ReesyncAI

REESYNC AI LTD · CONFIDENTIAL (SAMPLE)

AWS Data Platform Rescue Audit

Sample report — Acme Utilities Ltd

PREPARED BY JAMES REES, REESYNC AI LTD

JUNE 2026 · VERSION 1.1

SAMPLE · FICTIONAL CLIENT · SYNTHETIC DATA

About this sample. Acme Utilities Ltd is a fictional company. Every figure, finding and quote in this document is synthetic, built to show prospective clients exactly what a real Rescue Audit report looks like: its structure, its evidence standard, its scoring and its recommendations. A real report follows this format precisely, populated with your evidence. Estimates in a real report are stated as ranges and never guaranteed.

DOCUMENT CONTROL

Engagement record

Engagement	AWS Data Platform Rescue Audit (fixed price)
Sponsor	Head of Data, Acme Utilities Ltd
Audit window	12–23 May 2026 (10 working days)
Access route	Cross-account read-only IAM role (external ID protected) plus Tableau questionnaire and screenshots
Scope	One AWS organisation, 2 accounts (prod, dev), eu-west-2; Tableau Server on EC2
Out of scope	Data content review, penetration testing, formal compliance certification, environment changes
Interviews	Head of Data; Senior Data Engineer; Finance BI Lead

SECTION 1

Executive summary

Acme Utilities runs a genuinely useful AWS data platform. It feeds the dashboards the business runs on, and the team keeping it alive is doing so under real pressure. This audit is not a criticism of that team. It is a map of where the platform is leaking money, where it is fragile, and what to fix first.

THE HEADLINE NUMBERS

£14,200/mo
current platform spend
up ~45% in 12 months without a matching rise in workload

£3,100–3,600/mo
addressable cost identified (£37,000–43,000/year)
~£2,300/month is high confidence, evidenced in billing and configuration

2.2 / 5
overall health score — **At Risk**
weakest: observability 1.4, cost 1.8, maintainability 1.9

21
findings: 2 Critical, 6 High, 10 Medium, 3 Low
nine qualify as quick wins deliverable inside 30 days

All savings figures in this report are stated as ±30% ranges derived from observed evidence; they are estimates, not guarantees.

The three things that matter most

First, the business finds out about failures before the data team does.

There is not a single CloudWatch alarm on any data platform component. The nightly finance pipeline failed 11 times in the last 90 days, and on every occasion it was Finance who noticed, on average 4.2 hours after the failure (REL-01, OBS-01). This is the root cause of the erosion of trust in reporting, and it is fixable in days.

Second, the platform pays repeatedly for work it has already done.

Nine of the 25 Glue jobs reprocess full history on every run because the raw data layer is unpartitioned. Combined with ad hoc Athena queries scanning the same unpartitioned data, this accounts for roughly £1,700/month of avoidable spend (COST-01, COST-04), and it also explains why the overnight window keeps stretching.

Third, the platform depends on one person.

A single senior engineer holds the working knowledge of the estate; 19 of 25 jobs have no runbook, and the 2023 contractor handover document no longer reflects reality (MAINT-01). This is the largest unquantified risk in this report.

RECOMMENDED NEXT STEP

The 30/60/90-day plan in Section 7 sequences the work. The first 30 days require no procurement and no architectural change: an alerting baseline, S3 lifecycle policies, key rotation, decommissioning dead pipelines and partitioning the three heaviest datasets. Delivered together, these remove an estimated **£900–£1,200/month** of cost and, more importantly, end the era of finding out about failures from Finance.

SECTION 2

Platform health scorecard

Scores run from 1 (critical weakness) to 5 (exemplary). Bands: 1.0–1.9 Critical, 2.0–2.9 At Risk, 3.0–3.9 Stable but inefficient, 4.0–4.9 Good, 5.0 Exemplary. The overall score is the unweighted mean.

