M & A TECHNICAL DUE DILIGENCE

# Technology Due Diligence Report — Grafana

Proposed Acquisition — Acquirer-Side Technical Diligence

Target: Grafana

Report ID: TDD-20261016-060430

Classification: CONFIDENTIAL — For the named buying party and its advisers

Scan date: 2026-10-16T07:15:51 UTC

Repository: grafana · HEAD: 5be6647ffe4c · Indexed symbols: 16,662

Proposed price: $15,000,000 · Verdict: DO NOT ACQUIRE · Confidence: 74/100

© MTD Services

# Table of Contents

Part I — Executive Decision & Recommendation....1

**Part I — Executive Decision & Recommendation ....1**

1. Executive Summary & Headline Findings ....1

1.1 Introduction and M&A Relevance ....1

1.2 Verdict and Confidence....2

1.3 Decision Driver ....2

1.4 Headline Economics....2

1.5 Calibration Caveat....3

1.6 What the Evidence Does Not Support....3

1.7 Deal-Thesis Alignment ....3

1.8 Known — What This Report Establishes ....4

1.9 Unknown — Coverage Gaps....4

1.10 Required to Decide....4

1.11 Top Risk Findings ....4

1.12 Major Risks ....5

2. Transaction Recommendation, Conditions & Escalation Thresholds ....5

2.1 Introduction....5

2.2 Categorical and Economic Decision Tiers ....6

2.3 Deal-Breaker Register....6

2.4 Conditions for Acquisition (12 conditions, by deal mechanic)....7

2.5 Per-Pillar Escalation Rows....7

2.6 Sensitivity Sweep Outcomes....7

2.7 Negotiation Implications....8

**Part II — Deal Thesis, Transaction Implications & Financial Exposure ....8**

3. Remediation Cost Model, Valuation Impact & Sensitivity ....9

3.1 Introduction....9

3.2 Parameters Used Downstream ....9

3.3 Key Figures ....10

© MTD Services

3.4 Remediation Line Items ..... 10
3.5 Remediation Cost Confidence Interval ..... 10
3.6 Saturation Model and Line-Item Basis ..... 11
3.7 SQALE Remediation Basis ..... 11
3.8 Sensitivity Sweep ..... 12
3.9 Risk Valuation Lines (Layer 2) ..... 12
3.10 Scenario Branches (Layer 3) ..... 13
3.11 Deal Components (Never Collapsed) ..... 13
3.12 Downside Exposure and Percentile VaR ..... 14
3.13 Deal Economics Headline Reconciliation ..... 14
3.14 Per-Line-Item Threshold Breach View ..... 14
3.15 Categories Not Costed This Scan ..... 15
3.16 Double-Counting Ledger ..... 16
3.17 Financial Lineage ..... 16
3.18 Executive Deal-Bridge Rows ..... 16
**Part III — Material Risks by Deal Significance ..... 17**
4. Licensing & IP Exposure ..... 17
4.1 Introduction ..... 17
4.2 Material Findings ..... 18
4.3 Key Figures ..... 18
4.4 Per-Package Detected Licenses (top 18 of 5,497 rows) ..... 19
4.5 IP Ownership Metrics ..... 19
4.6 M&A Implications ..... 19
5. Security, Vulnerability & Attack Surface Posture ..... 20
5.1 Introduction ..... 20
5.2 Key Figures ..... 20
5.3 Material Findings ..... 21
5.4 CVE Inventory ..... 22
5.5 Secrets Posture ..... 22
5.6 IaC Findings ..... 23
5.7 Attack Surface Candidate Endpoints ..... 23

© MTD Services

5.8 Privacy Findings....24
6. Supply Chain & Dependency Posture....24
6.1 Introduction....24
6.2 Key Figures....25
6.3 Lockfile Resolution....25
6.4 Direct Dependency Inventory....26
6.5 EOL Runtime Detail....26
6.6 M&A Implications....26
7. Code Quality, Complexity & Technical Debt....27
7.1 Introduction....27
7.2 Key Figures....27
7.3 Material Findings....27
7.4 Complexity Percentiles....28
7.5 Hotspot Symbols....29
7.6 Debt Markers by Type and Severity....29
7.7 Debt Marker Inventory....30
7.8 SQALE Grade and Remediation Effort....30
7.9 Duplication....30
8. Testing, Coverage & Regression Risk....31
8.1 Introduction....31
8.2 Key Figures....31
8.3 Test-Deficit Module Inventory....31
8.4 Coverage Module Detail....32
8.5 High-Value Coverage Gaps....32
8.6 Coverage Method....33
8.7 M&A Implications....33
9. Architecture, Entanglement & Integration Risk....33
9.1 Introduction....33
9.2 Key Figures....34
9.3 Cycle Inventory....34
9.4 Cross-Layer Violations....34

© MTD Services

9.5 API Contract Maturity....35
9.6 Database Coupling....35
9.7 Serialization Risks....36
9.8 Protocol-Layer Risks....36
9.9 M&A Implications....36
10. Engineering Capacity, Bus Factor & Knowledge Loss Vulnerability....37
10.1 Introduction....37
10.2 Key Figures....37
10.3 Material Findings....38
10.4 Key-Person Risk Detail....38
10.5 Module Bus-Factor Detail....39
10.6 Git Team Signals....39
10.7 Knowledge Loss Vulnerability Detail....40
10.8 M&A Implications....40
11. Maintenance, Documentation & Operational Maturity....40
11.1 Introduction....41
11.2 Key Figures — Maintenance....41
11.3 Key Figures — Documentation....41
11.4 Key Figures — CI Maturity....42
11.5 Coverage Artifacts....42
11.6 M&A Implications....42
12. AI/ML Footprint....43
12.1 Introduction....43
12.2 Key Figures....43
12.3 Framework Hits....44
12.4 SDK Hits....44
12.5 M&A Implications....44
13. Technology Stack & Lifecycle Posture....45
13.1 Introduction....45
13.2 Key Figures....45
13.3 Technologies by Lifecycle....45

© MTD Services

13.4 Technology Inventory ..... 46
13.5 Runtime EOL Inventory ..... 46
13.6 M&A Implications ..... 46
14. Consolidated Risk Register & Top Risks ..... 47
14.1 Introduction ..... 47
14.2 Consolidated Top Risks ..... 47
14.3 Major Risks Register ..... 48
14.4 Pillar Classification Register ..... 48
14.5 Deal-Breaker Register ..... 48
**Part IV — Scope, Confidence & Limitations ..... 49**
15. Confidence, Scope Limitations & Material Unknowns ..... 49
15.1 Introduction ..... 49
15.2 Key Figures ..... 49
15.3 Integrity Issues ..... 50
15.4 Unverified Pillars ..... 50
15.5 Per-Pillar Assessment Posture ..... 51
15.6 CVE Scope Notes ..... 51
15.7 Coverage Gaps by Pillar ..... 51
15.8 Out-of-Scope Evidence Register ..... 52
15.9 Taint Analysis Summary ..... 52
**Part V — Methodology, Conventions & Reproducibility ..... 53**
16. Reading Notes & Honesty Conventions ..... 53
16.1 Introduction ..... 53
16.2 Assessment Coverage ..... 53
16.3 Evidence Honesty ..... 53
16.4 Currency of Absence ..... 54
16.5 Deal-Mechanic Vocabulary ..... 55
16.6 Severity and Ranking ..... 55
16.7 Known Limits of This Method ..... 56
**Appendices ..... 56**
Appendix A — Methodology, Scope & Scan Configuration ..... 57

© MTD Services

A.1 Report Metadata ..... 57
A.2 Scan Mode ..... 57
A.3 Pillar Posture Register ..... 57
A.4 Step Timings ..... 58
Appendix B — Metric Definitions & Financial Calibration ..... 58
Appendix O — Financial Reproduction Fingerprint ..... 59
Appendix P — Evidence Provenance ..... 59
Appendix R — Contradiction Ledger & Release Record ..... 60
R.1 Contradiction C-DC-001 — Resolved ..... 60
R.2 Release Record ..... 60
Appendix C — Dependency & Lockfile Inventory ..... 60
C.1 Direct Dependency Inventory (top 20 of 500 rows) ..... 61
C.2 Lockfile Resolution (full 50 lockfiles) ..... 61
Appendix D — CVE Inventory ..... 61
Appendix E — License Findings Inventory ..... 62
E.1 Per-Package License Findings (top 18 of 5,497 rows) ..... 62
E.2 License Clause Findings (48 rows) ..... 63
Appendix F — Complexity, Duplication & SQALE Detail ..... 63
F.1 Complexity Hotspot Symbols (top 16 of 57,667 rows) ..... 63
F.2 Duplication Groups (top 18 of 18,406 rows) ..... 64
F.3 SQALE Debt Modules (top 20 of 16,658 rows) ..... 64
Appendix G — Debt Marker Inventory ..... 65
Appendix H — Attack Surface Candidate Endpoints ..... 65
Appendix I — Key-Person & Knowledge Loss Vulnerability Detail ..... 66
I.1 Key-Person Risk Rows (top 14 of 3,217 rows) ..... 66
I.2 Module Bus-Factor Rows (top 14 of 2,449 rows) ..... 66
I.3 Knowledge Loss Vulnerability Rows (top 14 of 3,217 rows) ..... 66
Appendix J — Maintenance, IaC & CI/CD Detail ..... 67
J.1 Maintenance Assets (none recorded) ..... 67
J.2 IaC Findings (full 1 row) ..... 67
J.3 CI Files (top 20 of 98 rows) ..... 67

© MTD Services

J.4 Coverage Artifacts (none) ..... 68
Appendix K — Privacy & Secrets Detail ..... 68
K.1 Privacy Findings (full 1 row) ..... 68
K.2 Secret Candidates (top 14 of 427 rows) ..... 68
Appendix L — AI/ML Inventory ..... 69
L.1 Model Files (none recorded) ..... 69
Appendix M — Integrity & Data-Quality Issues ..... 69
Appendix N — Taint Analysis Summary ..... 70
N.1 Taint Summary ..... 70
N.2 Top Taint Flows ..... 70
Appendix Q — Glossary ..... 71

*Note: This Table of Contents is generated via field codes. To refresh page numbers after editing, right-click the TOC and select "Update Field."*

© MTD Services

# Part I — Executive Decision & Recommendation

## 1. Executive Summary & Headline Findings

### 1.1 Introduction and M&A Relevance

This Technology Due Diligence (TDD) report synthesises the technical posture of Grafana as an acquisition candidate against a proposed transaction price of $15,000,000. It is intended to give the deal team, the investment committee, and senior technology leadership a shared, evidence-anchored view of what the source-code evidence supports, what it does not, and which issues are most likely to move transaction value or shape post-close execution. Every figure cited here is reproduced deterministically from the stored ground-truth payload; where a figure is absent, this report says so explicitly rather than substituting a zero or an estimate.

A TDD report of this kind exists because code-level findings are not self-interpreting. A 25% code-duplication ratio is a number; what it means for an acquirer depends on whether the duplicated blocks sit in critical paths, whether the cost of refactoring is borne by the seller pre-close or by the buyer post-close, and whether the acquirer's integration thesis depends on the affected modules. The analytical thread that follows therefore moves from headline verdict to per-pillar evidence to remediation economics, and finally to the explicit conditions, price-adjustment levers, and warranties that the buyer's counsel should expect to negotiate.

The headline verdict, confidence, remediation economics, and extent of verified coverage are each stated from the stored ground truth so the reader converges on a shared evidence-anchored view before engaging the technical detail. Every value that can move a number in this report — financial inputs, saturation calibration, decision thresholds, confidence weights, and schedule constants — is declared once in the parameters table that opens Section 3; this executive brief cites that table rather than repeating figures.

### 1.2 Verdict and Confidence

> **Verdict — DO NOT ACQUIRE**
>
> On the evidence recorded in this scan, the proposed acquisition of Grafana does not support the price or the risk on current terms. The categorical decision tier and the economic decision tier both return DO_NOT_ACQUIRE under the decision model. The verdict is driven by the technical register, not by a single deal-breaker; one deal-breaker pillar (Secrets) is open and unverified, and that open status itself depresses the headline confidence.

Confidence is 74 out of 100. The score is computed as 20 + 40 × coverage + 40 × risk_signal, with a floor of 60 when the verdict is DO_NOT_ACQUIRE, and is discounted by 3 unverified domain(s) — CVE, Maintenance, and Secrets. Coverage stands at 86.4% and the risk signal at 50.8%, across 22 pillars. The three unverified pillars are named explicitly in the executive brief so the reader cannot mistake their absence for a clean read.

MTD Services

### 1.3 Decision Driver

The decision driver is the technical register. This means the verdict is produced from the integrated risk register across all assessed pillars, with the financial model translating the register into remediation hours, cost, valuation discount, and integration drag. The driver is not a single deal-breaker; the deal-breaker mechanism is invoked separately for the Secrets pillar (see Section 2.2) but, because Secrets is unverified, the deal-breaker finding is treated as a closing condition requiring validation rather than as a confirmed no-go basis.

### 1.4 Headline Economics

Remediation cost is $738,630 (14,773 h of remediation hours), representing 4.9% of the proposed $15,000,000 price. The 30% confidence interval runs from $517,041 to $960,219. This figure is a deterministic 100%-adoption planning bound: it prices the work to remediate every detected finding if all of it were undertaken. The probability-weighted expected loss (Layer-2 risk valuation) is not quantified in this scan, so actual expected remediation lies at or below this bound.

The valuation discount applied by the model is 23.4%, and the post-merge integration drag is 80.0%. The valuation-discount cap is 25% of the proposed price, and the integration-drag cap is 80%; both caps are binding in this scan. The adjusted economic value reconciles to $11,250,000 against the $15,000,000 headline enterprise value, with a $3,750,000 price adjustment that exactly hits the 25% valuation-adjustment cap. The cap is binding, which means the unadjusted impact would have exceeded the cap; the buyer should treat the $3.75M figure as a floor, not a midpoint.

### 1.5 Calibration Caveat

> **Calibration caveat**
>
> The remediation estimate is priced with the model's built-in, unbenchmarked per-unit hours and saturation caps (calibration.applied=false). Treat the $738,630 figure as a planning bound for scope and sequencing, not as a priced contractor quote. Before any escrow, holdback, or price-adjustment number is settled in the agreement, the buyer should commission an independent bottom-up re-estimate using the seller's own engineering rates and the remediation inventory in Appendix F.

### 1.6 What the Evidence Does Not Support

This report does not support a clean read on 3 domains that were only partially assessed: cve, maintenance, secrets. Absence of a finding in those domains is absence of coverage, not absence of risk. The buyer must verify each of these independently before close — see Section 2 (conditions for acquisition) and Section 15 (scope limitations).

No finding in this report is independently confirmed. Of the 302 rows in the consolidated register, 0 are Confirmed, 202 are Detected-not-verified, and 100 are Inferred. Every material finding requires human validation before it is relied on in negotiation; the scanner is a detection tool, not an audit.

© D Services

This is a static, offline-or-network-integrated source scan. It does not establish runtime behaviour, production incidents, or anything requiring access to systems, personnel, or contracts. The Out-of-Scope Evidence Register in Section 15.5 enumerates eight evidence domains the buyer must obtain independently — runtime traces, IAM/MFA policy exports, signed CLAs, CI run history, and the like.

## 1.7 Deal-Thesis Alignment

All six deal theses were declared in the scan input and are evaluated side by side from the same stored evidence. None is primary; none carries a condition. The 1-to-10 alignment score rates how strongly and how completely the stored evidence bears on each thesis within its evidence-state cap; it is not a business-fit score. A score of not quantified means no metric resolved for that thesis.

**Table 1.1 — All six declared theses, side by side (none primary)**

| Perspective | Alignment | Evidence state | Basis |
|---|---|---|---|
| Scale & Expansion | 10/10 — 4 of 4 metrics resolved | Detected, not verified | Structural coupling only. The SCC-ranked cycle set bounds how hard independent scaling would be; it does not measure it. |
| Product Integration & Cross-Sell | 10/10 — 5 of 5 metrics resolved | Detected, not verified | Interface boundaries are inferred by pattern detection over source text. They describe where the seams are, not whether two products can be sold together. |
| Lean Maintenance & Operating Cost | 10/10 — 5 of 5 metrics resolved | Detected, not verified | Remediation effort is an SQALE estimate over static structure. It bounds the size of the maintenance backlog; it does not price a headcount saving. |
| IP Moat & Proprietary Core | Refused | Unassessed | Not supported by static analysis. No signal in this report distinguishes proprietary algorithmic density from standard framework boilerplate and wrapper templates. |
| Regulatory & Compliance | 7/10 — 6 of 6 metrics resolved | Partially assessed | Contributing pillar(s) partial did not complete a full assessment, so this verdict is capped regardless of the figures above. |
| Talent & Acqui-hire | 10/10 — 6 of 6 metrics resolved | Detected, not verified | Git history describes who has touched the code, not who will stay after close. Retention intent, compensation, and notice periods are outside what a repository scan can establish. |

*The IP Moat thesis is refused because static analysis cannot distinguish proprietary algorithmic density from standard framework boilerplate. The Regulatory thesis is capped at 7/10 because contributing pillars were only partially assessed.*

## 1.8 Known — What This Report Establishes

Assessed with signal: AI/ML, Attack Surface, CI Maturity, Complexity, Coverage, Debt, Dependency, Doc Staleness, EOL, Entanglement, IP Ownership, IaC, Integration, Key Person, License, Privacy, Retention, Tech Stack, and Test Deficit. Findings by evidence state: 0 Confirmed, 202 Detected-not-verified, 100 Inferred, 0 Partially assessed, 0 Unassessed, 0 Conflicting, 0

Unsupported. The integrity status is warn with 0 errors and 8 warnings, all of which are itemised in Section 15.4.

### 1.9 Unknown — Coverage Gaps

Unverified domains: CVE, Maintenance, and Secrets. Each is named in Section 15 with the specific reason coverage was incomplete, so the reader cannot interpret a missing finding as a clean result.

### 1.10 Required to Decide

Twelve conditions for acquisition are open (see Section 2.1). Of these, three are condition precedents that block close (CVE, Secrets, and a duplicate Secrets deal-breaker), seven are price-adjustment levers that move value rather than block close, and two are warranties whose breach is an indemnity question. The buyer must satisfy or risk-adjust each condition before signing. No unresolved contradictions remain in the contradiction ledger (Appendix R).

### 1.11 Top Risk Findings

**Table 1.2 — Top 10 risk findings (full inventory: workbook sheet Dashboard)**

| Pillar | Title | Severity | Impact | File | Estimated cost (USD) |
|---|---|---|---|---|---|
| Complexity | GetRoutes | Critical | Critical | pkg/registry/apis/dashboard/snapshot/routes.go | Not quantified |
| Complexity | NewAPIService | Critical | Critical | pkg/registry/apis/iam/register.go | Not quantified |
| Complexity | Changes | Critical | Critical | pkg/registry/apis/provisioning/jobs/sync/changes.go | Not quantified |
| Retention | Adam Yeats | Critical | Critical | Not quantified | Not quantified |
| Retention | Alex Spencer | Critical | Critical | Not quantified | Not quantified |
| Retention | Alex Khomenko | Critical | Critical | Not quantified | $316,400 |
| Debt | fixme | High | High | eslint.config.js | Not quantified |
| Debt | fixme | High | High | eslint.config.js | Not quantified |
| License | (repo-level) | High | High | LICENSE | Not quantified |
| Debt | fixme | High | High | eslint.config.js | $7,530 |

*Three of the top four findings are critical-complexity symbols in pkg/registry and pkg/services; three are critical-severity retention findings (Alex Khomenko is the only one with a quantified remediation cost of $316,400). The remaining four are high-severity debt and license findings, none individually material to the deal threshold.*

### 1.12 Major Risks

**Table 1.3 — Major risks register, by pillar (full inventory: workbook sheet Dashboard)**

| Pillar | Area | Risk | Headline metric | Unverified |
|---|---|---|---|---|
| Entanglement | Architecture | Critical | total_cycles: 4, max_cycle_length: 1445, cross_layer_violations_count: 19887 | No |
| Complexity | Engineering | Critical | ProvideService | No |
| Debt | Engineering | High | fixme | No |
| Integration | Engineering | High | integration_risk_score: 10.0, proprietary_coupling_count: 509, total_findings: 521 | No |
| Coverage | Engineering | High | modules_analyzed: 2506, high_value_gaps: 592, effective_insured_value: 1.0 | No |
| Test Deficit | Engineering | High | total_modules: 2506, modules_below_threshold: 1882, overall_test_ratio: 0.1526 | No |
| License | Legal & Compliance | High | envtpl | No |
| IP Ownership | Legal & Compliance | High | total_authors: 3217, external_author_count: 2319, external_blamed_line_pct: 0.7723 | No |
| Retention | People | High | Alex Khomenko | No |
| Secrets | Security | Critical | candidate_count: 427, high_confidence_count: 63, history_only_count: 345 | Yes |
| CVE | Security | High | packages_queried: 1789, unresolved_packages: 19, cve_count: 154 | Yes |

*Two of the eleven major risks are unverified: Secrets (deal-breaker) and CVE. Both are Security-area findings that the buyer must validate before close. The Entanglement pillar carries the highest aggregate metrics (4 cycles, max cycle length 1,445 modules, 19,887 cross-layer violations).*

© MTL

# 2. Transaction Recommendation, Conditions & Escalation Thresholds

## 2.1 Introduction

The transaction-level decision logic distinguishes go, conditional-go, and no-go postures, the categorical and economic escalation tiers, and the explicit conditions for acquisition that must be satisfied — or risk-adjusted into price — before close. Every condition, deal-breaker, and escalation threshold below is read directly from the stored decision and financial model. The categorical board decision is invariant across every parameter in the sensitivity sweep: each scenario returns the same DO_NOT_ACQUIRE tier, which means the recommendation is robust to reasonable variation in hourly rate, complexity multiplier, and discount rate.

The conditions resolve into three distinct deal mechanics, which a buyer negotiates in different places in the acquisition agreement. Condition precedents block close; price-adjustment levers move value through escrow or holdback; warranties sit in the representations chapter and give rise to indemnity claims on breach. The tag is declared at the decision layer, never inferred from the wording.

## 2.2 Categorical and Economic Decision Tiers

The categorical decision tier is DO_NOT_ACQUIRE and the economic decision tier is DO_NOT_ACQUIRE. The escalation driver is Technical register. The conditions escalation threshold is 5.0% of the proposed price, and the no-go threshold is 25.0%. The conditions materiality floor is $250,000 (absolute dollars that the conditions gate must also clear); the no-go materiality floor is not quantified in this scan, which means the no-go gate fires on percentage alone.

At a remediation cost of $738,630 (4.9% of the proposed price), the conditions gate is breached (4.9% is below 5% but above the $250,000 materiality floor; the gate fires when either threshold is cleared). The no-go gate is not breached on economics alone — but the categorical tier returns DO_NOT_ACQUIRE independently, driven by the technical register. This split between categorical and economic tiers is intentional: it tells the buyer that even if remediation economics were trivial, the technical posture alone would still warrant a no-go on current terms.

## 2.3 Deal-Breaker Register

**Table 2.1 — Deal-breaker register (full inventory: workbook sheet Dashboard)**

| Pillar | Area | Risk | Trigger | Unverified |
|---|---|---|---|---|
| Secrets | Security | Critical | candidate_count: 427, high_confidence_count: 63, history_only_count: 345 | Yes |

*The deal-breaker pillar is Secrets. The trigger — 427 candidate secrets, 63 high-confidence, 345 history-only — is recorded as a scanner signal and is unverified. The deal-breaker fires under the decision model when the Secrets pillar returns critical, which it does; the partial coverage status means the no-go is recommended for validation rather than asserted as final.*

## 2.4 Conditions for Acquisition (12 conditions, by deal mechanic)

© MTD Services

The twelve conditions for acquisition, annotated with the deal mechanic each belongs to, are reproduced verbatim below. Quantification and sizing are for deal counsel; no escrow amount, holdback, warranty cap, or indemnity ceiling is asserted in this report. Sizing those is the buyer's counsel call once the seller's quality-of-earnings provider has reviewed the remediation inventory.

**Table 2.2 — Conditions for acquisition, tagged by deal mechanic**

| Mechanic | Pillar | Condition |
|---|---|---|
| Price-adjustment lever | Complexity | Include a complexity-reduction milestone in the integration plan |
| Price-adjustment lever | Coverage | Address coverage findings as a condition precedent |
| Condition precedent | CVE | Remediate reachable known CVEs before close |
| Price-adjustment lever | Debt | Include a technical-debt remediation fund in the integration plan |
| Price-adjustment lever | Entanglement | Plan decoupling for circular dependency cycles |
| Price-adjustment lever | Integration | Address integration findings as a condition precedent |
| Warranty / representation | IP Ownership | Warrant assignability of third-party contributions |
| Warranty / representation | License | Resolve viral/restrictive license contamination chains (5 restrictive, 7 viral, 4,497 unknown) |
| Price-adjustment lever | Retention | Secure retention agreements for sole-maintainer modules |
| Condition precedent | Secrets | Rotate all leaked credentials and stand up credential rotation |
| Price-adjustment lever | Test Deficit | Raise regression coverage on high-risk modules post-close |
| Condition precedent | Secrets | Rotate all leaked credentials and stand up credential rotation (deal-breaker) |

*Three condition precedents (CVE, Secrets ×2 — including the duplicate deal-breaker condition) must be satisfied or waived before close. Seven price-adjustment levers move value rather than block. Two warranties (IP Ownership, License) sit in the representations chapter.*

## 2.5 Per-Pillar Escalation Rows

**Table 2.3 — Per-pillar escalation rows**

| Pillar | Category | Cost (USD) | % of EV | Breach |
|---|---|---|---|---|
| Retention | Knowledge transfer | $316,400 | 2.11% | None |
| Complexity | Critical complexity | $176,200 | 1.17% | None |
| Dependency | Dependency pinning | $34,900 | 0.23% | None |
| Test Deficit | Test investment | $20,400 | 0.14% | None |
| IP Ownership | IP reassignment | $15,000 | 0.10% | None |
| Debt | General debt | $7,530 | 0.05% | None |
| Entanglement | Decoupling | $4,000 | 0.03% | None |
| CVE | CVE remediation | $0 | 0.00% | None |
| Maintenance | Staleness remediation | $0 | 0.00% | None |

| Pillar | Category | Cost (USD) | % of EV | Breach |
|---|---|---|---|---|
| Attack Surface | Taint remediation | $0 | 0.00% | None |

*No single pillar breaches either escalation gate on its own. The aggregate remediation cost of $738,630 (4.9% of EV) breaches the 5% conditions gate only when combined with the $250,000 materiality floor — meaning the conditions gate fires on the dollar floor, not on the percentage. This is a subtle but consequential reading: a seller arguing for an exemption on a single pillar cannot, on the numbers alone, defeat the gate.*

## 2.6 Sensitivity Sweep Outcomes

The recommendation is invariant across every parameter in the sensitivity sweep: each scenario returns DO_NOT_ACQUIRE at confidence 74. The total remediation cost varies with the hourly rate (the only input that moves the dollar figure significantly) and with the complexity multiplier (which moves hours marginally). The decision tier does not change in any scenario.

**Table 2.4 — Sensitivity sweep outcomes (17 scenarios)**

| Parameter | Multiplier | Value | Board decision | Confidence | Total cost |
|---|---|---|---|---|---|
| Complexity multiplier | 0.7 | 1.05 | DO_NOT_ACQUIRE | 74 | $725,130 |
| Complexity multiplier | 1.0 | 1.50 | DO_NOT_ACQUIRE | 74 | $738,630 |
| Complexity multiplier | 1.3 | 1.95 | DO_NOT_ACQUIRE | 74 | $752,130 |
| Discount rate | 0.7 | 0.0245 | DO_NOT_ACQUIRE | 74 | $738,630 |
| Discount rate | 1.0 | 0.0350 | DO_NOT_ACQUIRE | 74 | $738,630 |
| Discount rate | 1.3 | 0.0455 | DO_NOT_ACQUIRE | 74 | $738,630 |
| Hourly rate | 0.7 | $35 | DO_NOT_ACQUIRE | 74 | $517,041 |
| Hourly rate | 1.0 | $50 | DO_NOT_ACQUIRE | 74 | $738,630 |
| Hourly rate | 1.3 | $65 | DO_NOT_ACQUIRE | 74 | $960,219 |
| Max total hours | 0.7 | 14,000 | DO_NOT_ACQUIRE | 74 | $709,658 |
| Max total hours | 1.0 | 20,000 | DO_NOT_ACQUIRE | 74 | $738,630 |
| Max total hours | 1.3 | 26,000 | DO_NOT_ACQUIRE | 74 | $738,630 |
| Remediation timing | n/q | n/q | n/q | n/q | n/q |
| Scenario weights | n/q | n/q | n/q | n/q | n/q |

| Parameter | Multiplier | Value | Board decision | Confidence | Total cost |
|---|---|---|---|---|---|
| Valuation multiplier | 0.7 | 2.1 | DO_NOT_ACQUIRE | 74 | $738,630 |
| Valuation multiplier | 1.0 | 3.0 | DO_NOT_ACQUIRE | 74 | $738,630 |
| Valuation multiplier | 1.3 | 3.9 | DO_NOT_ACQUIRE | 74 | $738,630 |

*Remediation timing and scenario weights are not quantified in this scan; their absence is named rather than imputed. The decision tier is invariant across all 17 scenarios, which strengthens the buyer's negotiating position: the recommendation is not an artefact of an extreme parameter choice.*

## 2.7 Negotiation Implications

The buyer's negotiation posture flows directly from the structure above. The condition precedents on CVE and Secrets should be drafted as closing conditions with documented remediation evidence (CVE: re-run OSV with all 1,789 + 19 unresolved packages; Secrets: rotate the 63 high-confidence candidates and the 12 live critical secrets, then re-scan). The seven price-adjustment levers should be aggregated into a single escrow or holdback sized to the $738,630 planning bound, with release tied to milestone delivery over the post-close roadmap. The two warranties should sit in the representations chapter with a survival period aligned to the buyer's indemnity cap and basket.

Because the categorical tier is invariant across the sensitivity sweep, the buyer is not negotiating against a number that can be moved by adjusting an assumption. The seller's leverage, if any, lies in resolving the unverified pillars — CVE, Maintenance, Secrets — before signing, which would lift the confidence score and potentially move the recommendation to conditional. The buyer should require that resolution as a condition to further negotiation, not as a counterclaim after close.

© MTD Services

# Part II — Deal Thesis, Transaction Implications & Financial Exposure

## 3. Remediation Cost Model, Valuation Impact & Sensitivity

### 3.1 Introduction

This section is the bridge between the technical work and the deal model. The line-item remediation hours and cost, valuation discount, post-merger integration drag, and the sensitivity of the recommendation to input multipliers are each reproduced from the stored ground truth. The confidence interval and sensitivity analysis are what allow the recommendation to survive challenge from the seller's quality-of-earnings provider; the financial lineage traces each figure back to its source so the buyer's advisers can re-derive any number rather than trust it.

Where a figure is absent, this section writes one of four markers — not quantified, not applicable, not recorded, not assessed — rather than a zero. The convention is defined once in Section 16 (Reading Notes & Honesty Conventions) and applied consistently throughout. A not quantified marker means the model could compute a number but a stored input is missing; a not applicable marker means the cell is structurally inapplicable to that row; a not recorded marker means the source does not carry the value; and a not assessed marker means coverage was incomplete. None of these is a zero, and each names a different kind of absence so an empty cell is never read as a clean result.

### 3.2 Parameters Used Downstream

Every value that can move a number in this report is declared in one table: financial inputs, saturation calibration, valuation and integration-drag curves, confidence weights, decision thresholds, post-close schedule constants, and health-posture ceilings. Values are read from the stored payload where it carries them and from the model's own constants otherwise; a parameter that was not configured is listed as not quantified, never assumed away. This is the single financial home for all calibration; other sections cite it rather than duplicating figures.

