

Valepath

Technology Due Diligence Report

Illustrative Sample

Target: Lexmarq Ltd (fictitious)

Sector: Vertical SaaS — Legal Technology (workflow + contract AI)

Headquarters: London, UK

Revenue (LTM ARR): £22.0M · Growth: 42% YoY

Employees: 95 (48 in technology)

Prepared for: Thames Partners III (fictitious, £300M Fund III)

Date: April 2026

Assessment period: 12 working days

ILLUSTRATIVE SAMPLE — Prepared to demonstrate the Valepath deliverable format. Target company, buyer, findings, and financial figures are entirely fabricated.

1. Executive Summary

IC-ready · skim layer (this page) + supporting detail (next)

RECOMMENDATION: GO WITH CONDITIONS

Proceed with (a) **price adjustment –£3.0M from ask** reflecting security remediation cost and ARR-at-risk; (b) **£1.5M escrow over 18 months** tied to SOC 2 Type II and AI provenance audit; (c) **enhanced IP and cybersecurity reps & warranties** (see §14); (d) **founder-CTO retention** with knowledge-transfer milestones. Full deal-structuring package in §14.

INVESTMENT THESIS IMPACT

Lexmarq's technology is a **net enabler** of the platform thesis but carries **material security and IP-governance gaps** warranting deal-structure protection. The proprietary LLM-assisted contract review module is a credible moat (estimated 18–24 month lead over second-tier competitors) and supports both the pricing-power and bolt-on hypotheses. Absence of SOC 2 Type II, incomplete EU AI Act readiness, and undocumented model-training data provenance create concentrated execution and regulatory risks in the first 12 months post-close.

COMPOSITE RISK SCORE (VS PEER MEDIAN)

Architecture & Infrastructure		4/5	+0.5
Team & Organization		3/5	–0.2
Security & Compliance		2/5	–1.3
Data & Analytics		4/5	+0.4
AI & Proprietary Technology		3/5	+0.1
Cost Structure		3/5	–0.3
Delivery Capability		4/5	+0.6
Strategic & Bolt-on Readiness		3/5	–0.1
Composite		3.25/5	–0.1

Scoring: 1 Critical · 2 Significant · 3 Moderate · 4 Good · 5 Excellent. Rubric sources: industry-standard maturity frameworks (DORA, NIST CSF, OWASP). Peer comparisons indicative.

TOP 3 RISKS — QUANTIFIED

1. Security & compliance remediation

SOC 2 Type II + GDPR DSR automation + pentest backlog. **£320K · 0–6 months. ARR at contractual risk: £3.1M (14%).**

2. AI IP contamination risk

Undocumented training-data provenance on core LLM fine-tune. **Audit + remediation £180K · 0–6 months. Downside scenario: £1.1M / 9 months.**

3. Multi-tenancy refactor for bolt-on

Hard-coded customer boundaries prevent efficient integration of acquired targets. **£680K · 3–12 months. Blocks £8M of bolt-on synergies if deferred.**

COST-TO-FIX TOTAL

£2.10M total over 24 months — ~2.3% of implied EV; within thesis remediation envelope. Phased detail in §13.

1. Executive Summary — Detail

Supporting detail to the recommendation on the previous page.

DEAL-BREAKERS & RED FLAGS

#	ISSUE	SEVERITY	STRUCTURE IMPLICATION
1	AI model training data provenance undocumented — potential IP contamination from scraped third-party legal corpora	RED	IP rep + escrow + 90-day audit
2	No SOC 2 Type II; three DACH customers (14% ARR) have contractual right to terminate if not achieved by Q4 2026	RED	Reprice + escrow milestone
3	Key-person concentration: Head of AI (founder-CTO) holds 70% of undocumented model knowledge	AMBER	Retention + knowledge-transfer plan
4	Multi-tenant architecture has hard-coded customer boundaries — impedes bolt-on integration thesis	AMBER	6-month refactor, costed

No absolute deal-breakers identified. All issues tractable with structure and remediation capital.

