Anyway Ventures

PRIVATE EQUITY · DEAL TEAMS · OPERATING PARTNERS · 2026 EDITION

AI disruption risk *scorecard*.

One scoring model for how exposed a business is to AI, built to work on a single target in exclusivity and across a portfolio of 30 holdings. Eight factors with a 0 to 3 rubric, explicit weights, a revenue-at-risk method, capacity to respond, and rules that turn the score into defend, adapt, accelerate or exit early.

PDF GUIDE | EXCEL SCORECARD

AUTHOR

Brannan Coady, Anyway Ventures

EDITION

Version 1.0 · October 2026

ONLINE, WITH DOWNLOADS

anywayventures.com/resources/ai-disruption-risk-scorecard

CONTACT

brannan@anywayventures.com

Free to use, adapt and share with attribution · Not legal, tax or investment advice

CONTENTS

In this guide

1. Executive summary
2. How to use this scorecard
3. How this model relates to the frameworks it replaces
4. The scorecard at a glance
5. The scoring rubric
6. Weights
7. Revenue at risk
8. Capacity to respond and time to exit
9. Classification rules
10. Sector heat map by business model
11. Where the factors come from
12. How to run a portfolio scan
13. The data request
14. The management session question set
15. The one-page output per company
16. What to tell boards and LPs
17. Worked examples
18. Common mistakes
19. Glossary
20. Sources and further reading
21. About this resource

About Anyway Ventures

WHO IT IS FOR

PE deal teams, operating partners, investment committees, portfolio CEOs and CFOs, investor relations

ALSO AVAILABLE

PDF guide, Excel scorecard at anywayventures.com/resources/ai-disruption-risk-scorecard

KEY TAKEAWAYS

The short version

1. The AI disruption risk scorecard scores eight exposure factors from 0 (protected) to 3 (exposed) and weights them into an exposure index from 0 to 100, for one target or a whole portfolio.

2. In AI disruption scoring, three factors carry half the default weight: revenue on tasks AI does acceptably today, pricing per hour, seat or FTE, and margin earned from labour arbitrage.

3. Revenue at risk from AI is revenue times exposed share times compression over the horizon times the share passed to customers, with flow-through converting it to EBITDA at risk.

4. Portfolio companies are classed defend, adapt, accelerate or exit early by rule, with exit early for an AI exposure index of 71 or more plus low capacity or under two years to exit.

5. In 2026, EY found around 60% of GPs had increased diligence on AI disruption risk, and Livingstone recorded BPO multiples resetting from around 12 to 13 times EBITDA to around 5 to 7 times.

This is a complete scoring model for AI disruption risk: how much of what a business sells AI can now do, what protects it, how much revenue is at risk over the hold, and what the fund should do about it. It is written for private equity deal teams scoring a single target, operating partners running a portfolio AI risk review across 10 to 30 holdings, and the investment committees and boards who have to act on the answer. The Excel companion does the arithmetic: factor dropdowns, weights, a revenue-at-risk calculator, automatic classification and a portfolio roll-up.

1. Executive summary

1. **One model, two depths.** The scorecard joins the eight-factor grid from our AI disruption risk assessment, the seven-factor portfolio triage and the replicability test into a single scale. Eight exposure factors, each scored 0 (protected) to 3 (exposed), are weighted into an exposure index from 0 to 100. A portfolio scan scores them from a data request and one interview. A deal scores the same factors from a data room and a replicability test.
2. **Three factors carry most of the threat.** Revenue on tasks AI does acceptably today, pricing per hour, seat or FTE, and margin earned from labour arbitrage carry half the default weight. Data, regulation and switching costs are the moat. AI-native competition says whether the threat is already in the pipeline.
3. **A single 3 can matter more than the total.** A 3 on replicability, pricing or AI-native competition raises an alert that blocks a defend classification, whatever the index says.
4. **The number that goes in the model is revenue at risk.** For each revenue line: revenue × exposed share × compression over the horizon × share passed to customers. Then apply flow-through to get EBITDA at risk. In the illustrative compliance consultancy below it is about 11% of revenue but over 40% of EBITDA, because the revenue falls at close to full margin if the team stays the same size.
5. **Exposure is not the decision.** Capacity to respond (AI delivery track record, funding headroom, board backing, change capacity) and time to exit sit beside the index. The same exposure is an adapt case with a funded team and four years left, and an exit-early case with a stretched balance sheet and 18 months.
6. **Four classes, applied by rule.** Exit early when the index is 71 or more and capacity is low or exit is under two years away. Accelerate when the company owns data, distribution or the automation itself and can act on it. Defend when the index is 35 or less, the moat is strong and no alert is raised. Everything else is adapt. A review that puts every company in adapt has been softened somewhere.
7. **The market is already pricing this.** In 2026 EY found around 60% of GPs had increased diligence on AI disruption risk; Livingstone records BPO multiples resetting from around 12 to 13 times EBITDA to around 5 to 7 times; Bain reports software marks in PE portfolios down roughly 8%. The exit buyer will run some version of this test, so run it first.
8. **The output is one page per company.** Index, band, alert, revenue at risk, class, two sentences of reasoning and three actions with owners and dates. If an investment committee member can't read it in three minutes, it's too long.

~60%

of GPs increased diligence on AI disruption risk. EY PE Pulse, Q1 2026

5 to 7x

EBITDA for BPO, down from around 12 to 13 times. Livingstone, 2026

~8%

fall in software valuations inside PE portfolios, Q1 marks. Bain, 2026

12 to 14

software categories Apollo ranks by AI susceptibility. Bloomberg, 2026

2. How to use this scorecard

Who it is for. Deal teams use it on a target in exclusivity. Operating partners use it across the portfolio to decide where to spend attention and capital. Portfolio CEOs and CFOs use it a year or more before a sale to see what a buyer will see. Investor relations use the portfolio output, and the method behind it, for LP reporting.

When to use it.

MOMENT	DEPTH	TYPICAL TIME
Screening a new opportunity	Desk score from the IM and public material, factors 1, 2, 7 and 8 only	Two to four hours
Exclusivity on a target	Full scorecard with replicability test, revenue at risk and a pricing view	Three to five days as a focused strand
Portfolio review	Full scorecard per company from a standard data request and one management session	Four to six weeks for 10 to 30 holdings
Six-monthly refresh	Re-score factors 1, 2 and 7, update revenue at risk and actions	Half a day per company
Exit preparation	Full re-score, evidence pack built to buyer standard	One to two weeks

Inputs. Revenue by service line and by billing basis, gross margin by line, headcount by function, retention cohorts, top 20 customers, contract terms, the data inventory, competitor names, AI tool spend and a sample of real deliverables. The full list is in the data request.