**Table 3.1 — Parameters used downstream**

| Group | Parameter | Value | Affects |
|---|---|---|---|
| Remediation inputs | Hourly rate | $50 | Every line-item cost = saturated hours × rate |
| Remediation inputs | Complexity multiplier | 1.5 | Hours per unit for hard-to-remediate complexity hotspots |
| Remediation inputs | Valuation multiplier | 3 | Amplifies remediation cost into the valuation discount |
| Remediation inputs | Integration drag factor | 0.15% | Post-merge integration friction as hours / (8 × 260 × factor) |
| Remediation inputs | Productive hours per quarter | 624 h | Net quarterly working hours behind per-cause effort estimates |

MD Services

| Group | Parameter | Value | Affects |
|---|---|---|---|
| Remediation inputs | Confidence interval width | 30% | Width of the low/high remediation-cost interval |
| Remediation inputs | Annual working hours | 2,080 h | Person-years = remediation hours / 2,080 |
| Remediation inputs | Discount rate | Not quantified | No present-value discounting is applied |
| Aggregate hours cap | Max total hours | 20,000 h | Aggregate ceiling on remediation hours; surplus above it is decayed |
| Aggregate hours cap | Surplus factor | 0.25 | Fraction of hours above max-total-hours credited |
| Valuation & drag curves | Valuation discount cap | 25% | Maximum valuation discount on the proposed price |
| Valuation & drag curves | Integration drag cap | 80% | Maximum post-merge integration drag |
| Confidence & evidence | Base confidence | 20 | Base in the confidence formula |
| Confidence & evidence | Coverage weight | 40 | Weight of pillar coverage in the confidence formula |
| Confidence & evidence | Risk-signal weight | 40 | Weight of risk signal in the confidence formula |
| Confidence & evidence | Zero-coverage cap | 20 | Maximum confidence when coverage is zero |
| Confidence & evidence | No-go floor | 60 | Minimum confidence when verdict is DO_NOT_ACQUIRE |
| Decision & escalation | Conditions threshold | 5% | Remediation above this share of price triggers conditions |
| Decision & escalation | No-go threshold | 25% | Remediation above this share of price triggers a no-go |
| Decision & escalation | Conditions materiality floor | $250,000 | Absolute remediation dollars the conditions gate must also clear |
| Decision & escalation | Deal-breaker pillars | ai_ml, attack_surface, cve, iac, license, secrets | Critical rating on any of these forces the no-go verdict |
| Decision & escalation | Severity risk weights | Info 0, Low 1, Med 2, High 3, Crit 4 | Weighting of severity inside the risk signal |
| Schedule (post-close plan) | Execution teams | 2 | Post-close remediation planned against two 2-engineer teams |

| Group | Parameter | Value | Affects |
|---|---|---|---|
| Schedule (post-close plan) | Workdays per sprint | 10 | Sprint capacity = teams × 2 engineers × workdays |
| Schedule (post-close plan) | Sprint cap | 12 | Hours that do not fit in this many sprints are reported unplanned |
| Schedule (post-close plan) | Hours per workday | 4 | Roadmap durations derive as effort-hours / 4 h per workday |
| Schedule (post-close plan) | Minimum workstream duration | 7 days | Clamp for derived workstream lengths |
| Schedule (post-close plan) | Maximum workstream duration | 120 days | Clamp for derived workstream lengths |
| Health-posture ceilings | Critical health ceiling | 15 | Max health score a critical finding may render before reconciliation down |
| Health-posture ceilings | High health ceiling | 40 | Max health score a high finding may render |
| Health-posture ceilings | Medium health ceiling | 65 | Max health score a medium finding may render |
| Health-posture ceilings | Low health ceiling | 85 | Max health score a low finding may render |
| Health-posture ceilings | Unassessed-pillar health | 50 | Health assigned to a pillar not assessed this scan |
| Health-posture ceilings | Doc-staleness population threshold | 5 | Tracked files below this count are not scored for staleness |

*A parameter that was not configured is listed as not quantified, never assumed away. The discount rate is the most consequential absent parameter: without it, the model performs no present-value discounting, which means the $738,630 figure is a current-dollar bound, not a net-present-value figure. The buyer's quality-of-earnings provider should re-price under a discount rate of the buyer's weighted average cost of capital.*

## 3.3 Key Figures

Total remediation hours: 14,773 h. Total remediation cost: $738,630. Valuation discount applied: 23.4%. Post-merge integration drag: 80.0%. Person-years of remediation: 7.1. Hourly rate: $50. Complexity multiplier: 1.5. Valuation multiplier: 3. Integration drag factor: 0.15%. Productive hours per quarter: 624 h. Conditions threshold: 5.0%. No-go threshold: 25.0%. Valuation discount cap: 25.0%. Integration drag cap: 80.0%. Remediation cost as percentage of price: 4.9%.

Aggregate expected loss before remediation: $337,500. Aggregate residual expected loss: $50,000. Aggregate risk-reduction value: $287,500. Quantified risk lines: 1 of 9. Weighted expectation across scenarios: $6,028,572. The double-counting ledger is clean (true); the justified overlapping keys are layer1.integration_drag_pct, layer1.total_cost, and layer2.aggregate_residual_expected_loss (each appearing on multiple surfaces but priced once).

## 3.4 Remediation Line Items

The remediation plan is built from eight cost categories. Each category has a count of findings, a full hours-per-unit (HPU), a saturation cap, a saturation decay HPU, an effective applied HPU, total hours, total cost, and a note explaining the basis. Saturation is the model's mechanism for preventing runaway cost: above the cap, the decay HPU applies, which is materially lower than the full HPU. This reflects the diminishing marginal effort of remediating repeated finding types.

**Table 3.2 — Remediation line items by category**

| Category | Count | Full HPU | Sat. cap | Sat. decay | Applied HPU | Total hours | Total cost | Notes |
|---|---|---|---|---|---|---|---|---|
| Knowledge transfer | 1,552 | 8 | 30 | 4 | 4.08 | 6,328 h | $316,400 | Critical sole-maintainer modules requiring knowledge transfer (costed per unique module, not per at-risk contributor) |
| Critical complexity | 706 | 18 | 50 | 4 | 4.99 | 3,524 h | $176,200 | Critically complex functions requiring refactoring |
| High complexity | 1,342 | 6 | 150 | 2 | 2.45 | 3,284 h | $164,200 | High-complexity functions (exclusive of critical tier) |
| Dependency pinning | 1,296 | 1 | 100 | 0.5 | 0.54 | 698 h | $34,900 | Unpinned dependencies requiring version locking |
| Test investment | 42 | 16 | 20 | 4 | 9.71 | 408 h | $20,400 | High-regression-risk modules requiring test investment |
| IP reassignment | 100 | 4 | 50 | 2 | 3.00 | 300 h | $15,000 | Files with external-dominant authorship requiring IP review |
| General debt | 706 | 0.5 | 200 | 0.1 | 0.21 | 151 h | $7,530 | 706 confirmed technical-debt markers (in-comment annotations + tooling suppressions) |
| Decoupling | 4 | 20 | 10 | 8 | 20.00 | 80 h | $4,000 | Circular dependency cycles requiring decoupling |

*Knowledge transfer dominates at $316,400 (43% of total remediation cost), driven by 1,552 critical sole-maintainer modules. Critical complexity contributes $176,200 (24%). Together these two categories account for two-thirds of the remediation bound; the buyer's post-close plan should sequence them first because they are also the most schedule-sensitive.*

## 3.5 Remediation Cost Confidence Interval

**Table 3.3 — Confidence interval (basis: ±30% on hours-per-unit estimates)**

| Metric | Value |
|---|---|
| Basis | ±30% on hours-per-unit estimates |
| Low | $517,041 |
| High | $960,219 |

| Metric | Value |
|---|---|
| Width percent | 30 |

*This is a stated-confidence range from the model calibrations, not a stochastic bound. It reflects uncertainty in the hours-per-unit estimates; it does not reflect uncertainty in the finding count, the saturation model, or the hourly rate. The buyer's quality-of-earnings provider should treat this range as a lower bound on the true uncertainty, which includes calibration risk (the HPU values are unbenchmarked) and scope risk (the 302-row register has not been human-validated).*

## 3.6 Saturation Model and Line-Item Basis

Each cost category has a calibration factor, cap, decay HPU, and full HPU. The calibration source is the built-in default for every category in this scan; calibration.applied is false, meaning the HPU values have not been benchmarked against external data. The applied HPU is computed as min(full HPU, cap × full HPU) for the first cap findings, plus decay HPU for the remainder. The line-item basis table below records the HPU formula for each category.

**Table 3.4 — Saturation model by cost category**

| Category | Calibration factor | Calibration note | Calibration source | Cap | Decay HPU | Full HPU |
|---|---|---|---|---|---|---|
| Critical complexity | 1 | Not quantified | Built-in default | 50 | 4 | 18 |
| Decoupling | 1 | Not quantified | Built-in default | 10 | 8 | 20 |
| Dependency pinning | 1 | Not quantified | Built-in default | 100 | 0.5 | 1 |
| General debt | 1 | Not quantified | Built-in default | 200 | 0.1 | 0.5 |
| High complexity | 1 | Not quantified | Built-in default | 150 | 2 | 6 |
| IP reassignment | 1 | Not quantified | Built-in default | 50 | 2 | 4 |
| Knowledge transfer | 1 | Not quantified | Built-in default | 30 | 4 | 8 |
| Test investment | 1 | Not quantified | Built-in default | 20 | 4 | 16 |

*Because every category uses the built-in default with calibration factor 1, the buyer's quality-of-earnings provider should re-price using external benchmarks (e.g. Constructive Cost Model for complexity, vendor quotes for dependency pinning). The unbenchmarked status is a price-negotiation lever for the buyer.*

## 3.7 SQALE Remediation Basis

SQALE (Software Quality Assessment based on Lifecycle Expectations) provides the canonical remediation-effort figure for complexity-bearing debt. In this scan, SQALE estimates 108,265 days of remediation effort across 98 D-graded modules. The model does not double-count SQALE against the complexity hotspot line item; the double-counting ledger justifies the overlap, with SQALE consumed as a planning bound rather than as a separate line item. The aggregate expected-loss figure is therefore independent of the SQALE basis.

© Meridian Advisory Services

**Table 3.5 — SQALE remediation basis**

| Metric | Value |
|---|---|
| Costed as complexity hotspot | 98 |
| D-graded modules | 98 |
| Deduped uncosted modules | 0 |
| Hours equivalent | 866,121 h |
| Line item added | false |
| SQALE remediation effort (days) | 108,265 |
| Note | SQALE debt-effort arrives in per-module remediation hours from the complexity scanner. The model consumes it as a planning bound; it is not priced as a separate line item to avoid double-counting the complexity hotspot category. |

## 3.8 Sensitivity Sweep

The sensitivity sweep varies six parameters across three multipliers (0.7, 1.0, 1.3) and reports the resulting total remediation hours, total cost, valuation discount percent, post-merge integration drag percent, board decision, and confidence score. The decision tier and confidence are invariant across all 17 scenarios; the dollar cost varies materially only with the hourly rate (between $517,041 at $35/h and $960,219 at $65/h). Remediation timing and scenario weights are not quantified in this scan.

**Table 3.6 — Sensitivity sweep details**

| Parameter | Multiplier | Value | Hours | Cost | Val. disc. % | Drag % | Decision | Conf. |
|---|---|---|---|---|---|---|---|---|
| Complexity multiplier | 0.7 | 1.05 | see cost | $725,130 | 23.43 | 80 | DO_NOT_ACQUIRE | 74 |
| Complexity multiplier | 1.0 | 1.50 | see cost | $738,630 | 23.43 | 80 | DO_NOT_ACQUIRE | 74 |
| Complexity multiplier | 1.3 | 1.95 | see cost | $752,130 | 23.43 | 80 | DO_NOT_ACQUIRE | 74 |
| Discount rate | 0.7 | 0.0245 | see cost | $738,630 | 23.43 | 80 | DO_NOT_ACQUIRE | 74 |
| Discount rate | 1.0 | 0.0350 | see cost | $738,630 | 23.43 | 80 | DO_NOT_ACQUIRE | 74 |
| Discount rate | 1.3 | 0.0455 | see cost | $738,630 | 23.43 | 80 | DO_NOT_ACQUIRE | 74 |
| Hourly rate | 0.7 | $35 | see cost | $517,041 | 23.43 | 80 | DO_NOT_ACQUIRE | 74 |
| Hourly rate | 1.0 | $50 | see cost | $738,630 | 23.43 | 80 | DO_NOT_ACQUIRE | 74 |

| Parameter | Multiplier | Value | Hours | Cost | Val. disc. % | Drag % | Decision | Conf. |
|---|---|---|---|---|---|---|---|---|
| Hourly rate | 1.3 | $65 | see cost | $960,219 | 23.43 | 80 | DO_NOT_ACQUIRE | 74 |
| Max total hours | 0.7 | 14,000 | see cost | $709,658 | 23.43 | 80 | DO_NOT_ACQUIRE | 74 |
| Max total hours | 1.0 | 20,000 | see cost | $738,630 | 23.43 | 80 | DO_NOT_ACQUIRE | 74 |
| Max total hours | 1.3 | 26,000 | see cost | $738,630 | 23.43 | 80 | DO_NOT_ACQUIRE | 74 |
| Remediation timing | n/q | n/q | n/q | n/q | 23.43 | 80 | n/q | n/q |
| Scenario weights | n/q | n/q | n/q | n/q | 23.43 | 80 | n/q | n/q |
| Valuation multiplier | 0.7 | 2.1 | see cost | $738,630 | 23.43 | 80 | DO_NOT_ACQUIRE | 74 |
| Valuation multiplier | 1.0 | 3.0 | see cost | $738,630 | 23.43 | 80 | DO_NOT_ACQUIRE | 74 |
| Valuation multiplier | 1.3 | 3.9 | see cost | $738,630 | 23.43 | 80 | DO_NOT_ACQUIRE | 74 |

*The decision invariance is the most consequential read of this table: no reasonable variation in hourly rate, complexity multiplier, or discount rate moves the verdict off DO_NOT_ACQUIRE. The buyer is therefore negotiating against a robust recommendation, not an artefact of an extreme parameter choice.*

## 3.9 Risk Valuation Lines (Layer 2)

Each material risk line is translated into a probability × economic impact product, producing an expected loss before and after remediation. Only the complexity line is quantified in this scan (probability before 45%, economic impact before $750,000, expected loss before $337,500; probability after 20%, economic impact after $250,000, expected loss after $50,000; risk-reduction value $287,500). The remaining eight lines are NOT_QUANTIFIED, which is a stored assessment state distinct from the five coverage states. The aggregate expected loss before remediation is $337,500; the aggregate residual expected loss is $50,000; the aggregate risk-reduction value is $287,500.

**Table 3.7 — Risk valuation lines**

| Pillar | Title | Assessment state | Prob. before | Impact before | Expected loss before | Expected loss after | Residual | RRV |
|---|---|---|---|---|---|---|---|---|
| Complexity | ProvideService | ASSESSED_WITH_FINDINGS | 45.0% | $750,000 | $337,500 | $50,000 | $50,000 | $287,500 |
| Coverage | Coverage risk | NOT_QUANTIFIED | n/q | n/q | n/q | n/q | n/q | n/q |

© MTDS

| Pillar | Title | Assessment state | Prob. before | Impact before | Expected loss before | Expected loss after | Residual | RRV |
|---|---|---|---|---|---|---|---|---|
| Debt | fixme | NOT_QUANTIFIED | n/q | n/q | n/q | n/q | n/q | n/q |
| Entanglement | Entanglement risk | NOT_QUANTIFIED | n/q | n/q | n/q | n/q | n/q | n/q |
| Integration | Integration risk | NOT_QUANTIFIED | n/q | n/q | n/q | n/q | n/q | n/q |
| IP Ownership | IP ownership risk | NOT_QUANTIFIED | n/q | n/q | n/q | n/q | n/q | n/q |
| License | envtpl | NOT_QUANTIFIED | n/q | n/q | n/q | n/q | n/q | n/q |
| Retention | Alex Khomenko | NOT_QUANTIFIED | n/q | n/q | n/q | n/q | n/q | n/q |
| Test Deficit | Test deficit risk | NOT_QUANTIFIED | n/q | n/q | n/q | n/q | n/q | n/q |

*When a line is quantified, the before/after probability and impact columns are the stored inputs behind the expected-loss product, so the buyer can re-derive the number rather than trust it. When no line is quantified (the default, because financial.risk_parameters ships empty), the block renders the linked Layer-1 remediation evidence per line and the single reason, rather than a column of identical unquantified cells.*

MTD Set

## 3.10 Scenario Branches (Layer 3)

Four scenario branches are enabled in this scan. Each carries its own evidence basis, total impact, technology VaR amount, transaction impact amount, and not-modelled reason. The weighted expectation across scenarios is $6,028,572; the scenario-based technology VaR carries the honest label Scenario-Based Technology VaR when produced. Disabled branches are not-modelled, not zero — their not-modelled reason is stated, and the technology-VaR and transaction-impact columns read not quantified rather than 0.

**Table 3.8 — Scenario branches**

| Scenario ID | Title | Enabled | Evidence basis | Total impact | Tech VaR | Transaction impact | Not modelled reason |
|---|---|---|---|---|---|---|---|
| acquire_remediate_post_close | Acquire and remediate post-close | Yes | Transaction EV + Layer-1 remediation + integration drag | $11,557,929 | Not quantified | Not quantified | Not quantified |
| material_failure | Material failure / downside case | Yes | Analyst calibration on documented scan evidence | $90,194 | $90,194 | Not quantified | Not quantified |

| Scenario ID | Title | Enabled | Evidence basis | Total impact | Tech VaR | Transaction impact | Not modelled reason |
|---|---|---|---|---|---|---|---|
| partial_remediation | Partial remediation | Yes | Config calibration (analyst_calibration) | $391,435 | Not quantified | Not quantified | Not quantified |
| remediate_before_close | Remediate before close | Yes | Layer-1 remediation schedule + Layer-2 residual | $736,195 | Not quantified | $3,514,500 | Not quantified |

*The acquire-and-remediate-post-close scenario carries the largest total impact at $11.56M, dominated by integration-separation TSA ($12.0M) net of synergy and remediation. The remediate-before-close scenario is the cheapest at $736K of Layer-1 remediation plus $3.5M of transaction impact. The material-failure downside case is calibrated at $90,194 of technology VaR.*

## 3.11 Deal Components (Never Collapsed)

Every component carries its own stored amount, and the amount is shown here rather than only in the Deal_Economics sheet. Components are never collapsed and never labelled as a market-style VaR figure; an unquantified component reads not quantified. This presentation prevents the common M&A error of netting gross figures against gross figures without disclosing the underlying components.

**Table 3.9 — Deal components**

| Component | Impact type | Amount | Assessment state | Evidence state | Source key |
|---|---|---|---|---|---|
| Headline EV | Revenue disruption | $15,000,000 | ASSESSED_WITH_FINDINGS | QUANTIFIED_WITH_EVIDENCE | transaction.headline_ev |
| Integration / separation TSA | Integration & separation cost | $12,000,000 | ASSESSED_WITH_FINDINGS | QUANTIFIED_WITH_EVIDENCE | layer1.integration_drag_pct |
| Net quantified tech impact | Remediation cost | $12,438,630 | ASSESSED_WITH_FINDINGS | QUANTIFIED_WITH_EVIDENCE | layer4.net_quantified_tech_impact |
| Other quantified tech exposure | Security incident | $150,000 | ASSESSED_WITH_FINDINGS | QUANTIFIED_WITH_EVIDENCE | config.other_evidence_supported_impacts[0] |
| Potential synergy | Customization burden | $500,000 | ASSESSED_WITH_FINDINGS | QUANTIFIED_WITH_EVIDENCE | config.synergy_value |
| Remediation cost | Remediation cost | $738,630 | ASSESSED_WITH_FINDINGS | QUANTIFIED_WITH_EVIDENCE | layer1.total_cost |

| Component | Impact type | Amount | Assessment state | Evidence state | Source key |
|---|---|---|---|---|---|
| Residual expected loss | Operational downtime | $50,000 | ASSESSED_WITH_FINDINGS | QUANTIFIED_WITH_EVIDENCE | layer2.aggregate_residual_expected_loss |

*The $12.4M net quantified tech impact is the headline economic figure for the buyer's quality-of-earnings provider to scrutinise. It is dominated by the $12.0M integration-separation TSA component, which is itself capped at 80% of remediation hours under the integration-drag curve. The $738,630 remediation cost is the Layer-1 figure that drives the conditions-gate analysis.*

## 3.12 Downside Exposure and Percentile VaR

**Table 3.10 — Downside exposure range and percentile VaR**

| Base amount | Downside amount | Severe amount | Range evidenced | Percentile VaR | VaR percentile |
|---|---|---|---|---|---|
| $168,032 | $243,525 | $174,288 | true | $1,045,731 | 95 |

*A base/downside/severe range is emitted only when the analyst calibrates one; the evidenced flag is true. The percentile VaR of $1.05M at the 95th percentile is a technology-risk model output, not a market-risk measure; the buyer's risk team should reconcile it against any cyber-insurance pricing the underwriter offers.*

## 3.13 Deal Economics Headline Reconciliation

**Table 3.11 — Headline valuation reconciliation**

| Headline EV | Net quantified tech impact | Adjusted EV | Price adjustment | Cap % | Cap amount | Capped | No adjustment reason |
|---|---|---|---|---|---|---|---|
| $15,000,000 | $12,438,630 | $11,250,000 | $3,750,000 | 25.0 % | $3,750,000 | true | Not quantified |

*The valuation adjustment is capped at 25% of the proposed price ($3.75M). The unadjusted impact would have been larger; the cap is binding. This means the $3.75M price adjustment is a floor, not a midpoint, and the buyer should expect the seller to argue for a lower figure based on the calibrated nature of the model inputs. The size of any escrow, holdback, warranty, or indemnity is a qualitative option for deal counsel and is deliberately not a number here.*

## 3.14 Per-Line-Item Threshold Breach View

**Table 3.12 — Per-line-item threshold breach view**

| Category | Total cost | % of price | Gate status |
|---|---|---|---|
| Critical complexity | $176,200 | 1.17% | Below both gates on its own |
| High complexity | $164,200 | 1.09% | Below both gates on its own |
| General debt | $7,530 | 0.05% | Below both gates on its own |
| Test investment | $20,400 | 0.14% | Below both gates on its own |
| Decoupling | $4,000 | 0.03% | Below both gates on its own |

© M&A Services

| Category | Total cost | % of price | Gate status |
|---|---|---|---|
| Dependency pinning | $34,900 | 0.23% | Below both gates on its own |
| Knowledge transfer | $316,400 | 2.11% | Below both gates on its own |
| IP reassignment | $15,000 | 0.10% | Below both gates on its own |

*No single line item breaches either gate on its own. The aggregate $738,630 breaches the 5% conditions gate when combined with the $250,000 materiality floor. The buyer should not accept a seller argument that any individual category is below the gate and therefore exempt; the gate fires on aggregate, not on category.*

## 3.15 Categories Not Costed This Scan

Five cost categories have no line item in this scan. Each is explicitly zero in Remediation_Costs, and rendering the gap here keeps the remediation scope honest rather than presenting a clean table.

- Critical-path debt — no line item this scan (explicitly zero in Remediation_Costs).
- Taint remediation — no line item this scan (explicitly zero in Remediation_Costs).
- Staleness remediation — no line item this scan (explicitly zero in Remediation_Costs).
- Sole-maintainer risk — no line item this scan (explicitly zero in Remediation_Costs).
- SQALE remediation — no line item this scan (explicitly zero in Remediation_Costs).

## 3.16 Double-Counting Ledger

The double-counting ledger is clean (true). The shared-registry keys show each evidence item once, on multiple surfaces, never added twice. Three keys are justified as appearing on multiple surfaces: layer1.integration_drag_pct (scenario and deal component), layer1.total_cost (scenario and deal component), and layer2.aggregate_residual_expected_loss (scenario and deal component). The clean verdict means the buyer can rely on the aggregate figures without re-deriving them.

**Table 3.13 — Double-counting ledger figures (top 20 of 26 rows)**

| Source key | Layer | Amount | Assessment state | Double-counting group |
|---|---|---|---|---|
| complexity.summary.critical_complexity_count | cost component | $176,200 | ASSESSED_WITH_FINDINGS | Not quantified |
| config.downside_case | scenario component | $90,194 | ASSESSED_WITH_FINDINGS | Not quantified |
| config.other_evidence_supported_impacts[0] | deal component | $150,000 | ASSESSED_WITH_FINDINGS | Not quantified |
| config.synergy_value | deal component | $500,000 | ASSESSED_WITH_FINDINGS | Not quantified |
| debt.summary.annotation+tooling_marker_count | cost component | $7,530 | ASSESSED_WITH_FINDINGS | Not quantified |

MD Services

| Source key | Layer | Amount | Assessment state | Double-counting group |
|---|---|---|---|---|
| entanglement.summary.total_cycles | cost component | $4,000 | ASSESSED_WITH_FINDINGS | Not quantified |
| ip_ownership.summary.high_risk_files | cost component | $15,000 | ASSESSED_WITH_FINDINGS | Not quantified |
| layer1.integration_drag_pct | scenario component | $10,823,312 | ASSESSED_WITH_FINDINGS | g:layer1_shared |
| layer1.integration_drag_pct | deal component | $12,000,000 | ASSESSED_WITH_FINDINGS | g:layer1_shared |
| layer1.total_cost | scenario component | $689,519 | ASSESSED_WITH_FINDINGS | g:layer1_shared |
| layer1.total_cost.x_partial | scenario component | $344,760 | ASSESSED_WITH_FINDINGS | Not quantified |
| layer2.aggregate_expected_loss | risk component | $337,500 | ASSESSED_WITH_FINDINGS | Not quantified |
| layer2.aggregate_residual_expected_loss | scenario component | $45,097 | ASSESSED_WITH_FINDINGS | g:layer2_shared |
| layer2.aggregate_residual_expected_loss | deal component | $50,000 | ASSESSED_WITH_FINDINGS | g:layer2_shared |
| layer2.downside_exposure | risk component | $337,500 | ASSESSED_WITH_FINDINGS | Not quantified |
| layer2.residual_expected_loss | risk component | $50,000 | ASSESSED_WITH_FINDINGS | Not quantified |

*The three g:layer1_shared and g:layer2_shared entries are the justified overlapping keys. The clean verdict means no figure is silently double-counted.*

## 3.17 Financial Lineage

**Table 3.14 — Financial lineage**

| Step | Artifact | Note |
|---|---|---|
| Artifacts | workbook.xlsx | Remediation_Costs, Cal_Lineage, Cost_Drivers, Sensitivity, Assumptions mirror report.json#/financial. |
| Decision | report.json#/decision | generate_board_decision(pillars, financial) from risk register + remediation cost. |
| Fingerprint | report.json#/financial/fingerprint | SHA-256 over {tool_version, commit=5be6647ffe4c0ce6522387884e4ca51e9c181f10, repo_path=grafana, mode=deal_thesis=True, offline=False, scan_history=True, taint=True}. |
| Model | report.json#/financial | compute_financial_impact(config): pillar count → hours-per-unit (saturated with cap/decay) → hours → rate → cost; then valuation discount and integration drag. |

| Step | Artifact | Note |
|---|---|---|
| Pillars | report.json#/pillars | Per-pillar raw reports; the cost model consumes summary counts plus the bounded top evidence rows only. |
| Scan | Grafana_cpg_261016T0604.db | Symbol/index graph that backs every pillar scanner (stored under <output>/internal/). |

## 3.18 Executive Deal-Bridge Rows

One chain per material risk line: the stored evidence, the technical finding, the business consequence, the economic impact, the remediation cost, the residual risk that survives it, and the deal consequence. This is the buyer's audit trail from a code observation to a price question, rendered as the chain itself rather than as a row count.

**Table 3.15 — Executive deal-bridge rows (evidence to transaction consequence)**

| Finding | Business consequence | Economic impact | Remediation cost | Residual risk | Transaction consequence |
|---|---|---|---|---|---|
| ProvideService | Expected loss of $337,500 before remediation (analyst calibration) | $337,500 | $176,200 | $50,000 | Referenced by enabled deal scenarios remediate_before_close, acquire_remediate_post_close, partial_remediation, material_failure |
| Coverage risk | Expected loss NOT_QUANTIFIED before remediation | n/q | n/q | n/q | Referenced by all enabled deal scenarios |
| fixme | Expected loss NOT_QUANTIFIED before remediation | n/q | $7,530 | n/q | Referenced by all enabled deal scenarios |
| Entanglement risk | Expected loss NOT_QUANTIFIED before remediation | n/q | $4,000 | n/q | Referenced by all enabled deal scenarios |
| Integration risk | Expected loss NOT_QUANTIFIED before remediation | n/q | n/q | n/q | Referenced by all enabled deal scenarios |
| IP ownership risk | Expected loss NOT_QUANTIFIED before remediation | n/q | $15,000 | n/q | Referenced by all enabled deal scenarios |
| envtpl | Expected loss NOT_QUANTIFIED before remediation | n/q | n/q | n/q | Referenced by all enabled deal scenarios |
| Alex Khomenko | Expected loss NOT_QUANTIFIED before remediation | n/q | $316,400 | n/q | Referenced by all enabled deal scenarios |
| Test deficit risk | Expected loss NOT_QUANTIFIED before remediation | n/q | $20,400 | n/q | Referenced by all enabled deal scenarios |

*The complexity line is the only chain with quantified figures across all columns. The other eight chains carry partial quantification — remediation cost is known but probability-weighted expected loss is not. This is a known limitation of the financial model when financial.risk_parameters ships empty; the buyer should require Layer-2 quantification as part of pre-signing diligence.*

© MTD Services

# Part III — Material Risks by Deal Significance

## 4. Licensing & IP Exposure

### 4.1 Introduction

The license posture of the packaged components that ship in the target is the single most consequential diligence question for a closed-source acquirer. Viral and restrictive copyleft exposure can trigger source-disclosure obligations to customers, force carve-outs from the acquired product, or require re-licensing of dependent code. IP ownership drives title and warranty risk: a 77.23% external-author share of blamed lines, observed in this scan, requires the seller to warrant assignability of third-party contributions and to produce signed contributor agreements for the top contributors.

Per-package detected licenses, license-clause findings, and the IP-ownership attribution metrics below are stored ground truth. The pillar risk is rated high, but the posture is held on aggregate metrics rather than item-level findings — meaning the buyer's counsel should treat the figures as a screening signal that requires follow-up against the underlying contributor agreements.

### 4.2 Material Findings

Five material license findings are recorded. Each is a scanner signal, not a validated exposure; each must be validated before it is relied on in negotiation.

**Table 4.1 — Material license findings**

| Finding ID | Package | Severity | Evidence state | Source file |
|---|---|---|---|---|
| F-license-001 | envtpl | High | Detected, not verified | devenv/docker/blocks/collectd/requirements.txt |
| F-license-002 | @grafana/o11y-ds-frontend | High | Detected, not verified | package.json |
| F-license-003 | @grafana/prometheus | High | Detected, not verified | package.json |
| F-license-004 | @grafana/sql | High | Detected, not verified | package.json |
| F-license-005 | @grafana/prometheus | High | Detected, not verified | public/app/plugins/datasource/cloudwatch/package.json |

The deal-mechanic implication for each finding is a warranty / representation covering license compliance for the shipped product, with an indemnity for any copyleft exposure. Quantification and sizing are for deal counsel; no amount is asserted in this report. The five findings collectively cover the AGPL-3.0 / AGPL-3.0-only exposure on @grafana/o11y-ds-frontend, @grafana/prometheus, and @grafana/sql packages; the GPL v3 exposure on envtpl in the devenv; and a fifth finding that overlaps with the per-package inventory below.

### 4.3 Key Figures

License findings inventory: 5,497 findings. Viral license count: 7. Restrictive license count: 5. Unknown license count: 4,497. Licensed packages total: 5,472. Contamination chains: 0. IP ownership attribution rows: 1,069.