BAR COLOUR = BAND — Critical 1.0–1.9 | At Risk 2.0–2.9 | Stable but inefficient 3.0–3.9 | Good 4.0–4.9 | Exemplary 5.0

Cost optimisation	1.8	Significant avoidable spend; no cost ownership, tagging or review cycle
Reliability	2.1	Recurrent silent failures; no retries on key jobs; duplicates from manual reruns
Operational excellence	2.6	Schedules exist but change goes straight to prod; thin runbooks
Security and governance	2.7	No exposure found, but stale keys, admin sprawl and weak separation
Observability	1.4	No alarms anywhere; failures detected by consumers
Data engineering hygiene	2.4	No data quality checks; dead pipelines still running; schema drift breaks jobs
BI and Tableau readiness	2.5	Slow executive dashboards; refresh storms; conflicting KPI definitions
Maintainability	1.9	Single point of human failure; documentation badly out of date

Overall 2.2 / 5 AT RISK

What “At Risk” means in practice: the platform works on a normal day, but it has little margin. Failures are absorbed by individual effort rather than by design, costs drift upward unchallenged, and a single resignation would put reporting continuity in genuine doubt.

SECTION 3

Platform overview

Estate summary as observed, May 2026.

ACCOUNTS

prod and dev in one AWS organisation, eu-west-2. **No separate test account**; engineers work directly against prod data.

STORAGE

S3 data lake, **38 TB** current objects plus **21 TB** of non-current versions on two versioned buckets with no expiry rules. Raw zone is JSON, unpartitioned. Curated zone partially Parquet.

PROCESSING

25 Glue jobs (PySpark), 14 triggers, 6 crawlers. One Step Functions state machine wraps the nightly finance pipeline (4 jobs). 12 Lambda functions handle ingestion and glue logic. EventBridge holds 13 schedules.

WAREHOUSE

Redshift, **4-node ra3.xlplus, running 24/7**. CPU utilisation under 15% outside 06:00–18:00 weekdays.

QUERY

Athena, two workgroups, no per-workgroup scan limits.

BI

Tableau Server on a single EC2 m5.2xlarge. **120 licensed users**, roughly 85 weekly active. 220 published workbooks. 31 extract refreshes scheduled between 08:00 and 10:00 weekdays.

MONITORING

CloudWatch logs present (several groups set to never expire); **zero alarms**; no platform health dashboard.

SPEND

~£14,200/month across the platform accounts (breakdown in Section 6).

SECTION 4

Findings in detail

Nine findings carry the most weight and are written up in full below. All 21 appear in the register in Section 5. Each finding follows the same schema; field definitions are in Appendix B.

CRITICAL QUICK WIN

REL-01 Nightly finance pipeline fails silently and is discovered by Finance

SERVICE Step Functions, AWS Glue | CATEGORY Reliability | EFFORT S (2 days) | CONFIDENCE High

PACKAGE 30-day quick-win sprint

EVIDENCE

Job run history (90 days): 11 failed executions of the nightly state machine. Two of four wrapped Glue jobs have no retry configuration. CloudWatch shows no alarm on the state machine or jobs. Interview corroboration: Finance BI Lead reports discovering stale data “at the 9am meeting, roughly monthly”. Mean time from failure to detection: 4.2 hours, detected by consumers in 11 of 11 cases. (EV-014, EV-015, EV-031, INT-03)

BUSINESS IMPACT

Finance periodically presents or nearly presents stale numbers; each incident costs an estimated half day of engineering rework and measurable credibility with the board. This is the single largest driver of distrust in reporting.

LIKELY CAUSE

Pipeline built without failure-path design; no alerting culture established after contractor handover.

RECOMMENDATION

Add retry with backoff to the two unprotected jobs; add a failure catch state publishing to SNS; route to email and Teams; add a freshness check that blocks the dashboard data source and flags “data incomplete” rather than silently serving stale data.

Saving / risk reduction: risk reduction primarily; eliminates the recurring trust incident. Around £400/month of rework effort recovered.

