

AWS WELL-ARCHITECTED FRAMEWORK REVIEW · SAMPLE DELIVERABLE

Workload Definition, Current-State Architecture, and Evidence

Week 1 of 6: Workload and evidence · prepared for Tidepost, Inc.

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

February 27, 2026 · AWS Well-Architected Framework Review: The Remediation (six weeks)

SAMPLE DELIVERABLE

This is a sample deliverable. Tidepost, Inc. is a fictional company: its people, products, users, accounts, investors, and figures are invented for this sample, and any resemblance to a real company, person, product, or publication 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 18, 2026.

Contents

- Summary 2
- The workload 3
 - The lens 4
 - What is in scope, and what is not 4
- How the Review reaches the workload 5
- The architecture as found 5
 - Engines and support dates..... 7
- The components, path by path 7
 - Serving 7
 - Upload and moderation 8
 - Ranking pipeline 9
 - Telemetry 9
 - Network 10
- What the Review read 10
- The week-one sessions 11
- What week two does 12

Summary

The Brief is defined as a workload in Tidepost’s own AWS Well-Architected Tool, its architecture is written down as found, and the evidence is in hand. Nothing is scored yet. Dev Raman and Kofi Brandvold: this is the first of the Review’s reports. It records what we reviewed, what we read to review it, and what happens next, so that the 57 questions we answer next week start from a workload both of us have agreed on.

The Review runs six weeks, and this week it reviews what exists. The 23 components below are what Tidepost runs today, with no changes proposed yet. Next week we answer every question and bring the scorecard to the sponsor readout on March 6.

Where the Review stands on February 27, 2026. The counts of single-zone and console-built components describe the estate; they are not findings and are not scored.

Week one at a glance	Result
Workload	the Brief, in tp-prod, us-east-1
Defined in the Tool	February 23, 2026
Lens that is answered and counted	AWS Well-Architected Framework (57 questions, six pillars)
Components described	23 in 5 paths
Components in one Availability Zone	10
Components built in the console	14
Components defined in code (CDK)	9
Sessions held	4, 5 hours with the engineers
Evidence sources read	7
Questions answered or scored	None; they are answered in week two
Consultant hours this week	20

The workload

The workload is the Brief: how a contributor's upload becomes a moderated, ranked item on a reader's timeline. It is the product Tidepost sells. Everything a contributor uploads, every moderation decision, every ranked timeline, and every playback runs through it, which makes it the one workload worth reviewing end to end. Tidepost's Series A conditions and the 18 newsrooms whose vendor reviews ask about it both come back to how this workload behaves when something fails.

We defined it in the Tool on February 23, 2026, the day of the kickoff, with Kofi Brandvold as the Review owner.

The workload's properties as entered in the Tool. The Tool takes the account's 12-digit ID, shown here by its last four digits; it is tp-prod.

Property in the Tool	Value
Name	the Brief
Description	Upload, moderation, embedding, ranking, timeline, and playback for the Tidepost apps
Review owner	Kofi Brandvold, Head of Platform
Environment	Production
AWS account	...4417
AWS Region	us-east-1
Lens	AWS Well-Architected Framework

The lens

One lens is answered and counted, and one is context. The AWS Well-Architected Framework lens has 57 questions across the six pillars: Operational Excellence, Security, Reliability, Performance Efficiency, Cost Optimization, and Sustainability. It is the lens the engagement's guarantee is measured against.

The lenses the Review uses. Only the first is attached to the workload in the Tool, answered, scored, and counted toward the gate; the second is read as context.

Lens	QuestionsRole
AWS Well-Architected Framework	57The gate
Machine Learning Lens	NotContext only; its questions are not answered in the Tool answeredor counted

We read the Machine Learning Lens while we learn how the ranking and moderation models are built and served. We do not answer its questions in the Tool, and nothing in it is counted.

What is in scope, and what is not

In scope is the path from upload to playback: upload, moderation, embedding, ranking, timeline, and playback. That is every component in the architecture below.