TOP 3 VALUE CREATION OPPORTUNITIES — QUANTIFIED

1. AI-premium pricing tier Proprietary contract review module supports tiered pricing. +12% ARR over 18 months (~£2.6M) with no material incremental COGS.
2. Cloud & license rationalization AWS right-sizing + SaaS audit. £480K annual run-rate savings , achievable in 0–9 months.
3. DACH expansion acceleration Product already localized; sales motion and data residency remaining. +£4.5M ARR over 24 months if DACH team doubled.

COST-TO-FIX — PHASED

TIMEFRAME	INVESTMENT	RATIONALE
0–3 months (quick wins)	£320K	SOC 2 readiness, pentest, AI provenance audit, DSR automation
3–12 months	£1,200K	Multi-tenant refactor, EU AI Act compliance, cloud cost program, hiring buffer
12–24 months	£580K	Observability maturity, billing service rewrite, DR/BCP hardening
Total	£2,100K	~2.3% of implied EV; within thesis remediation envelope

Initiative-level breakdown in §13.

2. First 100 Days — Operational Plan

WEEK-BY-WEEK ACTION PLAN

WINDOW	WORKSTREAM	ACTION	OWNER	EXPECTED OUTCOME
Days 1–14	Security	SOC 2 Type II readiness assessment; schedule pentest	CTO + incoming Head of Security	Gap list, remediation plan, auditor engaged
Days 1–14	AI IP	Commission AI model provenance audit (training data, licensing, contamination)	External counsel + CTO	Audit engaged; data-room of training sources assembled
Days 1–21	Retention	Founder–CTO retention agreement; AI team equity refresh; successor-track Head of AI designated	Chair + CHRO	Signed retention; 9-month knowledge-transfer plan
Days 14–30	Cost	AWS cost review (RI commitment, unused resources, right-sizing)	VP Engineering	£120K+ annualized savings identified
Days 14–30	Cost	SaaS license audit across 47 tools; consolidate identity + collab	COO	£100K+ annualized savings identified
Days 30–60	Compliance	DSR automation deployed; GDPR Art. 15/17/20 SLAs met	DPO + Platform team	GDPR risk closed; deletion + portability automated
Days 30–60	Customer	Proactive outreach to 3 DACH accounts (14% ARR) with SOC 2 roadmap	CRO	Retention risk neutralized
Days 45–75	AI governance	AI model cards, training data registry, EU AI Act classification documented	Head of AI + DPO	EU AI Act baseline; IP rep-backed documentation
Days 60–90	Architecture	Multi-tenancy refactor design complete; PoC for tenant-boundary abstraction	Principal Architect	Engineering plan for H2; bolt-on cost re-estimated
Days 60–90	Strategic	Bolt-on target pipeline validated (3 candidates, tech-compatibility screens)	Corp Dev + Valepath	Tech-vetted pipeline for H2 LOIs
Days 90–100	Governance	Technology steering committee established; quarterly tech review rhythm	Chair + CTO	Operating cadence live; KPI dashboard

100-DAY KPIS

KPI	BASELINE (CLOSE)	100-DAY TARGET	12-MONTH TARGET
SOC 2 Type II readiness (% controls)	~40%	75%	Report issued
AI training-data provenance documented (% of corpus)	<20%	100%	Auto-lineage maintained
AWS monthly spend (£)	94K	82K	68K
Pentest high/critical findings open	12 (est.)	<3	0

KPI	BASELINE (CLOSE)	100-DAY TARGET	12-MONTH TARGET
DSR SLA compliance rate	60%	95%	99%
Key-person risk (roles with <1 successor)	4	2	0

GOVERNANCE CADENCE

- **Weekly** (first 90 days): CTO + deal team operating review
- **Bi-weekly**: Security remediation standup
- **Monthly**: Technology steering committee (Chair, CEO, CTO, Valepath advisor, deal partner)
- **Quarterly** (post-100 days): Technology KPI review at Board

3. Company Technology Profile

FIELD	DETAIL
Company	Lexmarq Ltd (fictitious)
Sector	Vertical SaaS — Legal Technology (workflow + contract AI)
Headquarters	London, UK
Revenue (LTM ARR)	£22.0M
Growth (YoY)	42% (FY25); 35% guided (FY26)
Employees	95 total (48 engineering, product, data, AI)
Primary customers	Mid-market UK + DACH law firms (340 firms, avg £65K ACV)
Primary product	Matter management + AI contract review platform
Cloud provider	AWS (eu-west-2 primary, eu-central-1 secondary)
Annual technology spend	£4.8M (22% of revenue)