CRITICAL QUICK WIN

OBS-01 No alerting exists on any data platform component

SERVICE CloudWatch (platform-wide) | CATEGORY Observability | EFFORT S (2–3 days) | CONFIDENCE High

PACKAGE 30-day quick-win sprint

EVIDENCE

CloudWatch alarm inventory: zero alarms in both accounts. No log metric filters. No dashboard for platform health. (EV-030, EV-031)

BUSINESS IMPACT

Every failure mode in this report is currently detected by humans downstream, if at all. The team operates blind; the business absorbs the consequences.

LIKELY CAUSE

Monitoring was deferred during build and never revisited.

RECOMMENDATION

Establish an alerting baseline: Glue job failure alarms, Step Functions execution failure alarms, Lambda error and throttle alarms, Redshift disk and queue alarms, S3 ingestion freshness checks. One SNS topic, two channels (email, Teams). Add a one-page CloudWatch dashboard as the platform’s front door.

Saving / risk reduction: risk reduction; converts every silent failure in this report into a same-hour alert.

HIGH | QUICK WIN · PARTIAL

COST-01 Nine Glue jobs reprocess full history on every run

SERVICE AWS Glue, S3 | CATEGORY Cost optimisation | EFFORT M (5–8 days, top three datasets) | CONFIDENCE High | PACKAGE Remediation Sprint, weeks 1–2

EVIDENCE

Job run metadata: 9 of 25 jobs read the full raw zone on each execution (no partition pushdown possible; raw zone unpartitioned JSON). Combined DPU consumption of these jobs is 78% of total Glue spend. Billing data: Glue ~£4,150/month. Comparable incremental design (observed in the 3 jobs already using bookmarks against partitioned Parquet) consumes ~12% of equivalent volume. (EV-006, EV-007, EV-002)

BUSINESS IMPACT

Roughly £1,420/month of avoidable DPU spend, an overnight window that grows with history, and rising failure exposure as runtimes stretch.

LIKELY CAUSE

Raw zone landed as unpartitioned JSON in the original build; incremental patterns never retrofitted.

RECOMMENDATION

Partition the three heaviest datasets by ingest date; convert raw-to-curated steps to write Parquet; enable job bookmarks; right-size DPU allocations after conversion.

Saving: £1,200–£1,600/month once the top three datasets are converted (±30%). Quick win on the top dataset only: 2 days, ~£550/month.

HIGH

COST-02 Redshift sized for a peak that rarely arrives

SERVICE Amazon Redshift | CATEGORY Cost optimisation | EFFORT M (3–5 days incl. measurement checkpoints) | CONFIDENCE Medium-High | QUICK WIN No (sequenced) | PACKAGE Remediation Sprint, weeks 2–3

EVIDENCE

4-node ra3.xlplus, 24/7, ~£2,640/month. CloudWatch: CPU below 15% outside 06:00–18:00 weekdays; WLM queueing only during the 08:00–10:00 extract window (see BI-02). No reserved instances or committed pricing. (EV-010, EV-011, EV-002)

BUSINESS IMPACT

An estimated £700–£900/month of spend buys headroom that is only consumed during a two-hour self-inflicted refresh storm.

LIKELY CAUSE

Cluster sized at build for an anticipated workload; never revisited. The morning queue pressure (BI-02) masks how oversized the cluster is the rest of the time.

RECOMMENDATION

First fix BI-02 (reschedule extracts), then re-measure for two weeks, then either resize to 2 nodes or evaluate Redshift Serverless against the observed pattern. Apply committed pricing only after right-sizing.

Saving: £700–£900/month (±30%), contingent on BI-02 landing first.

HIGH QUICK WIN

GOV-01 Long-lived access keys with administrator rights

SERVICE IAM | CATEGORY Security and governance | EFFORT S (2 days) | CONFIDENCE High

PACKAGE 30-day quick-win sprint

EVIDENCE

IAM credential report: 4 IAM users hold access keys older than 400 days; 2 of those users have AdministratorAccess attached and the keys are embedded in ETL scripts on the Tableau Server host. MFA absent on 3 console users. No key rotation policy. (EV-020, EV-021)

BUSINESS IMPACT

A leaked key from a script or host backup grants standing administrator access to the production estate. No exposure was observed, but the blast radius of this configuration is the whole platform.

