

AUDIT REPORT

GCP Cost & Performance Audit

Northwind Metrics (fictional)

● **SAMPLE:** synthetic data, illustrative only. Fictional company. Not a client engagement.

SCOPE

nw-prod, nw-data, nw-dev

REPORT VERSION

northwind-synthetic-2026-3b941c0373

ANALYSIS POLICY

0.1.0

DATA WINDOW

2026-07-21 to 2026-10-18 (UTC)

TEMPLATE / NARRATIVE

0.1.0 / templates-0.1.0 (no model)

GENERATED

2026-10-19T22:10:32Z

Recommendations only. Oculix does not access or change your systems. Estimates are not guaranteed savings.

Executive summary

! SAMPLE: synthetic data, illustrative only. Fictional company. Not a client engagement.

IDENTIFIED MONTHLY SAVINGS

\$708.78

\$8,505.36 a year

GROSS SPEND, LAST 30 DAYS

\$7,410.50

+8.1% vs \$6,852.85

FINDINGS

8

1 savings, 1 spend driver,
6 reliability or governance

WITHHELD FOR REVIEW

0

All findings published

This report covers nw-prod, nw-data and nw-dev from 2026-07-21 to 2026-10-18 (UTC), using 90 complete days of billing data. It identifies \$708.78 in estimated monthly savings (\$8,505.36 a year) from 1 finding, plus 1 spend driver and 6 reliability or governance items that carry no savings claim. Gross cost over the last 30 days was \$7,410.50 against \$6,852.85 in the 30 days before. The first decision is to review DEBUG logging in nw-prod (app-debug). Actions are recommendations only; you decide, test and implement.

Top three actions

1

F-001 DEBUG logging in nw-prod (app-debug)

Review whether DEBUG logging needs to stay on in production; if not, raise the log level or add an exclusion filter for DEBUG entries. Do not touch required audit logs.

2

F-002 Cloud Logging spend spike in nw-prod

Identify what changed on those dates (a release, a debug flag, a new source) and confirm it was intended; add an alert on daily cost for this service.

3

F-003 No monitoring coverage on nw-batch-1

Add an alert policy for the failure modes that matter and an uptime check where it is user-facing.

First 30 days: assign an owner to each top action and confirm the validation step before changing anything.


Methodology and coverage

Sources: billing export (daily cost by project, service and SKU), log volume by source, monitoring configuration, SLO configuration and the client's critical-journey list. Billing coverage: 90 complete days from 2026-07-21 to 2026-10-18 (14 required), no gaps. Currency: USD. Costs are gross (before credits).

Limitations: this slice does not analyze compute, disk, IP, snapshot, storage or commitment data, so it does not claim anything about them. Savings figures are estimates from observed cost and stated methods; spend drivers are observations, not savings. Alert policies counted: 5; metric references: 10 (a fee for these is not yet verified, so no dollar figure is given).

Spend by service (last 30 days vs previous 30)

■ Last 30 days □ Previous 30 days



Total \$7,410.50 vs \$6,852.85 (+8.1%)

Cloud Logging is highlighted because its cost rose by more than 50%. The two logging findings (F-001, F-002) explain it.

Findings register

ID	FINDING	SAVINGS / MO	OBSERVED	EFFORT	RISK	CONFIDENCE	SOURCE
F-001	DEBUG logging in nw-prod (app-debug)	\$708.78	n/a	S	● Low	Medium	LOG-0001, BQ-0544
F-002	Cloud Logging spend spike in nw-prod	n/a	\$555.22	S	● Low	High	BQ-0625, BQ-0634
F-003	No monitoring coverage on nw-batch-1	n/a	\$2,100.00	S	● Low	High	OBS-0004
F-004	No monitoring coverage on nw-web	n/a	\$410.00	S	● Low	High	OBS-0002
F-005	No monitoring coverage on nw-worker	n/a	\$230.00	S	● Low	High	OBS-0003
F-006	Cost with no team label in nw-data	n/a	\$1,573.81	S	● Low	High	BQ-0547, BQ-0548
F-007	No SLO for the search journey	n/a	n/a	M	● Low	High	JNY-0003
F-008	No SLO for the nightly export journey	n/a	n/a	M	● Low	High	JNY-0004

Savings findings are ranked by savings x confidence, then spend drivers, then reliability and governance items. Risk and confidence are shown as words and are never color-only. Observed amounts are not savings.