The Excel companion. The workbook has a blank **Scorecard** tab for one company (factor dropdowns, editable weights, weighted index, band, alert, capacity, revenue-at-risk calculator and classification), a **Portfolio** tab for up to 30 companies with a summary by class, a **Rubric** tab with every scoring definition, tabs for the **Data request** and the **Management session**, and filled-in examples for the three fictional companies in the worked examples. Every score and classification is a live formula: change a dropdown and the class changes.

Who should score. One small team, internal or external, applying the same standard to every company. When each deal team scores its own holdings, every company comes back resilient.

3. How this model relates to the frameworks it replaces

Three earlier pieces on this site set out parts of the method. This scorecard keeps their factors, thresholds and classes, puts them on one scale and adds explicit weights. Nothing in them is overturned.

EARLIER FRAMEWORK	WHAT IT CONTRIBUTES	HOW IT APPEARS HERE
AI disruption risk assessment: eight factors scored 1 to 5, where 5 is resilient	The deal-level evidence list, the revenue-at-risk calculation, the price and structure table, and the rule that upside is scored separately	Its seven threat factors map onto factors 1, 2, 4, 5, 6 and 8 and onto capacity to respond. Its factor 8 (cost base upside) stays separate as the upside score
Portfolio triage: seven factors scored 0 to 3, where 3 is exposed	The 0 to 3 scale, labour arbitrage and AI-native competition as factors, capacity to respond, time to exit, and the four classes	Factors 1 to 7 are its seven factors, unchanged. Its 0 to 21 total is still shown as the "triage score"

EARLIER FRAMEWORK	WHAT IT CONTRIBUTES	HOW IT APPEARS HERE
Replicability test: a timeboxed AI rebuild of the core product, or AI drafting of past deliverables	The evidence behind factor 1, and the distinction between the cost of the features and the cost of the customers	It sets the score for factor 1 and informs factors 4 and 6

Mapping the 1 to 5 deal grid onto the 0 to 3 scale

The deal grid scores resilience; the scorecard scores exposure. To convert a deal-grid score: exposure = (5 − resilience) × 0.75. So 5 becomes 0, 4 becomes 0.75, 3 becomes 1.5, 2 becomes 2.25 and 1 becomes 3. Round to the nearest whole number when you enter it in the scorecard.

DEAL GRID FACTOR (1 TO 5)	SCORECARD FACTOR (0 TO 3)
1 Replicability of the core service	1 Revenue on replicable tasks
2 Revenue from hours billed or seats	2 Pricing basis
3 Data and proprietary workflow	4 Proprietary data and workflow
4 Switching costs and embeddedness	6 Switching costs
5 Regulatory and trust barriers	5 Regulation and accountability
6 Distribution and relationship ownership	8 Distribution and demand capture
7 Management speed to adopt	Capacity to respond, sub-factor C1
8 Cost base upside (separate)	Upside score U (separate)

The deal grid's management factor moves out of the exposure score because management speed doesn't change how exposed the business is. It changes what can be done about it, which is the job of capacity to respond. That was already the logic of the triage.

Why the thresholds line up

The band boundaries were set so that each earlier threshold lands on one of them. With the factors scored evenly:

- A triage score of 7 or less (the triage's defend range) converts to an index of 33 or less, inside the **Low** band (0 to 35).
- A deal-grid threat total below 20 (the assessment's "high threat", where you price on post-compression revenue) converts to an index of roughly 54 or more, the start of the **High** band.
- A triage score of 15 or more (the triage's exit-early threshold) converts to an index of 71 or more, the **Severe** band.

Uneven score profiles and non-equal weights move individual companies a few points either way. That is the point of weighting, and the Excel workbook shows the triage score beside the index so you can see where they differ.

4. The scorecard at a glance

COMPONENT	WHAT IT MEASURES	SCALE	USED FOR
Exposure factors 1 to 8	How much of what customers pay for AI can do, and what protects it	0 to 3 each, 3 is most exposed	Exposure index
Weights	How much each factor matters for this type of business	Sum to 100	Exposure index
Exposure index	Weighted exposure	0 to 100	Band and classification
Single-factor alert	A 3 on factor 1, 2 or 7	Yes or none	Blocks defend
Moat strength	Average of factors 4, 5 and 6	0 to 3	Defend test
Revenue at risk	Annual revenue lost or repriced by the end of the horizon	£ and % of revenue	Price, plan, marks
Upside score U	Volume of repetitive work not yet automated	0 to 3, 3 is largest upside	100-day plan, kept out of the index
Capacity to respond	Ability to change pricing and delivery	Four sub-factors, 0 to 12, then high, medium or low	Classification
Time to exit	Years left in the hold	Years	Classification
Owns the shift	Whether the company owns data, distribution or the automation that AI makes more valuable	Yes or no, with evidence	Accelerate test

The index. Exposure index = Σ (weight × score ÷ 3), with weights summing to 100, rounded to a whole number. A company scoring 3 on every factor scores 100; one scoring 0 everywhere scores 0.

The bands.

BAND	INDEX	PLAIN MEANING
Low	0 to 35	Core value sits behind something AI doesn't easily remove. Watch the back office and sales channels
Moderate	36 to 53	Real exposure in parts of the business, usually pricing or one service line. A plan, not a crisis
High	54 to 70	A material share of revenue is on work AI is already doing. Price it, structure for it, fund the change
Severe	71 to 100	The business model itself is being repriced. The question is whether this owner can change it in time

5. The scoring rubric

Score every factor from evidence. Where the evidence is thin, record the score with low confidence and say what would change it. The Excel workbook has a confidence column for this.

Factor 1: revenue on replicable tasks

The question. What share of revenue comes from work that current AI tools do to an acceptable standard today, with light human review?

How to score it. For services, take ten anonymised past deliverables and the briefs that produced them. Have a competent operator run the briefs through current models with sensible prompting, then have a sector practitioner grade the output blind against the originals. For software, run the replicability test on the core workflow. Then map the result onto revenue: what share of billed effort, or of the product's value to customers, sits on work that passed?