TECHNOLOGY STACK

LAYER	SYSTEMS	NOTES
Frontend	React 18, TypeScript, Next.js 14	Modern; consistent
Backend	Python 3.11 (FastAPI), TypeScript (Node services)	Mixed but deliberate
Database	PostgreSQL 15, Redis 7, OpenSearch 2	Single-writer RDS cluster
Infrastructure	AWS (EKS, RDS, S3, Lambda, SageMaker)	Cloud-native; Terraform IaC
CI/CD	GitHub Actions, ArgoCD	GitOps
Monitoring	Datadog, Sentry	Tracing coverage 65%
Security	AWS GuardDuty, Snyk, 1Password, Okta SSO	Baseline; SIEM absent
AI / ML	Fine-tuned open-weight 70B-class LLM, frontier API (secondary fallback), SageMaker	Proprietary primary model

4. Architecture & Infrastructure

Score: 4/5 (Good) — vs peer median +0.5. AWS-native, well-factored; modest tech debt concentrated in legacy billing service.

CURRENT STATE

Lexmarq runs a modular-monolith Python backend with extracted microservices for AI inference, document ingestion, and audit logging. AWS-native with EKS, Terraform-managed infrastructure, GitOps deployment. Multi-region active-passive (15-minute RPO to eu-central-1). Scales to current load (2,400 req/s peak) with headroom to 3.5x.

TECHNICAL DEBT INVENTORY

ITEM	SEVERITY	REMEDIATION COST	TIMELINE
Legacy billing service (Rails 6, monolith-internal)	High	£240K	9 months
Hard-coded multi-tenancy boundaries (see §11)	High	£680K	12 months
Ingestion pipeline throughput ceiling (~50 docs/s)	Medium	£120K	6 months
Observability gaps (tracing coverage 65%)	Medium	£90K	4 months
Terraform modules lack versioning	Low	£30K	2 months
Total quantified technical debt		£1,160K	

SCALABILITY ASSESSMENT

Current capacity supports **3.5x current load** on the core platform. Constraint is the AI inference layer: proprietary 70B-class LLM on 4xg5.12xlarge caps at ~180 contract-review req/min. Scaling to 10x requires inference optimization (quantization, speculative decoding) — est. £200K engineering + £180K/year additional infra.

SINGLE POINTS OF FAILURE

SPOF	IMPACT	MITIGATION STATUS
RDS single-writer (Postgres primary)	Full write-path outage; 15-min RPO, 4-min RTO	Acceptable for stage; plan Aurora FY27
Fine-tuned model weights in single S3 bucket, no versioned backup	Model rollback impossible on bad fine-tune	Fix in 0–3 months (£15K)
Single payment gateway (Stripe)	Revenue collection outage	Acceptable; Stripe reliability strong

5. Team & Organization

Score: 3/5 (Moderate) — vs peer median -0.2. Strong core (60% senior+), but key-person risk on AI lead; hiring plan optimistic for London.

ENGINEERING TEAM COMPOSITION

ROLE	HEADCOUNT	IN-HOUSE	CONTRACTORS	NOTES
Frontend	8	8	0	
Backend	14	13	1	
Platform / DevOps	5	5	0	
Data / Analytics	4	4	0	
AI / ML	6	5	1	Research consultant 2d/wk
QA / SDET	3	3	0	
Engineering management	4	4	0	
Principal / Staff	4	4	0	
Total technology org	48	46	2	51% of total FTE

KEY-PERSON RISK

PERSON	ROLE	RISK	MITIGATION
Founder-CTO	Head of AI; 70% of model IP knowledge	High	Retention + 9-month knowledge transfer to successor-track hire
Principal Platform Engineer	Terraform + EKS keyholder; sole prod deployer	Medium	Cross-training underway; 2 backups partially trained
VP Engineering	Org architect; 4y tenure	Medium	Standard retention
Lead Data Engineer	Owns curated legal corpus pipeline	Medium	Documentation gap; 3-month sprint planned

HIRING PLAN ASSESSMENT

Plans to hire 14 engineers in FY26 (29% team growth). **Optimistic** given London market for senior AI/ML talent (time-to-hire median 4.2 months for staff+ AI roles). Realistic achievement: 9–11 hires. Buffer recommended.