Out of scope, so that a finding about one of these is never a finding about the Brief:

- **Staging.** It still runs in tp-prod beside production, but it is a copy used for testing, not part of the Brief. Where it shares production's network or account, the Review answers for production's use.
- **Developer instances.** They exist to build and test the Brief and serve no reader or contributor.
- **The data science sandbox and notebooks.** They explore data and train models; nothing a reader or contributor touches depends on them.

The supporting accounts the earlier engagements created (tp-log-archive for the organization trail and configuration history, tp-audit, tp-backup, and tp-management for billing data) are places we read evidence, not part of the workload.


How the Review reaches the workload

Vantalect reads through a read-only inspection role and writes nothing to Tidepost's production. The role can describe and list; it cannot change, and Tidepost can revoke it at any time.

- **Every change to the estate lands as code.** When a fix is agreed, it goes in as a pull request in Tidepost's tidepost-cdk repository (or in Tidepost's console where a fix is a setting with no code behind it yet). The owning engineer reviews it and Tidepost's pipeline ships it; a fix that is a runbook, a process, or a standard is written and adopted by its owner. Vantalect pairs on the change and reviews it; it never deploys.
- **One thing is written in Tidepost's account.** The workload and its answers, in Tidepost's own Well-Architected Tool. Tidepost owns the workload, and its milestones stay with Tidepost after the engagement.
- **Sessions are recorded with permission.** Each session below was recorded with the engineers' permission, so that an answer can be traced to what was said.

The architecture as found

The Brief runs on four request and data paths, and its stateful stores, the serving tiers, and the network each sit in one Availability Zone. Uploads, moderation, ranking, and serving are separate paths that meet at DynamoDB and the API; telemetry is a fourth that feeds ranking.



The Brief's current-state architecture on February 23, 2026. Boxes outlined in red are stores that sit in one Availability Zone with no standby.

Upload and moderation. A contributor uploads to the originals bucket with a presigned URL. An S3 event starts the Lambda orchestrators, which run every upload through several services. Video is moderated twice: frames sampled at one a second go to Rekognition image moderation, and the whole video also goes to Rekognition Video moderation. Audio goes to Transcribe and then Comprehend toxicity. A SageMaker classifier scores the rest, and the transcoder fleet writes the renditions readers play.

Ranking. Ranking is asynchronous. The embedding endpoint embeds each upload, features land in DynamoDB, and the API ranks candidates per request against a real-time SageMaker endpoint.

Serving. The API and web tiers run on EC2 behind Application Load Balancers, and the PostgreSQL 13 writer with two read replicas, ElastiCache for Redis OSS, and OpenSearch serve them. CloudFront delivers renditions.

Telemetry. The collector on EKS puts events on Kinesis, Firehose writes them to raw JSON in S3, and a Glue job and Athena turn them into the reports the ranking and product teams read.

Engines and support dates

Two engines are at the end of standard support, and the dates are facts the Review will weigh under the Reliability and Operational Excellence pillars.

Engine versions and their support dates on February 27, 2026. Support statements are as published by AWS.

Component	Version	Support
RDS for PostgreSQL 13 (production and staging)	PostgreSQL 13	Standard support ends February 28, 2026; RDS Extended Support bills from March 1, 2026
Amazon EKS, production and staging	Kubernetes 1.30	Standard support ended July 23, 2025; extended support until July 23, 2026
Amazon EKS, development	Kubernetes 1.33	Standard support to July 29, 2026
Amazon ElastiCache	Redis OSS 7.1	Supported
Amazon OpenSearch Service	OpenSearch 2.13	Supported

The components, path by path

23 components in 5 paths: 10 sit in one zone, 14 were built in the console, and 9 are in the CDK. Each table says what the component runs on, where it is placed, and whether Tidepost’s code defines it. The quantities are from the January 2026 bill, the most recent complete month.

Serving

8 components: 6 in one zone, 3 built in the console, 5 defined in the CDK.

The serving path. "Placement" is where the component runs; "Defined in" says whether tidepost-cdk creates it or someone built it in the console.