SCORE	WHAT IT LOOKS LIKE	EXAMPLES
0	Under 10% of revenue on replicable work. The core service is physical, on-site or judgement work carrying personal liability	Fire and security maintenance; environmental testing with lab analysis; residential care
1	10% to 25%. Some reporting, scheduling or documentation that AI handles, around a core that it doesn't	Building surveying (the report drafting is replicable, the inspection isn't); specialist engineering consultancy
2	25% to 50%. A substantial part of the work product is drafting, analysis, code or routine support that passed the blind test	Compliance consultancy; market research; vertical software whose core workflow was rebuilt in a few days
3	Over 50%. The output customers pay for is mostly a document, summary, translation, ticket resolution or code that AI produces acceptably	Translation and transcription; tier-one contact handling; content and SEO agencies; offshore development by the day

Traps. Management will say "our work needs judgement". Some does. The blind test settles it in a day. And don't score the whole firm on its best work: score the revenue mix, including the routine work that pays the bills.

Factor 2: pricing basis

The question. What share of revenue is priced per hour, day, FTE or seat, so that it shrinks when AI reduces the effort or the customer's headcount?

How to score it. Split the sales ledger by how revenue is actually billed, not how management describes it. "Value-based" fixed fees built up from an hours estimate and renegotiated each year count as hours.

SCORE	WHAT IT LOOKS LIKE	EXAMPLES
0	Under 10% priced on effort or seats. Priced per outcome, asset, visit, transaction value or multi-year contract	Statutory inspection contracts priced per site; software priced per claim settled or per property managed
1	10% to 40%	A managed services firm with some project days; SaaS with a platform fee plus a minority of seat revenue

SCORE	WHAT IT LOOKS LIKE	EXAMPLES
2	40% to 70%	An engineering consultancy with a mix of fixed-fee and time-and-materials work; SaaS mostly priced per user
3	Over 70% priced per hour, day, FTE or seat	A contact centre billed per agent; an IT services firm billing day rates; horizontal SaaS priced per knowledge-worker seat

Traps. Seat pricing looks safe because seats renew. It isn't safe if the customer's headcount in that role is what AI reduces. Apollo's chief economist noted in July 2026 that investors appear to be pricing lower demand for seat-based licences into application and vertical software.

Factor 3: labour-arbitrage model

The question. Does margin come from the product, asset or intellectual property, or from the gap between what staff cost and what they're billed at?

How to score it. Look at gross margin by line and at delivery headcount as a share of revenue. A business whose gross profit grows only when it adds billable people is a labour-arbitrage model. Offshore delivery sold on rate difference is the purest form.

SCORE	WHAT IT LOOKS LIKE	EXAMPLES
0	Margin comes from product, asset or IP. Delivery staff cost under 20% of revenue	Software with low services attach; asset-owning testing labs
1	Mixed. People deliver, but pricing reflects a product, method or asset	Managed services on owned tooling; field services with route density as the margin driver
2	Mostly people-driven margin, with some productised or recurring elements	Professional services with a retainer book; digital agencies with some productised offers
3	Margin is the spread between labour cost and bill rate. Delivery staff cost over 60% of revenue	Staff augmentation; BPO priced per FTE; offshore development centres

Why it is separate from factor 2. A firm can price on fixed fees and still depend on labour arbitrage, in which case AI is margin upside for it until competitors pass the saving on. A firm pricing per seat may have no labour arbitrage at all. The two factors move independently, and Livingstone's 2026 market update is explicit that the market is repricing "labor-arbitrage, FTE-billing business models, not the sector as a whole".

Factor 4: proprietary data and workflow

The question. Does the business own data and workflow that a competitor using the same AI models can't reach, and do its contracts let it use them?

How to score it. Ask for the data inventory and the contract clauses on data use. Check the rights, not the volume: years of client data the business may not reuse count for nothing.

SCORE	WHAT IT LOOKS LIKE	EXAMPLES
0	Owns years of structured data, generated or aggregated by its own operations, with clear rights to use it. A workflow customers run their process inside	A claims platform with 12 years of settled claims and rights to aggregated use; a testing lab's historical results database

SCORE	WHAT IT LOOKS LIKE	EXAMPLES
1	Some proprietary data or method, partially structured or with limited rights	A consultancy with a benchmarking database it built; SaaS whose data is customer-owned but whose configuration is deep
2	Methods and templates, but the work runs on the client's data with generic tools	Most professional services; agencies with a playbook
3	No proprietary data. Works on clients' data in clients' systems with generic tools	Staff augmentation; contact centres on the client's CRM; translation

Traps. "We have a lot of data" is not a moat if it's unstructured, unlabelled or contractually the client's. Also look at the other direction: Bain's 2025 SaaS analysis lists proprietary data depth among the indicators that make a workflow harder for AI agents to capture.

Factor 5: regulation and accountability

The question. Does someone accountable, licensed or insured have to put their name to the output?

How to score it. Ask who signs. List licences, accreditations, professional indemnity terms and regulatory obligations, and check whether they attach to the firm's output or only to its customers.

SCORE	WHAT IT LOOKS LIKE	EXAMPLES
0	Licensed activity or a named, legally required sign-off	Statutory audit; gas safety and fire certification; regulated clinical services
1	Accredited or insured activity that customers or their regulators require	ISO-accredited testing; FCA-regulated outsourced complaint handling; software in an audited regulated process
2	Professional norms and reputation, but no requirement	Management consultancy; most compliance advisory work that the client signs, not the firm
3	Anyone can do it and no liability attaches	Content production; data entry; first-line support

Note. Regulation protects revenue more than price. If a named professional has to sign, AI shortens the work but rarely removes the firm. The fee may still fall. Apollo's Rob Bittencourt pointed to healthcare as more protected because regulation slows adoption, according to Bloomberg's June 2026 report.

Factor 6: switching costs

The question. How hard and how expensive is it for a customer to leave, and does their behaviour prove it?

How to score it. Use retention cohorts, churn reasons and integration depth, not contract length. A five-year contract with termination for convenience on six months' notice is a six-month contract.

SCORE	WHAT IT LOOKS LIKE	EXAMPLES
0	Integrated into customer operations, multi-year, painful to unpick. Gross revenue retention 95% or more for three years	A system of record integrated with the customer's policy, payroll or ERP; statutory maintenance contracts across an estate
1	Sticky but separable. Gross revenue retention 90% to 95%	Managed services with annual renewal and a long tenure; SaaS with moderate integration

SCORE	WHAT IT LOOKS LIKE	EXAMPLES
2	Annual contracts, periodic re-tender, some splitting of work across suppliers. Retention 80% to 90%	Professional services on panels; agencies on retainers
3	Project work, re-tendered each time, easy to split or take in-house. Retention below 80% or not measurable	Project-based development; one-off research; staffing

Traps. Watch customers with capable IT teams. The competitor that matters may be the customer's own developers, who now find an internal tool cheaper than the renewal. Ask which customers have tried to build their own.