ORGANIZATION HEALTH SIGNALS

- Glassdoor: 4.3 (14 reviews; small n)
- GitHub commit cadence: consistent, no post-Series-A anomalies
- Voluntary turnover (FY25): 9% (peer median: 12%)
- Tenure distribution: median 2.1 years

6. Security & Compliance Posture

Score: 2/5 (Significant) — vs peer median -1.3. **MATERIAL GAP** — primary risk dimension.

COMPLIANCE STATUS

STANDARD	STATUS	GAP	REMEDIATION COST
SOC 2 Type II	Not achieved	Readiness ~40%; no auditor engaged	£120K + 6–9 months
GDPR	Partial	DSR automation incomplete; DPIA gaps on AI processing	£80K + 3 months
NIS2	Not in scope (below threshold)	Monitor; likely in-scope if DACH expands	£30K baseline if triggered
ISO 27001	Not pursued	Defer to SOC 2	—
EU AI Act	Incomplete	Contract-review product likely “limited risk”; docs absent	£140K + 6 months

VULNERABILITY ASSESSMENT

Last pentest: **August 2025 (8 months overdue)**. Last findings: 3 High, 7 Medium, 14 Low. All High remediated per internal records; Medium/Low backlog not verified. Fresh pentest commissioned in Day 1–14.

INCIDENT HISTORY

Last 12 months: 4 customer-reported incidents, MTTR 42 minutes. One P1 in Nov 2025 (Redis cluster failover failure, 38 min impact). Post-mortem quality adequate. No security incidents detected by GuardDuty.

IAM & ACCESS CONTROL

- **Authentication:** Okta SSO; MFA enforced (98% coverage; 2 exceptions flagged)
- **Privileged access:** AWS via SSO + role assumption; 14 users with production access
- **PAM:** No dedicated tooling; break-glass via 1Password, audited quarterly
- **Third-party access:** 3 contractors with limited read access; no offboarding gaps in spot check
- **Dormant accounts:** 4 identified in spot check (disabled but not deleted)

7. Data & Analytics Maturity

Score: 4/5 (Good) — vs peer median +0.4.

DATA QUALITY ASSESSMENT

Dimension	Score	Notes
Completeness	4/5	<2% missing in production tables
Accuracy	4/5	API-enforced validation; audit trails complete
Timeliness	5/5	Real-time; <1s freshness hot paths
Consistency	3/5	Two terminology drifts (contract vs. matter); cleanup scheduled

INTEGRATION HEALTH

Source	Target	Method	Latency	Reliability
Customer iManage	Lexmarq ingestion	API + webhook	<30s	99.6%
Customer Relativity	Lexmarq ingestion	Partner connector	<60s	99.4%
Lexmarq → Salesforce	Customer CRM	Outbound webhook	<5s	99.9%

DATA GOVERNANCE

- **Ownership:** Data Product Manager (hired Q3 2025) owns catalog
- **Lineage:** Partial (40% of critical flows in Atlan)
- **Retention:** 7-year default; customer-controlled deletion via DSR flow (automation incomplete — see §6)
- **Privacy:** EU residency enforced for DACH; UK in eu-west-2

8. AI & Proprietary Technology

Score: 3/5 (Moderate) — vs peer median +0.1. Genuine moat; IP documentation thin; EU AI Act exposure unquantified.

AI ROADMAP & MATURITY

Lexmarq operates a **fine-tuned open-weight 70B-class LLM** for contract clause extraction and risk flagging, hosted on-instance (SageMaker). Internal benchmark: **F1 = 0.87** on clause classification (peer benchmark: 0.78). Second-generation fine-tune in progress; target F1 = 0.92 by Q3 2026. Primary model on-instance; frontier API as secondary fallback. Additional frontier API integration planned Q3 2026.