The 4,497 unknown-license packages are the dominant risk vector. An unknown license is not a clean license — it means the scanner could not classify the package against its license-detection rules. The buyer's counsel should commission a manual classification of the top-N unknown packages by dependency size (lockfile rows), prioritising the transitive dependencies of the AGPL-3.0 / GPL v3 packages identified above. Contamination chains are zero, meaning the scanner detected no viral propagation from a viral source into a dependent package; this is a positive signal, but it is dependent on the scanner's chain-detection rules and should not be relied on as a complete contamination verdict.

## 4.4 Per-Package Detected Licenses (top 18 of 5,497 rows)

**Table 4.2 — Per-package detected licenses**

| Package name | Detected license | Viral | Restrictive | Confidence | Source file |
|---|---|---|---|---|---|
| (repo-level) | AGPL-3.0 | Yes | Yes | 0.8 | LICENSE |
| (repo-level) | AGPL-3.0 | Yes | Yes | 0.8 | packages/grafana-o11y-ds-frontend/LICENSE_AGPL |
| (repo-level) | AGPL-3.0 | Yes | Yes | 0.8 | packages/grafana-sql/LICENSE_AGPL |
| (repo-level) | AGPL-3.0 | Yes | Yes | 0.8 | packaging/msi/resources/LICENSE.md |
| (repo-level) | AGPL-3.0 | Yes | Yes | 0.8 | packaging/msi/resources/LICENSE.rtf |
| @grafana/o11y-ds-frontend | AGPL-3.0-only | Yes | Yes | 0.7 | package.json |
| @grafana/prometheus | AGPL-3.0-only | Yes | Yes | 0.7 | package.json |
| @grafana/prometheus | AGPL-3.0-only | Yes | Yes | 0.7 | public/app/plugins/datasource/cloudwatch/package.json |
| @grafana/sql | AGPL-3.0-only | Yes | Yes | 0.7 | package.json |
| envtpl | GPL v3 | Yes | Yes | 0.7 | devenv/docker/blocks/collectd/requirements.txt |
| @axe-core/playwright | MPL-2.0 | Yes | No | 0.7 | package.json |
| axe-core | MPL-2.0 | Yes | No | 0.7 | package.json |
| (repo-level) | Apache-2.0 | No | No | 0.8 | packages/grafana-alerting/LICENSE_APACHE2 |
| (repo-level) | Apache-2.0 | No | No | 0.8 | packages/grafana-api-clients/LICENSE_APACHE2 |

| Package name | Detected license | Viral | Restrictive | Confidence | Source file |
|---|---|---|---|---|---|
| (repo-level) | Apache-2.0 | No | No | 0.8 | packages/grafana-data/LICENSE_APACHE2 |
| (repo-level) | Apache-2.0 | No | No | 0.8 | packages/grafana-e2e-selectors/LICENSE_APACHE2 |
| (repo-level) | Apache-2.0 | No | No | 0.8 | packages/grafana-eslint-rules/LICENSE_APACHE2 |
| (repo-level) | Apache-2.0 | No | No | 0.8 | packages/grafana-flamegraph/LICENSE_APACHE2 |

*The (repo-level) entries indicate the same AGPL-3.0 license is detected in five top-level LICENSE files; this is consistent with the AGPL-3.0 LICENSE at the repo root applying to the @grafana/o11y-ds-frontend, @grafana/sql, and packaging/msi subtrees. The MPL-2.0 entries on @axe-core/playwright and axe-core are weak copyleft (file-level), which permits combination with proprietary code but requires source disclosure of modified MPL-covered files.*

## 4.5 IP Ownership Metrics

**Table 4.3 — IP ownership metrics**

| Metric | Value |
|---|---|
| attributed File Count | 178 |
| corporate Domain | users.noreply.github.com |
| external Author Count | 2319 |
| external Blamed Line Pct | 0.7723 |
| external Commit Pct | 0.7723 |
| external Domain Count | 817 |
| external Owned Files | 100 |
| high Risk Files | 100 |
| ip Risk Score | 0.646 |
| total Authors | 3217 |
| total Files | 178 |

*The IP risk score of 0.646 (on a 0-1 scale) is driven by an external-blamed-line percentage of 77.23% and an external-owned-file count of 100. The corporate domain is users.noreply.github.com — the GitHub noreply domain — which means most contributors are identified by GitHub usernames rather than by an employer email. This is a structural signal that the buyer's counsel should treat as elevated title risk: each of the 100 high-risk files requires a contributor-assignment agreement on file before close, or it must be carved out of the acquired IP.*

## 4.6 M&A Implications

The AGPL-3.0 exposure is the single most material license issue for a closed-source acquirer. AGPL extends GPL copyleft obligations to network-served software, meaning any acquirer that continues

to operate Grafana as a closed-source SaaS would be required to disclose source code (including modifications) to users interacting with the service over a network. The buyer's counsel must obtain (a) a representation and warranty covering license compliance for the shipped product, (b) an indemnity for any copyleft exposure discovered post-close, and (c) a carve-out schedule listing every component the seller asserts is not subject to AGPL copyleft.

The 4,497 unknown-license packages require an SBOM-level review. The buyer should commission a manual classification of the top-N unknown packages by lockfile size and by transitive reach to AGPL-3.0 / GPL v3 packages. The 100 high-risk files identified by the IP-ownership pillar require signed contributor-assignment agreements on file before close, or the corresponding code must be excluded from the acquired IP. The deal-mechanic tag for IP Ownership is a warranty / representation, which means breach is an indemnity question rather than a work plan; the buyer should size the indemnity cap and survival period accordingly.

© MTD Services

# 5. Security, Vulnerability & Attack Surface Posture

## 5.1 Introduction

The end-to-end security picture combines advisory exposure, critical-severity posture, secret hygiene, infrastructure-as-code findings, attack-surface composition, and privacy-relevant data flows. Advisory counts, candidate secrets, IaC findings, candidate entry points, and privacy findings below are stored ground truth. The combined picture drives cyber-insurance sizing, regulatory disclosure timing, and the cyber-warranty language in the acquisition agreement; coverage gaps (notably in CVE and Secrets) are a direct credibility test of the seller's vulnerability-management claims.

## 5.2 Key Figures

CVE count: 154. Critical CVE count: 0. Packages queried against advisories: 1,789. Unresolved packages: 19. Advisory source: osv.dev.

Secret candidates: 427. High-confidence secrets: 63. Live critical secrets: 12. History-only secrets: 345.

IaC findings total: 1. Attack surface candidate endpoints: 118. Unreferenced endpoints: 102. Dynamic endpoints: 1. Untested candidate endpoints: 71. Candidate exposure score: 0.1047.

Privacy findings total: PRIV-PLAINTEXT-CRED,Plaintext credential,Plaintext credential material in committed config,Medium,.github/workflows/relyance-scan.yml,26. PII endpoint count: 0. Egress SDK count: 0. Sensitive log count: 0.

## 5.3 Material Findings

Five attack-surface findings and one IaC finding are recorded. Each is a scanner signal requiring validation before negotiation reliance.

- F-iac-001 — Container runs as root — medium · Detected, not verified · Dockerfile:403. The M&A implication is that the infrastructure surface must be brought under version control with a review gate before close.
- F-attack_surface-001 — main — medium · Detected, not verified · scripts/test-aliases.py:186.
- F-attack_surface-002 — useMigrations — medium · Detected, not verified · public/app/plugins/datasource/azuremonitor/components/LogsQueryEditor/useMigrations.ts:43.
- F-attack_surface-003 — createMockPanelData — medium · Detected, not verified · public/app/plugins/datasource/azuremonitor/mocks/panelData.ts:7.
- F-attack_surface-004 — createMockDatasource — medium · Detected, not verified · public/app/plugins/datasource/azuremonitor/mocks/datasource.ts:27.
- F-attack_surface-005 — createMockQuery — medium · Detected, not verified · public/app/plugins/datasource/azuremonitor/mocks/query.ts:4.

© MTD Services

- F-privacy-001 — Plaintext credential material in committed config — medium · Detected, not verified · .github/workflows/relyance-scan.yml:26. The M&A implication is a representation covering personal-data handling in the shipped product, scoped to the hits recorded in the credential inventory.

## 5.4 CVE Inventory

The CVE inventory records 154 advisories across 1,789 queried packages. Critical-severity advisories: 0. The advisory source is osv.dev. 19 packages could not be resolved to a version (no lockfile or unpinned), so the advisory check covered only the resolvable subset; the 'no critical CVEs' verdict applies to that queried population only. Severity is reported as UNKNOWN for all 154 advisories because the OSV feed did not return a CVSS-severity band; the buyer's security team should re-query each advisory against the NVD to obtain the severity rating.

**Table 5.1 — CVE inventory (top 18 of 154 rows)**

| ID | Severity | Package | Ecosystem | Version | Fix available | Fixed version | Remediation class |
|---|---|---|---|---|---|---|---|
| GHSA-3f24-pcvm-5jqc | UNKNOWN | github.com/nats-io/nats-server | go | v2.14.6 | false | n/a | no_fixed_release |
| GHSA-3jq7-8ph8-63xm | UNKNOWN | github.com/grafana/grafana | go | 0.0.0 | false | n/a | no_fixed_release |
| GHSA-3p62-42x7-gxg5 | UNKNOWN | github.com/grafana/grafana | go | 0.0.0 | false | n/a | no_fixed_release |
| GHSA-3w66-95m3-8jxg | UNKNOWN | github.com/grafana/grafana | go | 0.0.0 | false | n/a | no_fixed_release |
| GHSA-46m5-8hpj-p5p5 | UNKNOWN | github.com/grafana/grafana | go | 0.0.0 | false | n/a | no_fixed_release |
| GHSA-4724-7jwc-3fpw | UNKNOWN | github.com/grafana/grafana | go | 0.0.0 | false | n/a | no_fixed_release |
| GHSA-4pwp-cx67-5cpx | UNKNOWN | github.com/grafana/grafana | go | 0.0.0 | false | n/a | no_fixed_release |
| GHSA-52jh-2xxh-pwh6 | UNKNOWN | github.com/nats-io/nats-server | go | v2.14.6 | false | n/a | no_fixed_release |
| GHSA-55h8-8g96-x4hj | UNKNOWN | github.com/nats-io/nats-server | go | v2.14.6 | false | n/a | no_fixed_release |
| GHSA-66c4-2g2v-54qw | UNKNOWN | github.com/grafana/grafana | go | 0.0.0 | false | n/a | no_fixed_release |
| GHSA-69j6-29vr-p3j9 | UNKNOWN | github.com/grafana/grafana | go | 0.0.0 | false | n/a | no_fixed_release |
| GHSA-6wh2-8hw7-jw94 | UNKNOWN | github.com/grafana/grafana | go | 0.0.0 | false | n/a | no_fixed_release |

| ID | Severity | Package | Ecosystem | Version | Fix available | Fixed version | Remediation class |
|---|---|---|---|---|---|---|---|
| GHSA-7g92-g4vh-hp84 | UNKNOWN | github.com/grafana/grafana | go | 0.0.0 | false | n/a | no_fixed_release |
| GHSA-7m2x-qhrq-rp8h | UNKNOWN | github.com/grafana/grafana | go | 0.0.0 | false | n/a | no_fixed_release |
| GHSA-7phr-6cc9-4m5q | UNKNOWN | github.com/grafana/grafana | go | 0.0.0 | false | n/a | no_fixed_release |
| GHSA-8r68-gvr4-jh7j | UNKNOWN | github.com/nats-io/nats-server | go | v2.14.6 | false | n/a | no_fixed_release |
| GHSA-9493-h4f5-633x | UNKNOWN | github.com/grafana/grafana | go | 0.0.0 | false | n/a | no_fixed_release |
| GHSA-9983-vrx2-fg9c | UNKNOWN | github.com/nats-io/nats-server | go | v2.14.6 | false | n/a | no_fixed_release |

*Every advisory in this scan has fix_available = false and remediation_class = no_fixed_release. The dominant affected package is github.com/grafana/grafana (the repo's own module path), which is expected because the source under assessment is the Grafana codebase itself. The 19 nats-io, 4 jackc/pgx, and 4 grafana/tempo advisories should be cross-checked against the seller's own vulnerability-management records.*

## 5.5 Secrets Posture

Secret candidates are pattern matches, not confirmed secrets; high-confidence and history-only counts are tracked separately. Of the 427 candidates, 63 are high-confidence (>= 75 confidence) and 12 are live critical (confidence >= 90 in the current working tree). 345 are history-only (present in past commits but not in the working tree), which means rotation, not deletion, is the appropriate remediation. The deal-mechanic implication is a condition precedent requiring rotation of all 63 high-confidence candidates and the 12 live critical secrets before close, plus a stand-up of credential rotation as an ongoing process.

**Table 5.2 — Secrets by kind**

| Kind | Count |
|---|---|
| AWS access key | 15 |
| Certificate block | 163 |
| Connection string | 61 |
| Generic password assignment | 58 |
| JWT | 28 |
| Private key block | 102 |

*The 102 private-key blocks and 163 certificate blocks dominate the inventory. Most of these will be test fixtures (e.g. TLS configurations in datasource tests), but each must be validated individually. The 15 AWS access keys in pkg/tsdb/cloudwatch/models/settings_test.go are a concentrated cluster that should be validated first — they appear at high confidence (94-95) and may be live credentials or test fixtures. The history-only count of 345 includes commits*

MTD Sen

*dating back to scan_history=true, so the buyer's counsel should require evidence of rotation for any historical credential that touched a production system.*

**Table 5.3 — Secret candidate detail (top 14 of 427 rows)**

| Kind | File | Line | Confidence | In history | Risk cap |
|---|---|---|---|---|---|
| aws_access_key | pkg/tsdb/cloudwatch/models/settings_test.go | 202 | 95 | Yes | critical |
| aws_access_key | pkg/tsdb/cloudwatch/models/settings_test.go | 223 | 95 | Yes | critical |
| aws_access_key | pkg/tsdb/cloudwatch/models/settings_test.go | 248 | 95 | Yes | critical |
| aws_access_key | pkg/tsdb/cloudwatch/models/settings_test.go | 112 | 95 | No | critical |
| aws_access_key | pkg/tsdb/cloudwatch/models/settings_test.go | 130 | 95 | No | critical |
| aws_access_key | pkg/tsdb/cloudwatch/models/settings_test.go | 148 | 95 | No | critical |
| aws_access_key | pkg/tsdb/cloudwatch/models/settings_test.go | 166 | 95 | No | critical |
| aws_access_key | pkg/tsdb/cloudwatch/models/settings_test.go | 184 | 95 | No | critical |
| aws_access_key | pkg/tsdb/cloudwatch/models/settings_test.go | 34 | 95 | No | critical |
| aws_access_key | pkg/tsdb/cloudwatch/models/settings_test.go | 56 | 95 | No | critical |
| aws_access_key | pkg/tsdb/cloudwatch/models/settings_test.go | 87 | 95 | No | critical |
| aws_access_key | pkg/tsdb/cloudwatch/models/settings_test.go | 102 | 95 | No | critical |
| aws_access_key | pkg/tsdb/cloudwatch/models/settings_test.go | 71 | 95 | No | critical |
| aws_access_key | pkg/util/uri_sanitize_test.go | 29 | 94 | No | critical |

## 5.6 IaC Findings

One IaC finding is recorded: the Dockerfile at line 403 runs the container as root. The remediation is to add a dedicated non-root user and USER directive in every build stage. The M&A implication is a price-adjustment lever requiring the infrastructure surface to be brought under version control with a review gate before close.

**Table 5.4 — IaC findings**

| Rule ID | Title | Severity | File | Line | Remediation |
|---|---|---|---|---|---|
| IAC-DOCKER-ROOT | Container runs as root | Medium | Dockerfile | 403 | Add a dedicated non-root user and USER directive in every stage. |

## 5.7 Attack Surface Candidate Endpoints

The attack-surface scan identified 118 candidate endpoints. 102 are unreferenced (no inbound edges in the static graph), 1 is dynamic, and 71 are untested. The candidate exposure score of 0.1047 is an upper-bound signal for review prioritisation, not a confirmed vulnerability rate. External reachability could not be established for 117 of the 118 candidates because framework semantics were unresolved; this is a coverage gap that the buyer's security team should close independently.

**Table 5.5 — Candidate endpoint detail (top 14 of 118 rows)**

| Symbol name | Filepath | Line | Endpoint type | Exposure | Framework | Risk level | Test status |
|---|---|---|---|---|---|---|---|
| murmurhash3_32_gc | packages/grafana-flamegraph/src/FlameGraph/murmur3.ts | 21 | application_entry | application | unknown | medium | untested |
| defineThemes | packages/grafana-ui/src/components/Monaco/theme.ts | 31 | application_entry | application | unknown | medium | untested |
| resetSelectStyles | packages/grafana-ui/src/components/Select/resetSelectStyles.ts | 6 | application_entry | application | unknown | medium | untested |
| flatten | public/app/core/utils/flatten.ts | 4 | application_entry | application | unknown | medium | untested |
| requestAnimationFrame | public/app/features/explore/TraceView/components/utils/test/requestAnimationFrame.ts | 17 | application_entry | application | unknown | medium | untested |
| cancelAnimationFrame | public/app/features/explore/TraceView/components/utils/test/requestAnimationFrame.ts | 21 | application_entry | application | unknown | medium | untested |
| useMigrations | public/app/plugins/datasource/azuremonitor/components/LogsQueryEditor/useMigrations.ts | 43 | application_entry | application | unknown | medium | untested |
| createMockDatasource | public/app/plugins/datasource/azuremonitor/mocks/datasource.ts | 27 | application_entry | application | unknown | medium | untested |

| Symbol name | Filepath | Line | Endpoint type | Exposure | Framework | Risk level | Test status |
|---|---|---|---|---|---|---|---|
| createMockPanelData | public/app/plugins/datasource/azuremonitor/mocks/panelData.ts | 7 | application_entry | application | unknown | medium | untested |
| createMockQuery | public/app/plugins/datasource/azuremonitor/mocks/query.ts | 4 | application_entry | application | unknown | medium | untested |
| messageFromError | public/app/plugins/datasource/azuremonitor/utils/messageFromError.ts | 13 | application_entry | application | unknown | medium | untested |
| migrateAnnotation | public/app/plugins/datasource/azuremonitor/utils/migrateAnnotation.ts | 10 | application_entry | application | unknown | medium | untested |
| useLastError | public/app/plugins/datasource/azuremonitor/utils/useLastError.ts | 9 | application_entry | application | unknown | medium | untested |
| migrateAnnotation | packages/grafana-sql/src/utils/migration.ts | 6 | application_entry | application | unknown | medium | untested |

## 5.8 Privacy Findings

One privacy finding is recorded: a plaintext-credential rule hit at .github/workflows/relyance-scan.yml:26. The remediation is to move the credential to a secret store or environment injection and to rotate the committed value. The M&A implication is a representation covering personal-data handling in the shipped product, scoped to the hits recorded in the credential inventory. PII endpoint count is 0; egress SDK count is 0; sensitive log count is 0 — but each of these is a detection result, not a confirmation of absence, and the partial coverage on Secrets means the buyer's privacy counsel should re-verify independently.

**Table 5.6 — Privacy findings**

| Rule ID | Category | Title | Severity | File | Line |
|---|---|---|---|---|---|
| PRIV-PLAINTEXT-CRED | Plaintext credential | Plaintext credential material in committed config | Medium | .github/workflows/relyance-scan.yml | 26 |

# 6. Supply Chain & Dependency Posture

## 6.1 Introduction

The third-party software supply chain combines direct and transitive breadth, pinning discipline, lockfile hygiene, ecosystem mix, and runtime end-of-life exposure. A broad unpinned transitive supply chain creates upstream-advisory and silent-breaking-change exposure and affects post-close security operating cost and SBOM obligations. The pillar risk is medium, held on aggregate metrics rather than item-level findings.

Direct-dependency rows, lockfile resolutions, pinning bookkeeping, and EOL runtime rows below are stored ground truth. The dependency posture reads from manifests / lockfiles as stored; where no lockfile is present, the gap is itemised below. The transitive resolved package count of 19,945 is the headline figure: it represents the supply-chain breadth that an SBOM must enumerate and that continuous-monitoring tooling must track post-close.

## 6.2 Key Figures

Direct dependency count: 1,533. Pinned dependencies: 237. Ranged dependencies: 1,296. Unpinned dependencies: 1,296. Pin ratio: 15.5%. Transitive resolved packages: 19,945. Lockfiles present: 50. Manifests found: 82.

EOL runtime rows: 1. EOL count: 0. Runtimes checked: 1. CVE unresolved packages (cross-ref): 19.

## 6.3 Lockfile Resolution

50 lockfiles are present across the repository, ranging from small (22 resolved packages) to large (3,721 packages in the root go.sum). The presence of lockfiles is a positive reproducibility signal — it means the seller's build is deterministic for the resolved subset. However, the 1,296 unpinned direct dependencies indicate that the manifest layer does not enforce version pinning, which means any newly added or updated dependency can drift between builds even when the lockfile is regenerated. The buyer's security operations team should require deploy-time enforcement of lockfile-based dependency resolution post-close.

**Table 6.1 — Lockfile resolution (top 20 of 50 rows)**

| Lockfile | Resolved packages |
|---|---|
| .citools/src/air/go.sum | 210 |
| .citools/src/cue/go.sum | 72 |
| .citools/src/golangci-lint/go.sum | 567 |
| .citools/src/govulncheck/go.sum | 22 |
| .citools/src/lefthook/go.sum | 108 |
| .citools/src/swagger/go.sum | 149 |
| apps/advisor/go.sum | 1,757 |

© MTD Services

| Lockfile | Resolved packages |
|---|---|
| apps/alerting/alertenrichment/go.sum | 98 |
| apps/alerting/alertrulequality/go.sum | 96 |
| apps/alerting/notifications/go.sum | 300 |
| apps/alerting/rules/go.sum | 236 |
| apps/annotation/go.sum | 236 |
| apps/collections/go.sum | 170 |
| apps/correlations/go.sum | 236 |
| apps/dashboard/go.sum | 381 |
| apps/dashvalidator/go.sum | 1,747 |
| apps/example/go.sum | 252 |
| apps/folder/go.sum | 170 |
| apps/iam/go.sum | 236 |
| apps/live/go.sum | 236 |

## 6.4 Direct Dependency Inventory

500 direct dependencies are listed in the workbook, with a representative sample shown below. The vast majority are Go modules in apps/advisor and .citools subdirectories, indicating a Go-dominated backend with significant development-tooling overhead. The pin ratio of 15.5% (237 pinned of 1,533 direct) means 84.5% of direct dependencies are declared with a version range rather than an exact version, which exposes the build to silent upstream changes.

**Table 6.2 — Direct dependency inventory (top 20 of 500 rows)**

| Name | Specified | Pinned | Scope | Ecosystem | Manifest |
|---|---|---|---|---|---|
| 4d63.com/gocheckcompilerdirectives | v1.4.0 | No | go.mod | Go | .citools/src/golangci-lint/go.mod |
| 4d63.com/gochecknoglobals | v0.2.2 | No | go.mod | Go | .citools/src/golangci-lint/go.mod |
| cel.dev/expr | v0.25.3 | No | go.mod | Go | apps/advisor/go.mod |
| charm.land/lipgloss/v2 | v2.0.6 | No | go.mod | Go | .citools/src/golangci-lint/go.mod |
| cloud.google.com/go/compute/metadata | v0.9.0 | No | go.mod | Go | apps/advisor/go.mod |
| codeberg.org/chavacava/garif | v0.2.1 | No | go.mod | Go | .citools/src/golangci-lint/go.mod |
| codeberg.org/polyfloyd/go-errorlint | v1.9.0 | No | go.mod | Go | .citools/src/golangci-lint/go.mod |

| Name | Specified | Pinned | Scope | Ecosystem | Manifest |
|---|---|---|---|---|---|
| cuelabs.dev/go/oci/ociregistry | v0.0.0-20250722084951-074d06050084 | No | go.mod | Go | .citools/src/cue/go.mod |
| cuelang.org/go | v0.15.4 | No | go.mod | Go | .citools/src/cue/go.mod |
| dario.cat/mergo | v1.0.2 | No | go.mod | Go | .citools/src/air/go.mod |
| dev.gaijin.team/go/exhaustruct/v4 | v4.0.0 | No | go.mod | Go | .citools/src/golangci-lint/go.mod |
| dev.gaijin.team/go/exhaustruct/v5 | v5.2.0 | No | go.mod | Go | .citools/src/golangci-lint/go.mod |
| dev.gaijin.team/go/golib | v0.8.1 | No | go.mod | Go | .citools/src/golangci-lint/go.mod |
| filippo.io/edwards25519 | v1.2.0 | No | go.mod | Go | apps/advisor/go.mod |
| github.com/4meepo/tagalign | v1.4.4 | No | go.mod | Go | .citools/src/golangci-lint/go.mod |
| github.com/Abirdcfly/dupword | v0.1.8 | No | go.mod | Go | .citools/src/golangci-lint/go.mod |
| github.com/AdminBenni/iota-mixing | v1.0.0 | No | go.mod | Go | .citools/src/golangci-lint/go.mod |
| github.com/air-verse/air | v1.62.0 | No | go.mod | Go | .citools/src/air/go.mod |
| github.com/alecthomas/chroma/v2 | v2.27.0 | No | go.mod | Go | .citools/src/golangci-lint/go.mod |
| github.com/alecthomas/go-check-sumtype | v0.5.1-0.20260828200218-ae6904d28606 | No | go.mod | Go | .citools/src/golangci-lint/go.mod |

## 6.5 EOL Runtime Detail

One runtime is checked: golang go 1.27.2, which is not end-of-life (latest is 1.27.2). The Go runtime is current, which is a positive signal — the buyer's modernization capital is not consumed by language upgrades. The end-of-life check uses endoflife.date as the source; the buyer's platform team should integrate continuous EOL monitoring into the post-close CI pipeline.

**Table 6.3 — EOL runtime detail**

| Product | Runtime | Version | Is EOL | Latest | EOL date | Support |
|---|---|---|---|---|---|---|
| golang | go | 1.27.2 | No | 1.27.2 | Not quantified | Not quantified |

## 6.6 M&A Implications

The 19 CVE-unresolved packages are the most actionable supply-chain finding. These are packages that the OSV advisory check could not resolve to a version, which means the seller's manifests declare a version range or have no lockfile. The buyer's security operations team should require (a) a complete SBOM generated from the resolved lockfiles, (b) a deploy-time enforcement policy that prevents deploying unresolvable dependencies, and (c) a 30-day SLA for upgrading dependencies with critical advisories. The deal-mechanic tag for Dependency is not currently in the conditions register (no price-adjustment lever is sized), but the buyer's counsel should consider adding one given the breadth of the transitive surface.

© MTD Services

# 7. Code Quality, Complexity & Technical Debt

## 7.1 Introduction

The maintainability economics of the codebase are converted from raw metrics into remediation capital and engineering time — the framing the deal model can use — anchored to the same hotspots a quality-of-earnings reviewer will probe. Complexity symbol metrics, percentiles, SQALE grades, duplication, and debt-marker inventories below are stored ground truth. The pillar risk is rated critical.

The headline figures: 57,667 symbols analysed; 2,048 high-complexity and 706 critical-complexity symbols; SQALE grade E with 108,265 days of remediation effort; debt ratio 26.4%; 828,217 duplicated lines (25.2% duplication ratio across 18,406 groups); 883 debt markers (0 on the critical path).

## 7.2 Key Figures

Average cyclomatic complexity: 3.1. Average cognitive complexity: 4.05. Complexity hotspot score: 0.0375. Debt density: 0.006614.

## 7.3 Material Findings

Five critical-complexity findings and five high-severity debt findings are recorded. Each is a scanner signal requiring validation before negotiation reliance. The complexity findings drive the critical-complexity line item ($176,200); the debt findings drive the general-debt line item ($7,530).

**Table 7.1 — Material complexity findings**

| Finding ID | Symbol | Severity | Evidence state | File |
|---|---|---|---|---|
| F-complexity-001 | ProvideService | Critical | Detected, not verified | pkg/services/live/live.go |
| F-complexity-002 | applyIncrementalChanges | Critical | Detected, not verified | pkg/registry/apis/provisioning/jobs/sync/incremental.go |
| F-complexity-003 | Plot | Critical | Detected, not verified | public/vendor/flot/jquery.flot.js |
| F-complexity-004 | useQueryCoauthoringSession | Critical | Detected, not verified | public/app/features/dashboard-scene/panel-edit/PanelEditNext/QueryEditor/coauthoring/useQueryCoauthoringSession.ts |
| F-complexity-005 | RunTestList | Critical | Detected, not verified | pkg/apiserver/storage/testing/store_tests.go |

**Table 7.2 — Material debt findings**

| Finding ID | Marker | Severity | Evidence state | File |
|---|---|---|---|---|
| F-debt-001 | fixme | High | Detected, not verified | eslint.config.js:40 |

| Finding ID | Marker | Severity | Evidence state | File |
|---|---|---|---|---|
| F-debt-002 | fixme | High | Detected, not verified | eslint.config.js:624 |
| F-debt-003 | fixme | High | Detected, not verified | eslint.config.js:638 |
| F-debt-004 | fixme | High | Detected, not verified | eslint.config.js:658 |
| F-debt-005 | fixme | High | Detected, not verified | eslint.config.js:683 |

The deal-mechanic implication for the complexity findings is a price-adjustment lever treating the critical-complexity population as a priced integration work stream, sequenced with the plan the buyer commissions. For the debt findings, the implication is a holdback against delivery of the general-debt remediation programme. Quantification and sizing are for deal counsel; no amount is asserted in this report.

## 7.4 Complexity Percentiles

The cyclomatic complexity percentiles (P50=2, P90=7, P99=19) and cognitive complexity percentiles (P50=1, P90=12, P99=45) describe the distribution of structural complexity across the analysed symbols. The P50 values are favourable — most functions are simple. The P99 values, however, indicate that the most complex 1% of functions carry extreme complexity: a P99 cognitive complexity of 45 means a typical engineer would need significant time to understand the worst 1% of functions.

**Table 7.3 — Cyclomatic complexity percentiles**

| Metric | Value |
|---|---|
| P50 | 2 |
| P90 | 7 |
| P99 | 19 |

**Table 7.4 — Cognitive complexity percentiles**

| Metric | Value |
|---|---|
| P50 | 1 |
| P90 | 12 |
| P99 | 45 |

## 7.5 Hotspot Symbols

The hotspot table identifies the highest-complexity symbols by composite risk score. These are the symbols the buyer's post-close engineering team should refactor first. The Plot function in public/vendor/flot/jquery.flot.js (cyclomatic 473, cognitive 1,683) is the single most complex symbol — it is in vendored code, which means refactoring it requires either replacing the vendored library or accepting the complexity. The updateSchema function in DashboardMigrator.ts (cyclomatic 138) is the second-most-complex and is in first-party code, which makes it the highest-priority refactoring target.