LIKELY CAUSE

Scripts written quickly against user keys during the original build; roles never retrofitted.

RECOMMENDATION

Replace script credentials with an instance profile role on the Tableau host; scope a least-privilege policy for ETL actions; rotate and then delete the user keys; enforce MFA on all console users.

Risk reduction: major; removes standing admin credentials from disk.

HIGH

BI-01 Executive dashboard takes 38 seconds to load

SERVICE Tableau, Amazon Redshift | CATEGORY BI and Tableau readiness | EFFORT M (4–6 days) | CONFIDENCE High

QUICK WIN No | PACKAGE Remediation Sprint, weeks 1–2

EVIDENCE

Tableau Server admin views (screenshots provided): P95 load time 38s on the executive pack, P50 22s. Workbook inspection: live connection to Redshift, 14 quick filters, two cross-database joins, no extract. Load times peak during the 08:00–10:00 window, correlating with WLM queueing (EV-040, EV-041, EV-011). Interview: “people export to Excel because the dashboard is too slow” (INT-03).

BUSINESS IMPACT

The flagship report is effectively unusable at the moment of peak demand, driving exports to Excel, which in turn feeds the conflicting-numbers problem (BI-03).

LIKELY CAUSE

Live-connection design against a queued warehouse, compounded by filter and join complexity.

RECOMMENDATION

Convert to an extract refreshed off-peak; collapse quick filters to parameter-driven views; pre-join in the curated layer; target sub-8-second P95, which the existing hardware supports comfortably once redesigned.

Value: restores the primary reporting surface; indirect cost saving via reduced Redshift pressure.

MEDIUM QUICK WIN

BI-02 Thirty-one extract refreshes collide in a two-hour morning window

SERVICE Tableau, Amazon Redshift | CATEGORY BI and Tableau readiness | EFFORT S (1–2 days) | CONFIDENCE High | PACKAGE 30-day quick-win sprint

EVIDENCE

Tableau schedule export: 31 of 38 extract refreshes scheduled 08:00–10:00 weekdays, directly overlapping peak interactive use and the tail of the overnight load. Redshift WLM queue depth peaks in the same window (EV-042, EV-011).

BUSINESS IMPACT

Self-inflicted contention that slows dashboards (BI-01), queues analyst queries, and makes the warehouse look undersized when it is the opposite (COST-02).

LIKELY CAUSE

Default schedule choices accumulated workbook by workbook; nobody owns the refresh calendar.

RECOMMENDATION

Move heavy refreshes to 04:30–06:30 after the overnight load completes; stagger the remainder; assign ownership of the refresh calendar; document the dependency on pipeline completion (links REL-01 freshness check).

Value: unblocks BI-01 and COST-02; immediate experience improvement at zero infrastructure cost.

HIGH

MAINT-01 The platform depends on one engineer

SERVICE Platform-wide | CATEGORY Maintainability | EFFORT M (intermittent over 4 weeks, ~6 days total) | CONFIDENCE High | QUICK WIN No | PACKAGE Remediation Sprint, weeks 1–4 (running thread)

EVIDENCE

Runbook inventory: 19 of 25 Glue jobs have no runbook. The contractor handover document (2023) references 9 jobs that no longer exist and omits 11 that do. Interviews: both the sponsor and the engineer independently identified the same single person as the only holder of end-to-end knowledge (INT-01, INT-02, EV-050).

BUSINESS IMPACT

A resignation, illness or extended absence puts reporting continuity in doubt within days. This is the largest unquantified risk on the platform and it compounds every other finding.

LIKELY CAUSE

Knowledge concentrated through attrition; documentation never funded as work.

RECOMMENDATION

Run a four-week documentation sprint: one-page runbook per job (purpose, schedule, inputs, outputs, failure modes, rerun procedure), an estate map, and a recorded walkthrough. Pair this with the alerting baseline so operational knowledge lives in systems, not memory.

Risk reduction: removes the single largest continuity risk; halves onboarding time for any future hire or contractor.