Factor 7: AI-native competition

The question. Is the threat theoretical, or is it in the pipeline now?

How to score it. Scan the category for funded AI-native entrants, ask for lost-deal reasons for the last four quarters, and ask which customers have asked for price cuts citing AI. This is the factor that moves fastest, so refresh it every quarter for adapt and exit-early companies.

SCORE	WHAT IT LOOKS LIKE	EXAMPLES
0	No credible AI-native entrant in the category, no customer pressure citing AI	Most physical and site-based services
1	Entrants exist, but aren't winning in the company's segment yet. Occasional customer questions	Vertical software in a niche with few funded entrants
2	Funded entrants are winning deals, customers are piloting alternatives, or price pressure citing AI has appeared at renewal	Compliance and legal drafting; marketing content; tier-two support
3	The company is losing deals or renewals to AI-led alternatives this year, or customers are building in-house	Translation; tier-one contact handling; transcription; simple data and reporting tools

Weighting in conversation. A company scoring 3 here is losing business this quarter whatever the total says. That is why a 3 on factor 7 raises an alert.

Factor 8: distribution and demand capture

The question. Does the business own how customers find and choose it, or does it rent that from search, aggregators and price-led tenders?

How to score it. Ask how the last 20 customers arrived. Split new business by source, and check framework agreements, referral partnerships and channel arrangements.

SCORE	WHAT IT LOOKS LIKE	EXAMPLES
0	Owns relationships, frameworks or referral channels; most new work comes from existing customers or named partners	Panel positions with insurers; long-standing public framework agreements
1	Mostly relationship-led, with some competitive tendering	Specialist consultancies with referral networks
2	A large share won through open tenders on price, or through a few intermediaries	Facilities services on public tenders; agencies winning through pitch lists

SCORE	WHAT IT LOOKS LIKE	EXAMPLES
3	Wins through search, comparison sites, aggregators or marketplaces	Online legal and accounting services; lead-generation-led trades; app-store-dependent products

Why it matters. A business that wins through search or price comparison is exposed twice: AI cuts the cost of the work and changes how buyers find suppliers. There is not yet reliable data on how fast AI answer engines are shifting demand away from search in UK B2B markets, so score this from the company's own source-of-business data rather than from market claims.

The upside score (kept separate)

The question. How much repetitive, measurable work in the company's own delivery and back office has not yet been automated?

SCORE	WHAT IT LOOKS LIKE
0	Little repetitive work; already automated where it makes sense
1	Some repetitive work in one or two functions
2	Material repetitive volume in delivery or back office, partly sized
3	A large, measurable volume of repetitive work (scheduling, quoting, reporting, invoicing, first-line queries) not yet automated

Score it separately because mixing threat and upside lets a large cost-saving opportunity disguise a revenue line that is about to shrink. Size it by timing a sample of real work in two or three processes, using the method in measuring AI ROI. Upside feeds the 100-day plan and the AI value creation plan; it never reduces the exposure index.

6. Weights

The default weights reflect what the market is repricing hardest: the work itself, the way it is priced and where margin comes from.

#	FACTOR	DEFAULT WEIGHT	WHY
1	Revenue on replicable tasks	20	The direct measure of what AI can do instead
2	Pricing basis	15	Decides whether AI efficiency becomes lost revenue or retained margin
3	Labour-arbitrage model	15	The model the market is repricing most visibly in 2026
4	Proprietary data and workflow	10	The strongest moat, but rarer than claimed
5	Regulation and accountability	10	Protects revenue, less so price
6	Switching costs	10	Sets how fast exposure turns into churn
7	AI-native competition	12	Says whether the threat is live; also raises an alert at 3
8	Distribution and demand capture	8	Second-order for most B2B firms, first-order for search-led ones

#	FACTOR	DEFAULT WEIGHT	WHY
	Total	**100**	

Weight presets by business model

Change weights only for a reason you can write in one line, and change them for a whole class of business, never for one company. Otherwise weighting becomes a way to move a favoured company into a better band.

FACTOR	SERVICES AND OUTSOURCING	SOFTWARE AND SAAS	PHYSICAL AND LICENSED SERVICES
1 Revenue on replicable tasks	20	20	15
2 Pricing basis	15	15	10
3 Labour-arbitrage model	20	5	10
4 Proprietary data and workflow	10	15	10
5 Regulation and accountability	10	10	15
6 Switching costs	10	15	10
7 AI-native competition	10	15	10
8 Distribution and demand capture	5	5	20

The software preset shifts weight from labour arbitrage to data, switching costs and competition, which is where software exposure is decided: when the features can be rebuilt, what stops customers leaving is data, integration and trust. The physical preset shifts weight to regulation and distribution, because for these businesses the core service is rarely replicable and the exposure sits in how work is won and how much of it is reporting. Use the default for mixed portfolios so companies are comparable.

The single-factor alert

A 3 on factor 1, 2 or 7 raises an alert. An alerted company can't be classed as defend, whatever its index. The reason is simple: a business that scores well everywhere except pricing has a pricing problem, and one losing deals to AI-native rivals has a live problem, and in both cases the average hides it.

7. Revenue at risk

The score tells you how exposed the business is. Revenue at risk tells you how much that is worth, and it is the number that goes into the model, the price and the marks.

The formula

For each revenue line:

Revenue at risk = revenue × exposed share × compression × pass-through

TERM	DEFINITION	WHERE IT COMES FROM
Revenue	Current annual revenue of the line	Revenue split by service line and billing basis
Exposed share	The share of the line's revenue that the mechanism applies to	The factor 1 test, the factor 2 ledger split, contested segments from factor 7
Compression	How much of that exposed revenue's effort, seats or volume disappears by the end of the horizon	Productivity evidence, the replicability test, customer behaviour
Pass-through	The share of the reduction that customers capture through lower fees, fewer seats or lost volume	Pricing basis, contract terms, procurement behaviour

Then: **EBITDA at risk = total revenue at risk × flow-through**, where flow-through is the share of lost revenue that falls to EBITDA if the cost base doesn't change.

Three mechanisms, one formula

MECHANISM	EXPOSED SHARE IS	COMPRESSION IS	PASS-THROUGH IS
Effort compression (hours, days, FTEs)	Share of billed effort on work that passed the replicability test	Reduction in effort on that work over the horizon	Share of the saving customers take through fees or procurement
Seat compression	Share of seat revenue tied to roles at customers that AI reduces	Expected fall in those seats	Share of lost seat revenue not recovered by moving to usage or outcome pricing
Share loss	Share of revenue in segments where AI-native rivals or in-house builds are winning	Expected share of that revenue lost over the horizon	100%, because lost customers pay nothing

Setting the horizon

Use the longer of the remaining hold and three years. The exit buyer prices the next three years whatever the hold has left, so a company 18 months from sale still needs a three-year view.