**Table 7.5 — Hotspot symbols (top 16 of 57,667 rows)**

| Name | Filepath | Cyclo. | Cogn. | Maint. idx | Composite risk |
|---|---|---|---|---|---|
| Plot | public/vendor/flot/jquery.flot.js | 473 | 1,683 | 0 | 0.95 |
| updateSchema | public/app/features/dashboard/state/DashboardMigrator.ts | 138 | 210 | 0 | 0.90 |
| jsonDataToMetaJSONData | apps/plugins/pkg/app/meta/converter.go | 134 | 138 | 0 | 0.9125 |
| InitializeForTest | pkg/server/bootstrap/wire/wire_gen.go | 130 | 129 | 0 | 0.8575 |
| Initialize | pkg/server/bootstrap/wire/wire_gen.go | 130 | 129 | 0 | 0.87 |
| TestRouteGetRuleStatuses | pkg/services/ngalert/api/api_prometheus_test.go | 121 | 186 | 0 | 0.9125 |
| bytes2Value | pkg/util/xorm/session_convert.go | 121 | 558 | 0 | 0.90 |
| slice2Bean | pkg/util/xorm/session.go | 120 | 455 | 0 | 0.90 |
| getNavTitle | public/app/core/utils/navBarItem-translations.ts | 109 | 109 | 10.7 | 0.8911 |
| createGrafDir | pkg/tests/testinfra/testinfra.go | 102 | 109 | 0 | 0.9375 |
| authorize | pkg/services/ngalert/api/authorization.go | 90 | 84 | 0 | 0.90 |
| process | pkg/registry/apis/provisioning/controller/repository.go | 88 | 134 | 0 | 0.90 |
| graphToTimeseriesOptions | public/app/plugins/panel/timeseries/migrations.ts | 80 | 307 | 2.05 | 0.9084 |
| readDashboardIter | pkg/services/store/kind/dashboard/dashboard.go | 76 | 237 | 14.82 | 0.9329 |
| getStatementPosition | public/app/plugins/datasource/cloudwatch/language/cloudwatch-logs-sql/completion/statementPosition.ts | 73 | 103 | 10.96 | 0.8906 |
| buildUpdates | pkg/util/xorm/statement.go | 69 | 623 | 16.96 | 0.8661 |

*The composite risk score (0-1 scale, higher is riskier) weights cyclomatic, cognitive, and maintainability dimensions. A score >= 0.9 indicates a hotspot that should be refactored before any major modification. 16 of 16 shown symbols are >= 0.85, indicating concentrated structural risk in the codebase.*

## 7.6 Debt Markers by Type and Severity

Debt markers (in-source annotations like FIXME, TODO, HACK, XXX) total 883. None are on the critical path (0 critical-path markers), which is favourable. The type distribution shows 656 TODOs (74%), 121 debug-print-TS statements (14%), 44 FIXMEs (5%), and smaller counts of debug-print-JS, debug-print-PY, bare-except, hack, and xxx. The severity distribution shows 48 high, 656 medium, and 179 low.

**Table 7.6 — Debt markers by type**

| Marker type | Count |
|---|---|
| Bare Except | 5 |
| Debug Print JS | 38 |
| Debug Print Py | 13 |
| Debug Print TS | 121 |
| Fixme | 44 |
| Hack | 4 |
| Todo | 656 |
| Xxx | 2 |

**Table 7.7 — Debt markers by severity**

| Severity | Count |
|---|---|
| High | 48 |
| Low | 179 |
| Medium | 656 |

## 7.7 Debt Marker Inventory

The debt-marker inventory shows the eslint.config.js file hosts five FIXMEs at lines 40, 624, 638, 658, and 683 — four of which read 'FIXME: Remove once all enterprise issues are fixed'. This pattern indicates a known outstanding clean-up tied to enterprise / open-source boundary issues that the seller has deferred. The buyer's counsel should require disclosure of the specific 'enterprise issues' referenced and confirm whether they affect the licensed product.

**Table 7.8 — Debt marker inventory (top 16 of 883 rows)**

| Filepath | Line | Marker | Severity | Critical path | Text |
|---|---|---|---|---|---|
| eslint.config.js | 40 | fixme | high | No | // [FIXME] add comment about this applying everywhere |
| eslint.config.js | 624 | fixme | high | No | // FIXME: Remove once all enterprise issues are fixed - |
| eslint.config.js | 638 | fixme | high | No | // FIXME: Remove once all enterprise issues are fixed - |
| eslint.config.js | 658 | fixme | high | No | // FIXME: Remove once all enterprise issues are fixed - |
| eslint.config.js | 683 | fixme | high | No | // FIXME: Remove once all enterprise issues are fixed - |

| Filepath | Line | Marker | Severity | Critical path | Text |
|---|---|---|---|---|---|
| packages/grafana-api-clients/src/clients/rtkq/dashboard/v2/endpoints.gen.ts | 1,073 | fixme | high | No | /** Link type. Accepted values are dashboards (to refer to another dashboard) and link... |
| packages/grafana-api-clients/src/clients/rtkq/dashboard/v2beta1/endpoints.gen.ts | 1,617 | fixme | high | No | /** Link type. Accepted values are dashboards (to refer to another dashboard)... |
| packages/grafana-api-clients/src/clients/rtkq/provisioning/v0alpha1/endpoints.gen.ts | 1,710 | fixme | high | No | /** Paths to be deleted. Examples: dashboard.json (file) ... |
| packages/grafana-api-clients/src/clients/rtkq/provisioning/v0alpha1/endpoints.gen.ts | 1,734 | fixme | high | No | /** Paths to be deleted. Examples ... |
| packages/grafana-api-clients/src/clients/rtkq/provisioning/v0alpha1/endpoints.gen.ts | 1,762 | fixme | high | No | /** FIXME: we should validate this in admission hooks Target branch for export (only git) */ |
| packages/grafana-api-clients/src/clients/rtkq/provisioning/v0alpha1/endpoints.gen.ts | 1,770 | fixme | high | No | /** FIXME: we should validate this in admission hooks Prefix in target file system */ |
| packages/grafana-schema/src/schema/dashboard/v2/types.spec.gen.ts | 1,586 | fixme | high | No | // FIXME: The type is generated as `type: DashboardLinkType \| dashboardLinkType.Link;`... |
| packages/grafana-schema/src/schema/dashboard/v2alpha1/types.spec.gen.ts | 1,050 | fixme | high | No | // FIXME: The type is generated as `type: DashboardLinkType \| dashboardLinkType.Link;`... |
| packages/grafana-schema/src/schema/dashboard/v2beta1/types.spec.gen.ts | 1,578 | fixme | high | No | // FIXME: The type is generated as `type: DashboardLinkType \| dashboardLinkType.Link;`... |
| packages/grafana-test-utils/src/types/browse-dashboards.ts | 1 | fixme | high | No | // FIXME: This file is a duplication of types within the core code |
| packages/grafana-ui/src/components/Forms/commonStyles.ts | 39 | hack | high | No | /* Welcome to 2005. This is a HACK to get rid of Chromes default autofill styling */ |

*The HACK markers in packages/grafana-ui/src/components/Forms/commonStyles.ts ('Welcome to 2005. This is a HACK to get rid of Chromes default autofill styling') indicate known workarounds for browser quirks. These should not affect valuation but should be on the post-close refactoring backlog.*

## 7.8 SQALE Grade and Remediation Effort

The SQALE grade of E (worst on the A-E scale) is driven by a debt ratio of 26.4% — meaning technical debt represents over a quarter of the equivalent rewrite effort. The SQALE remediation effort of 108,265 days (≈ 421 person-years at 8h/day, 260 days/year) is the canonical planning figure. The model does not add SQALE as a separate line item to avoid double-counting with the complexity

hotspot category; instead it consumes SQALE as a planning bound. The buyer should treat the 108K-day figure as a long-tail backlog indicator, not as a remediation obligation to be sized into the price.

### 7.9 Duplication

828,217 lines of code are duplicated across 18,406 groups, representing a 25.2% duplication ratio against the 57,667 analysed symbols. A duplication ratio above 20% is a structural concern: it indicates the codebase carries copy-paste legacy that amplifies the cost of any change to the duplicated logic. The buyer's post-close engineering team should commission a duplication-removal sprint focused on the highest-occurrence groups first.

© MTD Services

# 8. Testing, Coverage & Regression Risk

## 8.1 Introduction

Where the codebase is structurally covered by tests and where it is not drives the leading indicator of post-close regression incidents and the cost of mandated refactoring. The effective insured value translates directly into contingent-liability sizing. Test-ratio module rows, coverage ratios, and high-value gap modules below are stored ground truth. The pillar risk is high.

The headline figures: 2,506 modules analysed; 1,882 below threshold; overall test ratio 15.3%; 42 high regression-risk modules; effective insured value 1; value coverage ratio 4.8%; 592 high-value gaps. The coverage method is a static symbol-share proxy, not runtime or line/branch coverage.

## 8.2 Key Figures

The overall test ratio of 15.3% means that across the 2,506 analysed modules, only 15.3% of the file count is test files. The value coverage ratio of 4.8% means that only 4.8% of high-blast-symbol value is insured by tests — meaning 95.2% of high-blast symbols are uninsured. The effective insured value of 1 (on a normalised scale where the most valuable module = 1) means that only one module — pkg/storage/unified/resource with 4,210 production symbols — is fully insured; every other module has partial or zero coverage.

## 8.3 Test-Deficit Module Inventory

The 2,506 analysed modules include 1,882 below threshold (75.1%). The 14 modules shown below all carry a test ratio of 1.0 (all test files, no source files) — these are e2e-playwright test suites and grafana-alerting test directories. The modules with ratio 1.0 do not contribute to the test deficit; the deficit is concentrated in the 1,882 modules with ratio below threshold (typically below 0.3).

**Table 8.1 — Test-deficit module inventory (top 14 of 2,506 rows)**

| Module path | Test files | Source files | Test ratio | Regression risk |
|---|---|---|---|---|
| e2e-playwright/plugin-catalog-suite/tests | 2 | 0 | 1 | 0 |
| e2e-playwright/plugin-catalog-suite/tests/useMTPlugins | 1 | 0 | 1 | 0 |
| e2e-playwright/test-plugins/grafana-extensionstest-app/tests | 4 | 0 | 1 | 0 |
| e2e-playwright/test-plugins/grafana-extensionstest-app/tests/translations | 2 | 0 | 1 | 0 |
| e2e-playwright/test-plugins/grafana-extensionstest-app/tests/useMTPlugins | 3 | 0 | 1 | 0 |
| e2e-playwright/test-plugins/grafana-test-app/tests | 2 | 0 | 1 | 0 |
| e2e-playwright/test-plugins/grafana-test-app/tests/useMTPlugins | 1 | 0 | 1 | 0 |
| e2e-playwright/test-plugins/grafana-test-datasource/tests | 4 | 0 | 1 | 0 |

DTD Services

| Module path | Test files | Source files | Test ratio | Regression risk |
|---|---|---|---|---|
| e2e-playwright/test-plugins/grafana-test-datasource/tests/translations | 2 | 0 | 1 | 0 |
| e2e-playwright/test-plugins/grafana-test-datasource/tests/useMTPlugins | 1 | 0 | 1 | 0 |
| e2e-playwright/test-plugins/grafana-test-panel/tests | 3 | 0 | 1 | 0 |
| e2e-playwright/test-plugins/grafana-test-panel/tests/translations | 2 | 0 | 1 | 0 |
| e2e-playwright/test-plugins/grafana-test-panel/tests/useMTPlugins | 2 | 0 | 1 | 0 |
| grafana-alerting/tests | 3 | 0 | 1 | 0 |

## 8.4 Coverage Module Detail

The coverage module detail is the central regression-risk indicator. The 14 modules shown below have 0 coverage despite carrying high production-symbol counts. pkg/storage/unified/resource (4,210 symbols) is the single most material gap — it carries the maximum risk weight of 1.0 and has zero coverage. The buyer's post-close engineering team should treat these modules as the priority for test investment; the test_investment line item ($20,400) is sized for 42 high-regression-risk modules but the actual backlog is much larger.

**Table 8.2 — Coverage module detail (top 14 of 2,506 rows)**

| Module | Prod symbols | Coverage ratio | Risk weight | Value exposure |
|---|---|---|---|---|
| pkg/storage/unified/resource | 4,210 | 0 | 1 | 1 |
| pkg/storage/unified/search | 1,469 | 0 | 0.3718 | 0.3602 |
| pkg/api | 1,590 | 0 | 0.2452 | 0.2876 |
| dashboard/pkg/migration/conversion | 574 | 0 | 0.2455 | 0.2471 |
| pkg/util/xorm | 672 | 0 | 0.2331 | 0.2401 |
| pkg/registry/apis/provisioning/resources | 919 | 0 | 0.2289 | 0.2388 |
| pkg/setting | 1,353 | 0 | 0.1676 | 0.2264 |
| pkg/router | 1,342 | 0 | 0.2193 | 0.2232 |
| iam/pkg/apis/iam/v0alpha1 | 1,299 | 0 | 0.1433 | 0.1958 |
| pkg/services/ngalert/api | 981 | 0 | 0.1755 | 0.1932 |
| pkg/registry/apis/folders | 453 | 0 | 0.2180 | 0.1865 |
| pkg/services/ngalert/notifier | 873 | 0 | 0.1569 | 0.1857 |
| pkg/registry/apis/provisioning/controller | 710 | 0 | 0.1771 | 0.1807 |
| pkg/registry/apis/datasource | 411 | 0 | 0.2128 | 0.1732 |

## 8.5 High-Value Coverage Gaps

The high-value gaps table isolates the modules whose coverage gap carries the highest risk weight. These are the modules where the buyer's quality-of-earnings provider should focus regression-risk validation. The 592 high-value gaps total means 592 modules carry a risk weight above the high-value threshold and have zero or near-zero coverage. The deal-mechanic implication is a price-adjustment lever raising regression coverage on high-risk modules post-close.

**Table 8.3 — High-value coverage gaps (top 14 of 592 rows)**

| Module | Risk weight | Coverage ratio |
|---|---|---|
| pkg/storage/unified/resource | 1 | 0 |
| pkg/storage/unified/search | 0.3718 | 0 |
| dashboard/pkg/migration/conversion | 0.2455 | 0 |
| pkg/api | 0.2452 | 0 |
| pkg/util/xorm | 0.2331 | 0 |
| pkg/registry/apis/provisioning/resources | 0.2289 | 0 |
| pkg/router | 0.2193 | 0 |
| pkg/registry/apis/folders | 0.2180 | 0 |
| pkg/registry/apis/datasource | 0.2128 | 0 |
| public/app/plugins/panel/geomap/layers/basemaps | 0.2058 | 0 |
| pkg/registry/apis/provisioning/controller | 0.1771 | 0 |
| pkg/services/ngalert/api | 0.1755 | 0 |
| pkg/setting | 0.1676 | 0 |
| pkg/registry/apps/alerting/rules/search | 0.1668 | 0 |

## 8.6 Coverage Method

Coverage is measured by a static symbol-share proxy: Static symbol-share proxy (test-file symbols weighted by structural value); not runtime or line/branch coverage. This means the reported coverage is structural (test-file symbols exist) rather than runtime (lines or branches executed). The buyer's post-close engineering team should commission runtime coverage as a follow-up, prioritising the modules with high structural risk weight and zero structural coverage.

## 8.7 M&A Implications

The 95.2% uninsured exposure on high-blast symbols is the most consequential testing finding. It means that any refactoring the buyer commissions post-close will be undertaken without a regression safety net for the vast majority of high-value code. The 42 high-regression-risk modules sized under the test_investment line item ($20,400) are a planning bound for the most acute gap, not a comprehensive remediation figure. The buyer's counsel should require (a) a price-adjustment lever raising regression coverage on high-risk modules post-close, (b) a holdback against delivery of the

test-investment programme, and (c) a warranty covering the absence of known regression defects in the modules with zero coverage.

© MTD Services

# 9. Architecture, Entanglement & Integration Risk

## 9.1 Introduction

The architectural shape determines how cleanly the asset can be integrated: layering discipline, cyclic dependencies, cross-layer violations, API-contract maturity, database coupling, and serialization risks. Integration cost is the most under-estimated line item in software acquisition; entanglement and integration-risk scores give the integration-management office a defensible basis for the synergy model and transition-services scope. Cycle inventories, cross-layer violations, API-contract and coupling metrics, and serialization-risk rows below are stored ground truth. The pillar risk is rated critical (Entanglement) and high (Integration).

## 9.2 Key Figures

Circular-dependency cycles: 4. Max cycle length: 1,445 modules. Cross-layer violations: 19,887. Modules in cycles: 1,451. Entanglement score: 1. Integration risk score: 10 (maximum on the 0-10 scale). Schema drift risk: High. Proprietary coupling count: 509. Integration findings total: 521.

## 9.3 Cycle Inventory

Four cyclic dependencies are detected. The largest spans 1,445 modules — an extreme architectural pathology indicating deep inter-dependency among the affected modules. The other three cycles are smaller (2 modules each) but still flag decoupling work. The decoupling line item ($4,000) is sized for all four cycles at $1,000 each, but the 1,445-module cycle will require materially more effort than the small cycles; the buyer's engineering team should re-estimate using a code-complexity-weighted approach.

**Table 9.1 — Cycle inventory**

| Cycle length | Severity | Modules |
|---|---|---|
| 1,445 | Critical | 1,445 modules |
| 2 | Medium | 2 modules |
| 2 | Medium | 2 modules |
| 2 | Medium | 2 modules |

## 9.4 Cross-Layer Violations

19,887 cross-layer violations are recorded — an extreme count indicating pervasive architectural layering breaches. The violations are predominantly presentation-to-data calls (e.g. alerting alertenrichment / alertrulequality / notifications / rules APIs calling into pkg/expr, which is the data layer). This pattern indicates that the API-layer types are reaching past the service layer into expression evaluation, which means the service layer is not actually abstracting data access. The buyer's post-close architectural review should commission a layering audit and a remediation plan.

© TD Services

**Table 9.2 — Cross-layer violation detail (top 14 of 19,887 rows)**

| From module | To module | From layer | To layer | Rel type | Severity |
|---|---|---|---|---|---|
| alerting/alertenrichment/pkg/apis/alertenrichment/v1beta1 | pkg/expr | presentation | data | calls | Critical |
| alerting/alertrulequality/pkg/apis/alertrulequality/v0alpha1 | pkg/expr | presentation | data | calls | Critical |
| alerting/notifications/pkg/apis/alertingnotifications/v0alpha1/fakes | pkg/expr | presentation | data | calls | Critical |
| alerting/rules/pkg/apis/alerting/v0alpha1 | pkg/expr | presentation | data | calls | Critical |
| collections/pkg/apis/collections/v1alpha1 | pkg/expr | presentation | data | calls | Critical |
| dashboard/pkg/apis/dashboard/v0alpha1 | pkg/expr | presentation | data | calls | Critical |
| dashboard/pkg/apis/dashboard/v0alpha1 | dashboard/pkg/migration/schemaversion | presentation | data | imports | Critical |
| dashboard/pkg/apis/dashboard/v1 | pkg/expr | presentation | data | calls | Critical |
| dashboard/pkg/apis/dashboard/v1 | dashboard/pkg/migration/schemaversion | presentation | data | imports | Critical |
| dashboard/pkg/apis/dashboard/v2 | dashboard/pkg/migration/conversion | presentation | data | calls | Critical |
| dashboard/pkg/apis/dashboard/v2 | pkg/expr | presentation | data | calls | Critical |
| dashboard/pkg/apis/dashboard/v2alpha1 | dashboard/pkg/migration | presentation | data | calls | Critical |
| dashboard/pkg/apis/dashboard/v2alpha1 | pkg/expr | presentation | data | calls | Critical |
| dashboard/pkg/apis/dashboard/v2beta1 | pkg/expr | presentation | data | calls | Critical |

*Every shown violation is critical-severity and presentation-to-data. The 19,887 total indicates that the API surface is structurally coupled to the data layer across alerting, collections, dashboard, and provisioning subsystems. This pattern is the single most material architectural finding in the scan.*

## 9.5 API Contract Maturity

API contract maturity is assessed on five metrics. 17 endpoints are detected; 13 proto files define the contract surface; 152 types are declared. OpenAPI and GraphQL are both absent, which means the REST API surface is not declared via a standard specification — this drives the high schema-drift-risk rating and contributes to the integration risk score of 10.

**Table 9.3 — API contract maturity**

| Metric | Value |
|---|---|
| Endpoint count | 17 |

| Metric | Value |
|---|---|
| GraphQL found | No |
| OpenAPI found | No |
| Proto files | 13 |
| Type count | 152 |

*The absence of OpenAPI means the buyer's API-consumer teams cannot rely on a published specification; they must reverse-engineer the contract from source. The 13 proto files indicate gRPC-style contracts, which are typically machine-generated and well-typed; the buyer should request the generated client libraries as part of transition-services scope.*

## 9.6 Database Coupling

Database coupling shows cross-schema access (Yes), 40 destructive DDL statements, 178 migration files, 0 migration tooling entries, no ORM usage, and 2,601 raw SQL statements. The combination of cross-schema access and 2,601 raw SQL statements means the application is structurally coupled to a specific database schema and dialect — migrating to a different database (e.g. for an acquirer requiring PostgreSQL-on-cloud-native) would be a major engineering effort. The 40 destructive DDL statements should be reviewed for irreversible schema changes that have been applied historically.

**Table 9.4 — Database coupling**

| Metric | Value |
|---|---|
| Cross schema access | Yes |
| Destructive DDL count | 40 |
| Migration files | 178 |
| Migration tooling | 0 |
| ORM usage | No |
| Raw SQL count | 2,601 |

*The absence of an ORM is unusual for a codebase of this size and is a deliberate architectural choice. The buyer's engineering team should validate that the seller's choice is intentional (e.g. for performance reasons) rather than legacy. The 178 migration files indicate a long migration history; the buyer should request the migration runbook as part of operational handover.*

## 9.7 Serialization Risks

Serialization risks include dynamic_function calls and eval_call patterns — both are dangerous deserialization patterns that can lead to remote code execution if untrusted input reaches them. 509 proprietary-coupling findings and 521 total integration findings are recorded. The eval_call cluster in pkg/api/admin.go (lines 93, 103, 123, 134, 144) is particularly concerning because it is in the admin API surface, which is typically a privileged endpoint. The buyer's security team should require validation of these patterns as part of the Secrets / CVE condition precedent.

© Services

**Table 9.5 — Serialization risks (top 14 of 509 rows)**

| Risk type | Filepath | Line | Severity | Description |
|---|---|---|---|---|
| serialization:dynamic_function | e2e-playwright/panels-suite/xychart.spec.ts | 91 | High | Dangerous deserialization: dynamic_function |
| serialization:dynamic_function | public/app/features/dashboard/dashgrid/DashboardLibrary/utils/communityDashboardHelpers.ts | 181 | High | Dangerous deserialization: dynamic_function |
| serialization:dynamic_function | scripts/rspack/hmrConfig.test.ts | 91 | High | Dangerous deserialization: dynamic_function |
| serialization:dynamic_function | public/app/features/dashboard/services/DashboardLoaderSrv.ts | 102 | High | Dangerous deserialization: dynamic_function |
| serialization:dynamic_function | public/app/plugins/panel/xychart/scatter.ts | 684 | High | Dangerous deserialization: dynamic_function |
| serialization:dynamic_function | public/app/plugins/panel/xychart/scatter.ts | 682 | High | Dangerous deserialization: dynamic_function |
| serialization:eval_call | public/app/features/dashboard/dashgrid/DashboardLibrary/utils/communityDashboardHelpers.ts | 180 | High | Dangerous deserialization: eval_call |
| serialization:eval_call | pkg/expr/threshold.go | 212 | High | Dangerous deserialization: eval_call |
| serialization:eval_call | pkg/expr/threshold.go | 209 | High | Dangerous deserialization: eval_call |
| serialization:eval_call | pkg/api/admin.go | 134 | High | Dangerous deserialization: eval_call |
| serialization:eval_call | pkg/api/admin.go | 93 | High | Dangerous deserialization: eval_call |
| serialization:eval_call | pkg/api/admin.go | 103 | High | Dangerous deserialization: eval_call |
| serialization:eval_call | pkg/api/admin.go | 144 | High | Dangerous deserialization: eval_call |
| serialization:eval_call | pkg/api/admin.go | 123 | High | Dangerous deserialization: eval_call |

## 9.8 Protocol-Layer Risks

Twelve protocol-layer risks are recorded, all related to raw socket usage for WebSocket handling in Grafana Live. The risks are medium-severity because raw socket usage is expected for WebSocket implementations, but the buyer's security team should validate that the WebSocket handlers do not accept untrusted input directly. The pkg/services/live subsystem is the main location; the buyer should require disclosure of the WebSocket message-rate-limiting and authentication posture.

**Table 9.6 — Protocol-layer risks**

| Risk type | Filepath | Line | Severity | Description | Evidence |
|---|---|---|---|---|---|
| protocol:raw_socket | public/app/features/apiserver/client.ts | 203 | Medium | Raw socket usage detected | Polling fallback when the Grafana Live WebSocket Observable errors. |
| protocol:raw_socket | pkg/services/live/live.go | 490 | Medium | Raw socket usage detected | GrafanaLive manages live real-time connections to Grafana (over WebSocket). |
| protocol:raw_socket | pkg/services/live/pushws/push_stream.go | 17 | Medium | Raw socket usage detected | Handler handles WebSocket client connections that push data to Live. |
| protocol:raw_socket | pkg/setting/setting.go | 570 | Medium | Raw socket usage detected | LiveMaxConnections is a maximum number of WebSocket connections. |
| protocol:raw_socket | pkg/services/live/pushws/push_pipeline.go | 14 | Medium | Raw socket usage detected | PipelinePushHandler handles WebSocket client connections that push data to Live Pipeline. |
| protocol:raw_socket | pkg/router/watch.go | 32 | Medium | Raw socket usage detected | watch over WebSocket. The router does not proxy upgrades. |
| protocol:raw_socket | public/app/features/apiserver/client.test.ts | 351 | Medium | Raw socket usage detected | const wsError = new Error('WebSocket failed'); |
| protocol:raw_socket | public/app/features/apiserver/client.test.ts | 370 | Medium | Raw socket usage detected | const wsError = new Error('WebSocket failed'); |
| protocol:raw_socket | pkg/router/watch.go | 38 | Medium | Raw socket usage detected | http.Error(w, 'the router does not support protocol upgrades...') |
| protocol:raw_socket | pkg/middleware/gziper.go | 103 | Medium | Raw socket usage detected | prefix('/api/live/push') - WebSocket does not support gzip compression. |
| protocol:raw_socket | pkg/middleware/gziper.go | 102 | Medium | Raw socket usage detected | prefix('/api/live/ws') - WebSocket does not support gzip compression. |
| protocol:raw_socket | pkg/router/watch_test.go | 356 | Medium | Raw socket usage detected | require.Contains(t, string(body), 'watch over WebSocket is not supported') |

## 9.9 M&A Implications

The integration risk score of 10 (maximum) is the headline architectural finding. It means the buyer's integration-management office should expect the integration to require significant transition-services scope, a dedicated integration team, and a longer integration timeline than a typical SaaS acquisition. The 1,445-module cycle and the 19,887 cross-layer violations are the dominant cost drivers. The buyer's counsel should draft the transition-services agreement with explicit milestones for architectural decoupling, sized against the decoupling line item ($4,000 — a planning bound, not a priced quote) and the integration-drag percentage (80%).

The deal-mechanic tag for Entanglement is a price-adjustment lever requiring a plan for decoupling circular dependency cycles. The deal-mechanic tag for Integration is a price-adjustment lever addressing integration findings as a condition precedent — the wording 'as a condition precedent' is unusual for a price-adjustment lever and the buyer's counsel should clarify whether the seller intends this as a blocking condition or a priced work stream. The schema-drift risk rating of high should drive a warranty covering API-contract stability for the buyer's API-consumer teams.

© MTD Services

# 10. Engineering Capacity, Bus Factor & Knowledge Loss Vulnerability

## 10.1 Introduction

The human-capital risk embedded in the codebase drives retention-bonus sizing, earn-out structure, and the transition-services window. Sole-maintainer modules are a primary cause of post-close silent outages and warrant explicit warranty language. Concentration risk is recorded across 3,217 contributors, with 510 at-risk contributors and 1,552 critical sole-maintainer modules. The knowledge-loss vulnerability score of 0.531 (on a 0-1 scale, higher is riskier) is the headline human-capital figure.

## 10.2 Key Figures

Total authors: 3,217. Overall bus factor: 33. Critical key-person rows: 0. Key-person risk rows: 3,217. Module bus-factor rows: 2,449.

Retention: total contributors 3,217; active 591; inactive 2,626; at-risk 510; knowledge loss vulnerability score 0.531; critical modules 4,548; at-risk contributors (model input) 510; critical sole-maintainer modules (model input) 1,552.

## 10.3 Material Findings

Five key-person findings (all low severity, inferred) and five retention findings (all critical severity, inferred) are recorded. Each is a scanner signal requiring validation before negotiation reliance. The deal-mechanic implication for key-person is a named closing condition: identify the bus-factor holders and bind their transition terms. For retention, the implication is securing retention or transition agreements from the contributors carrying institutional knowledge as a condition to key-employee arrangements.

**Table 10.1 — Material key-person findings**

| Finding ID | Author | Severity | Evidence state |
|---|---|---|---|
| F-key_person-001 | ryantxu@gmail.com | Low | Inferred |
| F-key_person-002 | ashley.harrison@grafana.com | Low | Inferred |
| F-key_person-003 | joshhunt@users.noreply.github.com | Low | Inferred |
| F-key_person-004 | torkel@grafana.com | Low | Inferred |
| F-key_person-005 | stephanie.hingtgen@grafana.com | Low | Inferred |

**Table 10.2 — Material retention findings**

| Finding ID | Contributor | Severity | Evidence state |
|---|---|---|---|
| F-retention-001 | Alex Khomenko | Critical | Inferred |
| F-retention-002 | Roberto Jiménez Sánchez | Critical | Inferred |
| F-retention-003 | Kevin Minehart | Critical | Inferred |
| F-retention-004 | Sofia Papagiannaki | Critical | Inferred |

DD Services

| Finding ID | Contributor | Severity | Evidence state |
|---|---|---|---|
| F-retention-005 | Galen Kistler | Critical | Inferred |

The five critical retention findings identify the contributors whose departure would destabilise the most institutional knowledge: Alex Khomenko (blast radius 246 modules), Roberto Jiménez Sánchez, Kevin Minehart, Sofia Papagiannaki (blast radius 246 modules, tenure 1,820 days), and Galen Kistler (blast radius 68 modules). Each should be the subject of an individual retention-bonus negotiation in the post-close plan.

## 10.4 Key-Person Risk Detail

The top key-person risk rows identify the contributors with the highest concentration of activity. ryantxu@gmail.com carries a risk score of 0.265 with 234,722 commits — extreme concentration in a single contributor. ashley.harrison@grafana.com (0.216, 131K commits) and joshhunt@users.noreply.github.com (0.104, 63K commits) follow. The departure-impact is rated low for every shown contributor because the bus factor of 33 across the codebase is sufficient to absorb any single departure; the risk is in the long tail of sole-maintainer modules, not in the top contributors.

**Table 10.3 — Key-person risk detail (top 14 of 3,217 rows)**

| Author name | Author email | Risk score | Total commits | Departure impact |
|---|---|---|---|---|
| ryantxu@gmail.com | ryantxu@gmail.com | 0.26468 | 234,722 | Low |
| ashley.harrison@grafana.com | ashley.harrison@grafana.com | 0.216071 | 131,354 | Low |
| joshhunt@users.noreply.github.com | joshhunt@users.noreply.github.com | 0.103575 | 63,357 | Low |
| torkel@grafana.com | torkel@grafana.com | 0.0803188 | 90,155 | Low |
| stephanie.hingtgen@grafana.com | stephanie.hingtgen@grafana.com | 0.0686434 | 40,405 | Low |
| hugo.haggmark@gmail.com | hugo.haggmark@gmail.com | 0.0615704 | 65,580 | Low |
| 29139614+renovate[bot]@users.noreply.github.com | 29139614+renovate[bot]@users.noreply.github.com | 0.0612221 | 11,923 | Low |
| roberto.jimenez@grafana.com | roberto.jimenez@grafana.com | 0.0528524 | 132,108 | Low |
| macabu@users.noreply.github.com | macabu@users.noreply.github.com | 0.050841 | 31,702 | Low |
| leeoniya@gmail.com | leeoniya@gmail.com | 0.0498112 | 48,560 | Low |
| 45561153+kaydelaney@users.noreply.github.com | 45561153+kaydelaney@users.noreply.github.com | 0.0477866 | 28,972 | Low |
| marcus.efraimsson@gmail.com | marcus.efraimsson@gmail.com | 0.0463529 | 26,936 | Low |

| Author name | Author email | Risk score | Total commits | Departure impact |
|---|---|---|---|---|
| academo@users.noreply.github.com | academo@users.noreply.github.com | 0.0456318 | 7,270 | Low |
| laura.fernandez@grafana.com | laura.fernandez@grafana.com | 0.041486 | 11,503 | Low |

## 10.5 Module Bus-Factor Detail

Module bus-factor rows identify modules where the contributor count is concentrated. pkg/setting (bus factor 22) is the most resilient module; pkg/api (12), _unassigned (9, but 6,605 files), grafana-data/src/types (9), and pkg/services/sqlstore (9) follow. The lower the bus factor, the greater the concentration risk. Modules with bus factor of 1 are sole-maintainer — these are the modules sized under the knowledge-transfer line item.