Component	Runs on	Placement	Defined in
API tier	6 c6i.4xlarge, On-Demand, always on	One zone	The CDK
Web tier	2 m5.2xlarge, On-Demand, always on	One zone	The CDK
EKS production node group (event collector and timeline services)	9 m5.4xlarge, On-Demand, always on	One zone	The CDK
Ranking endpoint (SageMaker real-time)	2 ml.g5.2xlarge, about 9% GPU	Placed by SageMaker	The CDK
PostgreSQL 13 writer and two read replicas	1 db.r6g.4xlarge; 2 db.r6g.2xlarge	One zone, no standby	The console
ElastiCache for Redis OSS	2 cache.r6g.4xlarge	One zone, no replicas	The console
OpenSearch Service	3 r6g.2xlarge.search, 3 TB gp3	One zone, no dedicated masters	The console
Renditions to viewers (CloudFront)	400 TB a month	Global	The CDK

Upload and moderation

8 components: 2 in one zone, 8 built in the console, 0 defined in the CDK.

The upload and moderation path. "Placement" is where the component runs; "Defined in" says whether tidepost-cdk creates it or someone built it in the console.

Component	Runs on	Placement	Defined in
Originals bucket (presigned uploads)	Not billed separately	Regional	The console
Moderation orchestrators (Lambda)	20 million requests, 25 million GB-seconds	Regional	The console
Frame sampler (one frame a second)	3 c6i.2xlarge, On-Demand, always on	One zone	The console
Rekognition image moderation	10.53 million images a month	Regional	The console
Rekognition Video moderation	157,500 minutes a month	Regional	The console
Transcribe, then Comprehend toxicity	157,500 minutes a month; 4.71 million units a month	Regional	The console
Moderation classifier (SageMaker)	2 ml.g5.xlarge	Placed by SageMaker	The console
Transcoder fleet	6 g4dn.2xlarge, On-Demand, always on	One zone	The console

Ranking pipeline

2 components: 0 in one zone, 1 built in the console, 1 defined in the CDK.

The ranking pipeline path. "Placement" is where the component runs; "Defined in" says whether tidepost-cdk creates it or someone built it in the console.

Component	Runs on	Placement	Defined in
Embedding endpoint (SageMaker asynchronous)	2 ml.g5.12xlarge, one GPU of four used	Placed by SageMaker	The CDK
DynamoDB features	3 billion write and 15 billion read request units; 2 TB	Regional, no point-in-time recovery	The console

Telemetry

3 components: 0 in one zone, 0 built in the console, 3 defined in the CDK.

The telemetry path. "Placement" is where the component runs; "Defined in" says whether tidepost-cdk creates it or someone built it in the console.

Component	Runs on	Placement	Defined in
Kinesis Data Streams	100 shards	Regional	The CDK
Firehose to raw telemetry JSON in S3	6,809 GB	Regional	The CDK
Glue hourly ETL and Athena	720 DPU-hours; 300 TB scanned	Regional	The CDK

Network

2 components: 2 in one zone, 2 built in the console, 0 defined in the CDK.

The network path. "Placement" is where the component runs; "Defined in" says whether tidepost-cdk creates it or someone built it in the console.

Component	Runs on	Placement	Defined in
VPC, subnets, and NAT	1 NAT gateway, 55 TB processed	One zone	The console
Application Load Balancers	3 load balancers, 60 LCUs on average	One zone's subnets	The console

These are descriptions, not scores. How a one-zone store or a console-built resource affects an answer is decided question by question in week two, with the engineers who run it.

What the Review read

7 sources, all read-only, and each one used for a stated purpose. An answer in week two cites the source it came from, so that every question can be traced to evidence rather than to an opinion.

The evidence base for the Review, with the access used for each source.