Choosing compression assumptions

Don't take compression rates from vendor claims. The controlled evidence is narrower and more useful:

- In a randomised experiment published in Science, Noy and Zhang found ChatGPT cut the time professionals took on mid-level writing tasks by 40% and raised quality by 18% (2023).
- Brynjolfsson, Li and Raymond found an AI assistant raised customer support agents' issues resolved per hour by 14% on average, and by 34% for novice workers (NBER, 2023; published in the Quarterly Journal of Economics, 2025).
- Dell'Acqua and colleagues found BCG consultants using GPT-4 completed 12.2% more tasks, 25.1% faster, at higher quality, on tasks inside AI's capability frontier, and were 19 percentage points less likely to reach a correct answer on a task outside it (Harvard Business School working paper, 2023).
- METR's 2025 randomised trial found experienced open-source developers took 19% longer with AI tools, and its February 2026 update said they were probably sped up by then but described its own newer evidence on the size of the effect as very weak.

These studies used 2023 to 2025 tools, and capability has moved since. What they establish is that effort compression on suitable tasks is real and sizeable, uneven across tasks, and much smaller on work outside the frontier. A working range for effort compression on work that passed the factor 1 test, over a three-year horizon, is 30% to 60%. Use the low end where review and liability slow adoption, the high end where customers are already running the tools themselves. Then test the result against what the company's own timesheets and the replicability test show.

Choosing pass-through

Pass-through depends mostly on factor 2 and factor 6.

SITUATION	TYPICAL PASS-THROUGH
Time and materials or per-FTE pricing, sophisticated procurement	70% to 90%
Time and materials, relationship-led, some pricing power	50% to 70%
Fixed fee renewed annually against an hours estimate	40% to 60%
Genuine outcome or per-transaction pricing	0% to 30%
Lost customers (share loss)	100%

These ranges are judgement, not measured market data. Write down the reason for each choice and test it in the management session: ask which customers have already asked for AI-related price cuts and what was agreed.

Flow-through to EBITDA

MODEL	TYPICAL FLOW-THROUGH	WHY
Professional services keeping the same team	70% to 90%	Lost fees fall almost straight to EBITDA until headcount changes
BPO or contact centre priced per FTE	25% to 40%	Lost revenue takes the agent cost with it; the lost margin is the gross margin plus unabsorbed overhead

MODEL	TYPICAL FLOW-THROUGH	WHY
Software	80% to 95%	Little cost goes with lost seats

Worked arithmetic

A compliance consultancy bills £7.2m on time and materials. Timesheet sampling shows 40% of billed hours, £2.88m, go on drafting and review that passed the blind test. Assume AI halves the effort on that work over three years (compression 50%), and clients capture 60% of the saving (pass-through 60%). Revenue at risk is £2.88m × 50% × 60% = £0.86m. This is the same arithmetic as the compliance case in our AI disruption risk assessment, and the Excel calculator reproduces it.

Run a low and a high case by changing compression and pass-through, not by applying a percentage to the answer. The Excel tab gives a quick range with multipliers, but the scenario that goes to committee should state its assumptions.

> **Tip.** Revenue at risk is a gross figure before the company acts. Show it beside a mitigated figure that assumes the actions in the plan happen on time. The difference between the two is the value of the plan, and it is the number the board should hold management to.

8. Capacity to respond and time to exit

Exposure says how big the problem is. Capacity and time say whether this owner, with this team, can fix it before it shows up in the exit price.

Capacity to respond

Score four sub-factors from 0 (weak) to 3 (strong).

SUB-FACTOR	0	1	2	3
C1 AI delivery track record	No AI in delivery, no owner, no budget	Licences bought, scattered personal use, no measured results	One or two tools in live delivery with usage data	Someone in the team has shipped AI into operations, with measured before-and-after results and a named owner
C2 Funding headroom	No cash or covenant headroom; lender would resist new spend	Headroom only by cutting elsewhere	Can fund a change programme from cash flow	Cash, facility or shareholder commitment to fund pricing and delivery change for 12 to 18 months
C3 Board and shareholder backing	Board resists or hasn't discussed it	Discussed, no decision or budget	Decision taken, budget agreed, reported quarterly	A named executive owns the change with board-level targets on pricing mix and cost per unit
C4 Change capacity	Management fully absorbed by another crisis; data and systems not fit to automate	Bandwidth limited; systems fragmented	Bandwidth exists; systems workable with some clean-up	Spare leadership bandwidth, clean data, modern systems, a record of delivering change

Add them up: **9 to 12 is high, 5 to 8 medium, 0 to 4 low.** Score C1 from evidence: interview below board level and ask for 12 months of AI vendor invoices and usage data. The CEO will say the right things; a team lead showing you a tool they use every week is evidence.

Time to exit

Record the years left in the expected hold. Under two years is short: there is no time to change a business model and show the results in the numbers before a buyer looks. For deals, record the planned hold. Where the answer is under two years, the exit readiness checklist sets out what can still be fixed in the time.

Owns the shift

Record yes only with evidence that the company owns at least one of:

- proprietary, structured data with rights to use it, which becomes more valuable as AI makes it usable;
- the automation itself, sold to others or embedded in a platform customers pay for;
- a distribution channel an AI-native entrant would take years to build.

Livingstone describes the logic in its note on Capgemini paying 13.5 times EBITDA for WNS: buyers pay premium prices for platforms "that own the automation rather than compete with it".

9. Classification rules

Each company gets exactly one class. The rules are applied in order, and the first that fits wins.

ORDER	CLASS	RULE	WHAT THE FUND DOES
1	Exit early	Index 71 or more **and** (capacity low **or** under two years to exit)	Bring forward a sale while buyers still separate this asset from its sector, or restructure the plan and reset the mark
2	Accelerate	Owns the shift **and** capacity medium or high	Invest more: product, buy-and-build, earlier capex; consider whether this becomes the platform
3	Defend	Index 35 or less **and** moat strength 1 or less **and** no alert	Protect the moat, take the back-office upside, don't overpay for add-ons on an "AI-proof" story
4	Adapt	Everything else	Fund a specific change: reprice exposed lines, put AI into delivery, move seats to usage or outcome pricing

These are the triage classes and thresholds, written as rules. Three cases need a note.