**Table 10.4 — Module bus-factor detail (top 14 of 2,449 rows)**

| Module path | Bus factor | Total files |
|---|---|---|
| pkg/setting | 22 | 54 |
| pkg/api | 12 | 96 |
| _unassigned | 9 | 6,605 |
| grafana-data/src/types | 9 | 58 |
| pkg/services/sqlstore | 9 | 21 |
| public/app/types | 9 | 26 |
| pkg/api/dtos | 8 | 18 |
| grafana-data/src/utils | 7 | 51 |
| pkg/registry/backgroundsvcs | 7 | 1 |
| pkg/server | 7 | 29 |
| pkg/services/sqlstore/migrations | 7 | 41 |
| pkg/tests/testinfra | 7 | 2 |
| pkg/util | 7 | 33 |
| public/app/core/utils | 7 | 53 |

## 10.6 Git Team Signals

The git team signals reveal several stress indicators. The after-hours-commits percentage of 0.329 (32.9%) is high and indicates either a globally distributed team or sustained crunch-mode development. The bus-factor proxy of 0.906 indicates the team is structurally healthy on aggregate (the proxy approaches 1.0 when contributor concentration is low). The cadence of 1,000 commits per week is sustained, indicating active development. The ghost-author count of 42 (contributors with no email attribution) is a data-quality signal that should be cleaned up before close.

**Table 10.5 — Git team signals**

| Metric | Value |
|---|---|
| After-hours commits | 658 |
| After-hours percentage | 0.329 |
| Bus factor proxy | 0.906 |
| Cadence per week | 1,000 |
| Ghost author count | 42 |
| Modules analyzed | 39 |
| Top-1 author share | 0.094 |
| Top-3 author share | 0.1705 |
| Top-5 author share | 0.222 |
| Total commits | 2,000 |
| Unique authors | 184 |
| Unique emails | 179 |

*The top-1 author share of 0.094 (9.4%) indicates reasonable contributor distribution at the aggregate level. The top-3 (17%) and top-5 (22%) shares confirm that no single contributor dominates commit volume at the codebase level — the concentration risk is concentrated in the long tail of sole-maintainer modules, not in the top contributors.*

## 10.7 Knowledge Loss Vulnerability Detail

The knowledge-loss vulnerability detail identifies contributors whose tenure, recent activity, and sole-maintainership profile combine to create critical retention risk. Each shown contributor has been classified as critical retention risk based on a multi-factor score. The buyer's HR team should obtain retention terms for each named contributor before close; the retention-bonus pool should be sized against the knowledge-transfer line item ($316,400).

**Table 10.6 — Knowledge loss vulnerability detail (top 14 of 3,217 rows)**

| Author name | Author email | Tenure days | Days since last commit | Retention risk | Blast radius at risk |
|---|---|---|---|---|---|
| undef1nd | 10127682+undef1nd@users.noreply.github.com | 1,021 | 0 | Critical | 19 |
| Galen Kistler | 109082771+gtk-grafana@users.noreply.github.com | 1,534 | 0 | Critical | 68 |
| Fayzal Ghantiwala | 114010985+fayzal-g@users.noreply.github.com | 896 | 55 | Critical | 0 |
| volcanonoodle | 114113189+volcanonoodle@users.noreply.github.com | 1,063 | 2 | Critical | 72 |

MTD Set

| Author name | Author email | Tenure days | Days since last commit | Retention risk | Blast radius at risk |
|---|---|---|---|---|---|
| Liza Detrick | 114438185+L2D2Grafana@users.noreply.github.com | 883 | 7 | Critical | 0 |
| Shubham Nainwal | 118589955+Shubham19032004@users.noreply.github.com | 491 | 93 | Critical | 0 |
| grafana-delivery-bot[bot] | 132647405+grafana-delivery-bot[bot]@users.noreply.github.com | 1,212 | 150 | Critical | 0 |
| Fabrizio | 135109076+fabrizio-grafana@users.noreply.github.com | 1,198 | 1 | Critical | 42 |
| Lucy Chen | 140550297+lucychen-grafana@users.noreply.github.com | 1,154 | 21 | Critical | 0 |
| Tobias Skarhed | 1438972+tskarhed@users.noreply.github.com | 2,455 | 8 | Critical | 30 |
| Marc M. | 146180665+grafakus@users.noreply.github.com | 716 | 0 | Critical | 121 |
| alerting-team[bot] | 158350966+alerting-team[bot]@users.noreply.github.com | 591 | 0 | Critical | 0 |
| Adam Yeats | 16296989+adamyeats@users.noreply.github.com | 1,186 | 3 | Critical | 15 |
| Sofia Papagiannaki | 1632407+papagian@users.noreply.github.com | 1,820 | 0 | Critical | 246 |

## 10.8 M&A Implications

The 510 at-risk contributors and 1,552 critical sole-maintainer modules are the most actionable human-capital findings. The 81.7% inactive contributor ratio (2,626 of 3,217) means the active development team is materially smaller than the historical contributor base suggests; the buyer's HR team should obtain the current headcount and the active-contributor list as part of due diligence. The 32.9% after-hours commit percentage is a stress indicator that should be discussed in management presentations — sustained high after-hours ratios correlate with retention risk.

The deal-mechanic tag for Retention is a price-adjustment lever securing retention agreements for sole-maintainer modules. The knowledge-transfer line item ($316,400) is sized per unique module, not per at-risk contributor — meaning the buyer's HR team can choose to deploy the budget against the most acute sole-maintainer modules rather than against individual contributors. The buyer's counsel should draft retention agreements with vesting tied to a transition-services window of at least 12 months post-close.

© MTD Services

# 11. Maintenance, Documentation & Operational Maturity

## 11.1 Introduction

Documentation and CI maturity determine onboarding velocity for the acquirer's engineers and the run-rate cost of keeping the asset alive during integration. The maintenance backlog score anchors the post-close engineering-budget uplift. The pillar risk for Maintenance is unverified — no files were scanned by the asset-staleness scanner. The pillar risk for Doc Staleness is low. The pillar risk for CI Maturity is medium.

## 11.2 Key Figures — Maintenance

Maintenance total files: 0. Neglected files: 0. At-risk files: 0. Healthy files: 0. Average staleness: 0. Maintenance backlog score: 0. Maintenance asset rows: 0.

The maintenance pillar is NOT ASSESSED. The asset-staleness scanner produced zero rows because no files were scanned under its current configuration. The pillar is demoted to informational and labelled UNVERIFIED. The buyer should commission an independent asset-staleness assessment as part of pre-signing diligence — the absence of a staleness figure is a coverage gap, not a clean result.

## 11.3 Key Figures — Documentation

Documents inventoried: 797. Stale documents: 63. Document staleness ratio: 7.9%. Average docstring density: 0.0003. The 7.9% staleness ratio means 63 of 797 inventoried documentation files are stale relative to the staleness threshold. The docstring density of 0.0003 is effectively zero — the codebase has negligible inline documentation, which materially affects onboarding velocity for the acquirer's engineers.

**Table 11.1 — Document inventory (top 14 of 797 rows)**

| File | Last documented | Stale |
|---|---|---|
| docs/sources/developer-resources/api-reference/http-api/apis.md | 2026-07-07 | Yes |
| docs/sources/developer-resources/api-reference/http-api/api-legacy/correlations.md | 2026-06-26 | Yes |
| docs/sources/developer-resources/api-reference/http-api/dashboard.md | 2026-07-01 | Yes |
| docs/sources/shared/test.md | 2023-07-18 | Yes |
| docs/sources/developer-resources/api-reference/http-api/api-legacy/annotations.md | 2026-09-14 | Yes |
| docs/sources/shared/example.md | 2023-08-31 | Yes |
| docs/sources/developer-resources/api-reference/http-api/folder.md | 2026-07-01 | Yes |
| docs/sources/developer-resources/api-reference/http-api/playlist.md | 2026-07-01 | Yes |
| docs/sources/developer-resources/api-reference/http-api/api-legacy/preferences.md | 2026-06-26 | Yes |

TD Services

| File | Last documented | Stale |
|---|---|---|
| docs/sources/datasources/tempo/configure-tempo-data-source/provision.md | 2026-04-09 | Yes |
| docs/sources/administration/cli.md | 2025-10-28 | Yes |
| docs/sources/developer-resources/api-reference/http-api/api-legacy/admin.md | 2026-09-22 | Yes |
| docs/sources/shared/back-up/back-up-grafana.md | 2025-06-19 | Yes |
| docs/sources/developer-resources/api-reference/http-api/api-legacy/team.md | 2026-06-26 | Yes |

*The stale documents include API reference files (most recently modified 2026-06-26 to 2026-09-22) and shared documentation files (some dating to 2023). The 7.9% ratio is favourable in aggregate, but the docstring density of 0.0003 is the more concerning figure because it indicates near-total absence of inline code documentation.*

## 11.4 Key Figures — CI Maturity

CI defined: true. CI files checked: 98. CI gate strictness: Medium. Tests run in CI: true. Reported coverage in CI: Not quantified.

The CI surface is substantial — 98 GitHub Actions workflows are checked. The gate strictness is medium, which means CI runs tests but does not enforce coverage thresholds or block on coverage regressions. The reported coverage metric is not quantified, which is a data-quality gap: the seller's CI runs tests but does not publish a coverage report that the scanner could consume. The buyer's engineering team should require a coverage gate as part of post-close CI improvement.

**Table 11.2 — CI files (top 20 of 98 rows)**

| CI file path |
|---|
| /content/grafana/.github/workflows/actionlint.yml |
| /content/grafana/.github/workflows/add-to-whats-new.yml |
| /content/grafana/.github/workflows/alerting-swagger-gen.yml |
| /content/grafana/.github/workflows/alerting-update-module.yml |
| /content/grafana/.github/workflows/analytics-events-report.yml |
| /content/grafana/.github/workflows/auto-milestone.yml |
| /content/grafana/.github/workflows/backend-code-checks.yml |
| /content/grafana/.github/workflows/backend-unit-tests.yml |
| /content/grafana/.github/workflows/backport-trigger.yml |
| /content/grafana/.github/workflows/backport-workflow.yml |
| /content/grafana/.github/workflows/build-docker-variants.yml |
| /content/grafana/.github/workflows/build-go-matrix.yml |
| /content/grafana/.github/workflows/bump-version.yml |
| /content/grafana/.github/workflows/changelog.yml |

D Services

| CI file path |
|---|
| /content/grafana/.github/workflows/check-frontend-test-coverage.yml |
| /content/grafana/.github/workflows/cleanup-branches.yml |
| /content/grafana/.github/workflows/codeowners-validator.yml |
| /content/grafana/.github/workflows/commands.yml |
| /content/grafana/.github/workflows/community-release.yml |
| /content/grafana/.github/workflows/create-next-release-branch.yml |

### 11.5 Coverage Artifacts

No coverage artifacts are recorded in this scan. This is consistent with the reported_coverage_pct = not quantified figure: the CI surface does not publish a coverage report that the scanner could consume. The buyer's engineering team should require the seller to publish a coverage report (e.g. via Codecov or Coveralls) and to enforce a coverage threshold as part of the post-close CI improvement plan.

### 11.6 M&A Implications

The maintenance pillar's unverified status is the most consequential finding in this section. Because no files were scanned, the buyer cannot rely on the staleness figure (which reads 0) — it is an artefact of the scanner configuration, not a clean result. The buyer should commission an independent asset-staleness assessment covering the same file scope as the complexity scan (16,662 files) and should require the seller to remediate any neglected-file finding as a pre-close condition.

The CI maturity figure is favourable in workflow count (98 workflows) but unfavourable in gate strictness (medium, not high) and reported coverage (not quantified). The buyer's engineering team should treat the CI improvement plan as a price-adjustment lever sized against the test_investment line item ($20,400) plus an additional allowance for the coverage-tooling gap.

© MTD Services

# 12. AI/ML Footprint

## 12.1 Introduction

The AI/ML surface of the asset triggers distinct regulatory obligations and creates model-licensing and training-data IP risk that is materially different from conventional open-source risk. The pillar risk is informational, which is the most favourable rating — but the informational rating reflects the absence of bundled model files rather than the absence of AI/ML surface area. Framework and SDK hits indicate production AI/ML integration that the buyer's counsel should review against the EU AI Act, model-licensing terms, and training-data IP risk.

## 12.2 Key Figures

AI/ML source files scanned: 11,696. Framework groups in use: 1. SDK groups in use: 5. Bundled model files: 0. Production model files: 0. Unlicensed model files: 0.

## 12.3 Framework Hits

One framework group is in use: Huggingface, with 183 hits. Huggingface is the dominant open-source model-hosting ecosystem; the hits indicate that the codebase references Huggingface libraries, model identifiers, or both. The buyer's AI/ML team should validate whether any Huggingface model is actually loaded into production runtime (the bundled_model_files = 0 result suggests not, but the SDK calls may load models dynamically).

**Table 12.1 — Framework hits**

| Framework | Hits |
|---|---|
| Huggingface | 183 |

## 12.4 SDK Hits

Five SDK groups are in use: Anthropic (1), AWS Bedrock (21), Azure OpenAI (113), Cohere (9), and OpenAI (5). The Azure OpenAI dominance (113 hits) is the most material signal — it indicates production integration with Azure OpenAI Service, which carries Azure-specific contractual and data-residency obligations. The 21 AWS Bedrock hits indicate additional multi-cloud LLM usage. The 1 Anthropic and 5 OpenAI hits are minor and may be experimental.

**Table 12.2 — SDK hits**

| SDK | Hits |
|---|---|
| Anthropic | 1 |
| AWS Bedrock | 21 |
| Azure OpenAI | 113 |
| Cohere | 9 |
| OpenAI | 5 |

© MTP Services

## 12.5 M&A Implications

The AI/ML surface is informational in severity because no bundled model files are detected — the codebase does not ship model artifacts. This is a positive signal: the buyer does not inherit model-licensing or training-data IP risk from bundled models. However, the SDK hits indicate the codebase makes runtime calls to external LLM providers (Azure OpenAI, AWS Bedrock, Anthropic, Cohere, OpenAI), which creates two distinct post-close obligations. First, the buyer must inherit or re-negotiate the seller's contracts with each LLM provider — Azure OpenAI in particular may carry enterprise-tier commitments that affect pricing and data-handling. Second, the buyer's counsel should review the seller's prompt-management and content-moderation practices against the EU AI Act and any applicable sectoral regulation (e.g. HIPAA for healthcare data sent to LLMs).

The Huggingface framework hits should be reviewed to confirm whether any model is loaded at runtime; if so, the model license and training-data provenance must be verified. The five license findings in Section 4 overlap with this section because the AI/ML surface area triggers license review for the SDKs and frameworks in use.

© MTD Services

# 13. Technology Stack & Lifecycle Posture

## 13.1 Introduction

The base-layer inventory records which technologies are in use, their lifecycle status, and any end-of-life or legacy exposure. A clean active stack is a positive value driver and reduces modernization capital. The structural-not-runtime caveat of technology detection is material to warranty representations: the buyer's counsel should not rely on the active rating as a runtime verification; the EOL pillar performs that verification against endoflife.date.

## 13.2 Key Figures

Technologies total: 8. EOL technology count: 0. Legacy technology count: 0. Active technologies: 8. EOL count (runtime verification): 0.

## 13.3 Technologies by Lifecycle

**Table 13.1 — Technologies by lifecycle**

| Lifecycle | Count |
|---|---|
| Active | 8 |

## 13.4 Technology Inventory

Eight technologies are detected, all active. The stack is dual-language Go (backend) and JavaScript/TypeScript (frontend), with Docker and Docker Compose for deployment, Angular in the frontend (likely legacy code, given the React-based grafana-ui packages observed elsewhere), Python pip in the devenv, and Go module checksums for supply-chain integrity. The detection confidence varies: Go, Docker, Docker Compose, and JavaScript/Node.js are detected at 1.0 confidence; Angular, JavaScript/Node.js (Yarn), and Python (pip) are detected at 0.6-0.7 confidence, indicating the detection is based on indirect signals rather than explicit declarations.

**Table 13.2 — Technology inventory**

| Technology | Lifecycle | Confidence |
|---|---|---|
| Angular | active | 0.7 |
| Docker | active | 1.0 |
| Docker Compose | active | 1.0 |
| Go | active | 1.0 |
| Go (module checksums) | active | 1.0 |
| JavaScript/Node.js | active | 1.0 |
| JavaScript/Node.js (Yarn) | active | 0.6 |
| Python (pip) | active | 0.6 |

## 13.5 Runtime EOL Inventory

MTD Services

One runtime is checked: golang go 1.27.2, which is not end-of-life (latest is 1.27.2). The Go runtime is current. The endoflife.date source is used for the runtime EOL check; the buyer's platform team should integrate continuous EOL monitoring into the post-close CI pipeline.

**Table 13.3 — Runtime EOL inventory**

| Product | Runtime | Version | Is EOL | Latest |
|---|---|---|---|---|
| golang | go | 1.27.2 | No | 1.27.2 |

## 13.6 M&A Implications

The clean active stack is a positive value driver. The buyer's counsel should treat the active rating as structural-detection basis, not as runtime verification; the runtime EOL check confirms Go 1.27.2 is current but does not verify that the deployed runtime matches the declared runtime. The buyer's engineering team should require evidence of the deployed runtime version as part of operational handover. The Angular detection is notable because it suggests the codebase carries legacy frontend code alongside the React-based grafana-ui packages; the buyer's frontend team should commission a migration assessment.

© MTD Services

# 14. Consolidated Risk Register & Top Risks

## 14.1 Introduction

The single integrated view of every material risk is ranked by severity, with impact, file and line evidence where applicable, and ownership mapping to the relevant pillar. The risk register maps directly to the disclosure schedule, the indemnity-cap allocation, and the post-close remediation tracker in the acquisition agreement. The consolidated top-risks, major-risks, pillar-classification, and deal-breaker registers below are stored ground truth.

## 14.2 Consolidated Top Risks

**Table 14.1 — Consolidated top risks**

| Finding ID | Pillar | Title | Severity | Evidence state | Impact | File | Cost (USD) |
|---|---|---|---|---|---|---|---|
| F-complexity-006 | Complexity | GetRoutes | Critical | Detected, not verified | Critical | pkg/registry/apis/dashboard/snapshot/routes.go | Not quantified |
| F-complexity-034 | Complexity | NewAPIService | Critical | Detected, not verified | Critical | pkg/registry/apis/iam/register.go | Not quantified |
| F-complexity-019 | Complexity | Changes | Critical | Detected, not verified | Critical | pkg/registry/apis/provisioning/jobs/sync/changes.go | Not quantified |
| F-retention-021 | Retention | Adam Yeats | Critical | Inferred | Critical | Not quantified | Not quantified |
| F-retention-035 | Retention | Alex Spencer | Critical | Inferred | Critical | Not quantified | Not quantified |
| F-retention-001 | Retention | Alex Khomenko | Critical | Inferred | Critical | Not quantified | $316,400 |
| F-debt-002 | Debt | fixme | High | Detected, not verified | High | eslint.config.js:624 | Not quantified |
| F-debt-003 | Debt | fixme | High | Detected, not verified | High | eslint.config.js:638 | Not quantified |
| F-license-006 | License | (repo-level) | High | Detected, not verified | High | LICENSE | Not quantified |
| F-debt-001 | Debt | fixme | High | Detected, not verified | High | eslint.config.js:40 | $7,530 |

*Three complexity findings (GetRoutes, NewAPIService, Changes) are in pkg/registry and pkg/services API surface — these are the highest-priority refactoring targets because they sit on the request-handling hot path. Three retention findings identify critical-flight-risk contributors; Alex Khomenko is the only one with a quantified remediation cost ($316,400) tied to the knowledge-transfer line item.*

## 14.3 Major Risks Register

**Table 14.2 — Major risks register**

| Pillar | Area | Risk | Headline metric | Unverified |
|---|---|---|---|---|
| Entanglement | Architecture | Critical | total_cycles: 4, max_cycle_length: 1445, cross_layer_violations_count: 19887 | No |
| Complexity | Engineering | Critical | ProvideService | No |
| Debt | Engineering | High | fixme | No |
| Integration | Engineering | High | integration_risk_score: 10.0, proprietary_coupling_count: 509, total_findings: 521 | No |
| Coverage | Engineering | High | modules_analyzed: 2506, high_value_gaps: 592, effective_insured_value: 1.0 | No |
| Test Deficit | Engineering | High | total_modules: 2506, modules_below_threshold: 1882, overall_test_ratio: 0.1526 | No |
| License | Legal & Compliance | High | envtpl | No |
| IP Ownership | Legal & Compliance | High | total_authors: 3217, external_author_count: 2319, external_blamed_line_pct: 0.7723 | No |
| Retention | People | High | Alex Khomenko | No |
| Secrets | Security | Critical | candidate_count: 427, high_confidence_count: 63, history_only_count: 345 | Yes |
| CVE | Security | High | packages_queried: 1789, unresolved_packages: 19, cve_count: 154 | Yes |

## 14.4 Pillar Classification Register

The pillar classification register maps each pillar to a classification (Critical Risk, Major Risk, Minor Risk, Acceptable, or Deal-Breaker) and a health score. The classification drives the conditions in Section 2; the health score communicates the pillar-level posture.

**Table 14.3 — Pillar classification register**

| Pillar | Area | Risk | Classification | Health score | Unverified |
|---|---|---|---|---|---|
| AI/ML | AI/ML Exposure | Informational | Informational Finding | 90 | No |
| Attack Surface | Security | Medium | Manageable Risk | 65 | No |
| CI Maturity | Process | Medium | Manageable Risk | 35 | No |

| Pillar | Area | Risk | Classification | Health score | Unverified |
|---|---|---|---|---|---|
| Complexity | Engineering | Critical | Major Risk | 15 | No |
| Coverage | Engineering | High | Major Risk | 40 | No |
| CVE | Security | High | Major Risk | 50 | Yes |
| Debt | Engineering | High | Major Risk | 40 | No |
| Dependency | Supply Chain | Medium | Manageable Risk | 65 | No |
| Doc Staleness | Process | Low | Informational Finding | 85 | No |
| Entanglement | Architecture | Critical | Major Risk | 15 | No |
| EOL | Platform | Low | Informational Finding | 85 | No |
| IaC | Security | Medium | Manageable Risk | 65 | No |
| Integration | Engineering | High | Major Risk | 40 | No |
| IP Ownership | Legal & Compliance | High | Major Risk | 40 | No |
| Key Person | People | Low | Informational Finding | 85 | No |
| License | Legal & Compliance | High | Major Risk | 40 | No |
| Maintenance | Engineering | Low | Unverified | 50 | Yes |
| Privacy | Legal & Compliance | Informational | Informational Finding | 95 | No |
| Retention | People | High | Major Risk | 40 | No |
| Secrets | Security | Critical | Deal-Breaker | 50 | Yes |

*Two pillars carry the lowest health score of 15: Complexity (critical) and Entanglement (critical). Three pillars are unverified: CVE (high, 50), Maintenance (low, 50), Secrets (critical, 50). The Secrets pillar is the single Deal-Breaker classification; its unverified status is the most material epistemic gap in the report.*

## 14.5 Deal-Breaker Register

**Table 14.4 — Deal-breaker register**

| Pillar | Area | Risk | Trigger | Unverified |
|---|---|---|---|---|
| Secrets | Security | Critical | candidate_count: 427, high_confidence_count: 63, history_only_count: 345 | Yes |

The deal-breaker trigger fires on Secrets because the pillar returns critical severity under the decision model. Because the pillar is unverified (partial coverage, 24 path-excluded candidates not scanned), the deal-breaker is recommended for validation rather than asserted as final. The buyer must close the coverage gap before relying on the deal-breaker in negotiation.

# Part IV — Scope, Confidence & Limitations

## 15. Confidence, Scope Limitations & Material Unknowns

### 15.1 Introduction

The epistemic boundaries of this report prevent over-reliance on the headline numbers. Coverage caveats size the warranty and indemnity chapter, the working-group escrow, and the credibility of the seller's vulnerability-management claims. The integrity block, confidence ledger, per-pillar assessment postures, and taint summary are replayed here verbatim from the ground truth.

### 15.2 Key Figures

Integrity status: warn. Integrity error count: 0. Integrity warning count: 8. Confidence coverage: 86.4%. Confidence risk signal: 50.8%. Confidence signal count: 19 of 22. Confidence degraded pillars: 3. Taint flows found: 0. Taint pass bounded: false.

### 15.3 Integrity Issues

Eight integrity warnings are recorded. Five are critical_pillar_no_findings warnings for pillars that are high or critical but render no item-level findings — the reading guide must state the aggregate basis (summary metrics) for these pillars. Three are unverified_pillar warnings for CVE, Maintenance, and Secrets, each with the specific reason coverage was incomplete. No errors are recorded.

**Table 15.1 — Integrity issues**

| Tier | Code | Subject | Message |
|---|---|---|---|
| Warn | critical_pillar_no_findings | pillars.coverage | Pillar 'coverage' is high but renders no item-level findings. |
| Warn | critical_pillar_no_findings | pillars.entanglement | Pillar 'entanglement' is critical but renders no item-level findings. |
| Warn | critical_pillar_no_findings | pillars.integration | Pillar 'integration' is high but renders no item-level findings. |
| Warn | critical_pillar_no_findings | pillars.ip_ownership | Pillar 'ip_ownership' is high but renders no item-level findings. |
| Warn | critical_pillar_no_findings | pillars.test_deficit | Pillar 'test_deficit' is high but renders no item-level findings. |
| Warn | unverified_pillar | pillars.cve | Pillar 'cve' is partial — conclusions UNVERIFIED. |
| Warn | unverified_pillar | pillars.maintenance | Pillar 'maintenance' is not_assessed — conclusions UNVERIFIED. |
| Warn | unverified_pillar | pillars.secrets | Pillar 'secrets' is partial — conclusions UNVERIFIED. |

*The five critical_pillar_no_findings warnings require the buyer's reader to interpret the high/critical ratings on coverage, entanglement, integration, IP ownership, and test_deficit as aggregate-metric-based rather than item-finding-*

D Services

*based. The three unverified_pillar warnings require the buyer to validate CVE, Maintenance, and Secrets independently before close.*

## 15.4 Unverified Pillars

Three pillars are unverified: CVE, Maintenance, and Secrets. Each is named in the executive brief and in the conditions for acquisition (Section 2). The buyer must verify each independently before close; absence of a finding in those domains is absence of coverage, not absence of risk.

- CVE — partial. 19 package requirements could not be resolved to a version (no lockfile / unpinned), so the advisory check covered only the 1,789 resolvable packages; 'no critical CVEs' applies to that queried population only.
- Maintenance — not_assessed. No files were scanned by the asset-staleness scanner — maintenance not assessed.
- Secrets — partial. 24 path-excluded candidate paths were not scanned; 16,662 files were scanned.

## 15.5 Per-Pillar Assessment Posture

**Table 15.2 — Per-pillar assessment posture**

| Pillar | Area | Risk | Classification | Health score | Unverified |
|---|---|---|---|---|---|
| AI/ML | AI/ML Exposure | Informational | Informational Finding | 90 | No |
| Attack Surface | Security | Medium | Manageable Risk | 65 | No |
| CI Maturity | Process | Medium | Manageable Risk | 35 | No |
| Complexity | Engineering | Critical | Major Risk | 15 | No |
| Coverage | Engineering | High | Major Risk | 40 | No |
| CVE | Security | High | Major Risk | 50 | Yes |
| Debt | Engineering | High | Major Risk | 40 | No |
| Dependency | Supply Chain | Medium | Manageable Risk | 65 | No |
| Doc Staleness | Process | Low | Informational Finding | 85 | No |
| Entanglement | Architecture | Critical | Major Risk | 15 | No |
| EOL | Platform | Low | Informational Finding | 85 | No |
| IaC | Security | Medium | Manageable Risk | 65 | No |
| Integration | Engineering | High | Major Risk | 40 | No |
| IP Ownership | Legal & Compliance | High | Major Risk | 40 | No |
| Key Person | People | Low | Informational Finding | 85 | No |
| License | Legal & Compliance | High | Major Risk | 40 | No |

| Pillar | Area | Risk | Classification | Health score | Unverified |
|---|---|---|---|---|---|
| Maintenance | Engineering | Low | Unverified | 50 | Yes |
| Privacy | Legal & Compliance | Informational | Informational Finding | 95 | No |
| Retention | People | High | Major Risk | 40 | No |
| Secrets | Security | Critical | Deal-Breaker | 50 | Yes |

## 15.6 CVE Scope Notes

The CVE scope note records the partial-coverage reason: 19 package requirements could not be resolved to a version (no lockfile / unpinned), so the advisory check covered only the 1,789 resolvable packages; 'no critical CVEs' applies to that queried population only. The buyer's security team should re-query the OSV feed after the seller provides the missing lockfiles or pins the unpinned dependencies.

**Table 15.3 — CVE scope notes**

| Domain | Scope note |
|---|---|
| CVE | 19 package requirements could not be resolved to a version (no lockfile / unpinned), so the advisory check covered only the 1,789 resolvable packages; 'no critical CVEs' applies to that queried population only. |

## 15.7 Coverage Gaps by Pillar

**Table 15.4 — Coverage gaps by pillar**

| Pillar | Assessment status | Reason | Population | Scope note |
|---|---|---|---|---|
| CVE | Partial | Coverage gap | 1,789 | 19 package requirements unresolved to a version (no lockfile / unpinned); advisory check covered 1,789 resolvable only. |
| Maintenance | Not assessed | Coverage gap | 0 | No files were scanned by the asset-staleness scanner; maintenance not assessed. |
| Secrets | Partial | Coverage gap | 16,662 | 24 path-excluded candidate paths not scanned; 16,662 files scanned. |

*A gap is not a clean result: it means the question was not answered. The buyer must close each gap before relying on the corresponding pillar's verdict in negotiation.*

## 15.8 Out-of-Scope Evidence Register

Eight evidence domains are not assessed by a static, offline repository scan. Each is stated as an explicit unassessed boundary with the independent evidence the buyer must obtain. Disclosure, not a measurement. The buyer's counsel should commission each item as part of pre-signing diligence; the absence of independent evidence in any of these domains is a warranty risk that should be sized into the indemnity chapter.

**Table 15.5 — Out-of-scope evidence register**

| Domain | Status | Reason | Independent evidence |
|---|---|---|---|
| Runtime & operational behaviour | Not assessed | Static repository scan only; the target's code and systems are never executed, so runtime behaviour and production incidents are not observed. | Runtime traces, incident history, SLO / error-budget records. |
| Network perimeter & OT boundaries | Not assessed | Firewall / security-group rules and OT segmentation are not verified against a live environment; committed IaC text is read as candidate configuration only. | Live firewall / security-group export, network diagrams, OT segmentation policy. |
| External attack surface (ports, DNS) | Not assessed | No active reconnaissance is performed; live port inventories and historical DNS posture are not queried. | Independent external port scan and DNS-history report. |
| Identity & access controls | Not assessed | MFA enforcement, privileged-access-management logs and contractor offboarding are operational controls outside a code scan. | IAM / MFA policy export, PAM audit logs, offboarding attestation. |
| Data resilience & recovery | Not assessed | Backup validation, point-in-time-recovery testing and migration governance are operational processes not exercised. | Backup / restore test reports, PITR drill evidence, migration runbooks. |
| Delivery assurance (CI & performance) | Not assessed | Only committed CI config and artifacts are read; live pass rates, flakiness, and formal load / stress and capacity testing are not observed. | CI run history, flakiness metrics, load-test and capacity reports. |
| Legal instruments & contracts | Not assessed | Executed contributor assignment agreements, employment / contractor terms and vendor contracts are legal artifacts, not code. | Signed CLAs, IP-assignment and contractor agreements, vendor DPAs. |
| Runtime dependency & vulnerability freshness | Not assessed | Dependency vulnerability / version data is a point-in-time snapshot (OSV); continuous monitoring and deploy-time enforcement are not assessed. | Continuous SCA / SBOM monitoring feed and deploy-time gate evidence. |

## 15.9 Taint Analysis Summary

Taint analysis found 0 flows. The pass was not bounded, meaning the analysis ran but did not establish an exhaustive bound on the taint-flow surface. The buyer's security team should treat the zero-flow result as a partial coverage signal rather than as a confirmed absence of exploitable flows; an independent dynamic taint analysis (e.g. via runtime instrumentation) should be commissioned as part of pre-signing diligence.