Source	How we read it	Used for
Live configuration	The read-only inspection role in tp-prod: describe and list calls on every service the workload uses, and the AWS Config history in tp-log-archive	Zone placement, encryption, network paths, scaling settings, backups, quotas
Telemetry	CloudWatch metrics and logs through the same role	Utilization, latency, error rates, throttling, and which alarms exist
Billing	Cost Explorer and the Cost and Usage Report data export in tp-management, read-only	What each component costs and which components run around the clock
Infrastructure code	The tidepost-cdk repository, read in GitHub	What is defined in code and what was built in the console
Service quotas	Service Quotas and the Rekognition throttling metrics	Which limits the workload meets at peak
Earlier engagements	The security engagement's findings and the Multi-Account Foundation's account map	What was already fixed, so the Review does not score it twice
Sessions	The four week-one sessions with the engineers who run the workload, recorded with their permission	Targets, history, and anything the configuration cannot show

The week-one sessions

We held 4 sessions with the people who run the Brief, 5 hours in all. Each session had a topic, so that the right engineers were in the room and nobody sat through a pillar that was not theirs.

The sessions held in week one.

Date	Session	Minutes	Who	Covered
February 23	Kickoff	60	Dev Raman, Kofi Brandvold, Yuki Sorensen, Lior Haddad, Amara Delacroix, Marcus Oyelaran	The workload's boundary, the recovery and service-level targets to review against, access, and the two-week schedule
February 24	Architecture and reliability session	90	Kofi Brandvold, Marcus Oyelaran, Yuki Sorensen, Lior Haddad, Amara Delacroix	The upload, moderation, ranking, and telemetry paths; data stores; failure modes; backups; quotas; and scaling
February 25	Security and operations session	90	Sasha Ilyin, Kofi Brandvold, Marcus Oyelaran	Workload identity, data classification, the pipeline, telemetry and alarms, and the incident history
February 26	Performance, cost, and sustainability session	60	Kofi Brandvold, Yuki Sorensen, Lior Haddad, Gwen Tallis	Instance and endpoint sizing against utilization, storage and retention, and how the bill is reviewed

Recovery objectives. Tidepost had no written recovery objectives at the kickoff. The Review records that under REL13 in week two and proposes objectives for Tidepost to write down; it does not treat any number as an existing target.

What week two does

Next week we answer all 57 questions, score them, and bring the result to the sponsor readout on March 6. In the Tool, each question is answered by selecting the best practices that are in place; the Tool counts the rest as risk. We record the evidence for each answer in the Tool's notes so that nothing is a bare opinion.

- **March 5, HRI and roadmap review (90 minutes).** Every HRI and MRI with its evidence, the proposed fix, and the sequence. With Kofi Brandvold, Marcus Oyelaran, Yuki Sorensen, Lior Haddad, Amara Delacroix, Sasha Ilyin, Gwen Tallis.
- **March 6, Sponsor readout (60 minutes).** The six-pillar scorecard, the agreed HRI register, the roadmap, and milestone one. With Dev Raman, Noor Castellanos, Kofi Brandvold, Gwen Tallis.

By March 5 you will have:

- The six-pillar scorecard, with every question's risk and evidence.
- The register of high-risk issues (HRIs), each with a proposed fix, an owner, and a week.
- The register of medium-risk issues (MRIs), kept as an improvement plan outside the guarantee.
- The answers for all 57 questions, pillar by pillar, as an appendix.

The answers are saved as **milestone one** at the readout on March 6. The guarantee is measured against that baseline: every agreed HRI closed, or accepted in writing, and verified in milestone two.

What we need from Tidepost. The engineers named on the week-two agenda available for the March 5 session, and the owner of each data store and each tier ready to confirm what the Review recorded about it.

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

AWS, Amazon Web Services, AWS Graviton, AWS Inferentia, and the AWS service names are trademarks of Amazon.com, Inc. or its affiliates. PostgreSQL is a registered trademark of the PostgreSQL Community Association of Canada. Kubernetes is a registered trademark of The Linux Foundation. Redis is a registered trademark of Redis Ltd. Valkey is a trademark of the Linux Foundation. OpenSearch is a registered trademark of Amazon.com, Inc. SOC 2 is a registered trademark of the American Institute of Certified Public Accountants. CIS Benchmarks is a trademark of the Center for Internet Security, Inc. GitHub is a trademark of GitHub, Inc. Jupyter is a trademark of LF Charities. Vantalect is not affiliated with or endorsed by any of them. Tidepost, Inc. is fictional.