use it"; \
continue; }
        hold=$(aws ec2 describe-snapshots --snapshot-ids "$resource" --query "Snapshots[0].Tags[? \
Key=='Retain-Until'].Value | [0]" --output text)
        [[ "$hold" == "None" || "$hold" < "$(date -u +%F)" ]] || { log "SKIP $id $resource is \
held until $hold"; continue; }
        [[ "$MODE" == "execute" ]] || { log "DRY $id delete $resource"; continue; }
        aws ec2 delete-snapshot --snapshot-id "$resource" && log "OK $id deleted $resource" ;;
    esac
done < "$PLAN"

```

Snapshots are incremental, so deleting one frees only the blocks no other snapshot references. That is why the W-4 estimate is each snapshot's billed cost in the Cost and Usage Report's resource lines, not its volume size, and why W-4 is rated Medium. Snapshots in validated accounts are never in the input.

W-3 releases one of the five p4d.24xlarge instances. The ML team keeps p4d instances running because capacity is hard to get back, so the release keeps the way back: the training AMI must be current before the instance goes, and a relaunch from it works if capacity allows.

BASH

```

aws ec2 describe-images --image-ids <training-ami-id> --query 'Images[0].CreationDate' --output \
text # confirm it is current
aws ec2 terminate-instances --instance-ids <instance- \
id> # the one idle 84% of hours

```

SageMaker notebooks: auto-stop after 60 idle minutes

W-2 is a lifecycle configuration. The on-start script checks the notebook's tags, and a notebook tagged autostop-exempt=true is left alone so a long training run can finish. Everything else gets a cron job that runs the idle-stop script from the AWS sample lifecycle configurations, which the landing-zone repository vendors and stages in S3.

BASH

```
cat > on-start.sh <<'SH'
#!/bin/bash
set -e
IDLE_SECONDS=$(( 60 * 60 ))
NB_ARN=$(jq -r '.ResourceArn' /opt/ml/metadata/resource-metadata.json)
EXEMPT=$(aws sagemaker list-tags --resource-arn "$NB_ARN" --query "Tags[?Key=='autostop-exempt'].Value | [0]" --output text)
if [[ "$EXEMPT" != "true" ]]; then
    aws s3 cp s3://<artifact-bucket>/sagemaker/autostop.py /home/ec2-user/autostop.py
    (crontab -l 2>/dev/null; echo "*/5 * * * * /usr/bin/python3 /home/ec2-user/autostop.py --time $IDLE_SECONDS --ignore-connections") | crontab -
fi
SH
aws sagemaker create-notebook-instance-lifecycle-config --notebook-instance-lifecycle-config-name idle-autostop --on-start Content="$(base64 -w0 on-start.sh)"
# for each notebook: stop it, attach the configuration, start it
aws sagemaker stop-notebook-instance --notebook-instance-name <name>
aws sagemaker wait notebook-instance-stopped --notebook-instance-name <name>
aws sagemaker update-notebook-instance --notebook-instance-name <name> --lifecycle-config-name idle-autostop
aws sagemaker start-notebook-instance --notebook-instance-name <name>
```

The four ml.g5.12xlarge notebooks ran around the clock for scientists who use them in bursts; kernel activity shows each busy about 22% of hours, and the estimate prices the hours that remain. The rollback is to remove the lifecycle configuration.

Schedules, rightsizing, and databases

S-2 puts the development accounts on 7 AM to 7 PM Eastern on weekdays (260 of 730 hours a month) with Instance Scheduler on AWS: a schedule with its period, and a Schedule tag on each host. S-1 uses the same scheduler for 20 of the 60 visualization workstations and adds an idle-stop alarm on the CloudWatch agent's GPU utilization; the other 40 stay on because scientists run overnight structure renders on them.

YAML

```
# Instance Scheduler on AWS: the schedule and its period (a CloudFormation custom resource, abbreviated)
DevWeekdaySchedule:
  Type: Custom::ServiceInstanceSchedule
  Properties:
    ServiceToken: !Ref InstanceSchedulerServiceToken
    NoStackPrefix: "True"
    Name: dev-weekday-12h
    Description: 7 AM to 7 PM Eastern on weekdays
    Timezone: America/New_York
    Periods:
      - Description: weekday-12h
        BeginTime: "07:00"
        EndTime: "19:00"
        WeekDays: mon-fri
```

R-1 is a launch template change with a rolling instance refresh, so the data platform services are replaced a few at a time. D-1 turns Multi-AZ off on six development PostgreSQL databases, which need no standby. Both have an instant way back.

BASH