30/60/90-day plan

0-30 DAYS

6 ACTIONS

- **F-001** Review whether DEBUG logging needs to stay on in production; if not, raise the log level or add an exclusion filter for DEBUG entries. Do not touch required audit logs.
Check: Ingested bytes for this log fall after the change and no alert or metric depended on it.
- **F-002** Identify what changed on those dates (a release, a debug flag, a new source) and confirm it was intended; add an alert on daily cost for this service.
Check: Daily cost for this service returns to the baseline after the change is understood.
- **F-003** Add an alert policy for the failure modes that matter and an uptime check where it is user-facing.
Check: The resource shows at least one enabled alert policy (and an uptime check where applicable).
- **F-004** Add an alert policy for the failure modes that matter and an uptime check where it is user-facing.
Check: The resource shows at least one enabled alert policy (and an uptime check where applicable).
- **F-005** Add an alert policy for the failure modes that matter and an uptime check where it is user-facing.
Check: The resource shows at least one enabled alert policy (and an uptime check where applicable).
- **F-006** Label the resources that drive this cost with an owning team so spend can be assigned.
Check: The unlabeled share falls below 10% in next month's export.

31-60 DAYS

2 ACTIONS

- **F-007** Define an availability or latency SLO for this journey and alert on its burn rate.
Check: An SLO exists for the service behind this journey.
- **F-008** Define an availability or latency SLO for this journey and alert on its burn rate.
Check: An SLO exists for the service behind this journey.

61-90 DAYS

NONE

No items.

Finding details

F-001

DEBUG logging in nw-prod (app-debug)

Category cost


Priority 1

Effort S

Risk Low

Confidence Medium

First observed 2026-09-19



Observed. app-debug ingested 1433.6 GiB of DEBUG-severity logs in 30 days, 67% of the project's Cloud Logging cost. This is an identified monthly saving of \$708.78, an estimate, not a guarantee (confidence Medium).

Recommended action. Review whether DEBUG logging needs to stay on in production; if not, raise the log level or add an exclusion filter for DEBUG entries. Do not touch required audit logs.

Verification. Ingested bytes for this log fall after the change and no alert or metric depended on it.

Calculation. DEBUG bytes x observed rate, where observed rate = 30-day Cloud Logging cost / total ingested GiB for the project. Assumes all DEBUG volume could be excluded; confirm first. Sources: LOG-0001, BQ-0544, BQ-0553, BQ-0562, BQ-0571, BQ-0580.

F-002

Cloud Logging spend spike in nw-prod

Category cost

Priority 2

Effort S

Risk Low

Confidence High

First observed 2026-09-28

Observed. Cloud Logging in nw-prod ran above its normal daily cost from 2026-09-28 to 2026-10-01 (4 days); excess over the \$17.66 median day was \$555.22. The observed amount is excess spend during the spike (one-off, not a monthly figure).

Recommended action. Identify what changed on those dates (a release, a debug flag, a new source) and confirm it was intended; add an alert on daily cost for this service.

Verification. Daily cost for this service returns to the baseline after the change is understood.

Calculation. Daily gross cost by project and service; days above 1.5x the median day and more than \$50 above it are grouped into runs; excess = sum of (day cost - median day). Sources: BQ-0625, BQ-0634, BQ-0643, BQ-0652.

F-003

No monitoring coverage on nw-batch-1

Category **reliability**Priority **3**Effort **S**Risk **Low**Confidence **High**First observed **2026-10-18**

Observed. nw-batch-1 (compute_instance) has no alert policy. It costs about \$2,100.00 a month. The observed amount is the 30-day cost of the unmonitored resource.