Severe exposure with capacity and time is adapt, not exit early. The triage says adapt covers companies above 14 "with strong capacity and time". The Excel marks these "Adapt (severe)" so the board knows the plan is large and urgent.

High exposure with low capacity or under two years is an exit-early watch. The rules keep it in adapt, because the index hasn't reached the severe band, but the workbook flags it. Revisit it at the next refresh.

Exit early is a judgement, not an instruction. In some categories the market has already priced the risk, and selling a BPO business into a 5 to 7 times market simply crystallises the loss. The case is strongest where the asset still trades on old assumptions, often because it sits in a niche buyers don't yet associate with AI, and where the honest plan says the company can't change fast enough. Where that isn't true, move it to adapt with a smaller, cheaper plan and reset expectations.

What the band does to a deal

For a target rather than a holding, the band feeds the price and structure table from the AI disruption risk assessment.

BAND	PRICE	STRUCTURE	100-DAY PLAN
Low	Price on the business plan; no premium for being "AI-proof"	Standard terms	Capture the upside score
Moderate	Haircut growth in exposed lines by the revenue-at-risk figure	Earn-out or deferred consideration linked to pricing migration or retention	Reprice exposed lines, put AI into delivery
High or severe	Price on post-compression revenue, or walk	Larger deferred element; management rollover weighted to delivering the pricing change	A funded, sequenced plan to change the business model, owned by a named executive

Livingstone reports earn-outs in IT services increasingly keyed to services gross-profit retention, at 10% to 25% of consideration, with seller rollover of 15% to 35%. Those are reasonable reference points for structuring a high-band services deal.

10. Sector heat map by business model

The heat map is where a review starts, not where it ends. Every row contains companies in every band; the scorecard exists to tell them apart. Our sector guides for software and SaaS, professional services and field services and TIC apply it to those models. Typical bands below are judgement from applying the rubric to each model's usual profile, with the 2026 market evidence beside them.

BUSINESS MODEL	TYPICAL BAND	2026 MARKET EVIDENCE	WHAT SEPARATES THE RESILIENT ONES
BPO, contact centres, transaction processing	Severe	Livingstone: multiples reset from around 12 to 13 times EBITDA to around 5 to 7 times "as agentic AI erodes transaction-based revenue"	Owning the automation (Capgemini paid 13.5 times EBITDA for WNS); regulated processes; outcome pricing
Digital engineering, staff augmentation, offshore development	Severe	Livingstone: digital engineering median EV/EBITDA fell from above 23 times (December 2024) to around 7.4 times as AI coding tools undercut specialised capacity	Managed services contracts; domain depth; outcome-based pricing
Content, SEO, translation and creative production agencies	High to severe	Microsoft Research ranks interpreters and translators, writers and editors among the occupations most applicable to generative AI (2025)	Strategy and relationship ownership; performance fees; proprietary audience data
Seat-priced horizontal application software	High	Bain: near 30% fall in listed software valuations in February 2026, PE software marks down roughly 8%; Apollo: application software lagging security and infrastructure	System-of-record status; proprietary data; usage or outcome pricing
Document-heavy professional services (compliance, paralegal, bookkeeping, research)	High	DfE ranked management consultants, financial managers and accountants among the most AI-exposed occupations (2023); Skills England puts 70% of UK workers in occupations with AI-exposed tasks (2026)	Regulated sign-off; liability; fixed-fee pricing; long relationships
Vertical software with proprietary data	Moderate, wide range	Apollo's framework reportedly treats healthcare as more protected by regulation and data visualisation as highly exposed; Bain's SaaS work rates proprietary data and certification barriers as defences	Data rights; regulated workflow; integration with systems of record
Managed IT services, managed security, infrastructure software	Low to moderate	Livingstone lists contracted recurring revenue and managed security among premium attributes; Apollo notes security and infrastructure software outperforming	Multi-year contracts; owned tooling; certifications
Regulated professions with named sign-off (audit, surveying, actuarial, clinical)	Moderate	Exposure research puts much of the analytical work in high-exposure occupations, while licensing keeps the firm in the loop	Who signs; PI insurance; regulator relationships

BUSINESS MODEL	TYPICAL BAND	2026 MARKET EVIDENCE	WHAT SEPARATES THE RESILIENT ONES
Physical, licensed, on-site services (field services, testing and inspection, care, construction services)	Low	Bain: buyers showed "greater interest in companies with physical or labor-intensive components"; ONS: AI use at 13% in construction against 58% in information and communication (2026)	Already resilient; the exposure is in back-office reporting and how work is won
Businesses that win demand through search, comparison or aggregators	Moderate to high on factor 8	No reliable 2026 UK data yet on AI answer engines displacing B2B search; score from the company's own source-of-business data	Owned relationships, frameworks and referral channels

The heat map supports what investors are already doing. EY's Q1 2026 Private Equity Pulse found technology fell from roughly 30% of global PE deployment by value in 2025 to 12% in the first quarter of 2026, with nearly two-thirds of GPs taking a more targeted approach. Bloomberg reported in June 2026 that Apollo assesses every new software opportunity for AI disruption risk, using a framework dividing software into 12 to 14 categories ranked by susceptibility, and has overlaid it on existing holdings. The same report said Ares had hired an outside consultant to review software investments in its largest listed private credit fund.

11. Where the factors come from

The factors are not invented from first principles. They draw on the best public exposure research and investor frameworks, adjusted for a question those sources don't answer: what happens to one company's revenue and margin.

SOURCE	WHAT IT MEASURES	WHAT THIS SCORECARD TAKES FROM IT	WHAT IT CAN'T TELL YOU
Eloundou, Manning, Mishkin and Rock, "GPTs are GPTs" (Science, 2024; preprint 2023)	Task-level exposure of US occupations to large language models: around 80% of workers could have at least 10% of tasks affected, and around 19% at least half	Task-level thinking for factor 1: score the share of revenue on exposed tasks, not whether a job is "exposed"	Exposure is not displacement, and occupations are not companies
IMF Staff Discussion Note, Gen-AI: Artificial Intelligence and the Future of Work (2024)	Around 40% of global employment and around 60% in advanced economies exposed, split by whether AI is likely to complement or substitute	The complement or substitute split, which in a company shows up as pass-through: who keeps the saving	Country and occupation averages
ILO, Generative AI and Jobs: a refined global index (2025)	One in four workers globally in occupations with some generative AI exposure; 3.3% in the highest category; clerical work most exposed	The gradient view: exposure comes in degrees, which is why factors use 0 to 3 rather than yes or no	Company-level pricing and moats
Felten, Raj and Seamans, language modelling exposure (2023)	Occupational and industry exposure; legal services and securities among the most exposed industries	Industry rows in the heat map	Within-industry variation, which is most of the variation that matters

