

Examiner's Assessment

Career Provenance Analysis — Sample Report

Candidate Reference	CX-2026-0147 (Anonymised)
Role Target	Staff Software Engineer — Platform Infrastructure
Assessment Type	Full Examiner's Assessment
Date Issued	June 19, 2026
Examiner	TechCredalyst Provenance Practice

EXECUTIVE SUMMARY

This assessment evaluates the experience claims of a candidate presented for a Staff Software Engineer role on a platform infrastructure team. The analysis applies TechCredalyst's seven-dimension Career Provenance Framework against the candidate's resume, supporting materials, and the target role's job description.

The candidate's experience claims present **elevated fabrication risk**. While technical vocabulary and surface-level narrative are largely plausible, several signal patterns common to AI-assisted or embellished career narratives were detected — most notably ownership diffusion and metric orphaning. These signals do not confirm fabrication on their own, but they materially increase the probability that claimed seniority and impact exceed demonstrated reality.

We recommend proceeding to a structured validation interview using the probe set in Section 6 before extending an offer. Do not rely on standard technical or behavioural interview performance alone — interviews validate capability, not historical ownership.

PROVENANCE SCORE™	TECHNICAL MATCH	RISK CLASSIFICATION
68 / 100	6.9 / 10	ELEVATED

Not a verdict. A structured risk assessment — to be confirmed through targeted validation.

SECTION 1

Seven-Dimension Breakdown

Each dimension is scored 1–10 based on signal analysis of the candidate's submitted materials. Lower scores indicate higher fabrication risk for that specific signal category.

Temporal Plausibility

6.4 /10

Seniority progression is broadly plausible for the 9-year stated tenure, though the jump to Staff-level ownership in the final role is compressed relative to typical IC growth pacing.

Stack Saturation

4.8 /10

Candidate claims production-level depth across 9 distinct technical domains (Kubernetes, Kafka, Go, Rust, Terraform, AWS, GCP, ML pipelines, gRPC). Genuine mastery typically produces depth asymmetry — this profile does not.

Metric Coherence

5.1 /10

Resume cites 'reduced infrastructure cost by 38%' and 'improved deployment frequency 4x' without baseline figures, measurement methodology, or named attribution to a specific initiative.

Ownership Signals

4.6 /10

'Led' appears 11 times across three roles with no accompanying org structure, headcount, or decision-authority context. Ownership language is present but not evidenced.

Narrative Arc

6.8 /10

Career progression from mid-level to Staff is logically sequenced and consistent with the claimed company trajectory. No major narrative discontinuities detected.

Linguistic Register

4.3 /10

Vocabulary register is unusually consistent across all three employers and a 6-year span. Genuine career documents typically accumulate stylistic drift; this document does not.

Constraint Presence


5.0 /10

One instance of a stated tradeoff was found. Resource constraints, rollback events, or failure modes are otherwise absent from a multi-year infrastructure migration narrative.

SECTION 2

Provenance Risk Composition

The composite Provenance Score is not a flat average. Signal categories are weighted by how strongly they correlate with fabricated or embellished experience claims in prior pattern analysis.



Primary risk concentration: Ownership attribution and metric evidence. These two categories alone account for 55% of the composite Provenance Score, indicating the assessment's core concern is not technical credibility but verifiable accountability for the claimed outcomes.

SECTION 3

Fabrication Signal Flags

Three high-confidence signal patterns were detected, consistent with AI-assisted authoring or embellished career narratives. Flags identify patterns — they do not independently confirm fabrication.

VOCABULARY UNIFORMITY

Formal register remains statistically consistent across all three employers and a six-year span. Genuine career documents accumulate stylistic drift as writing habits, vocabulary, and tone evolve across roles and time — this document shows no such variation.

OWNERSHIP DIFFUSION

'Led' or equivalent ownership language appears 11 times without org structure, headcount, or accountability context in any instance. Ownership is asserted but not evidenced.

METRIC ORPHANING

Two precise impact figures ('38% cost reduction', '4x deployment frequency') are presented without baseline values, measurement methodology, or a named individual or team responsible for tracking the metric.

EXAMINER'S NOTE

Ownership language is present but diffuse across this candidate's narrative — 'led' recurs without org structure or decision authority context, and the two headline metrics lack causation chains. Taken individually, any one signal could be explained by resume-writing convention rather than fabrication. Taken together, the pattern materially increases the probability that claimed seniority and impact exceed demonstrated reality. We recommend validating ownership scope and metric attribution directly before proceeding.

SECTION 4

Role Alignment Context

This candidate was evaluated against the target role's job description. Role-specific context strengthens fabrication risk reasoning beyond resume-only analysis.

JD-Required Skill	Resume Evidence	Status
Kubernetes (production scale)	Claimed — 200-node cluster migration	Present
Terraform / IaC	Claimed — multi-region infrastructure	Present
Go or Rust (production)	Claimed — both, equivalent depth	Caution
Incident response / on-call leadership	Not addressed in resume	Missing
Cross-team architectural reviews	Not addressed in resume	Missing

Two role-critical responsibilities — incident response leadership and cross-team architectural review — are entirely absent from the resume despite being core to the Staff-level scope of this role. This gap, combined with the ownership diffusion pattern in Section 3, suggests the candidate's actual scope of authority may be narrower than the title implies.

SECTION 5

Recommended Action

- Do not proceed based on interview performance alone — interview signal and provenance signal are independent.
- Validate the candidate's actual decision-making authority in the claimed Kubernetes migration — specifically, who approved architectural changes.
- Request the baseline measurement and named owner for both cited impact metrics.
- Directly probe incident response and on-call experience, which is absent from the resume but required for this role.

SECTION 6

Validation Probe Pack

Five interview questions generated specifically from this candidate's claims and the detected signal patterns. Designed to surface the gap between claimed and demonstrated ownership — generic interview questions will not expose these patterns.

1. Your resume uses "led" 11 times across three roles. Pick the instance where you had the most direct decision-making authority — what specifically could only you decide, and who would have blocked a different choice?
2. You cite a 38% infrastructure cost reduction. What was the exact baseline before your work, who measured the outcome, and what broke or regressed during the implementation?
3. You claim production depth in both Go and Rust at equivalent levels. For whichever you consider weaker — describe a real production failure or constraint you personally navigated, not a tutorial-level scenario.
4. This role requires incident response leadership, which doesn't appear on your resume. Walk me through your most recent on-call incident — what was your specific role versus the rest of the team?
5. Describe the Kubernetes migration's most significant architectural disagreement. Who proposed the alternative, and why was your approach chosen over theirs?

FINAL EXAMINER RECOMMENDATION

**VALIDATE
FURTHER**

Do not extend an offer based on current interview performance alone. The candidate presents elevated provenance risk concentrated in ownership and metric attribution. Complete the validation probe set above before proceeding to offer stage. This assessment reflects forensic signal analysis and human professional judgment — it is not an automated hiring verdict.

This is a sample report generated to illustrate the structure and depth of a TechCredalyst Examiner's Assessment. Candidate details are fictionalised. Real assessments are reviewed by a human examiner before delivery and tailored to the specific resume and role submitted.