Recommended action. Add an alert policy for the failure modes that matter and an uptime check where it is user-facing.

Verification. The resource shows at least one enabled alert policy (and an uptime check where applicable).

Calculation. Resource configuration from the monitoring export; cost from the client's own figure if supplied. Sources: OBS-0004.

F-004

No monitoring coverage on nw-web

Category **reliability**Priority **4**Effort **S**Risk **Low**Confidence **High**First observed **2026-10-18**

Observed. nw-web (cloud_run) has no alert policy. It costs about \$410.00 a month. The observed amount is the 30-day cost of the unmonitored resource.

Recommended action. Add an alert policy for the failure modes that matter and an uptime check where it is user-facing.

Verification. The resource shows at least one enabled alert policy (and an uptime check where applicable).

Calculation. Resource configuration from the monitoring export; cost from the client's own figure if supplied. Sources: OBS-0002.

F-005

No monitoring coverage on nw-worker

Category **reliability**Priority **5**Effort **S**Risk **Low**Confidence **High**First observed **2026-10-18**

Observed. nw-worker (cloud_run) has no alert policy and no uptime check. It costs about \$230.00 a month. The observed amount is the 30-day cost of the unmonitored resource.

Recommended action. Add an alert policy for the failure modes that matter and an uptime check where it is user-facing.

Verification. The resource shows at least one enabled alert policy (and an uptime check where applicable).

Calculation. Resource configuration from the monitoring export; cost from the client's own figure if supplied. Sources: OBS-0003.

F-006

Cost with no team label in nw-data

Category **governance**Priority **6**Effort **S**Risk **Low**Confidence **High**First observed **2026-09-19**

Share of nw-data cost with no team label

100%

Observed. 100% of nw-data's cost in the last 30 days carries no team label. The observed amount is the 30-day cost with no team label.

Recommended action. Label the resources that drive this cost with an owning team so spend can be assigned.

Verification. The unlabeled share falls below 10% in next month's export.

Calculation. Gross cost for billing rows with an empty team label, last 30 days. Sources: BQ-0547, BQ-0548, BQ-0556, BQ-0557, BQ-0565, BQ-0566.

F-007

No SLO for the search journey

Category **reliability**Priority **7**Effort **M**Risk **Low**Confidence **High**First observed **2026-10-18**

Observed. The search journey runs on nw-web, which has no SLO defined.

Recommended action. Define an availability or latency SLO for this journey and alert on its burn rate.

Verification. An SLO exists for the service behind this journey.

Calculation. Client's critical-journey list compared with the SLO configuration export. Sources: JNY-0003.

F-008

No SLO for the nightly export journey

Category **reliability**

Priority **8**

Effort **M**

Risk **Low**

Confidence **High**

First observed **2026-10-18**

Observed. The nightly export journey runs on nw-worker, which has no SLO defined.

Recommended action. Define an availability or latency SLO for this journey and alert on its burn rate.

Verification. An SLO exists for the service behind this journey.

Calculation. Client's critical-journey list compared with the SLO configuration export. Sources: JNY-0004.

Appendix

Source reference index

Each reference (BQ = billing row, LOG = log source, OBS = monitoring resource, SLO, JNY = journey) is the row number in the file you supplied; the evidence is kept in your bundle under the retention terms.

Glossary

SKU	Stock - keeping unit: the line item Google bills for, such as a type of storage or compute time.
SLO	Service level objective: a target for reliability, such as 99.9% of requests succeed.
Gross cost	Cost before credits and discounts are applied.
Spend driver	A cost pattern worth understanding. It is an observation, not a savings estimate.
Uptime check	A probe that tests from outside whether a service responds.
Alert policy	A rule that notifies someone when a metric crosses a threshold.

Report northwind-synthetic-2026-3b941c0373 | template 0.1.0 | analysis policy 0.1.0 | generated 2026-10-19T22:10:32Z

Recommendations only. Oculix does not access or change your systems. Estimates are not guaranteed savings.