SOURCE	WHAT IT MEASURES	WHAT THIS SCORECARD TAKES FROM IT	WHAT IT CAN'T TELL YOU
Microsoft Research, Tomlinson and colleagues, AI applicability to occupations (2025)	Observed overlap between real Copilot use and occupational work activities	Evidence that exposure follows observed use, which is why factor 7 asks what customers are actually doing	Whether the overlap replaces or assists
McKinsey Global Institute, The economic potential of generative AI (2023)	Technical potential to automate activities absorbing 60% to 70% of employees' time	A ceiling, not a forecast: the reason compression assumptions stay well below technical potential	Adoption speed in a specific customer base
Bain, Will agentic AI disrupt SaaS? (2025)	Two axes for software workflows: automation potential and AI penetration potential, with indicators including proprietary data depth, switching friction and certification barriers	The split between threat (factors 1 to 3) and moat (factors 4 to 6)	Services businesses
J.P. Morgan Asset Management, Disruption decoded (2026)	A loan-level scoring framework for private credit; four channels from AI disruption to impairment: revenue erosion, margin compression, valuation compression and refinancing	The insistence on timelines relative to maturity, which here is time to exit and the horizon rule; the point that pricing model context is needed to read margins	Equity upside and accelerate cases
Apollo, as reported by Bloomberg (2026)	12 to 14 software categories ranked by susceptibility, applied to new deals and the existing book	The portfolio overlay: score the whole book on one basis, then decide where to reposition	The category list is not public

The gap these sources leave is the one this scorecard fills. Occupational studies measure jobs, not revenue lines. Consultancy frameworks are mostly written for software. Credit frameworks look at downside only. None of the public ones combine threat, moat, pricing, capacity and time into a decision with rules a committee can check.

12. How to run a portfolio scan

The process fits four to six weeks for 10 to 30 companies, with a light ask of each management team. It follows the portfolio triage and adds the scoring detail.

WEEK	STEP	OUTPUT	WHO
0	Agree scope, owner and weights. Send one standard data request to every company in the same format	Signed-off method; data requests out	Operating partner, scoring team
1 to 2	Desk score every company from the data. Hold a 60 to 90 minute management session with the CEO and whoever runs delivery or product	Draft scores with confidence ratings	Scoring team
2 to 3	Deep tests on the five to eight highest-index companies and every accelerate candidate: replicability test, competitor scan, four quarters of renewal and pricing data	Revised scores; revenue-at-risk ranges	Scoring team, sector practitioners
4	Calibration: every company on one page, scored side by side, with the operating team and deal partners	Consistent scores; draft classes	Operating partner, deal partners
4 to 5	Management challenge: share each draft with the CEO and CFO, allow a week to correct it with data	Final scores	Portfolio management teams
5 to 6	Classification, three actions per company, portfolio view and LP narrative	One-pagers; portfolio summary by class	Scoring team, IR

Running calibration

Calibration is where the scores become consistent, and it decides whether the review is worth anything.

- ☐ Put every company's eight scores on one screen, sorted by index, using the Portfolio tab.
- ☐ For each factor, read the evidence line aloud for the highest and lowest scorer. If the difference isn't obvious from the evidence, one of them is wrong.
- ☐ Challenge every defend: what is the share of revenue on hourly billing, and what did the replicability test show?
- ☐ Challenge every accelerate: what exactly does the company own, and do the contracts give it the rights?
- ☐ Check the distribution of classes. If 18 of 20 holdings are adapt, the scoring is being softened, usually by deal teams scoring their own companies.
- ☐ Record each change and the reason, so the next refresh starts from the same standard.

Refresh cadence

Re-score factors 1, 2 and 7 every six months, with a lighter quarterly check on adapt and exit-early companies focused on competitor activity and pricing at renewal. Do a full re-score when a company starts exit preparation, makes a large acquisition or changes its pricing model.

13. The data request

One request, the same for every company, answerable by a CFO in half a day. For a deal, the same list goes into the data room request, with the deliverable sample added. The AI due diligence questions cover the target's own AI claims alongside it.

#	ITEM	FORMAT	FEEDS
1	Revenue for the last two financial years and current year to date, split by service line or product	Spreadsheet	Factors 1 and 2; revenue at risk
2	The same revenue split by billing basis: hourly or daily, per FTE, per seat, fixed fee, per transaction or outcome, subscription	Spreadsheet	Factor 2
3	Gross margin by service line or product	Spreadsheet	Factor 3; flow-through
4	Headcount by function, with delivery staff cost as a share of revenue	Spreadsheet	Factor 3; upside score
5	Top 20 customers with tenure, revenue, contract end date and termination terms	Spreadsheet	Factor 6
6	Gross and net revenue retention by annual cohort for three years, with churn reasons	Spreadsheet	Factor 6
7	Data inventory: what data the company holds, how structured it is, and the contract clauses on its use	Document	Factor 4
8	Licences, accreditations, professional indemnity terms and regulatory obligations attaching to the company's output	Document	Factor 5
9	New business for the last two years by source (referral, existing customer, tender, search, aggregator, partner)	Spreadsheet	Factor 8
10	The five competitors management worries about most, and lost-deal reasons for four quarters	Document	Factor 7
11	Customer requests for price reductions or scope changes citing AI, and what was agreed	Document	Factor 7; pass-through
12	Twelve months of AI tool and vendor spend, with usage data where available	Spreadsheet	Capacity C1
13	Current cash, facility headroom and covenant position	Summary	Capacity C2
14	Ten anonymised past deliverables with their briefs (services), or product documentation and a trial account (software)	Files	Factor 1 replicability test
15	Process volumes and time-per-unit for the three largest back-office processes	Spreadsheet	Upside score

Tick-box version for sending:

- ☐ Revenue by line and by billing basis, two years plus year to date
- ☐ Gross margin by line; headcount by function
- ☐ Top 20 customers with tenure, contract end and termination terms
- ☐ Retention cohorts with churn reasons
- ☐ Data inventory with data-use clauses
- ☐ Licences, accreditations and PI terms

- [ ] New business by source
- [ ] Competitors, lost-deal reasons, AI-related price requests
- [ ] AI tool spend and usage
- [ ] Cash, headroom and covenants
- [ ] Deliverable sample or product access
- [ ] Back-office process volumes

14. The management session question set

Sixty to ninety minutes with the CEO and whoever runs delivery or product. Send the topics in advance, not the questions. The first two questions in the competition block move scores more than anything in the data pack.