MEDIUM QUICK WIN

COST-03 No S3 lifecycle management and unbounded version history

SERVICE Amazon S3 | CATEGORY Cost optimisation | EFFORT S (1 day) | CONFIDENCE High

PACKAGE 30-day quick-win sprint

EVIDENCE

Bucket inventory: 38 TB current objects, 71% untouched for 90+ days, all in Standard. Two versioned buckets hold 21 TB of non-current versions with no expiry rule. No lifecycle policies exist on any bucket. (EV-001, EV-002)

BUSINESS IMPACT

Roughly £410/month of avoidable storage spend, growing monthly as versions accumulate.

LIKELY CAUSE

Versioning enabled defensively during the build; lifecycle design never followed.

RECOMMENDATION

Apply Intelligent-Tiering to the raw and archive zones; add non-current version expiry (30 days) on the two versioned buckets; tag buckets with owners.

Saving: £350–£470/month (±30%).

SECTION 5

Findings register — all 21

QW = quick win. Effort: S ≤2 days, M 3–10 days, L >10 days. Savings are monthly estimates (±30%); "risk" denotes findings whose value is risk reduction rather than direct cost.

ID	TITLE	SEV	SERVICE	SAVING / RISK	EFFORT	QW	CONF
REL-01	Nightly finance pipeline fails silently	CRITICAL	Step Functions, Glue	risk + ~£400	S	Yes	High
OBS-01	No alerting on any platform component	CRITICAL	CloudWatch	risk	S	Yes	High
COST-01	Nine Glue jobs reprocess full history	HIGH	Glue, S3	£1,200–1,600	M	Part	High
COST-02	Redshift oversized for observed workload	HIGH	Redshift	£700–900	M	No	Med-High
GOV-01	Long-lived admin access keys in scripts	HIGH	IAM	risk	S	Yes	High
BI-01	Executive dashboard P95 load 38s	HIGH	Tableau, Redshift	trust	M	No	High
REL-02	Ingestion Lambda drops payloads on throttle, no DLQ	HIGH	Lambda	risk	S	Yes	High
MAINT-01	Single-person dependency, no runbooks	HIGH	Platform	risk	M	No	High
COST-03	No S3 lifecycle, unbounded versions	MEDIUM	S3	£350–470	S	Yes	High
COST-04	Athena scanning unpartitioned raw JSON	MEDIUM	Athena	£230–330	M	Part	High
REL-03	Non-idempotent jobs create duplicates on rerun	MEDIUM	Glue, Redshift	trust	M	No	Med
OBS-02	Log retention unmanaged, no metric filters	MEDIUM	CloudWatch	£80–120	S	Yes	High
GOV-02	No dev/test/prod separation; change straight to prod	MEDIUM	Org-wide	risk	L	No	High
GOV-03	No enforced KMS CMKs or TLS-only bucket policies; 2 buckets rely on account-level public block only	MEDIUM	S3, KMS	risk	S	Yes	Med
ENG-01	No data quality checks; schema drift broke jobs twice in 90 days	MEDIUM	Glue	risk	M	No	High
ENG-02	Seven dead pipelines still running	MEDIUM	Glue, EventBridge	£200–280	S	Yes	High
BI-02	31 extract refreshes collide 08:00–10:00	MEDIUM	Tableau, Redshift	enabler	S	Yes	High
BI-03	Two live definitions of "Jobs Completed" KPI	MEDIUM	Tableau	trust	S	No	High
COST-05	Orphaned EBS volumes, idle NAT gateways, stale snapshots	LOW	EC2, VPC	£180–260	S	Yes	High
GOV-04	CloudTrail lacks data events on sensitive buckets	LOW	CloudTrail	risk	S	Yes	Med
BI-04	47% of 220 workbooks unused in 90 days	LOW	Tableau	hygiene	S	Yes	High

SECTION 6

Cost opportunity summary

Current monthly platform spend: ~£14,200, up from ~£9,800 twelve months ago.