AI leverage in product: contract review = 38% of product usage minutes; cited by 71% of new customers as primary purchase driver.

INTELLECTUAL PROPERTY PORTFOLIO

IP ASSET	TYPE	PROTECTION	VALUE
Fine-tuned contract-review model weights	Trade secret	NDA + access controls	Core moat
Curated legal clause corpus (820K annotated clauses)	Trade secret + copyright	NDA + access controls	Supporting moat
Matter management workflow UI	Patent applications	2 filed UK/EPO; 1 provisional US	Defensive
Lexmarq wordmark	Trademark	Registered UK, EU, US	Standard

AI IP & TRAINING DATA GOVERNANCE — GAP

Training data provenance is undocumented

Initial fine-tune combined customer data (consent-gated), public-domain legal corpora, and — per informal engineering notes — at least one scraped commercial legal database with unclear licensing. Historical data pipeline does not retain source attribution.

- **Risk:** potential IP contamination; downstream liability if scraped corpus challenged
- **Mitigation:** 90-day AI provenance audit (external counsel + forensics); if contamination found, retrain from clean corpus (£1.1M, 9 months)

AI REGULATORY EXPOSURE (EU AI ACT)

- **Likely classification:** Limited risk (contract review is assistive, not determinative)
- **Required documentation:** model card, training data registry, use-case disclosure, human-oversight affirmations
- **Current status:** None documented
- **Remediation:** £140K + 6 months; £80K recurring annual compliance cost

COMPETITIVE MOAT DURABILITY

Moat is real but narrowing. Estimated **18–24 month window** before second-tier competitors reach parity via open-source fine-tunes. Value creation priority: capitalize via pricing tier (see §1) and data-flywheel acceleration.

9. Cost Structure

Score: 3/5 (Moderate) — vs peer median -0.3. Tech spend 22% of revenue vs peer 18–20%; overspend in cloud + SaaS.

CURRENT TECHNOLOGY SPEND

CATEGORY	ANNUAL (£)	% OF TOTAL	BENCHMARK	GAP
Cloud infrastructure (AWS)	1,130K	24%	18%	+22%
SaaS licenses	840K	18%	15%	+20%
Salaries (tech team)	2,420K	50%	55%	-9%
Contractors / outsourcing	180K	4%	5%	-20%
Hardware / devices	70K	1%	2%	-50%
Other	160K	3%	5%	-40%
Total	4,800K	100%		

OPTIMIZATION POTENTIAL

AREA	CURRENT	OPTIMIZED	SAVINGS	EFFORT
Cloud right-sizing (RDS, EKS, GPU inference)	£1,130K	£910K	£220K	Medium / 60d
SaaS license rationalization (47 → 30 tools)	£840K	£500K	£340K	Low / 30d
RI / Savings Plans commitment	£1,130K	£1,020K	£110K	Low / 14d
Total annualized savings			£670K	

Realistic 12-month captured savings: ~**£480K** (gradual ramp; growth offsets some gains).

BENCHMARK COMPARISON

AWS spend per-customer: £3,320/year vs peer median £2,400/year. Difference driven by AI inference (g5.12xlarge fleet, structurally necessary); remaining ~£220K overspend is addressable via right-sizing.

10. Delivery Capability

Score: 4/5 (Good) — vs peer median +0.6. DORA elite on 3 of 4 metrics; AI engineering leverage above median.

DORA METRICS

Metric	Current	Elite Benchmark	Rating
Deployment frequency	~23 deploys/day (prod)	Multiple per day	Elite
Lead time for changes	2.4 hours (median)	<1 day	Elite
Mean time to restore	38 minutes	<1 hour	Elite
Change failure rate	14.2%	0–15%	High (just within elite band)

AI Engineering Leverage

Metric	Current	Peer Median	Rating
Engineering seats with AI copilot licenses	44/48 (92%)	65%	Above median
% of PRs with AI-assisted code (self-reported)	68%	~45%	Above median
AI-assisted test generation	Adopted (40% new tests)	Emerging	Above median
AI code review in CI	Partial (Greptile, critical paths)	Low adoption	Above median
Internal LLM tooling for ops/runbooks	None	None (norm)	Parity