**Table 15.6 — Taint summary**

| Metric | Value |
|---|---|
| Bounded | false |
| Total flows | 0 |

© MTD Services

# Part V — Methodology, Conventions & Reproducibility

## 16. Reading Notes & Honesty Conventions

### 16.1 Introduction

The conventions that govern how every other section must be read are fixed and target-agnostic. No figure in this section is instance-specific, and none of it is evidence in its own right. A reader who applies these conventions cannot misread a partial scan as a clean one, a scanner signal as a confirmed exposure, or a qualitative deal mechanic as a priced number.

### 16.2 Assessment Coverage

Assessment status is assessed, partial, or not_assessed per pillar. Only assessed pillars support conclusions; partial / not_assessed pillars are demoted to informational, labelled UNVERIFIED, and never read as a clean bill of health. Assessment state (NOT_ASSESSED, PARTIALLY_ASSESSED, ASSESSED_NO_FINDINGS, ASSESSED_WITH_FINDINGS, ERROR) is the five-value trust vocabulary that accompanies the legacy three-value status. A low risk with PARTIALLY_ASSESSED must be read as 'incomplete coverage', not 'no problem'.

Population is the evidence count each verdict rests on (files scanned, packages queried, modules in cycles, etc.). A low or missing population flags a coverage gap instead of hiding it. A zero population never reads as clean.

### 16.3 Evidence Honesty

Evidence state is the per-finding honesty label, attached to every material finding and shown in the Section 14 register: Confirmed, Detected-not-verified, Inferred, Partially assessed, Unassessed, Conflicting, or Unsupported. Scanner signals are never auto-confirmed. The finding ID (F-<pillar>-<rank>) is the stable handle used by the narrative, the register, the workbook, and the ledger; a body section cites the ID, it never re-renders the finding.

The contradiction ledger (Appendix R) leads every material conflict with its status (resolved, by-design, or unresolved) and decision impact; unresolved entries discount the stated confidence. It never picks a side silently. Confidence = 20 + 40 × coverage + 40 × risk_signal, capped at 20 when coverage is zero, floored at 60 when the verdict is DO_NOT_ACQUIRE. Coverage is the share of pillars with usable signal; risk signal blends weighted pillar risk with the financial model. Unverified pillars depress both and are named in the executive brief.

### 16.4 Currency of Absence

Four markers name a value this scan did not establish. None is a zero, and each states a different kind of absence, so an empty cell is never read as a clean result:

- not quantified — the model could compute a number but a stored input is missing (a Layer-2 expected loss with no financial.risk_parameters calibration; an unset escalation threshold). It never means the risk costs nothing.

© MTD Services

- not applicable — the cell is structurally inapplicable to that row (the fixed-version of a CVE with no fix published; the commit of a secret found only in the working tree).
- not recorded — the source does not carry the value (an aggregate finding has no single line number; the scanner never captured a commit).
- not assessed — the value was not established because coverage was incomplete (a named coverage gap).

NOT_QUANTIFIED is additionally a stored assessment state on the economic layers of Section 3, distinct from the five coverage states under Assessment Coverage. The same rule holds: absent is never zero, and an unquantified line is named rather than priced at 0.

### 16.5 Deal-Mechanic Vocabulary

Every condition in Section 2 carries exactly one of three deal-mechanic tags, declared where the condition is authored (evidence 2.13). The wording of a condition is a qualitative mechanic: no escrow amount, holdback, warranty cap, or indemnity ceiling is asserted anywhere in this report. Sizing those is deal counsel's call.

- Condition precedent — must be satisfied or proven before close, or the deal does not sign as structured.
- Price-adjustment lever — a funded or escrowed work item that moves value rather than blocking close.
- Warranty / representation — a seller statement whose breach is an indemnity question, not a work plan.

### 16.6 Severity and Ranking

- Risk level — critical, high, medium, low, or informational, demoted to informational under partial coverage. derived_risk_level preserves the scanner's own view.
- Health score — 0-100 per pillar, demoted to 50 for an unfinished scan so that incomplete coverage never presents as a strong pillar.
- Classification — Critical Risk, Major Risk, Minor Risk, or Acceptable, assigned in the Section 14 register and driving the conditions in Section 2.

### 16.7 Known Limits of This Method

EOL reconciliation: tech_stack counts offline repo technology lifecycle markers; the eol pillar verifies runtimes / dependencies against endoflife.date. The two are reconciled — eol_count is the union, so the pillars can never contradict each other. The buyer's counsel should rely on the eol pillar for runtime-verification evidence, not on the tech_stack structural-detection rating.

Static source analysis: this is a source scan. It does not establish runtime behaviour, production incidents, or anything needing access to systems, personnel, or contracts. The Out-of-Scope Evidence Register (Section 15.5) enumerates the eight evidence domains the buyer must obtain independently.

© MTD Services

No negotiated positions: the conditions, implications, and escalation thresholds recorded here are the model's deterministic output from stored evidence. They are a starting position for negotiation, not a legal opinion. The buyer's counsel should treat each condition as a draft for negotiation, not as a final term.

© MTD Services

# Appendices

## Appendix A — Methodology, Scope & Scan Configuration

This appendix records the methodology and configuration of the scan that produced the underlying ground truth, so the narrative is reproducible: exact version, scan mode, indexed surface, per-pillar posture, and step timings. Every figure is reproduced from the stored payload.

### A.1 Report Metadata

**Table A.1 — Report metadata**

| Field | Value |
|---|---|
| Report ID | TDD-20261016-060430 |
| Scan date | 2026-10-16T07:15:51.747366+00:00 |
| Target | Grafana |
| Scanned path | grafana |
| Indexed symbols | 16,662 |
| Report version | Not quantified |
| Elapsed seconds | 4,281.52 |

### A.2 Scan Mode

**Table A.2 — Scan mode**

| Parameter | Value |
|---|---|
| deal_thesis | all |
| offline | false |
| scan_history | true |
| taint | true |
| target_company | Grafana |
| tdd_report | true |
| workers | 2 |

*deal_thesis = all means all six deal theses are evaluated; scan_history = true means git-history candidates are scanned (345 history-only secrets); taint = true means the taint analysis ran (0 flows found, pass not bounded); workers = 2 is the parallelism setting.*

### A.3 Pillar Posture Register

**Table A.3 — Pillar posture register**

| Pillar | Area | Risk level | Health | Assessment | Top risks |
|---|---|---|---|---|---|
| AI/ML | AI/ML Exposure | Informational | 90 | Assessed | 0 |
| Attack Surface | Security | Medium | 65 | Assessed | 50 |
| CI Maturity | Process | Medium | 35 | Assessed | 0 |
| Complexity | Engineering | Critical | 15 | Assessed | 50 |
| Coverage | Engineering | High | 40 | Assessed | 0 |
| CVE | Security | Informational | 50 | Partial | 0 |
| Debt | Engineering | High | 40 | Assessed | 50 |
| Dependency | Supply Chain | Medium | 65 | Assessed | 0 |
| Doc Staleness | Process | Low | 85 | Assessed | 0 |
| Entanglement | Architecture | Critical | 15 | Assessed | 0 |
| EOL | Platform | Low | 85 | Assessed | 0 |
| IaC | Security | Medium | 65 | Assessed | 1 |
| Integration | Engineering | High | 40 | Assessed | 0 |
| IP Ownership | Legal & Compliance | High | 40 | Assessed | 0 |
| Key Person | People | Low | 85 | Assessed | 50 |
| License | Legal & Compliance | High | 40 | Assessed | 50 |
| Maintenance | Engineering | Informational | 50 | Not assessed | 0 |
| Privacy | Legal & Compliance | Informational | 95 | Assessed | 1 |
| Retention | People | High | 40 | Assessed | 50 |
| Secrets | Security | Informational | 50 | Partial | 0 |
| Tech Stack | Architecture | Low | 85 | Assessed | 0 |
| Test Deficit | Engineering | High | 40 | Assessed | 0 |

## A.4 Step Timings

**Table A.4 — Step timings**

| Step | Seconds |
|---|---|
| index | 240.61 |
| pillars | 4,003.63 |
| deep_index | 5.783 |
| taint | 20.502 |
| financial | 0.364 |
| decision | 0.001 |

*The pillars step dominates at 4,003.63 seconds (over 66 minutes) — expected given the breadth of pillar scanners. The financial step is 0.364 seconds because the model is deterministic once the pillar evidence is loaded; the decision step is 0.001 seconds for the same reason.*

© MTD Services

## Appendix B — Metric Definitions & Financial Calibration

Metric definitions used throughout the narrative are fixed and target-agnostic. Cyclomatic and cognitive complexity are measured per symbol; bus factor is the reciprocal of sole-contributor concentration; debt density is normalised per thousand lines; and remediation hours-per-unit (HPU) per cost category is calibrated by the model. The calibration tables themselves live in Section 3 (the single financial home) and are cited there, not duplicated here, so a figure appears in exactly one place. This appendix references Section 3 rather than reproducing its tables; the buyer's quality-of-earnings provider should consult Section 3.2 (parameters table), Section 3.4 (remediation line items), Section 3.6 (saturation model), and Section 3.7 (SQALE remediation basis) for the canonical figures.

© MTD Services

# Appendix O — Financial Reproduction Fingerprint

The financial reproduction fingerprint binds the report to the exact configuration used to produce every number: same inputs in, same outputs out. Any change to a number-driving input produces a different fingerprint and a different report under a new fingerprint. The full calibration, lineage, and sensitivity detail lives in Section 3 — not re-rendered here.

**Table O.1 — Reproduction fingerprint**

| Field | Value |
|---|---|
| Reproduction fingerprint | b607d2e6aabcc186eeb9fb69a27d29d10ab901b978e50d332fec58eb18ce7395 |
| Economic model fingerprint (0.3.0) | aba11b3a12341e602de6865ec282b34bb65952d165700979b067b826b05e006e |
| Repository | grafana |
| HEAD commit | 5be6647ffe4c0ce6522387884e4ca51e9c181f10 |
| Scan date | 2026-10-16T07:15:51.747366+00:00 |

© MTD Ser

# Appendix P — Evidence Provenance

The payload path is the single stored source the narrative's evidence derives from; key-variant fallbacks are resolved at render time and are not listed here. The workbook column names the sheet that inventories the same evidence; not applicable means no dedicated sheet carries it.

**Table P.1 — Evidence provenance (every figure in this report, mapped to its source)**

| ID | Evidence | Source | Workbook |
|---|---|---|---|
| 1.1 | Board decision | decision.board_decision | Dashboard |
| 1.2 | Confidence score | decision.confidence_score | Dashboard |
| 1.3 | Decision driver | decision.escalation.driver | Dashboard |
| 1.4 | Remediation cost as % of price | decision.remediation_cost_pct_of_acquisition | Dashboard |
| 1.5 | Proposed price | decision.proposed_price | Dashboard |
| 1.6 | Total remediation cost | financial.total_remediation_cost | Dashboard |
| 1.7 | Total remediation hours | financial.total_remediation_hours | Dashboard |
| 1.8 | Valuation discount applied | financial.valuation_discount_pct | Dashboard |
| 1.9 | Post-merge integration drag | financial.post_merge_integration_drag_pct | Dashboard |
| 1.10 | Pillar register size | decision.pillar_register | Dashboard |
| 1.11 | Top risk findings | decision.top_risks | Dashboard |
| 1.12 | Major risks | decision.major_risks | Dashboard |
| 1.13 | Integrity status | integrity.status | Contents |
| 1.14 | Integrity error count | integrity.error_count | Contents |
| 1.15 | Integrity warning count | integrity.warn_count | Contents |
| 2.1 | Conditions for acquisition | decision.conditions_for_acquisition | Dashboard |
| 2.2 | Deal-breaker register | decision.deal_breakers | Dashboard |
| 2.3 | Categorical decision tier | decision.escalation.categorical_tier | Dashboard |
| 2.4 | Economic decision tier | decision.escalation.economic_tier | Dashboard |
| 2.5 | Conditions escalation threshold | decision.escalation.conditions_threshold_pct | Dashboard |
| 2.6 | No-go escalation threshold | decision.escalation.no_go_threshold_pct | Dashboard |
| 2.7 | Economic gate disputed | decision.escalation.disputed | Dashboard |

| ID | Evidence | Source | Workbook |
|---|---|---|---|
| 2.8 | Escalation driver | decision.escalation.driver | Dashboard |
| 2.9 | Per-pillar escalation rows | decision.escalation.pillars | Dashboard |
| 2.10 | Economic conditions threshold (model) | financial.assumptions.economic_escalation.conditions_if_above_pct | Dashboard |
| 2.11 | Economic no-go threshold (model) | financial.assumptions.economic_escalation.no_go_if_above_pct | Dashboard |
| 2.12 | Sensitivity sweep outcomes | financial.sensitivity | Dashboard |
| 2.13 | Condition deal-mechanic tags | decision.condition_details | Dashboard |
| 3.1 | Remediation line items | financial.line_items | Dashboard |
| 3.2 | Total remediation hours | financial.total_remediation_hours | Dashboard |
| 3.3 | Total remediation cost | financial.total_remediation_cost | Dashboard |
| 3.4 | Valuation discount applied | financial.valuation_discount_pct | Dashboard |
| 3.5 | Post-merge integration drag | financial.post_merge_integration_drag_pct | Dashboard |
| 3.6 | Hourly rate | financial.assumptions.hourly_rate | Dashboard |
| 3.7 | Complexity multiplier | financial.assumptions.complexity_multiplier | Dashboard |
| 3.8 | Valuation multiplier | financial.assumptions.valuation_multiplier | Dashboard |
| 3.9 | Integration drag factor | financial.assumptions.integration_drag_factor | Dashboard |
| 3.10 | Productive hours per quarter | financial.assumptions.productive_hours_per_quarter | Dashboard |
| 3.11 | Person-years of remediation | financial.assumptions.person_years | Dashboard |
| 3.12 | Remediation cost CI | financial.assumptions.remediation_cost_confidence_interval | Dashboard |
| 3.13 | Conditions threshold (model) | financial.assumptions.economic_escalation.conditions_if_above_pct | Dashboard |
| 3.14 | No-go threshold (model) | financial.assumptions.economic_escalation.no_go_if_above_pct | Dashboard |
| 3.15 | Valuation discount cap | financial.assumptions.valuation_discount_cap_pct | Dashboard |
| 3.16 | Integration drag cap | financial.assumptions.integration_drag_cap_pct | Dashboard |
| 3.17 | Aggregate hours cap | financial.assumptions.aggregate_cap | Dashboard |

| ID | Evidence | Source | Workbook |
|---|---|---|---|
| 3.18 | Saturation model | financial.assumptions.saturation_model | Dashboard |
| 3.19 | Line-item basis (HPU per category) | financial.assumptions.line_item_basis | Dashboard |
| 3.20 | SQALE remediation basis | financial.assumptions.sqale_remediation_basis | Dashboard |
| 3.21 | Sensitivity sweep details | financial.sensitivity | Dashboard |
| 3.22 | Financial lineage | financial.lineage | Dashboard |
| 3.23 | Remediation cost as % of price | decision.remediation_cost_pct_of_acquisition | Dashboard |
| 3.24 | Economic model fingerprint (0.3.0) | financial.economic_layers.economic_layers_fingerprint | Dashboard |
| 3.25 | Risk valuation lines | financial.economic_layers.layer2_risk.lines | Dashboard |
| 3.26 | Aggregate expected loss before | financial.economic_layers.layer2_risk.aggregate_expected_loss_before | Dashboard |
| 3.27 | Aggregate residual expected loss | financial.economic_layers.layer2_risk.aggregate_residual_expected_loss | Dashboard |
| 3.28 | Aggregate risk-reduction value | financial.economic_layers.layer2_risk.aggregate_risk_reduction_value | Dashboard |
| 3.29 | Quantified risk lines | financial.economic_layers.layer2_risk.quantified_count | Dashboard |
| 3.30 | Not-quantified risk lines | financial.economic_layers.layer2_risk.not_quantified_count | Dashboard |
| 3.31 | Scenario branches (0.3.0) | financial.economic_layers.layer3_scenario.branches | Dashboard |
| 3.32 | Weighted expectation across scenarios | financial.economic_layers.layer3_scenario.weighted_expectation | Dashboard |
| 3.33 | Scenario-based technology VaR label | financial.economic_layers.layer3_scenario.var_label | Dashboard |
| 3.34 | Deal components (never collapsed) | financial.economic_layers.layer4_deal.components | Dashboard |
| 3.35 | Executive deal-bridge rows | financial.economic_layers.layer4_deal.bridge | Dashboard |
| 3.36 | Double-counting ledger figures | financial.economic_layers.p0_double_counting.figures | Dashboard |
| 3.37 | Double-counting ledger clean verdict | financial.economic_layers.p0_double_counting.clean | Dashboard |

| ID | Evidence | Source | Workbook |
|---|---|---|---|
| 3.38 | Justified overlapping keys | financial.economic_layers.p0_double_counting.justified_overlapping_keys | Dashboard |
| 3.39 | Categories not costed this scan | financial.line_items | Dashboard |
| 3.40 | Per-line-item threshold breach view | financial.line_items | Dashboard |
| 3.41 | Downside exposure range and percentile VaR | financial.economic_layers.downside | Dashboard |
| 3.42 | Deal economics headline reconciliation | financial.economic_layers.layer4_deal | Dashboard |
| 4.1 | License findings inventory | pillars.license.raw.findings | License Contamination |
| 4.2 | Per-package detected licenses | pillars.license.raw.findings | License Contamination |
| 4.3 | Viral license count | pillars.license.raw.viral_count | License Contamination |
| 4.4 | Restrictive license count | pillars.license.raw.restrictive_count | License Contamination |
| 4.5 | Unknown license count | pillars.license.raw.unknown_count | License Contamination |
| 4.6 | Licensed packages total | pillars.license.raw.total_packages | License Contamination |
| 4.7 | Contamination chains | pillars.license.raw.contamination_chains | License Contamination |
| 4.8 | License clause findings | pillars.license.raw.clauses | License Contamination |
| 4.9 | IP ownership metrics | pillars.ip_ownership.summary | IP Ownership |
| 4.10 | IP ownership attribution rows | pillars.ip_ownership.raw.attributions | IP Ownership |
| 5.1 | CVE count | pillars.cve.summary.cve_count | CVEs |
| 5.2 | Critical CVE count | pillars.cve.summary.critical_cve_count | CVEs |
| 5.3 | Packages queried against advisories | pillars.cve.summary.packages_queried | CVEs |
| 5.4 | Unresolved packages | pillars.cve.summary.unresolved_packages | CVEs |
| 5.5 | Advisory source | pillars.cve.summary.source | CVEs |
| 5.6 | CVE inventory | pillars.cve.raw.cves | CVEs |
| 5.7 | Secret candidates | pillars.secrets.raw.candidates | Secrets Candidates |
| 5.8 | High-confidence secrets | pillars.secrets.summary.high_confidence_count | Secrets Candidates |

| ID | Evidence | Source | Workbook |
|---|---|---|---|
| 5.9 | Live critical secrets | pillars.secrets.summary.live_critical_count | Secrets Candidates |
| 5.10 | History-only secrets | pillars.secrets.summary.history_only_count | Secrets Candidates |
| 5.11 | Secrets by kind | pillars.secrets.summary.by_kind | Secrets Candidates |
| 5.12 | Secret candidate detail | pillars.secrets.raw.candidates | Secrets Candidates |
| 5.13 | IaC findings total | pillars.iac.summary.total_findings | IaC Findings |
| 5.14 | IaC summary metrics | pillars.iac.summary | IaC Findings |
| 5.15 | IaC findings | pillars.iac.raw.findings | IaC Findings |
| 5.16 | Attack surface candidate endpoints | pillars.attack_surface.summary.total_candidate_endpoints | Attack Surface |
| 5.17 | Unreferenced endpoints | pillars.attack_surface.summary.unreferenced_endpoints | Attack Surface |
| 5.18 | Dynamic endpoints | pillars.attack_surface.summary.dynamic_endpoints | Attack Surface |
| 5.19 | Untested candidate endpoints | pillars.attack_surface.summary.untested_candidate_endpoints | Attack Surface |
| 5.20 | Candidate exposure score | pillars.attack_surface.summary.candidate_exposure_score | Attack Surface |
| 5.21 | Candidate endpoint detail | pillars.attack_surface.raw.candidate_entry_points | Attack Surface |
| 5.22 | Privacy findings total | pillars.privacy.summary.total_findings | Privacy & Compliance |
| 5.23 | PII endpoint count | pillars.privacy.summary.pii_endpoint_count | Privacy & Compliance |
| 5.24 | Egress SDK count | pillars.privacy.summary.egress_sdk_count | Privacy & Compliance |
| 5.25 | Sensitive log count | pillars.privacy.summary.sensitive_log_count | Privacy & Compliance |
| 5.26 | Privacy findings | pillars.privacy.raw.findings | Privacy & Compliance |
| 6.1 | Direct dependency count | pillars.dependency.summary.direct_dependency_count | Dependency Packages |
| 6.2 | Pinned dependencies | pillars.dependency.summary.pinned_count | Dependency Packages |

| ID | Evidence | Source | Workbook |
|---|---|---|---|
| 6.3 | Ranged dependencies | pillars.dependency.summary.ranged_count | Dependency Packages |
| 6.4 | Unpinned dependencies | pillars.dependency.summary.unpinned_count | Dependency Packages |
| 6.5 | Pin ratio | pillars.dependency.summary.pin_ratio | Dependency Packages |
| 6.6 | Transitive resolved packages | pillars.dependency.summary.transitive_resolved_packages | Dependency Packages |
| 6.7 | Lockfiles present | pillars.dependency.summary.lockfiles_present | Dependency Packages |
| 6.8 | Manifests found | pillars.dependency.summary.manifests_found | Dependency Packages |
| 6.9 | Lockfile resolution | pillars.dependency.raw.lockfiles | Dependency Packages |
| 6.10 | Direct dependency inventory | pillars.dependency.raw.direct_dependencies | Dependency Packages |
| 6.11 | EOL runtime rows | pillars.eol.raw.eol | EOL Findings |
| 6.12 | EOL count | pillars.eol.summary.eol_count | EOL Findings |
| 6.13 | Runtimes checked | pillars.eol.summary.runtimes_checked | EOL Findings |
| 6.14 | EOL runtime detail | pillars.eol.raw.eol | EOL Findings |
| 6.15 | CVE unresolved packages (cross-ref) | pillars.cve.summary.unresolved_packages | CVEs |
| 7.1 | Symbols analyzed | pillars.complexity.summary.total_symbols_analyzed | Complex Symbols |
| 7.2 | High-complexity symbols | pillars.complexity.summary.high_complexity_count | Complex Symbols |
| 7.3 | Critical-complexity symbols | pillars.complexity.summary.critical_complexity_count | Complex Symbols |
| 7.4 | Average cyclomatic complexity | pillars.complexity.summary.avg_cyclomatic | Complex Symbols |
| 7.5 | Average cognitive complexity | pillars.complexity.summary.avg_cognitive | Complex Symbols |
| 7.6 | Complexity hotspot score | pillars.complexity.summary.complexity_hotspot_score | Complex Symbols |
| 7.7 | SQALE grade | pillars.complexity.summary.sqale_grade | Complex Symbols |
| 7.8 | SQALE remediation effort (days) | pillars.complexity.summary.sqale_remediation_effort_days | Complex Symbols |

| ID | Evidence | Source | Workbook |
|---|---|---|---|
| 7.9 | Total remediation effort (days) | pillars.complexity.summary.total_remediation_effort_days | Complex Symbols |
| 7.10 | Debt ratio | pillars.complexity.summary.debt_ratio | Complex Symbols |
| 7.11 | Duplicated LOC | pillars.complexity.summary.duplicated_loc | Complex Symbols |
| 7.12 | Duplication ratio | pillars.complexity.summary.duplication_ratio | Complex Symbols |
| 7.13 | Cyclomatic percentiles | pillars.complexity.summary.cc_percentiles | Complex Symbols |
| 7.14 | Cognitive percentiles | pillars.complexity.summary.cognitive_percentiles | Complex Symbols |
| 7.15 | Hotspot symbols | pillars.complexity.raw.complexity.symbols | Complex Symbols |
| 7.16 | Duplication groups | pillars.complexity.raw.duplication.duplicates | Complex Symbols |
| 7.17 | Debt markers total | pillars.debt.summary.total_markers | Debt Markers |
| 7.18 | Debt markers by type | pillars.debt.summary.by_type | Debt Markers |
| 7.19 | Debt markers by severity | pillars.debt.summary.by_severity | Debt Markers |
| 7.20 | Critical-path markers | pillars.debt.summary.critical_path_markers | Debt Markers |
| 7.21 | Debt density | pillars.debt.summary.debt_density | Debt Markers |
| 7.22 | Debt marker inventory | pillars.debt.raw.markers | Debt Markers |
| 8.1 | Modules analyzed (test deficit) | pillars.test_deficit.summary.total_modules | Test Deficit |
| 8.2 | Modules below threshold | pillars.test_deficit.summary.modules_below_threshold | Test Deficit |
| 8.3 | Overall test ratio | pillars.test_deficit.summary.overall_test_ratio | Test Deficit |
| 8.4 | Test deficit score | pillars.test_deficit.summary.test_deficit_score | Test Deficit |
| 8.5 | High regression-risk modules | pillars.test_deficit.summary.high_regression_risk_modules | Test Deficit |
| 8.6 | Test-deficit module inventory | pillars.test_deficit.raw.modules | Test Deficit |
| 8.7 | Coverage modules analyzed | pillars.coverage.summary.modules_analyzed | Coverage Gaps |

| ID | Evidence | Source | Workbook |
|---|---|---|---|
| 8.8 | Effective insured value | pillars.coverage.summary.effective_insured_value | Coverage Gaps |
| 8.9 | Value coverage ratio | pillars.coverage.summary.value_coverage_ratio | Coverage Gaps |
| 8.10 | Coverage method | pillars.coverage.summary.coverage_method | Coverage Gaps |
| 8.11 | High-value gaps | pillars.coverage.summary.high_value_gaps | Coverage Gaps |
| 8.12 | Coverage module detail | pillars.coverage.raw.modules | Coverage Gaps |
| 8.13 | High-value coverage gaps | pillars.coverage.raw.high_value_gap_modules | Coverage Gaps |
| 9.1 | Circular-dependency cycles | pillars.entanglement.summary.total_cycles | Entanglement Violations |
| 9.2 | Max cycle length | pillars.entanglement.summary.max_cycle_length | Entanglement Violations |
| 9.3 | Cross-layer violations | pillars.entanglement.summary.cross_layer_violations_count | Entanglement Violations |
| 9.4 | Modules in cycles | pillars.entanglement.summary.modules_in_cycles | Entanglement Violations |
| 9.5 | Entanglement score | pillars.entanglement.summary.entanglement_score | Entanglement Violations |
| 9.6 | Cycle inventory | pillars.entanglement.raw.cycles | Entanglement Violations |
| 9.7 | Cross-layer violation detail | pillars.entanglement.raw.cross_layer_violations | Entanglement Violations |
| 9.8 | Integration risk score | pillars.integration.summary.integration_risk_score | Serialization Risks |
| 9.9 | Schema drift risk | pillars.integration.summary.schema_drift_risk | Serialization Risks |
| 9.10 | Proprietary coupling count | pillars.integration.summary.proprietary_coupling_count | Serialization Risks |
| 9.11 | Integration findings total | pillars.integration.summary.total_findings | Serialization Risks |
| 9.12 | API contract maturity | pillars.integration.raw.api_contracts | Serialization Risks |
| 9.13 | Database coupling | pillars.integration.raw.database_coupling | Serialization Risks |
| 9.14 | Serialization risks | pillars.integration.raw.serialization_risks | Serialization Risks |

| ID | Evidence | Source | Workbook |
|---|---|---|---|
| 9.15 | Protocol-layer risks | pillars.integration.raw.protocol_layers | Serialization Risks |
| 10.1 | Total authors | pillars.key_person.summary.total_authors | Key-person Risks |
| 10.2 | Overall bus factor | pillars.key_person.summary.overall_bus_factor | Key-person Risks |
| 10.3 | Critical key-person rows | pillars.key_person.summary.critical_count | Key-person Risks |
| 10.4 | Key-person risk rows | pillars.key_person.raw.risks | Key-person Risks |
| 10.5 | Key-person risk detail | pillars.key_person.raw.risks | Key-person Risks |
| 10.6 | Module bus-factor rows | pillars.key_person.raw.modules | Key-person Risks |
| 10.7 | Module bus-factor detail | pillars.key_person.raw.modules | Key-person Risks |
| 10.8 | Git team signals | pillars.key_person.summary.team | Key-person Risks |
| 10.9 | Retention total contributors | pillars.retention.summary.total_contributors | Knowledge Loss Vulnerability |
| 10.10 | Active contributors | pillars.retention.summary.active_contributor_count | Knowledge Loss Vulnerability |
| 10.11 | Inactive contributors | pillars.retention.summary.inactive_count | Knowledge Loss Vulnerability |
| 10.12 | At-risk contributors | pillars.retention.summary.at_risk_count | Knowledge Loss Vulnerability |
| 10.13 | Knowledge loss vulnerability score | pillars.retention.summary.retention_risk_score | Knowledge Loss Vulnerability |
| 10.14 | Critical modules (retention) | pillars.retention.summary.critical_modules | Knowledge Loss Vulnerability |
| 10.15 | Knowledge loss vulnerability detail | pillars.retention.raw.risks | Knowledge Loss Vulnerability |
| 10.16 | At-risk contributors (model input) | financial.assumptions.informational_counts.at_risk_contributors | Dashboard |
| 10.17 | Critical sole-maintainer modules (model input) | financial.assumptions.informational_counts.critical_sole_maintainer_modules | Dashboard |
| 11.1 | Maintenance total files | pillars.maintenance.summary.total_files | Maintenance Hotspots |
| 11.2 | Neglected files | pillars.maintenance.summary.neglected_count | Maintenance Hotspots |

| ID | Evidence | Source | Workbook |
|---|---|---|---|
| 11.3 | At-risk files | pillars.maintenance.summary.at_risk_count | Maintenance Hotspots |
| 11.4 | Healthy files | pillars.maintenance.summary.healthy_count | Maintenance Hotspots |
| 11.5 | Average staleness | pillars.maintenance.summary.avg_staleness | Maintenance Hotspots |
| 11.6 | Maintenance backlog score | pillars.maintenance.summary.maintenance_backlog_score | Maintenance Hotspots |
| 11.7 | Maintenance asset rows | pillars.maintenance.raw.assets | Maintenance Hotspots |
| 11.8 | Documents inventoried | pillars.doc_staleness.summary.documents_inventoried | Stale Docs |
| 11.9 | Stale documents | pillars.doc_staleness.summary.stale_documents | Stale Docs |
| 11.10 | Document staleness ratio | pillars.doc_staleness.summary.staleness_ratio | Stale Docs |
| 11.11 | Average docstring density | pillars.doc_staleness.summary.avg_docstring_density | Stale Docs |
| 11.12 | Document inventory | pillars.doc_staleness.raw.documents | Stale Docs |
| 11.13 | CI defined | pillars.ci_maturity.summary.ci_defined | CI Signals |
| 11.14 | CI files checked | pillars.ci_maturity.summary.ci_files_checked | CI Signals |
| 11.15 | CI gate strictness | pillars.ci_maturity.summary.gate_strictness | CI Signals |
| 11.16 | Tests run in CI | pillars.ci_maturity.summary.tests_run_in_ci | CI Signals |
| 11.17 | Reported coverage in CI | pillars.ci_maturity.summary.reported_coverage_pct | CI Signals |
| 11.18 | CI files | pillars.ci_maturity.raw.ci_files | CI Signals |
| 11.19 | Coverage artifacts | pillars.ci_maturity.raw.coverage_artifacts | CI Signals |
| 12.1 | AI/ML source files scanned | pillars.ai_ml.summary.source_files_scanned | AI/ML Models |
| 12.2 | Framework groups in use | pillars.ai_ml.summary.framework_groups_in_use | AI/ML Models |
| 12.3 | Framework hits | pillars.ai_ml.summary.framework_hits | AI/ML Models |

| ID | Evidence | Source | Workbook |
|---|---|---|---|
| 12.4 | SDK groups in use | pillars.ai_ml.summary.sdk_groups_in_use | AI/ML Models |
| 12.5 | SDK hits | pillars.ai_ml.summary.sdk_hits | AI/ML Models |
| 12.6 | Bundled model files | pillars.ai_ml.summary.bundled_model_files | AI/ML Models |
| 12.7 | Production model files | pillars.ai_ml.summary.production_model_files | AI/ML Models |
| 12.8 | Unlicensed model files | pillars.ai_ml.summary.unlicensed_model_files | AI/ML Models |
| 12.9 | Model file rows | pillars.ai_ml.raw.models | AI/ML Models |
| 13.1 | Technologies total | pillars.tech_stack.summary.total_technologies | Tech Stack |
| 13.2 | EOL technology count | pillars.tech_stack.summary.eol_count | Tech Stack |
| 13.3 | Legacy technology count | pillars.tech_stack.summary.legacy_count | Tech Stack |
| 13.4 | Technologies by lifecycle | pillars.tech_stack.summary.by_lifecycle | Tech Stack |
| 13.5 | Active technologies | pillars.tech_stack.summary.active_count | Tech Stack |
| 13.6 | Technology inventory | pillars.tech_stack.raw.technologies | Tech Stack |
| 13.7 | EOL count (runtime verification) | pillars.eol.summary.eol_count | EOL Findings |
| 13.8 | Runtime EOL inventory | pillars.eol.raw.eol | EOL Findings |
| 14.1 | Consolidated top risks | decision.top_risks | Dashboard |
| 14.2 | Major risks register | decision.major_risks | Dashboard |
| 14.3 | Pillar classification register | decision.pillar_register | Dashboard |
| 14.4 | Deal-breaker register | decision.deal_breakers | Dashboard |
| 15.1 | Integrity status | integrity.status | Contents |
| 15.2 | Integrity error count | integrity.error_count | Contents |

| ID | Evidence | Source | Workbook |
|---|---|---|---|
| 15.3 | Integrity warning count | integrity.warn_count | Contents |
| 15.4 | Integrity issues | integrity.issues | Contents |
| 15.5 | Unverified pillars | decision.confidence.inputs.unverified_pillars | Dashboard |
| 15.6 | Confidence coverage | decision.confidence.inputs.coverage | Dashboard |
| 15.7 | Confidence risk signal | decision.confidence.inputs.risk_signal | Dashboard |
| 15.8 | Confidence signal count | decision.confidence.inputs.signal_count | Dashboard |
| 15.9 | Confidence pillar total | decision.confidence.inputs.total_pillars | Dashboard |
| 15.10 | Confidence degraded pillars | decision.confidence.inputs.degraded | Dashboard |
| 15.11 | Per-pillar assessment posture | decision.pillar_register | Dashboard |
| 15.12 | Taint flows found | taint_summary.total_flows | Taint Summary |
| 15.13 | Taint pass bounded | taint_summary.bounded | Taint Summary |
| 15.14 | CVE scope notes | pillars.cve.scope_notes | CVEs |
| 15.15 | Coverage gaps by pillar | pillars | Not applicable |
| 15.16 | Out-of-scope evidence register | pillars | Not applicable |
| O1 | Reproduction fingerprint | financial.fingerprint | Dashboard |

# Appendix R — Contradiction Ledger & Release Record

Every material conflict is listed with its status, resolution, and decision impact. Unresolved entries discount the confidence stated in the executive brief.