LINE	MONTHLY	NOTES
AWS Glue	**£4,150**	78% consumed by the 9 full-reprocess jobs
Redshift (4× ra3.xlplus)	**£2,640**	24/7, low off-peak utilisation
Dev/test account misc	**£1,460**	Mirrors prod patterns, including idle NAT
S3 (incl. 21 TB versions)	**£1,380**	No lifecycle policies
Other (orphaned EBS, snapshots, misc)	**£1,170**	Includes decommission debris
Tableau Server EC2 + EBS	**£1,120**	Single m5.2xlarge
Athena	**£870**	Unpartitioned scans
NAT and data transfer	**£640**	Two idle NAT gateways in dev
CloudWatch	**£460**	Several never-expire log groups
Lambda, Step Functions, EventBridge	**£310**	Reasonable
Total	**£14,200**	

IDENTIFIED ADDRESSABLE COST: £3,100–£3,600/MONTH (£37,000–£43,000/YEAR)

SOURCE FINDING	MONTHLY ESTIMATE	CONFIDENCE
COST-01 Glue incremental conversion	**£1,200–1,600**	High
COST-02 Redshift right-size	**£700–900**	Med-High
COST-03 S3 lifecycle and versions	**£350–470**	High
COST-04 Athena partitioning/Parquet	**£230–330**	High
ENG-02 Dead pipelines	**£200–280**	High
COST-05 + OBS-02 orphans and log retention	**£260–380**	High

High-confidence subtotal (excluding COST-02 pending re-measurement): **~£2,300/month**. All estimates are ±30%, derived only from observed billing and configuration data, and stated as ranges deliberately. They are not guarantees.

SECTION 7

30/60/90-day remediation plan

DAYS 1–30

Stop the bleeding, switch the lights on — no procurement, no architectural change

1. Alerting baseline across Glue, Step Functions, Lambda, Redshift; failure catch on the finance pipeline; freshness check on the finance data source **(OBS-01, REL-01)**
2. Rotate and remove long-lived keys; instance profile role on the Tableau host; MFA enforcement **(GOV-01)**
3. S3 lifecycle policies and version expiry **(COST-03)**; decommission the 7 dead pipelines **(ENG-02)**; clear orphaned EBS, snapshots and idle NAT **(COST-05)**; set log retention **(OBS-02)**
4. Reschedule the extract refresh calendar **(BI-02)**; add DLQ to the ingestion Lambda **(REL-02)**
5. Partition and convert the single heaviest dataset **(COST-01 first slice)**

Expected effect: £900–£1,200/month removed; every silent failure converted to a same-hour alert.

DAYS 31–60

Fix the structural cost and the flagship dashboard

6. Convert remaining heavy datasets to partitioned Parquet with bookmarks; right-size DPUs **(COST-01, COST-04)**
7. Rebuild the executive dashboard on an extract with redesigned filters **(BI-01)**
8. Add idempotency to the 6 rerun-unsafe jobs; dedupe the two affected fact tables **(REL-03)**
9. Begin the runbook sprint, one page per job **(MAINT-01)**
10. Re-measure Redshift for two clean weeks post BI-02 **(COST-02 preparation)**

DAYS 61–90

Right-size and govern

11. Redshift decision: resize to 2 nodes or move to Serverless based on measured profile; apply committed pricing after **(COST-02)**
12. KPI definition workshop and single source for “Jobs Completed”; retire the duplicate **(BI-03)**
13. Minimal viable environment separation: a test account for Glue changes with promotion via code **(GOV-02 first phase)**
14. Data quality checks on the three business-critical pipelines **(ENG-01)**; archive unused workbooks **(BI-04)**; CloudTrail data events on sensitive buckets **(GOV-04)**; KMS and TLS bucket policies **(GOV-03)**
15. Complete runbooks and estate map; recorded walkthrough **(MAINT-01)**