Development Practices

- Version control: ✓ (GitHub, all code)
- Code review: ✓ (2-reviewer required on main, enforced)
- Automated testing: partial (72% coverage backend, 54% frontend; integration tests thin)
- CI/CD pipeline: ✓ (GitHub Actions + ArgoCD)
- Infrastructure as code: ✓ (Terraform; modules unversioned — debt)
- Monitoring & alerting: partial (Datadog + Sentry; tracing coverage 65%)
- On-call rotation: ✓ (6-engineer rotation, weekly handoff)
- Post-mortem process: ✓ (blameless; Linear-tracked actions)

Backlog Health

- Age distribution: 68% <90d, 22% 90–365d, 10% >365d (peer norm: 15% >365d)
- Tech debt ratio: 18% of engineering time (healthy)
- Sprint predictability: 82% commitment hit rate

11. Strategic Alignment & Bolt-on Integration Readiness

Score: 3/5 (Moderate) — vs peer median -0.1.

TECHNOLOGY ROADMAP VS BUSINESS PLAN

Published FY26–FY28 plan: (a) DACH expansion +40% ARR, (b) 2–3 bolt-ons in 18 months, (c) AI premium tier launch FY26 Q4. Technology supports (a) and (c) well; **(b) is blocked by multi-tenancy architecture until refactored.**

BOLT-ON INTEGRATION READINESS

DIMENSION	ASSESSMENT	IMPLICATION
API surface	Strong; REST + webhooks, well-documented, versioned	Enables loose integration
Multi-tenancy model	Hard-coded tenant boundaries in 6 core services	Blocks efficient absorption
Data model	Clean domain model; legal entities reusable	Absorbable
Identity federation	Okta-based; federate acquired orgs <30 days	Ready
Billing & contracts	Legacy Rails billing; customer model tightly coupled	Refactor required before first bolt-on
Org / Conway	Two acquired-target scenarios modelled; team no prior M&A integration	Advisory support recommended

BOLT-ON INTEGRATION COST MODEL

INTEGRATION PATH	COST PER BOLT-ON	TIMELINE	NOTES
Current architecture (hard-coded)	£900K–1.4M	9–12 months	Unacceptable for thesis
Post-refactor (tenant-boundary abstraction)	£280K–420K	3–5 months	Supports thesis cadence

Multi-tenancy refactor (£680K, 3–12 months) is **gating investment** for bolt-on thesis. Pays back after bolt-on #1.

COMPETITIVE TECHNOLOGY MOAT

- **Differentiator:** fine-tuned contract AI (18–24 month lead, narrowing)
- **Commodity:** matter management workflow (feature parity with 5+ competitors)
- **Liability:** none identified

M&A EXIT READINESS (ACQUIRER PERSPECTIVE)

Strategic acquirer (Thomson Reuters, LexisNexis, listed LegalTech) could integrate Lexmarq in 12–18 months post-close. Primary attractors: AI capability + DACH footprint. Primary deductions: SOC 2 gap (immediately visible in diligence), AI training-data provenance question.

12. Risk Scorecard

Architecture		4/5	vs peer +0.5	
Team		3/5	vs peer -0.2	
Security		2/5	vs peer -1.3	← MATERIAL
Data		4/5	vs peer +0.4	
AI & Proprietary		3/5	vs peer +0.1	
Cost		3/5	vs peer -0.3	
Delivery		4/5	vs peer +0.6	
Strategic & Bolt-on		3/5	vs peer -0.1	
Composite		3.25/5	vs peer -0.1	

Strong technology foundation (Architecture, Data, Delivery above peer median), with one **material gap** (Security) driving the deal-structure recommendations in §14. AI & Proprietary score is deceptively moderate — the technology itself is above median, the governance is below; composite reflects the blend.