```
aws ec2 create-launch-template-version --launch-template-id <lt-id> --source-version '$Latest'
--launch-template-data '{"InstanceType":"m6i.2xlarge"}'
aws ec2 modify-launch-template --launch-template-id <lt-id> --default-version '$Latest'
aws autoscaling start-instance-refresh --auto-scaling-group-name <asg> --preferences
'{"MinHealthyPercentage":90,"InstanceWarmup":300}'

aws rds modify-db-instance --db-instance-identifier <dev-db> --no-multi-az --apply-
immediately      # D-1, one database at a time
```

gp2 to gp3 (E-1 and E-2a)

gp3 costs \$0.08 a GB-month against \$0.10 for gp2, with 3,000 IOPS and 125 MiB/s included at any size. The change is online with Elastic Volumes. The Analysis found the included values at or above every volume's gp2 baseline, so the estimate prices every volume at the gp3 defaults. A volume that needed more would be converted with `--iops` and `--throughput` set, up to gp3's limits of 80,000 IOPS and 2,000 MiB/s per volume, and priced at the extra rate. The ledger line comes first, so the old values are never lost, and a volume can be modified once every 6 hours, so a rollback waits for that. E-2a runs the same commands under the checklist in [E-2a](#), [L-2](#), and [L-3](#): [GxP standard changes](#).

BASH

```
vol=<volume-id>
read -r type size iops < <(aws ec2 describe-volumes --volume-ids "$vol" --query 'Volumes[0].
[VolumeType,Size,Iops]' --output text)
printf '%s\t%s\t%s\t%s\n' "$vol" "$type" "$size" "$iops" >> volume-ledger.tsv      # the way
back
aws ec2 modify-volume --volume-id "$vol" --volume-type gp3
aws ec2 describe-volumes-modifications --volume-ids "$vol" --query
'VolumesModifications[0].ModificationState' --output text      # modifying, optimizing, completed
```

The data lake: lifecycle rules and the tar job

A lifecycle configuration replaces the bucket's whole configuration, so every put-bucket-lifecycle-configuration starts from get-bucket-lifecycle-configuration and adds a rule; the previous document is kept as the way back. W-5 applies only to research buckets, and GxP buckets keep Object Lock and versioning. S3-2 moves cryo-EM raw movies to Glacier Flexible Retrieval at 180 days; a restore uses bulk retrieval, which is free per GB, or standard retrieval at \$0.01 a GB. S3-4 moves analysis results to Intelligent-Tiering.

JSON

```
{
  "Rules": [
    {
      "ID": "research-noncurrent-30d", "Status": "Enabled", "Filter": {"Prefix": ""},
      "NoncurrentVersionExpiration": {"NoncurrentDays": 30},
      "AbortIncompleteMultipartUpload": {"DaysAfterInitiation": 30}},
    {
      "ID": "cryoem-movies-flexible", "Status": "Enabled", "Filter": {"Prefix": "cryoem/raw/"},
      "Transitions": [{"Days": 180, "StorageClass": "GLACIER"}]},
    {
      "ID": "analysis-results-it", "Status": "Enabled", "Filter": {"Prefix": "results/"},
      "Transitions": [{"Days": 0, "StorageClass": "INTELLIGENT_TIERING"}]}
  ]
}
```

S3-1 is the one with a build step. Each sequencer run folder holds thousands of small files, and Deep Archive adds 40 KB of overhead per object (8 KB billed at Standard and 32 KB at Deep Archive) and charges \$0.05 per 1,000 transitions, so transitioning the files one by one would cost more than it saves. One Batch job per run tars it, writes a checksum manifest, verifies the tar by extracting it and comparing every file, and only then uploads it to Deep Archive. The estimate prices the move at the top Standard tier (\$0.021 a GB-month) against Deep Archive (\$0.00099).

BASH

```
#!/usr/bin/env bash
# tar-run.sh <bucket> <run-prefix>          one sequencer run folder per Batch job; scratch is
the container's ephemeral volume
set -euo pipefail
BUCKET="$1"; RUN="${2%/}"; WORK=$(mktemp -d)
aws s3 cp --recursive "s3://$BUCKET/$RUN/" "$WORK/src" --only-show-errors
( cd "$WORK/src" && find . -type f -print0 | sort -z | xargs -0 sha256sum ) > "$WORK/
manifest.sha256"
tar -C "$WORK/src" -cf "$WORK/run.tar" .
mkdir "$WORK/check" && tar -C "$WORK/check" -xf "$WORK/run.tar"
( cd "$WORK/check" && sha256sum --quiet -c "$WORK/manifest.sha256" )      # fails the job on
any mismatch
aws s3 cp "$WORK/run.tar" "s3://$BUCKET/archive/$RUN.tar" --storage-class DEEP_ARCHIVE --only-
show-errors
aws s3 cp "$WORK/manifest.sha256" "s3://$BUCKET/archive/$RUN.manifest.sha256" --only-show-
errors
printf '%s\tVERIFIED\t%s\n' "$(date -u +%FT%TZ)" "$RUN" >> verified-runs.tsv
```

The originals are not touched by this script. A second step reads `verified-runs.tsv` and the data owner's signed list, and deletes the original objects, every version, by version ID, only for runs that are in both, so no noncurrent copy of an original stays billed under the lake's 30-day noncurrent-version rule (W-5). S3 Inventory then shows the transitions and deletions complete; for the oldest runs that was Friday, March 13. To get a run back, request a standard retrieval of the tar from Deep Archive (about 12 hours), extract it, and compare against its manifest.

EFS, networking, logging, and Config

EF-1 moves files untouched for 30 days to Infrequent Access (\$0.30 to \$0.025 a GB-month), and back to Standard on first access. N-1a adds a gateway endpoint to each research VPC's route tables, so S3 traffic stops going through the NAT gateways at \$0.045 a GB.

BASH

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

YAML