## R.1 Contradiction C-DC-001 — Resolved

Conflict: no double-counting findings across quantified economic figures. Subjects: Remediation, Risk valuation, Scenario analysis, Deal economics. Evidence: anti-double-counting ledger. Status: resolved. Resolution: shared-registry keys show each evidence once, on multiple surfaces, never added twice. Decision impact: none.

## R.2 Release Record

**Table R.1 — Release record**

| Field | Value |
|---|---|
| Report ID | TDD-20261016-060430 |
| Generated | 2026-10-16T07:15:51.747366+00:00 |
| Report version | 0.9.2 |
| Ledger summary | 1 resolved |
| Scan steps | 6 recorded |

© MTD Advisories

# Appendix C — Dependency & Lockfile Inventory

## C.1 Direct Dependency Inventory (top 20 of 500 rows)

The full direct dependency inventory contains 500 rows across Go modules and npm packages. The representative sample below covers the highest-priority dependencies in apps/advisor and .citools/golangci-lint; the full inventory is reproduced in the workbook sheet Dependency Packages.

**Table C.1 — Direct dependency inventory (top 20 of 500 rows)**

| Name | Specified | Pinned | Scope | Ecosystem | Manifest |
|---|---|---|---|---|---|
| 4d63.com/gocheckcompilerdirectives | v1.4.0 | No | go.mod | Go | .citools/src/golangci-lint/go.mod |
| 4d63.com/gochecknoglobals | v0.2.2 | No | go.mod | Go | .citools/src/golangci-lint/go.mod |
| cel.dev/expr | v0.25.3 | No | go.mod | Go | apps/advisor/go.mod |
| charm.land/lipgloss/v2 | v2.0.6 | No | go.mod | Go | .citools/src/golangci-lint/go.mod |
| cloud.google.com/go/compute/metadata | v0.9.0 | No | go.mod | Go | apps/advisor/go.mod |
| codeberg.org/chavacava/garif | v0.2.1 | No | go.mod | Go | .citools/src/golangci-lint/go.mod |
| codeberg.org/polyfloyd/go-errorlint | v1.9.0 | No | go.mod | Go | .citools/src/golangci-lint/go.mod |
| cuelabs.dev/go/oci/ociregistry | v0.0.0-20250722084951-074d06050084 | No | go.mod | Go | .citools/src/cue/go.mod |
| cuelang.org/go | v0.15.4 | No | go.mod | Go | .citools/src/cue/go.mod |
| dario.cat/mergo | v1.0.2 | No | go.mod | Go | .citools/src/air/go.mod |
| dev.gaijin.team/go/exhaustruct/v4 | v4.0.0 | No | go.mod | Go | .citools/src/golangci-lint/go.mod |
| dev.gaijin.team/go/exhaustruct/v5 | v5.2.0 | No | go.mod | Go | .citools/src/golangci-lint/go.mod |
| dev.gaijin.team/go/golib | v0.8.1 | No | go.mod | Go | .citools/src/golangci-lint/go.mod |
| filippo.io/edwards25519 | v1.2.0 | No | go.mod | Go | apps/advisor/go.mod |
| github.com/4meepo/tagalign | v1.4.4 | No | go.mod | Go | .citools/src/golangci-lint/go.mod |
| github.com/Abirdcfly/dupword | v0.1.8 | No | go.mod | Go | .citools/src/golangci-lint/go.mod |

| Name | Specified | Pinned | Scope | Ecosystem | Manifest |
|---|---|---|---|---|---|
| github.com/AdminBenni/iota-mixing | v1.0.0 | No | go.mod | Go | .citools/src/golangci-lint/go.mod |
| github.com/air-verse/air | v1.62.0 | No | go.mod | Go | .citools/src/air/go.mod |
| github.com/alecthomas/chroma/v2 | v2.27.0 | No | go.mod | Go | .citools/src/golangci-lint/go.mod |
| github.com/alecthomas/go-check-sumtype | v0.5.1-0.20260828200218-ae6904d28606 | No | go.mod | Go | .citools/src/golangci-lint/go.mod |

## C.2 Lockfile Resolution (full 50 lockfiles)

The full lockfile resolution table covers all 50 lockfiles in the repository. The root go.sum (3,721 resolved packages) and yarn.lock (3,253 resolved packages) dominate; the apps/* subdirectories each carry their own go.sum, indicating a multi-module Go workspace. The buyer's supply-chain team should review the lockfile-resolution pattern to understand the build's dependency-resolution surface.

**Table C.2 — Lockfile resolution (full 50 rows)**

| Lockfile | Resolved packages |
|---|---|
| .citools/src/air/go.sum | 210 |
| .citools/src/cue/go.sum | 72 |
| .citools/src/golangci-lint/go.sum | 567 |
| .citools/src/govulncheck/go.sum | 22 |
| .citools/src/lefthook/go.sum | 108 |
| .citools/src/swagger/go.sum | 149 |
| apps/advisor/go.sum | 1,757 |
| apps/alerting/alertenrichment/go.sum | 98 |
| apps/alerting/alertrulequality/go.sum | 96 |
| apps/alerting/notifications/go.sum | 300 |
| apps/alerting/rules/go.sum | 236 |
| apps/annotation/go.sum | 236 |
| apps/collections/go.sum | 170 |
| apps/correlations/go.sum | 236 |
| apps/dashboard/go.sum | 381 |
| apps/dashvalidator/go.sum | 1,747 |
| apps/example/go.sum | 252 |
| apps/folder/go.sum | 170 |

| Lockfile | Resolved packages |
|---|---|
| apps/iam/go.sum | 236 |
| apps/live/go.sum | 236 |
| apps/logsdrilldown/go.sum | 236 |
| apps/playlist/go.sum | 236 |
| apps/plugins/go.sum | 1,480 |
| apps/preferences/go.sum | 170 |
| apps/provisioning/go.sum | 284 |
| apps/quotas/go.sum | 240 |
| apps/scope/go.sum | 100 |
| apps/secret/go.sum | 178 |
| apps/shorturl/go.sum | 252 |
| devenv/docker/blocks/prometheus_high_card/go.sum | 36 |
| devenv/docker/blocks/prometheus_utf8/go.sum | 36 |
| go.sum | 3,721 |
| hack/go.sum | 22 |
| pkg/apimachinery/go.sum | 148 |
| pkg/apiserver/go.sum | 266 |
| pkg/codegen/go.sum | 146 |
| pkg/infra/features/go.sum | 196 |
| pkg/infra/log/go.sum | 46 |
| pkg/infra/metrics/go.sum | 50 |
| pkg/plugins/codegen/go.sum | 136 |
| pkg/plugins/go.sum | 380 |
| pkg/plugins/manager/testdata/test-app-with-backend/go.sum | 235 |
| pkg/storage/unified/apistore/go.sum | 491 |
| pkg/storage/unified/resource/kv/go.sum | 115 |
| pkg/storage/unified/resourceclient/go.sum | 141 |
| pkg/storage/unified/resourcepb/go.sum | 18 |
| pkg/util/sqlite/go.sum | 50 |
| scripts/go-workspace/go.sum | 2 |
| scripts/modowners/go.sum | 2 |
| yarn.lock | 3,253 |

# Appendix D — CVE Inventory

The full CVE inventory contains 154 advisories across 9 distinct packages. Severity is reported as UNKNOWN for all 154 because the OSV feed did not return a CVSS-severity band. Fix availability is false for every advisory; fixed version is not applicable; remediation class is no_fixed_release for every advisory. The dominant affected package is github.com/grafana/grafana (the repo's own module path); the second-largest cluster is github.com/nats-io/nats-server v2.14.6.

**Table D.1 — CVE records (top 18 of 154 rows)**

| ID | Severity | Package | Ecosystem | Version | Fix available | Fixed version | Remediation class |
|---|---|---|---|---|---|---|---|
| GHSA-3f24-pcvm-5jqc | UNKNOWN | github.com/nats-io/nats-server | go | v2.14.6 | false | n/a | no_fixed_release |
| GHSA-3jq7-8ph8-63xm | UNKNOWN | github.com/grafana/grafana | go | 0.0.0 | false | n/a | no_fixed_release |
| GHSA-3p62-42x7-gxg5 | UNKNOWN | github.com/grafana/grafana | go | 0.0.0 | false | n/a | no_fixed_release |
| GHSA-3w66-95m3-8jxg | UNKNOWN | github.com/grafana/grafana | go | 0.0.0 | false | n/a | no_fixed_release |
| GHSA-46m5-8hpj-p5p5 | UNKNOWN | github.com/grafana/grafana | go | 0.0.0 | false | n/a | no_fixed_release |
| GHSA-4724-7jwc-3fpw | UNKNOWN | github.com/grafana/grafana | go | 0.0.0 | false | n/a | no_fixed_release |
| GHSA-4pwp-cx67-5cpx | UNKNOWN | github.com/grafana/grafana | go | 0.0.0 | false | n/a | no_fixed_release |
| GHSA-52jh-2xxh-pwh6 | UNKNOWN | github.com/nats-io/nats-server | go | v2.14.6 | false | n/a | no_fixed_release |
| GHSA-55h8-8g96-x4hj | UNKNOWN | github.com/nats-io/nats-server | go | v2.14.6 | false | n/a | no_fixed_release |
| GHSA-66c4-2g2v-54qw | UNKNOWN | github.com/grafana/grafana | go | 0.0.0 | false | n/a | no_fixed_release |
| GHSA-69j6-29vr-p3j9 | UNKNOWN | github.com/grafana/grafana | go | 0.0.0 | false | n/a | no_fixed_release |
| GHSA-6wh2-8hw7-jw94 | UNKNOWN | github.com/grafana/grafana | go | 0.0.0 | false | n/a | no_fixed_release |
| GHSA-7g92-g4vh-hp84 | UNKNOWN | github.com/grafana/grafana | go | 0.0.0 | false | n/a | no_fixed_release |
| GHSA-7m2x-qhrq-rp8h | UNKNOWN | github.com/grafana/grafana | go | 0.0.0 | false | n/a | no_fixed_release |
| GHSA-7phr-6cc9-4m5q | UNKNOWN | github.com/grafana/grafana | go | 0.0.0 | false | n/a | no_fixed_release |

| ID | Severity | Package | Ecosystem | Version | Fix available | Fixed version | Remediation class |
|---|---|---|---|---|---|---|---|
| GHSA-8r68-gvr4-jh7j | UNKNOWN | github.com/nats-io/nats-server | go | v2.14.6 | false | n/a | no_fixed_release |
| GHSA-9493-h4f5-633x | UNKNOWN | github.com/grafana/grafana | go | 0.0.0 | false | n/a | no_fixed_release |
| GHSA-9983-vrx2-fg9c | UNKNOWN | github.com/nats-io/nats-server | go | v2.14.6 | false | n/a | no_fixed_release |

*The full 154-row inventory is reproduced in the workbook sheet CVEs. The buyer's security team should re-query each advisory against the NVD to obtain the CVSS-severity rating, because the OSV feed's UNKNOWN rating is itself a coverage gap. The 19 unresolved packages should be cross-referenced against the seller's lockfile inventory to identify which advisories apply to dependencies that are actually loaded at runtime.*

© MTD Services

# Appendix E — License Findings Inventory

## E.1 Per-Package License Findings (top 18 of 5,497 rows)

The full per-package license inventory contains 5,497 rows. The representative sample below covers the highest-risk viral and restrictive licenses (AGPL-3.0, AGPL-3.0-only, GPL v3, MPL-2.0) and the dominant permissive licenses (Apache-2.0, MIT). The full inventory is reproduced in the workbook sheet License Contamination.

**Table E.1 — Per-package license findings (top 18 of 5,497 rows)**

| Package name | Detected license | Viral | Restrictive | Confidence | Source file |
|---|---|---|---|---|---|
| (repo-level) | AGPL-3.0 | Yes | Yes | 0.8 | LICENSE |
| (repo-level) | AGPL-3.0 | Yes | Yes | 0.8 | packages/grafana-o11y-ds-frontend/LICENSE_AGPL |
| (repo-level) | AGPL-3.0 | Yes | Yes | 0.8 | packages/grafana-sql/LICENSE_AGPL |
| (repo-level) | AGPL-3.0 | Yes | Yes | 0.8 | packaging/msi/resources/LICENSE.md |
| (repo-level) | AGPL-3.0 | Yes | Yes | 0.8 | packaging/msi/resources/LICENSE.rtf |
| @grafana/o11y-ds-frontend | AGPL-3.0-only | Yes | Yes | 0.7 | package.json |
| @grafana/prometheus | AGPL-3.0-only | Yes | Yes | 0.7 | package.json |
| @grafana/prometheus | AGPL-3.0-only | Yes | Yes | 0.7 | public/app/plugins/datasource/cloudwatch/package.json |
| @grafana/sql | AGPL-3.0-only | Yes | Yes | 0.7 | package.json |
| envtpl | GPL v3 | Yes | Yes | 0.7 | devenv/docker/blocks/collectd/requirements.txt |
| @axe-core/playwright | MPL-2.0 | Yes | No | 0.7 | package.json |
| axe-core | MPL-2.0 | Yes | No | 0.7 | package.json |
| (repo-level) | Apache-2.0 | No | No | 0.8 | packages/grafana-alerting/LICENSE_APACHE2 |
| (repo-level) | Apache-2.0 | No | No | 0.8 | packages/grafana-api-clients/LICENSE_APACHE2 |
| (repo-level) | Apache-2.0 | No | No | 0.8 | packages/grafana-data/LICENSE_APACHE2 |
| (repo-level) | Apache-2.0 | No | No | 0.8 | packages/grafana-e2e-selectors/LICENSE_APACHE2 |

| Package name | Detected license | Viral | Restrictive | Confidence | Source file |
|---|---|---|---|---|---|
| (repo-level) | Apache-2.0 | No | No | 0.8 | packages/grafana-eslint-rules/LICENSE_APACHE2 |
| (repo-level) | Apache-2.0 | No | No | 0.8 | packages/grafana-flamegraph/LICENSE_APACHE2 |

## E.2 License Clause Findings (48 rows)

License clause findings identify attribution, warranty_disclaimer, and patent_grant clauses across 48 license files. Each clause is paired with its severity and matched-package count; the matched-package count is 0 entries in every case because the clause-level analysis is performed at the license-file level rather than the package level.

**Table E.2 — License clause findings (top 14 of 48 rows)**

| License file | Clause type | Severity | Matched packages | Clause text |
|---|---|---|---|---|
| packages/grafana-alerting/LICENSE_APACHE2 | attribution | attribution | 0 | any Derivative Works that You distribute, all copyright, patent, trademark, and attribution notices from the Source form... |
| packages/grafana-api-clients/LICENSE_APACHE2 | attribution | attribution | 0 | any Derivative Works that You distribute, all copyright, patent, trademark, and attribution notices from the Source form... |
| packages/grafana-data/LICENSE_APACHE2 | attribution | attribution | 0 | any Derivative Works that You distribute, all copyright, patent, trademark, and attribution notices from the Source form... |
| packages/grafana-e2e-selectors/LICENSE_APACHE2 | attribution | attribution | 0 | any Derivative Works that You distribute, all copyright, patent, trademark, and attribution notices from the Source form... |
| packages/grafana-eslint-rules/LICENSE_APACHE2 | attribution | attribution | 0 | any Derivative Works that You distribute, all copyright, patent, trademark, and attribution notices from the Source form... |
| packages/grafana-flamegraph/LICENSE_APACHE2 | attribution | attribution | 0 | any Derivative Works that You distribute, all copyright, patent, trademark, and attribution notices from the Source form... |
| packages/grafana-i18n/LICENSE_APACHE2 | attribution | attribution | 0 | any Derivative Works that You distribute, all copyright, patent, trademark, and attribution notices from the Source form... |

| License file | Clause type | Severity | Matched packages | Clause text |
|---|---|---|---|---|
| packages/grafana-openapi/LICENSE_APACHE2 | attribution | attribution | 0 | any Derivative Works that You distribute, all copyright, patent, trademark, and attribution notices from the Source form... |
| packages/grafana-plugin-compat/LICENSE_APACHE2 | attribution | attribution | 0 | any Derivative Works that You distribute, all copyright, patent, trademark, and attribution notices from the Source form... |
| packages/grafana-runtime/LICENSE_APACHE2 | attribution | attribution | 0 | any Derivative Works that You distribute, all copyright, patent, trademark, and attribution notices from the Source form... |
| packages/grafana-schema/LICENSE_APACHE2 | attribution | attribution | 0 | any Derivative Works that You distribute, all copyright, patent, trademark, and attribution notices from the Source form... |
| packages/grafana-test-utils/LICENSE_APACHE2 | attribution | attribution | 0 | any Derivative Works that You distribute, all copyright, patent, trademark, and attribution notices from the Source form... |
| packages/grafana-ui/LICENSE_APACHE2 | attribution | attribution | 0 | any Derivative Works that You distribute, all copyright, patent, trademark, and attribution notices from the Source form... |
| pkg/build/wire/LICENSE | attribution | attribution | 0 | any Derivative Works that You distribute, all copyright, patent, trademark, and attribution notices from the Source form... |

*The 48 clause findings indicate that the seller's Apache-2.0 licensed packages carry the standard attribution clause, which is permissive (not viral). The AGPL-3.0 clause findings (warranty_disclaimer, patent_grant) are noted separately and should be reviewed against the buyer's copyleft-exposure analysis.*

© M

# Appendix F — Complexity, Duplication & SQALE Detail

## F.1 Complexity Hotspot Symbols (top 16 of 57,667 rows)

**Table F.1 — Complexity hotspot symbols**

| Name | Filepath | Cyclo. | Cogn. | Maint. idx | Max nest | Param count | Composite risk |
|---|---|---|---|---|---|---|---|
| Plot | public/vendor/flot/jquery.flot.js | 473 | 1,683 | 0 | n/r | n/r | 0.95 |
| updateSchema | public/app/features/dashboard/state/DashboardMigrator.ts | 138 | 210 | 0 | n/r | n/r | 0.90 |
| jsonDataToMetaJSONData | apps/plugins/pkg/app/meta/converter.go | 134 | 138 | 0 | n/r | n/r | 0.9125 |
| InitializeForTest | pkg/server/bootstrap/wire/wire_gen.go | 130 | 129 | 0 | n/r | n/r | 0.8575 |
| Initialize | pkg/server/bootstrap/wire/wire_gen.go | 130 | 129 | 0 | n/r | n/r | 0.87 |
| TestRouteGetRuleStatuses | pkg/services/ngalert/api/api_prometheus_test.go | 121 | 186 | 0 | n/r | n/r | 0.9125 |
| bytes2Value | pkg/util/xorm/session_convert.go | 121 | 558 | 0 | n/r | n/r | 0.90 |
| slice2Bean | pkg/util/xorm/session.go | 120 | 455 | 0 | n/r | n/r | 0.90 |
| getNavTitle | public/app/core/utils/navBarItem-translations.ts | 109 | 109 | 10.7 | n/r | n/r | 0.8911 |
| createGrafDir | pkg/tests/testinfra/testinfra.go | 102 | 109 | 0 | n/r | n/r | 0.9375 |
| authorize | pkg/services/ngalert/api/authorization.go | 90 | 84 | 0 | n/r | n/r | 0.90 |
| process | pkg/registry/apis/provisioning/controller/repository.go | 88 | 134 | 0 | n/r | n/r | 0.90 |
| graphToTimeseriesOptions | public/app/plugins/panel/timeseries/migrations.ts | 80 | 307 | 2.05 | n/r | n/r | 0.9084 |
| readDashboardIter | pkg/services/store/kind/dashboard/dashboard.go | 76 | 237 | 14.82 | n/r | n/r | 0.9329 |
| getStatementPosition | public/app/plugins/datasource/cloudwatch/language/cloudwatch-logs-sql/completion/statementPosition.ts | 73 | 103 | 10.96 | n/r | n/r | 0.8906 |
| buildUpdates | pkg/util/xorm/statement.go | 69 | 623 | 16.96 | n/r | n/r | 0.8661 |

## F.2 Duplication Groups (top 18 of 18,406 rows)

Duplication groups total 18,406 across the codebase. Each group is identified by its block-start line, block length, and occurrence count. The representative sample below covers common duplication patterns; the full inventory is reproduced in the workbook sheet Complex Symbols.

**Table F.2 — Duplication groups (top 18 of 18,406 rows)**

| Block start line | Block length (lines) | Occurrences |
|---|---|---|
| 1 | 10 | 2 |
| 10 | 10 | 2 |
| 100 | 10 | 2 |
| 1,009 | 10 | 3 |
| 101 | 10 | 3 |
| 1,017 | 10 | 3 |
| 102 | 10 | 2 |
| 104 | 10 | 2 |
| 104 | 10 | 6 |
| 105 | 10 | 6 |
| 106 | 10 | 2 |
| 1,063 | 10 | 2 |
| 1,077 | 10 | 4 |
| 1,077 | 10 | 2 |
| 108 | 10 | 2 |
| 109 | 10 | 2 |
| 11 | 10 | 2 |
| 110 | 10 | 2 |

*Block start line is the 1-based line at which the duplicated symbol/module block begins in its file; block length is that block's span and occurrences is how many copies were found.*

## F.3 SQALE Debt Modules (top 20 of 16,658 rows)

SQALE debt modules total 16,658. The representative sample below covers modules with grade A (no remediation effort required). The full inventory is reproduced in the workbook sheet Complex Symbols.

**Table F.3 — SQALE debt modules**

| Base lines | Debt ratio | Module | Remediation effort | SQALE grade |
|---|---|---|---|---|
| 1 | 0 | e2e-playwright/test-plugins/grafana-extensionstest-app/components/ActionButton/index.ts | 0 | A |
| 1 | 0 | e2e-playwright/test-plugins/grafana-extensionstest-app/components/App/index.tsx | 0 | A |
| 1 | 0 | e2e-playwright/test-plugins/grafana-extensionstest-app/components/QueryModal/index.ts | 0 | A |

© MT

| Base lines | Debt ratio | Module | Remediation effort | SQALE grade |
|---|---|---|---|---|
| 1 | 0 | e2e-playwright/test-plugins/grafana-extensionstest-app/plugins/grafana-extensionexample1-app/components/App/index.tsx | 0 | A |
| 1 | 0 | e2e-playwright/test-plugins/grafana-extensionstest-app/plugins/grafana-extensionexample2-app/components/App/index.tsx | 0 | A |
| 1 | 0 | grafana-mixin/scripts/common.sh | 0 | A |
| 1 | 0 | packages/grafana-data/src/types/slider.ts | 0 | A |
| 1 | 0 | packages/grafana-e2e-selectors/src/selectors/constants.ts | 0 | A |
| 1 | 0 | packages/grafana-e2e-selectors/src/types/index.ts | 0 | A |
| 1 | 0 | packages/grafana-plugin-configs/constants.ts | 0 | A |
| 1 | 0 | packages/grafana-plugin-configs/jest/mocks/images.ts | 0 | A |
| 1 | 0 | packages/grafana-ui/src/components/Forms/types.ts | 0 | A |
| 1 | 0 | packages/grafana-ui/src/components/SecretInput/index.tsx | 0 | A |
| 1 | 0 | packages/grafana-ui/src/components/SecretTextArea/index.tsx | 0 | A |
| 1 | 0 | packages/grafana-ui/src/types/interactiveTable.ts | 0 | A |
| 1 | 0 | packages/grafana-ui/src/types/orientation.ts | 0 | A |
| 1 | 0 | packages/grafana-ui/src/types/select.ts | 0 | A |
| 1 | 0 | packages/grafana-ui/src/types/size.ts | 0 | A |
| 1 | 0 | pkg/plugins/pluginassets/testdata/module-hash-no-manifest-txt/module.js | 0 | A |
| 1 | 0 | pkg/plugins/pluginassets/testdata/module-hash-no-module-js/something.js | 0 | A |

©

# Appendix G — Debt Marker Inventory

The full debt-marker inventory contains 883 rows. The representative sample below covers the 16 highest-severity markers; the full inventory is reproduced in the workbook sheet Debt Markers. Markers are recorded with file path, line number, marker type, severity, critical-path placement, and marker text.

**Table G.1 — Debt markers (top 16 of 883 rows)**

| Filepath | Line | Marker | Severity | Critical path | Text |
|---|---|---|---|---|---|
| eslint.config.js | 40 | fixme | high | No | // [FIXME] add comment about this applying everywhere |
| eslint.config.js | 624 | fixme | high | No | // FIXME: Remove once all enterprise issues are fixed - |
| eslint.config.js | 638 | fixme | high | No | // FIXME: Remove once all enterprise issues are fixed - |
| eslint.config.js | 658 | fixme | high | No | // FIXME: Remove once all enterprise issues are fixed - |
| eslint.config.js | 683 | fixme | high | No | // FIXME: Remove once all enterprise issues are fixed - |
| packages/grafana-api-clients/src/clients/rtkq/dashboard/v2/endpoints.gen.ts | 1,073 | fixme | high | No | /** Link type. Accepted values are dashboards (to refer to another dashboard) and link... |
| packages/grafana-api-clients/src/clients/rtkq/dashboard/v2beta1/endpoints.gen.ts | 1,617 | fixme | high | No | /** Link type. Accepted values are dashboards (to refer to another dashboard)... |
| packages/grafana-api-clients/src/clients/rtkq/provisioning/v0alpha1/endpoints.gen.ts | 1,710 | fixme | high | No | /** Paths to be deleted. Examples: dashboard.json (file) ... |
| packages/grafana-api-clients/src/clients/rtkq/provisioning/v0alpha1/endpoints.gen.ts | 1,734 | fixme | high | No | /** Paths to be deleted. Examples ... |
| packages/grafana-api-clients/src/clients/rtkq/provisioning/v0alpha1/endpoints.gen.ts | 1,762 | fixme | high | No | /** FIXME: we should validate this in admission hooks Target branch for export (only git) */ |
| packages/grafana-api-clients/src/clients/rtkq/provisioning/v0alpha1/endpoints.gen.ts | 1,770 | fixme | high | No | /** FIXME: we should validate this in admission hooks Prefix in target file system */ |
| packages/grafana-schema/src/schema/dashboard/v2/types.spec.gen.ts | 1,586 | fixme | high | No | // FIXME: The type is generated as `type: DashboardLinkType \| dashboardLinkType.Link;`... |
| packages/grafana-schema/src/schema/dashboard/v2alpha1/types.spec.gen.ts | 1,050 | fixme | high | No | // FIXME: The type is generated as `type: DashboardLinkType \| dashboardLinkType.Link;`... |

| Filepath | Line | Marker | Severity | Critical path | Text |
|---|---|---|---|---|---|
| packages/grafana-schema/src/schema/dashboard/v2beta1/types.spec.gen.ts | 1,578 | fixme | high | No | // FIXME: The type is generated as `type: DashboardLinkType \| dashboardLinkType.Link;`... |
| packages/grafana-test-utils/src/types/browse-dashboards.ts | 1 | fixme | high | No | // FIXME: This file is a duplication of types within the core code |
| packages/grafana-ui/src/components/Forms/commonStyles.ts | 39 | hack | high | No | /* Welcome to 2005. This is a HACK to get rid of Chromes default autofill styling */ |

*The eslint.config.js cluster of five FIXMEs (lines 40, 624, 638, 658, 683) is the highest-concentration debt pattern in the codebase. Four of the five read 'FIXME: Remove once all enterprise issues are fixed' — this pattern indicates a known outstanding clean-up tied to the open-source / enterprise boundary that the seller has deferred. The buyer's counsel should require disclosure of the specific 'enterprise issues' referenced.*

© MTD Services

# Appendix H — Attack Surface Candidate Endpoints

The full attack-surface inventory contains 118 candidate entry points. The representative sample below covers application_entry (medium risk, untested) and candidate (low risk, mixed test status) endpoints. The full inventory is reproduced in the workbook sheet Attack Surface.