13. Cost-to-Fix — Phased Remediation

PHASED REMEDIATION PLAN

#	INITIATIVE	DIMENSION	PRIORITY	COST	TIMELINE
1	SOC 2 Type II readiness + audit	Security	P1	£120K	0–9 mo
2	AI provenance audit + remediation	AI/IP	P1	£180K (base)	0–6 mo
3	Pentest + remediation backlog	Security	P1	£60K	0–3 mo
4	DSR automation + GDPR close-out	Security	P1	£80K	0–3 mo
5	EU AI Act documentation	AI/IP	P2	£140K	3–9 mo
6	Multi-tenancy refactor	Architecture	P2	£680K	3–12 mo
7	Cloud cost optimization program	Cost	P2	£60K (invest)	0–9 mo
8	SaaS license rationalization	Cost	P2	£20K (invest)	0–3 mo
9	Billing service rewrite	Architecture	P3	£240K	12–18 mo
10	Observability maturity	Delivery	P3	£90K	12–15 mo
11	DR/BCP hardening	Architecture	P3	£250K	12–24 mo
12	Key-person risk mitigation (hiring + KT)	Team	P1–P2	£180K	0–12 mo

SUMMARY

PHASE	INITIATIVES	TOTAL COST	TIMELINE	COMBINED ROI
P1: Quick wins	1(p), 3, 4, 8, 12(p)	£320K	0–3 mo	Protects £3.1M ARR + initiates £200K+ savings
P2: Foundation	1(c), 2, 5, 6, 7, 12(c)	£1,200K	3–12 mo	Unlocks bolt-on thesis (~£8M synergy capacity)
P3: Optimization	9, 10, 11	£580K	12–24 mo	Platform-grade operating profile
Total	12 initiatives	£2,100K	0–24 mo	~£11M+ value unlocked or protected

14. Deal Structuring Implications

PRICE ADJUSTMENT

Recommended **reduction from ask: £3.0M.**

COMPONENT	AMOUNT	RATIONALE
Security remediation (certain)	£320K	0–3 month P1 initiatives
AI IP remediation (expected-value)	£500K	Base £180K + 30% × £1.1M tail risk
ARR-at-risk (14% × 18mo × haircut)	£1,800K	3 DACH accounts with SOC 2 termination rights
EU AI Act future compliance	£220K	Certain near-term cost, not priced into ask
Multi-tenancy refactor (partial)	£160K	25% weighting; remainder is growth investment
Total price adjustment	£3,000K	~3.3% of typical mid-market SaaS EV

ESCROW / HOLDBACK

Recommended **£1.5M in escrow, 18 months**, milestone-released:

MILESTONE	RELEASED	DEADLINE
Pentest High/Critical findings closed	£250K	Month 3
AI provenance audit clean (or remediated)	£500K	Month 6
SOC 2 Type II report issued	£500K	Month 12
EU AI Act docs signed off by external counsel	£250K	Month 15

REPS & WARRANTIES (ENHANCED)

- IP — AI training data provenance:** seller represents data was lawfully obtained; surviving indemnity (24 months, capped at 100% of purchase price for third-party IP claims)
- Cybersecurity & data protection:** absence of unreported security incidents; GDPR compliance; carve-in for pentest backlog disclosure
- AI regulatory compliance:** no pending or threatened regulatory actions (EU AI Act, GDPR Art. 22, UK ICO)
- Data residency & customer contractual rights:** full disclosure of customer termination rights triggered by security gaps

EARNOUT DESIGN

Recommended **earnout cap £4.0M / 12 months**, tied to 100-day plan KPIs for post-close alignment:

- £2.0M — FY26 ARR growth ≥ 30% (YoY, constant currency)
- £1.5M — DACH segment ARR ≥ £4.0M by end FY26
- £0.5M — AI premium tier launch with ≥ 50 converted customers

MANAGEMENT RETENTION

- Founder-CTO:** 3-year lockup, equity refresh vesting 4 years, change-of-control acceleration, tied to 9-month KT milestones
- VP Engineering:** standard retention (2-year, equity refresh)
- AI team (6 engineers):** pooled retention bonus £480K vesting 24 months

SUMMARY DEAL PACKAGE

RECOMMENDED DEAL STRUCTURE

Price adjustment **–£3.0M** · Escrow **£1.5M / 18mo**, milestone-released · Enhanced reps (4) · Earnout cap **£4.0M / 12mo** KPI-tied · Founder-CTO 3y lockup + AI team retention pool