```
S3GatewayEndpoint:
  Type: AWS::EC2::VPCEndpoint
  Properties:
    VpcId: !Ref Vpc
    VpcEndpointType: Gateway
    ServiceName: !Sub com.amazonaws.${AWS::Region}.s3
    RouteTableIds: !Ref PrivateRouteTableIds
```

A log group's class is set when it is created, so L-1 points the Batch 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. L-2 sets 90-day retention on a positive list of research log groups, named by prefix and account and agreed with the owners; a group that is not on the list is never touched, so GxP audit logs and every other group keep their retention. It is a GxP standard change with notice to QA, run from the log archive account. L-3 records EC2 instances, network interfaces, and volumes daily instead of continuously in the three high-churn accounts (rc-genomics-prod, rc-genomics-dev, and rc-eks-prod; \$0.012 per item a day against \$0.003 for each change).

YAML

```
BatchJobLogs:
  Type: AWS::Logs::LogGroup
  Properties:
    LogGroupName: /batch/jobs/research
    LogGroupClass: INFREQUENT_ACCESS
    RetentionInDays: 30

ConfigRecorder:
  Type: AWS::Config::ConfigurationRecorder
  Properties:
    RoleARN: !GetAtt ConfigRole.Arn
    RecordingGroup: {AllSupported: true, IncludeGlobalResourceTypes: false}
    RecordingMode:
      RecordingFrequency: CONTINUOUS
      RecordingModeOverrides:
        - ResourceTypes: [AWS::EC2::Instance, AWS::EC2::NetworkInterface, AWS::EC2::Volume]
          RecordingFrequency: DAILY
          Description: Short-lived Batch and Karpen-ter instances and their interfaces and
volumes
```

BASH

```
# research-log-groups.tsv is the positive list: one "account<TAB>log-group-prefix" line per
research prefix, signed by the owners
while IFS=$'\t' read -r account prefix; do
  for group in $(aws logs describe-log-groups --profile "$account" --log-group-name-prefix
"$prefix" --query 'logGroups[].logGroupName' --output text); do
    aws logs put-retention-policy --profile "$account" --log-group-name "$group" --retention-
in-days 90 && echo "OK $account $group"
  done
done < research-log-groups.tsv
```

What moves to week 4

Week 4 runs from Monday, March 16, to Friday, March 20. The GMP campaign freezes the manufacturing systems in production from the first day, so everything in it is either non-GxP or, like S-3a on the GMP val environments (in no freeze), goes through the GxP change advisory board, which meets on Tuesdays, and then a release window.

What week 4 puts in place, in date order. The commitment decision is Rafael Quintero's and Adaeze Whitfield's.

ID	Planned change	Owner	Date	What happens	Net, a year
B-1	Genomics Batch: 40% of the work to Spot	Priya Okonkwo	Monday, March 16	Spot environments go live; per-genome cost tracked daily	\$238,111
F-1	Genomics Lustre resized to 48 TiB at 500 MB/s/TiB	Priya Okonkwo	Tuesday, March 17	A new file system takes over between pipeline batches; the old one is deleted only after a full pipeline batch has run on the new one, after a checksum comparison	\$255,040
C-1	Compute Savings Plan, one year, \$50 an hour	Rafael Quintero	Thursday, March 19	The commitment decision with Rafael Quintero and Adaeze Whitfield; bought that afternoon if approved	\$148,682
S-3a	GMP val environments on a weekday schedule	Daniel Achterberg	Thursday, March 19	The GxP change advisory board approves it on Tuesday, March 17; it executes in this release window, with a hold tag for executed protocols	\$87,540
Total					\$729,373

The thirteen changes that wait (\$1,192,737 a year net) keep the dates in the deferral calendar, which the table below repeats.

The deferral calendar, one set of dates for every report.

Changes	When	Waiting on
F-3, B-3	April 6	The next cryo-EM campaign
N-1b, S3-3	April 14	Approved change requests (S3-3 with a documented retrieval test)
S3-1b	April 20	The tar job's second pass
G-1	April 20	The molecular dynamics benchmark
S3-5	April 21	The disaster-recovery test
S-3b, E-2b	April 27	Topline: S-3b by change request, E-2b as a standard change
B-1b	May 11	Biomarker pipeline revalidation
R-3	May 18	The LIMS release, with operational qualification regression
R-4	June 15	The MES maintenance window after the periodic review
R-2	June 22	After the post-topline analyses, by change request

The week-4 report follows on Friday, March 20; it records what Rafael Quintero decided on Thursday, March 19. Week 5 is guardrails, and no savings by design.

© 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 Seqera Labs. Vantalect is not affiliated with or endorsed by any of them. Copperwick Therapeutics, Inc. is fictional.