QUICK WINS — DELIVERABLE IN THE FIRST 30 DAYS

FINDING	EFFORT	VALUE
OBS-01 alerting baseline	2–3 days	Ends silent failures
REL-01 retries + freshness check	2 days	Ends the Finance trust incident
GOV-01 key removal	2 days	Removes standing admin credentials
COST-03 lifecycle + versions	1 day	**£350–470/month**
ENG-02 dead pipelines	1 day	**£200–280/month**
COST-05 + OBS-02 cleanup	1 day	**£260–380/month**
BI-02 refresh calendar	1–2 days	Faster dashboards, calmer warehouse
REL-02 DLQ	1 day	Stops silent data loss
COST-01 first dataset	2 days	**~£550/month**

SECTION 8

Methodology and evidence basis

This audit followed the Reesync AI Rescue Audit methodology: metadata-only evidence collection through a read-only cross-account role, structured interviews with three stakeholders, analysis against eight scored dimensions, and a findings standard requiring every claim to trace to at least one evidence artefact. Fifty-three evidence artefacts were collected and are indexed in Appendix A. No changes were made to the environment. No table contents or business records were accessed. Cost estimates derive solely from observed billing and configuration data and are expressed as ranges at ±30%.

APPENDIX A · EVIDENCE REGISTER (FORMAT AND SAMPLE ROWS)

The full register accompanies the report as a separate file.

EVIDENCE ID	SOURCE	METHOD	COLLECTED	ARTEFACT	RELATED FINDINGS
EV-002	Cost Explorer	Read-only API export, 13-month view	13 May 2026	cost-export-2026-05.csv	COST-01..05, OBS-02
EV-006	Glue job run history	get_job_runs, 90 days, all jobs	13 May 2026	glue-runs.json	COST-01, REL-01, ENG-01
EV-011	Redshift CloudWatch metrics	CPU, queue depth, 30 days	14 May 2026	redshift-metrics.json	COST-02, BI-01, BI-02
EV-020	IAM credential report	generate_credential_report	14 May 2026	credential-report.csv	GOV-01
EV-042	Tableau schedule export	Admin screenshot set	15 May 2026	tableau-schedules.png	BI-02
INT-03	Finance BI Lead interview	Structured interview, 40 min	16 May 2026	int-03-notes.md	REL-01, BI-01, BI-03

APPENDIX B · FINDING SCHEMA AND SCALES

Every finding carries: ID; title; severity; affected service; category; evidence (with artefact references); business impact; likely cause; recommendation; implementation effort (S ≤2 days, M 3–10, L >10); estimated saving or risk reduction; confidence (High: direct billing or configuration evidence; Medium: partial evidence plus inference; Low: interview-only, reported as observation not estimate); quick-win flag (≤5 days, no procurement, reversible, high confidence); suggested remediation package.

Severity: **CRITICAL** = active or imminent material harm, act within days. **HIGH** = significant cost, risk or trust impact, act within 30 days. **MEDIUM** = meaningful inefficiency or risk, plan within 90 days. **LOW** = hygiene, address opportunistically.

APPENDIX C · SCORING BANDS

1.0–1.9 Critical weakness. 2.0–2.9 At Risk. 3.0–3.9 Stable but inefficient. 4.0–4.9 Good. 5.0 Exemplary. Dimension rubrics are defined in the Reesync AI audit methodology and applied identically across clients, so scores are comparable over time and against re-audit.

APPENDIX D · ASSUMPTIONS AND LIMITATIONS

Point-in-time assessment over a 10-day window. Metadata only; data content not reviewed. Tableau assessed via questionnaire, exports and screenshots. Estimates are ranges (±30%) from observed data, not guarantees. Not a penetration test or compliance certification. Recommendations assume current workload patterns continue.

APPENDIX E · ACCESS USED AND REVOCATION

Cross-account read-only IAM role (CloudFormation-deployed, external ID protected, least privilege, no write actions). Role deleted by the client on 26 May 2026, confirmed in CloudTrail.

ReesyncAI

NEXT STEP

Get this report for your platform

A fixed-price audit of your AWS data platform: £2,500 + VAT, two weeks, scored against the same eight dimensions, with a 30/60/90-day plan your team can execute alone.

Kill clause: if the first two days show the audit will not surface material value, you pay nothing.

reesyncai.com/aws-data-platform-rescue-audit

hello@reesyncai.com