Appendix — Methodology, Sources, Scope-Out

ASSESSMENT METHODOLOGY — VALEPATH 8-DIMENSION FRAMEWORK

- 1. **Architecture & Infrastructure** — topology, cloud maturity, technical debt, scalability, SPOFs
- 2. **Team & Organization** — composition, key-person risk, hiring plan, culture signals
- 3. **Security & Compliance** — vulnerability posture, certifications, IAM, incident history
- 4. **Data & Analytics** — quality, governance, integration, privacy
- 5. **AI & Proprietary Technology** — models, IP portfolio, training-data governance, regulatory exposure
- 6. **Cost Structure** — spend, optimization, benchmarking, cost-to-fix
- 7. **Delivery Capability** — DORA, AI engineering leverage, backlog, practices
- 8. **Strategic Alignment & Bolt-on Readiness** — roadmap fit, integration readiness, exit attractiveness

SCORING SCALE (WITH PEER-BENCHMARK ANCHORS)

SCORE	LABEL	DEFINITION	PEER ANCHOR
1	Critical	Emergency intervention; material risk to operations or thesis	Bottom decile
2	Significant	Major gap; 3–6 month remediation	Bottom quartile
3	Moderate	Functional but suboptimal; 6–12 month roadmap	Median
4	Good	Minor gaps; optimization opportunities	Top quartile
5	Excellent	Best-in-class for size/stage	Top decile

Scoring rubric: industry-standard maturity frameworks (DORA, NIST CSF, OWASP, EU AI Act readiness criteria). Peer comparisons drawn from Valepath working benchmark and publicly reported PE/CDD market data; methodology disclosed on request.

AI-AUGMENTED WORKFLOW DISCLOSURE

Valepath engagements use AI-assisted tooling for:

- Document ingestion and structured-signal extraction from data-room materials
- Source code static analysis (volume + pattern detection; human-validated findings)
- Financial model anomaly detection (cost data, vendor spend patterns)
- Cross-reference automation (flagging contradictions across interviews, documents, data)

Valepath does **not** delegate to AI:

- Scoring and judgment — every dimension score set by human principal
- Client interviews — conducted and interpreted by human principal
- Final risk ratings and deal-structuring recommendations — human judgment
- IP and legal interpretation — external counsel

Tooling stack disclosed on request under NDA.

DATA SOURCES (THIS ENGAGEMENT)

- Source code repository review (read-only, 23 repos)
- AWS console access (read-only, all accounts)
- CI/CD pipeline review (GitHub Actions + ArgoCD)
- Monitoring dashboards (Datadog, Sentry, GuardDuty)
- Architecture documentation (12 internal docs)

- Org chart and HRIS data
- Financial data (technology spend, 24-month history)
- Penetration test reports (last: August 2025)
- Incident logs (Linear + PagerDuty, 12 months)
- Product roadmap and backlog
- Stakeholder interviews (11 sessions): CEO, CTO, VP Engineering, CRO, DPO, CFO, Head of AI, Head of Product, Principal Architect, Head of Customer Success, External Counsel

WHAT WE DID NOT ASSESS (SCOPE-OUT)

This engagement did **not** include:

- Independent verification of AI model training data sources (recommended: separate forensic audit — see §14)
- Full source code security audit (sampled; SAST/DAST outputs reviewed)
- Customer contract legal review (counsel workstream)
- Financial audit or quality-of-earnings (finance DD)
- Commercial or market DD (commercial DD provider)
- HR/people DD beyond technology org
- Physical security or facilities assessment
- Independent validation of regulatory certifications beyond document review

DISCLAIMER

This report is prepared solely for Thames Partners III in connection with the proposed acquisition of Lexmarq Ltd. Findings and recommendations represent Valepath's professional judgment as of the report date. This report should not be relied upon by any other party without Valepath's prior written consent. Valepath is not a law firm and does not provide legal advice; items flagged for legal review should be directed to qualified counsel.

ILLUSTRATIVE SAMPLE — Prepared to demonstrate the Valepath deliverable format. Target company, buyer, findings, scores, and financial figures are entirely fabricated. No actual due diligence was conducted.