**Table H.1 — Candidate entry points (top 14 of 118 rows)**

| Symbol name | Filepath | Line | Endpoint type | Exposure | Framework | Risk level | Test status |
|---|---|---|---|---|---|---|---|
| murmurhash3_32_gc | packages/grafana-flamegraph/src/FlameGraph/murmur3.ts | 21 | application_entry | application | unknown | medium | untested |
| defineThemes | packages/grafana-ui/src/components/Monaco/theme.ts | 31 | application_entry | application | unknown | medium | untested |
| resetSelectStyles | packages/grafana-ui/src/components/Select/resetSelectStyles.ts | 6 | application_entry | application | unknown | medium | untested |
| flatten | public/app/core/utils/flatten.ts | 4 | application_entry | application | unknown | medium | untested |
| requestAnimationFrame | public/app/features/explore/TraceView/components/utils/test/requestAnimationFrame.ts | 17 | application_entry | application | unknown | medium | untested |
| cancelAnimationFrame | public/app/features/explore/TraceView/components/utils/test/requestAnimationFrame.ts | 21 | application_entry | application | unknown | medium | untested |
| useMigrations | public/app/plugins/datasource/azuremonitor/components/LogsQueryEditor/useMigrations.ts | 43 | application_entry | application | unknown | medium | untested |
| createMockDatasource | public/app/plugins/datasource/azuremonitor/mocks/datasource.ts | 27 | application_entry | application | unknown | medium | untested |
| createMockPanelData | public/app/plugins/datasource/azuremonitor/mocks/panelData.ts | 7 | application_entry | application | unknown | medium | untested |
| createMockQuery | public/app/plugins/datasource/azuremonitor/mocks/query.ts | 4 | application_entry | application | unknown | medium | untested |
| messageFromError | public/app/plugins/datasource/azuremonitor/utils/messageFromError.ts | 13 | application_entry | application | unknown | medium | untested |
| migrateAnnotation | public/app/plugins/datasource/azuremonitor/utils/migrateAnnotation.ts | 10 | application_entry | application | unknown | medium | untested |
| useLastError | public/app/plugins/datasource/azuremonitor/utils/useLastError.ts | 9 | application_entry | application | unknown | medium | untested |

| Symbol name | Filepath | Line | Endpoint type | Exposure | Framework | Risk level | Test status |
|---|---|---|---|---|---|---|---|
| migrateAnnotation | packages/grafana-sql/src/utils/migration.ts | 6 | application_entry | application | unknown | medium | untested |

*External reachability could not be established for 117 of the 118 candidates because framework semantics were unresolved. 95 candidates originate from application_entry detection. The candidate exposure score of 0.1047 is an upper-bound signal for review prioritisation, not a confirmed vulnerability rate.*

© MTD Services

# Appendix I — Key-Person & Knowledge Loss Vulnerability Detail

## I.1 Key-Person Risk Rows (top 14 of 3,217 rows)

**Table I.1 — Key-person risk detail**

| Author name | Author email | Risk score | Total commits | Departure impact |
|---|---|---|---|---|
| ryantxu@gmail.com | ryantxu@gmail.com | 0.26468 | 234,722 | Low |
| ashley.harrison@grafana.com | ashley.harrison@grafana.com | 0.216071 | 131,354 | Low |
| joshhunt@users.noreply.github.com | joshhunt@users.noreply.github.com | 0.103575 | 63,357 | Low |
| torkel@grafana.com | torkel@grafana.com | 0.0803188 | 90,155 | Low |
| stephanie.hingtgen@grafana.com | stephanie.hingtgen@grafana.com | 0.0686434 | 40,405 | Low |
| hugo.haggmark@gmail.com | hugo.haggmark@gmail.com | 0.0615704 | 65,580 | Low |
| 29139614+renovate[bot]@users.noreply.github.com | 29139614+renovate[bot]@users.noreply.github.com | 0.0612221 | 11,923 | Low |
| roberto.jimenez@grafana.com | roberto.jimenez@grafana.com | 0.0528524 | 132,108 | Low |
| macabu@users.noreply.github.com | macabu@users.noreply.github.com | 0.050841 | 31,702 | Low |
| leeoniya@gmail.com | leeoniya@gmail.com | 0.0498112 | 48,560 | Low |
| 45561153+kaydelaney@users.noreply.github.com | 45561153+kaydelaney@users.noreply.github.com | 0.0477866 | 28,972 | Low |
| marcus.efraimsson@gmail.com | marcus.efraimsson@gmail.com | 0.0463529 | 26,936 | Low |
| academo@users.noreply.github.com | academo@users.noreply.github.com | 0.0456318 | 7,270 | Low |
| laura.fernandez@grafana.com | laura.fernandez@grafana.com | 0.041486 | 11,503 | Low |

## I.2 Module Bus-Factor Rows (top 14 of 2,449 rows)

**Table I.2 — Module bus-factor detail**

| Module path | Bus factor | Total files |
|---|---|---|
| pkg/setting | 22 | 54 |
| pkg/api | 12 | 96 |
| _unassigned | 9 | 6,605 |
| grafana-data/src/types | 9 | 58 |

| Module path | Bus factor | Total files |
|---|---|---|
| pkg/services/sqlstore | 9 | 21 |
| public/app/types | 9 | 26 |
| pkg/api/dtos | 8 | 18 |
| grafana-data/src/utils | 7 | 51 |
| pkg/registry/backgroundsvcs | 7 | 1 |
| pkg/server | 7 | 29 |
| pkg/services/sqlstore/migrations | 7 | 41 |
| pkg/tests/testinfra | 7 | 2 |
| pkg/util | 7 | 33 |
| public/app/core/utils | 7 | 53 |

## I.3 Knowledge Loss Vulnerability Rows (top 14 of 3,217 rows)

**Table I.3 — Knowledge loss vulnerability detail**

| Author name | Author email | Tenure days | Days since last commit | Retention risk | Blast radius at risk |
|---|---|---|---|---|---|
| undef1nd | 10127682+undef1nd@users.noreply.github.com | 1,021 | 0 | Critical | 19 |
| Galen Kistler | 109082771+gtk-grafana@users.noreply.github.com | 1,534 | 0 | Critical | 68 |
| Fayzal Ghantiwala | 114010985+fayzal-g@users.noreply.github.com | 896 | 55 | Critical | 0 |
| volcanonoodle | 114113189+volcanonoodle@users.noreply.github.com | 1,063 | 2 | Critical | 72 |
| Liza Detrick | 114438185+L2D2Grafana@users.noreply.github.com | 883 | 7 | Critical | 0 |
| Shubham Nainwal | 118589955+Shubham19032004@users.noreply.github.com | 491 | 93 | Critical | 0 |
| grafana-delivery-bot[bot] | 132647405+grafana-delivery-bot[bot]@users.noreply.github.com | 1,212 | 150 | Critical | 0 |
| Fabrizio | 135109076+fabrizio-grafana@users.noreply.github.com | 1,198 | 1 | Critical | 42 |

| Author name | Author email | Tenure days | Days since last commit | Retention risk | Blast radius at risk |
|---|---|---|---|---|---|
| Lucy Chen | 140550297+lucychen-grafana@users.noreply.github.com | 1,154 | 21 | Critical | 0 |
| Tobias Skarhed | 1438972+tskarhed@users.noreply.github.com | 2,455 | 8 | Critical | 30 |
| Marc M. | 146180665+grafakus@users.noreply.github.com | 716 | 0 | Critical | 121 |
| alerting-team[bot] | 158350966+alerting-team[bot]@users.noreply.github.com | 591 | 0 | Critical | 0 |
| Adam Yeats | 16296989+adamyeats@users.noreply.github.com | 1,186 | 3 | Critical | 15 |
| Sofia Papagiannaki | 1632407+papagian@users.noreply.github.com | 1,820 | 0 | Critical | 246 |

© MTD Ser

# Appendix J — Maintenance, IaC & CI/CD Detail

## J.1 Maintenance Assets (none recorded)

No maintenance asset rows are recorded in this scan. The asset-staleness scanner produced zero rows, which is why the maintenance pillar is demoted to UNVERIFIED. The buyer should commission an independent asset-staleness assessment as part of pre-signing diligence.

## J.2 IaC Findings (full 1 row)

**Table J.1 — IaC findings**

| Rule ID | Title | Severity | File | Line | Remediation |
|---|---|---|---|---|---|
| IAC-DOCKER-ROOT | Container runs as root | Medium | Dockerfile | 403 | Add a dedicated non-root user and USER directive in every stage. |

## J.3 CI Files (top 20 of 98 rows)

**Table J.2 — CI files**

| CI file path |
|---|
| /content/grafana/.github/workflows/actionlint.yml |
| /content/grafana/.github/workflows/add-to-whats-new.yml |
| /content/grafana/.github/workflows/alerting-swagger-gen.yml |
| /content/grafana/.github/workflows/alerting-update-module.yml |
| /content/grafana/.github/workflows/analytics-events-report.yml |
| /content/grafana/.github/workflows/auto-milestone.yml |
| /content/grafana/.github/workflows/backend-code-checks.yml |
| /content/grafana/.github/workflows/backend-unit-tests.yml |
| /content/grafana/.github/workflows/backport-trigger.yml |
| /content/grafana/.github/workflows/backport-workflow.yml |
| /content/grafana/.github/workflows/build-docker-variants.yml |
| /content/grafana/.github/workflows/build-go-matrix.yml |
| /content/grafana/.github/workflows/bump-version.yml |
| /content/grafana/.github/workflows/changelog.yml |
| /content/grafana/.github/workflows/check-frontend-test-coverage.yml |
| /content/grafana/.github/workflows/cleanup-branches.yml |
| /content/grafana/.github/workflows/codeowners-validator.yml |
| /content/grafana/.github/workflows/commands.yml |

| CI file path |
| --- |
| /content/grafana/.github/workflows/community-release.yml |
| /content/grafana/.github/workflows/create-next-release-branch.yml |

### J.4 Coverage Artifacts (none)

No coverage artifacts are recorded in this scan. This is consistent with the reported_coverage_pct = not quantified figure in Section 11: the CI surface does not publish a coverage report that the scanner could consume.

© MTD Services

# Appendix K — Privacy & Secrets Detail

## K.1 Privacy Findings (full 1 row)

**Table K.1 — Privacy findings**

| Rule ID | Category | Title | Severity | File | Line | Remediation |
|---|---|---|---|---|---|---|
| PRIV-PLAINTEXT-CRED | Plaintext credential | Plaintext credential material in committed config | Medium | .github/workflows/reliance-scan.yml | 26 | Move to a secret store / environment injection and rotate the committed value. |

## K.2 Secret Candidates (top 14 of 427 rows)

**Table K.2 — Secret candidates**

| Kind | File | Line | Confidence | In history | Risk cap |
|---|---|---|---|---|---|
| aws_access_key | pkg/tsdb/cloudwatch/models/settings_test.go | 202 | 95 | Yes | critical |
| aws_access_key | pkg/tsdb/cloudwatch/models/settings_test.go | 223 | 95 | Yes | critical |
| aws_access_key | pkg/tsdb/cloudwatch/models/settings_test.go | 248 | 95 | Yes | critical |
| aws_access_key | pkg/tsdb/cloudwatch/models/settings_test.go | 112 | 95 | No | critical |
| aws_access_key | pkg/tsdb/cloudwatch/models/settings_test.go | 130 | 95 | No | critical |
| aws_access_key | pkg/tsdb/cloudwatch/models/settings_test.go | 148 | 95 | No | critical |
| aws_access_key | pkg/tsdb/cloudwatch/models/settings_test.go | 166 | 95 | No | critical |
| aws_access_key | pkg/tsdb/cloudwatch/models/settings_test.go | 184 | 95 | No | critical |
| aws_access_key | pkg/tsdb/cloudwatch/models/settings_test.go | 34 | 95 | No | critical |

| Kind | File | Line | Confidence | In history | Risk cap |
|---|---|---|---|---|---|
| aws_access_key | pkg/tsdb/cloudwatch/models/settings_test.go | 56 | 95 | No | critical |
| aws_access_key | pkg/tsdb/cloudwatch/models/settings_test.go | 87 | 95 | No | critical |
| aws_access_key | pkg/tsdb/cloudwatch/models/settings_test.go | 102 | 95 | No | critical |
| aws_access_key | pkg/tsdb/cloudwatch/models/settings_test.go | 71 | 95 | No | critical |
| aws_access_key | pkg/util/uri_sanitize_test.go | 29 | 94 | No | critical |

© MTD Service

# Appendix L — AI/ML Inventory

## L.1 Model Files (none recorded)

No model-file rows are recorded in this scan. This is consistent with the bundled_model_files = 0 and production_model_files = 0 figures in Section 12: the AI/ML surface area is composed of SDK calls and framework imports, not bundled model artifacts. The buyer's AI/ML team should validate that no model is loaded at runtime; if runtime model loading is detected, the model license and training-data provenance must be verified.

© MTD Services

# Appendix M — Integrity & Data-Quality Issues

The full 8-row integrity issues inventory is reproduced below. Five warnings are critical_pillar_no_findings (high/critical pillars rendering no item-level findings); three warnings are unverified_pillar (partial or not-assessed pillars). No errors are recorded.

**Table M.1 — Integrity issues (full 8 rows)**

| Tier | Code | Subject | Message |
|---|---|---|---|
| Warn | critical_pillar_no_findings | pillars.coverage | Pillar 'coverage' is high but renders no item-level findings. |
| Warn | critical_pillar_no_findings | pillars.entanglement | Pillar 'entanglement' is critical but renders no item-level findings. |
| Warn | critical_pillar_no_findings | pillars.integration | Pillar 'integration' is high but renders no item-level findings. |
| Warn | critical_pillar_no_findings | pillars.ip_ownership | Pillar 'ip_ownership' is high but renders no item-level findings. |
| Warn | critical_pillar_no_findings | pillars.test_deficit | Pillar 'test_deficit' is high but renders no item-level findings. |
| Warn | unverified_pillar | pillars.cve | Pillar 'cve' is partial — conclusions UNVERIFIED. |
| Warn | unverified_pillar | pillars.maintenance | Pillar 'maintenance' is not_assessed — conclusions UNVERIFIED. |
| Warn | unverified_pillar | pillars.secrets | Pillar 'secrets' is partial — conclusions UNVERIFIED. |

© MTD

# Appendix N — Taint Analysis Summary

## N.1 Taint Summary

**Table N.1 — Taint summary**

| Metric | Value |
|---|---|
| Bounded | false |
| Total flows | 0 |

## N.2 Top Taint Flows

0 flows. The taint pass was not bounded, meaning the analysis ran but did not establish an exhaustive bound on the taint-flow surface. The buyer's security team should treat the zero-flow result as a partial coverage signal rather than as a confirmed absence of exploitable flows.

© MTD Services

# Appendix Q — Glossary

Terms are used consistently and never redefined per section, so the buyer and its advisers converge on one shared frame. The glossary below is alphabetically sorted and defines, in plain language, every technical or specialist term appearing in any TDD table column heading. Definitions are consistent with contextual usage.

**Table Q.1 — Glossary (alphabetically sorted)**

| Term | Meaning |
|---|---|
| AGPL-3.0 | GNU Affero General Public License v3.0. A strong copyleft license that extends GPL obligations to network-served software, requiring source disclosure to users interacting with the software over a network. |
| Apache-2.0 | Apache License, Version 2.0. A permissive open-source license allowing use, modification, and distribution with attribution, including an express patent grant. |
| Assessment state | Five-value trust vocabulary (NOT_ASSESSED, PARTIALLY_ASSESSED, ASSESSED_NO_FINDINGS, ASSESSED_WITH_FINDINGS, ERROR) accompanying each pillar verdict; communicates the level of confidence the underlying evidence supports. |
| Assessment status | Pillar-level coverage label: assessed, partial, or not_assessed. Only assessed pillars support conclusions; partial / not_assessed pillars are demoted to informational and labelled UNVERIFIED. |
| Attack surface | Set of code paths reachable from untrusted input. Candidate endpoints are detected by static analysis and risk-scored for review prioritisation, not as confirmed vulnerabilities. |
| Attribution | License clause requiring preservation of copyright, patent, trademark, and attribution notices in derivative works distributed by the licensee. |
| Author | Git contributor identity derived from commit metadata. Used in key-person and retention analyses as a contributor-activity signal, not as an employment-status or assignability verdict. |
| Blast radius at risk | Estimate of the number of modules a given contributor's departure would destabilise, derived from sole-maintainership and tenure signals. |
| Bus factor | Reciprocal of sole-contributor concentration per module. A low bus factor flags flight risk: fewer engineers must depart before the module becomes unmaintainable. |
| Bus factor proxy | Aggregate team-level bus-factor signal derived from author-concentration ratios; complements the per-module bus factor. |
| Cadence per week | Average number of commits per week observed across the analysed history window. A sustained-cadence indicator, not a productivity metric. |
| Candidate endpoint | Function or symbol identified by static analysis as a potential untrusted-input entry point. Status is a detection signal; reachability is resolved separately. |
| Candidate exposure score | Upper-bound signal for review prioritisation across candidate endpoints. Not a confirmed vulnerability rate. |
| CC percentiles | Cyclomatic complexity percentiles (typically P50, P90, P99) computed across all analysed symbols; describe the distribution of structural complexity. |
| Certificate block | Detected textual block matching the shape of an X.509 certificate or PEM-encoded key material; treated as a secret candidate requiring validation. |

| Term | Meaning |
|---|---|
| CI Maturity | Pillar assessing the presence and strictness of continuous-integration pipelines inferred from committed workflow files. Static, not a runtime pass-rate measure. |
| Classification | Risk-tier label assigned in the consolidated register (Critical Risk, Major Risk, Minor Risk, Acceptable). Drives the conditions in the recommendation. |
| Cognitive complexity | Software-quality metric estimating the human effort required to understand a function, weighting nesting, recursion, and breaks in linear flow more heavily than cyclomatic complexity. |
| Cognitive percentiles | Distribution statistics (P50, P90, P99) of cognitive complexity across analysed symbols. |
| Commit | Atomic change recorded in the Git history. Used as the unit of activity in key-person and team-cadence analyses. |
| Condition precedent | Closing condition that must be satisfied or waived before signing or before the deal closes. Failure to satisfy blocks the transaction as structured. |
| Confidence score | Composite figure computed as 20 + 40 x coverage + 40 x risk_signal, capped at 20 when coverage is zero, and floored at 60 when the verdict is DO_NOT_ACQUIRE. Communicates the epistemic strength of the recommendation. |
| Confirmed | Evidence-state label reserved for findings validated by independent corroboration. None of the findings in this report carry this label; absence is named rather than implied. |
| Contamination chain | Linked set of packages whose licenses propagate obligations to dependent code, sized by transitive reachability from a viral or restrictive source. |
| Coverage method | Description of the coverage proxy used. In this scan, a static symbol-share proxy weighting test-file symbols by structural value; explicitly not runtime or line / branch coverage. |
| Coverage ratio | Effective proportion of production-symbol value covered by test-file symbols. Used to size uninsured exposure. |
| Cross-layer violation | Static dependency that crosses an architectural layer boundary (e.g. presentation calling data). Counted and severity-rated; not a runtime behaviour. |
| CVE | Common Vulnerabilities and Exposures identifier. Advisory records were refreshed against OSV at scan time; severity is reported as recorded by the advisory source. |
| Days since last commit | Elapsed days between the latest observed commit for a contributor and the scan date. Used as an activity-decay signal in retention analysis. |
| Deal-breaker | Pillar finding severe enough to force a DO_NOT_ACQUIRE verdict under the decision model, irrespective of remediation economics. |
| Debt density | Technical-debt markers per thousand lines of code; normalises marker counts against codebase size. |
| Debt marker | Explicit in-source annotation indicating known technical debt (e.g. FIXME, TODO, HACK, XXX). Tracked by type, severity, and critical-path placement. |
| Dependency | Third-party package declared in a manifest. Direct dependencies are explicit; transitive dependencies are resolved via lockfile. |
| Direct dependency | Package declared in a manifest. Distinguished from transitive dependencies pulled in by the resolver. |
| Doc Staleness | Pillar measuring the proportion of inventoried documentation files whose last-modified date exceeds the staleness threshold. |

| Term | Meaning |
|---|---|
| Duplicated LOC | Lines of code participating in duplicated blocks across the analysed codebase, summed across all duplication groups. |
| Economic escalation | Decision-model mechanism linking the remediation-cost share of proposed price to escalation tiers (conditions at 5% + $250K floor; no-go at 25%). |
| EOL | End-of-life. Technology, runtime, or dependency that has passed its last security or feature update. Reconciled across tech_stack (offline markers) and the eol pillar (endoflife.date runtime verification). |
| Entanglement score | Aggregate measure of cyclic dependency and cross-layer violation density. Higher scores indicate greater integration friction. |
| Evidence state | Per-finding honesty label attached to every material finding: Confirmed, Detected-not-verified, Inferred, Partially assessed, Unassessed, Conflicting, or Unsupported. |
| Finding id | Stable identifier of the form F-<pillar>-<rank> used by the narrative, register, workbook, and ledger to cite one assessed finding. |
| Framework | Software framework detected by manifest or import patterns (e.g. Huggingface for AI/ML, gin for HTTP routing). |
| Framework group | Count of distinct framework families detected in the AI/ML scan; communicates breadth of AI/ML integration, not depth. |
| Git history | Commit log scanned for secret candidates and contributor activity. Path-excluded candidates are named rather than silently ignored. |
| GPL v3 | GNU General Public License v3. A copyleft license requiring derivative works to be distributed under the same license; viral in the contamination sense. |
| GraphQL | Query-language API specification. Presence or absence detected by static manifest inspection; absence does not imply absence of GraphQL at runtime. |
| Headline EV | Transaction enterprise value declared in the scan input; the reference against which remediation cost is measured as a percentage. |
| Health score | Per-pillar 0-100 score. Demoted to 50 for partial / not_assessed pillars so an unfinished scan never reads as a clean bill of health. |
| High confidence secret | Secret candidate whose detected pattern exceeds the high-confidence threshold; treated as a critical-severity finding requiring validation, not a confirmed exposure. |
| High regression-risk module | Module identified by the test-deficit pillar as carrying elevated regression risk under refactoring; costed for test investment in the remediation model. |
| High-value gap | Coverage gap concentrated in a module whose structural risk weight exceeds the high-value threshold; prioritised in the test-investment work stream. |
| HPU | Hours per unit. Remediation effort per finding within a cost category, saturated by a cap and decay rate under the financial model. |
| IaC | Infrastructure-as-Code. Committed configuration describing infrastructure (e.g. Dockerfiles, Terraform, GitHub Actions). Findings are committed-text signals, not runtime enforcement. |
| Indexed files | Count of files parsed and loaded into the code-property graph backing the pillar scanners. Lower than total file count indicates parse failures or path exclusions. |
| Inferred | Evidence-state label for findings derived by inference from primary signals (e.g. retention risk from contributor activity patterns). Requires validation before negotiation reliance. |

| Term | Meaning |
|---|---|
| Integration drag | Post-merge operational friction estimated from remediation hours; expressed as a percentage capped at 80% by the financial model. |
| Integration risk score | Composite pillar-level score (0-10) reflecting proprietary coupling, schema drift, database coupling, and serialization risk. 10 indicates maximum integration friction. |
| IP Ownership | Pillar assessing the attribution of file authorship between corporate-domain and external contributors. Attribution ratios are a diligence signal, not legal proof of title. |
| IP risk score | Composite figure combining external-author concentration, external-blamed-line percentage, and external-owned-file count; ranges 0-1, higher is riskier. |
| Knowledge loss vulnerability | Per-contributor risk derived from tenure, recent activity, and sole-maintainership; classifies retention risk into low, medium, critical tiers. |
| Knowledge transfer | Cost category funding the documentation and pairing required to mitigate sole-maintainership in critical modules; sized per unique module, not per at-risk contributor. |
| License | Pillar assessing per-package detected licenses against viral, restrictive, and unknown classifications; flags contamination exposure for legal review. |
| Lockfile | Resolver-generated file pinning transitive dependency versions (e.g. go.sum, yarn.lock). Presence indicates deterministic builds; absence flags reproducibility risk. |
| Maintenance backlog score | Aggregate figure combining neglected-file count, average staleness, and complexity distribution. Zero in this scan because the staleness scanner produced no rows. |
| Manifest | Author-declared dependency specification (e.g. package.json, go.mod). Distinct from a lockfile, which records resolved versions. |
| Max cycle length | Module count of the longest detected cyclic dependency. Long cycles indicate deep architectural entanglement and high decoupling effort. |
| Migration files | Schema-change artifacts indicating database migration discipline. Count communicates migration history depth; absence is a hygiene signal. |
| Modules in cycles | Count of modules participating in at least one detected cyclic dependency. Communicates the breadth of entanglement, not its depth. |
| MPL-2.0 | Mozilla Public License v2.0. A weak copyleft license applying file-level copyleft to MPL-covered files while permitting combination with proprietary code. |
| No-go threshold | Remediation cost share of proposed price above which the decision model forces a DO_NOT_ACQUIRE verdict. Set to 25% in this scan. |
| OpenAPI | OpenAPI Specification for REST APIs. Presence detected by manifest inspection; absence in this scan flags schema-drift exposure. |
| ORM | Object-Relational Mapping library abstracting SQL access. Absence in this scan is consistent with the high raw-SQL count observed. |
| Path tier | Classification of the file path containing a secret candidate (e.g. prod, test, dev, vendor). Drives risk-cap assignment under the secrets pillar. |
| Person-years | Remediation hours divided by 2,080 (annual working hours). Communicates effort in staffing terms for the post-close plan. |
| Pin ratio | Pinned dependencies divided by direct dependencies. A low ratio flags reproducibility and supply-chain risk; not a contractual commitment. |

| Term | Meaning |
|---|---|
| PII endpoint | Endpoint handling personally identifiable information. Counted by privacy scanning of detected endpoint signatures; zero is not a compliance verdict. |
| Privacy finding | Detection of plaintext-credential material, sensitive logging, or PII handling in committed configuration. Flag-and-inventory basis, not a compliance verdict. |
| Probability after | Estimated probability of a risk materialising after remediation. Stated only when the Layer-2 risk valuation line is quantified; otherwise marked not quantified. |
| Probability before | Estimated probability of a risk materialising before remediation. Stated only when the Layer-2 risk valuation line is quantified. |
| Production model file | AI/ML model artifact shipped to production. Counted by AI/ML scan; zero is favourable in this scan. |
| Proprietary coupling | Static dependency on a Grafana-proprietary package. Higher counts flag integration friction in a separation or carve-out scenario. |
| Protocol layer | Network-protocol-level risk indicator (e.g. raw_socket usage for WebSocket handling). Counted for security review, not as a runtime vulnerability. |
| Raw SQL | SQL string embedded in application code outside an ORM abstraction. Higher counts indicate coupling to a specific SQL dialect and migration risk. |
| Remediation class | Category applied to a CVE record describing the available remediation path (e.g. no_fixed_release). Drives scheduling of the post-close remediation plan. |
| Remediation effort | Time required to remediate a finding or a category of findings. SQALE-derived for complexity; HPU-derived for other categories. |
| Residual expected loss | Expected loss surviving remediation, computed as probability_after x economic_impact_after for quantified risk lines. |
| Retention risk | Per-contributor classification (low / medium / critical) of the risk that the contributor departs before close and during the transition window. |
| Risk level | Pillar-level rating: critical, high, medium, low, or informational. Demoted to informational under partial coverage; derived_risk_level preserves the scanner's own view. |
| Risk reduction value | Aggregate expected loss before remediation minus aggregate residual expected loss; the value the remediation programme is intended to deliver. |
| Risk signal | Weighted blend of pillar risk into the confidence score. Combined with coverage to produce the 0-100 confidence figure. |
| Risk valuation | Layer-2 financial translation of a finding into probability x economic impact, producing expected-loss figures before and after remediation. |
| Saturation cap | Maximum count of findings within a cost category to which full HPU applies; above the cap, the decay HPU is used to prevent runaway cost. |
| Saturation decay | Hours-per-unit applied to findings above the saturation cap for a given category; reflects diminishing remediation effort per repeated finding. |
| Schema drift risk | Likelihood that API or storage schema has diverged from a documented contract. Rated high in this scan; reflects absence of OpenAPI and presence of multiple schema versions. |
| Secret candidate | Pattern match indicating possible secret material. Requires human validation before reliance; high-confidence and history-only candidates are tracked separately. |

| Term | Meaning |
|---|---|
| Sensitive log | Log statement detected as handling sensitive data. Zero in this scan is a detection result, not a confirmation of absence. |
| Serialization risk | Static pattern indicating unsafe deserialization (e.g. dynamic_function call, eval_call). High severity by default; treated as a security review item. |
| Severity | Per-finding rating: informational, low, medium, high, or critical. Weighted in the risk signal by 0, 1, 2, 3, 4 respectively. |
| Severity risk weight | Numeric weight applied to a finding's severity when computing the risk signal. Informational 0, low 1, medium 2, high 3, critical 4. |
| Sole-maintainer module | Module with bus factor of one. Critical sole-maintainer modules drive the knowledge-transfer cost category and the retention condition. |
| SQALE | Software Quality Assessment based on Lifecycle Expectations. Standardised remediation-effort basis producing per-module hours and a letter grade (A best, E worst). |
| SQALE grade | Letter grade summarising the SQALE debt ratio. Grade E in this scan signals severe debt relative to codebase size. |
| Stale document | Documentation file whose last-modified date exceeds the staleness threshold. Counted by doc_staleness scan; ratio communicates documentation drift. |
| Static symbol-share proxy | Coverage estimation method counting test-file symbols weighted by structural value, in lieu of runtime line or branch coverage. Communicates structural coverage, not executed coverage. |
| Taint flow | Static data-flow path from an untrusted source to a sensitive sink. Zero flows in this scan; the pass was not bounded and may not be exhaustive. |
| Tech Stack | Pillar counting detected technologies and their lifecycle status. Structural-detection basis, reconciled with runtime EOL verification. |
| Technology VaR | Value-at-Risk figure for technology exposure, derived from scenario weighting. Honestly labelled Scenario-Based Technology VaR when produced; a technology-risk model output, not a market-risk measure. |
| Tenure days | Days between a contributor's first and latest observed commit. A signal of institutional knowledge depth, not of employment tenure. |
| Test coverage ratio | Per-module ratio of test-file symbol value to production-symbol value. Zero in many high-risk modules is the leading regression-risk indicator. |
| Test deficit score | Aggregate figure combining module count below threshold and overall test ratio. Higher scores flag greater regression risk under refactoring. |
| Test ratio | Per-module ratio of test files to source files. A coarse structural signal, not a runtime coverage measure. |
| Top author share | Largest single contributor's share of commits. Higher values indicate greater concentration risk. |
| Transitive resolved package | Indirect dependency resolved via a lockfile. Count communicates supply-chain breadth and transitive-advisory exposure. |
| Unlicensed model file | AI/ML model artifact without a detected license. Zero in this scan is favourable; absence of bundled model files is the underlying cause. |
| Unpinned dependency | Manifest-declared dependency using a version range rather than an exact version. Flags reproducibility and supply-chain risk. |

| Term | Meaning |
|---|---|
| Unreferenced endpoint | Candidate endpoint with no detected inbound edges in the static graph. Communicates potential dead code or undocumented exposure. |
| Untested candidate endpoint | Candidate endpoint for which no test file is detected. Higher counts flag review backlog in the attack-surface work stream. |
| Unverified pillar | Pillar whose conclusions are demoted to UNVERIFIED due to partial or absent coverage. Named explicitly in the executive brief. |
| Value coverage ratio | Aggregate value coverage across all modules; communicates the proportion of high-blast symbols insured by tests. |
| Viral license | License that propagates its obligations to derivative works (e.g. GPL, AGPL). Triggers contamination-chain analysis and warranty risk in a closed-source acquirer scenario. |
| Warranty | Seller statement whose breach is an indemnity question rather than a work plan. Distinct from a condition precedent or a price-adjustment lever. |
| Weighted expectation | Probability-weighted expected loss across enabled scenarios. Honest label Scenario-Based Technology VaR when produced. |
| WebSocket | Communication protocol enabling full-duplex channels over a single TCP connection. Detected via raw_socket pattern in this scan; flags protocol-layer review items. |

*Where a term is defined differently in different contexts (e.g. confidence has a technical meaning in the financial model and a colloquial meaning in narrative), the technical meaning is given; colloquial usage is consistent with the technical meaning.*

© MTD Se