QUESTION	WHAT YOU'RE LISTENING FOR	FACTOR
Which pieces of work that customers pay for would you least want a customer to try doing with AI themselves? Why?	Whether management knows where the exposure is. Vague answers mean nobody has looked	1
Walk me through how a typical job is priced, from quote to invoice	Hours estimates hiding inside "fixed fees"	2
If your delivery team became 30% more productive tomorrow, who would keep the saving: you or the customer?	Pricing power, and whether management has thought about pass-through	2, revenue at risk
Where does your gross margin actually come from? What happens to it when you add a person?	Labour arbitrage	3
What data do you hold that a competitor couldn't get, and what do your contracts say about using it?	Rights, not volume	4
Who signs the output, and what happens to them if it's wrong?	Named accountability	5
Tell me about the last three customers you lost and the last three who reduced scope	Real switching costs and churn reasons	6
What have you lost to cheaper or AI-led competitors in the last year?	Factor 7 evidence. Expect "nothing" first; ask for lost-deal records	7
Which customers have asked for price cuts citing AI, and what did you agree?	Live pass-through	7, revenue at risk
Have any customers built their own version of what you do?	The in-house competitor	6, 7
How did your last 20 customers find you?	Owned or rented distribution	8
Show me an AI tool someone in delivery uses every week. What changed when they started?	Real adoption, measured	C1
If you needed to change your pricing model and retrain delivery over 12 months, how would you fund it?	Headroom and appetite	C2, C3
What else is taking leadership time this year?	Change capacity	C4

QUESTION	WHAT YOU'RE LISTENING FOR	FACTOR
Where is your team doing repetitive work you'd automate first? How much time does it take?	Upside, sized	U

> **Tip.** Interview at least one person below board level. A team lead's account of what tools are used, and what customers are asking, is usually more accurate than the CEO's.

15. The one-page output per company

Keep it to one page. The Scorecard tab in Excel prints to this layout.

SECTION	CONTENT
Header	Company, business model, revenue, EBITDA, years to exit, date, scorer
Score line	Exposure index and band; triage score (0 to 21); alert if any; moat strength
Factor evidence	Eight rows: score, one line of evidence, confidence
Revenue at risk	Gross and mitigated, as £ and % of revenue, with EBITDA at risk and the three main assumptions
Capacity and time	Capacity total and level; years to exit; owns the shift, with evidence
Class and reason	The class and two sentences explaining it
Actions	Three actions, each with an owner, a date and a measure
Upside	Upside score with the sized opportunity, if any

An investment committee member should be able to read it in three minutes and say whether they agree. If the evidence column is empty for any factor, the page isn't finished.

16. What to tell boards and LPs

Portfolio boards

Each board gets its own page and a conversation. Keep the portfolio league table inside the fund. The CEO of a company classed as exit early hears it from the deal partner first, with the reasoning, before anything changes in the board pack.

Then turn the class into standing agenda items with numbers.

CLASS	QUARTERLY BOARD MEASURES
Defend	Moat factors (retention, data rights, accreditations); back-office savings delivered against the upside score

CLASS	QUARTERLY BOARD MEASURES
Adapt	Share of revenue on fixed, usage or outcome pricing; cost per unit of delivery; win rates against named competitors; revenue at risk, gross and mitigated
Accelerate	Investment deployed against plan; revenue from AI-enabled products or data; new customers won from AI-native rivals
Exit early	Sale preparation milestones; the AI section of the information memorandum and the evidence pack behind it

The questions in what boards should ask about AI map directly onto these measures.

LPs

LPs are asking, and they expect the answer to separate managers. The Coller Capital Global Private Capital Barometer, Summer 2026 found 67% of LPs expect AI adoption to widen the gap between the best and worst performing funds.

What works:

1. **Method first.** Describe the eight factors, the weights and the classification rules, and say that one team applied them to every company.
2. **Portfolio by class as a share of NAV.** The Portfolio tab produces this table.
3. **Actions and dates against adapt and exit-early companies**, without naming companies if confidentiality requires.
4. **What changed in the marks, and why.**
5. **Upside only where the plan has a number and an owner.**

What doesn't work is "we have limited exposure to AI disruption". Every portfolio has some, if only in back offices and in how customers buy. Reporting that four companies sit in adapt with funded plans and dates is more credible than claiming a clean portfolio and being contradicted by the next valuation round.

17. Worked examples

Three illustrative examples, using fictional companies, with numbers chosen to show the method. They end in three different classes. All three are filled in on the example tabs of the Excel workbook, which uses the default weights. Each company here is described as a holding in a portfolio review; the same scoring applies to a target.

Illustrative example A: Calderbrook Compliance Ltd, a fictional consultancy (adapt)

Calderbrook writes policies, audit reports and regulatory submissions for financial services clients. Revenue £9.0m, EBITDA £1.8m. 80% of revenue, £7.2m, is billed on time and materials; £1.8m is fixed-fee retainers. Four years left in the hold. Its profile follows the compliance consultancy in our AI disruption risk assessment.

#	FACTOR	SCORE	EVIDENCE
1	Revenue on replicable tasks	2	Blind test: seven of ten policy and submission drafts rated "send with light edits". Timesheets put 40% of billed hours, £2.88m (32% of revenue), on that work
2	Pricing basis	3	80% time and materials
3	Labour arbitrage	2	Margin from people, but a small productised audit pack range
4	Data and workflow	2	Methods and templates; works on client data
5	Regulation and accountability	2	Clients want a named firm; no licence required
6	Switching costs	1	Gross revenue retention 92%; long relationships
7	AI-native competition	2	Two funded AI compliance tools winning smaller clients; two clients asked for fee cuts citing AI
8	Distribution	1	Referral-led

Index 66 (high). Triage score 14. Alert on factor 2 (pricing). Moat strength 1.67.

Revenue at risk (three-year horizon).

LINE	REVENUE	EXPOSED	COMPRESSION	PASS-THROUGH	AT RISK
Drafting and policy work (T&M)	£4.2m	55%	50%	60%	£0.69m
Audit fieldwork and reporting (T&M)	£3.0m	19%	50%	60%	£0.17m
Retainers, smaller clients (share loss)	£1.8m	20%	25%	100%	£0.09m
Total	**£9.0m**				**£0.95m (10.6%)**

At 80% flow-through, EBITDA at risk is £0.76m, 42% of current EBITDA. That is the figure the board needs to see: the revenue loss is modest, the margin loss isn't, because the fees fall and the team doesn't.

Capacity. C1 1 (Copilot licences, no measured use), C2 2, C3 3, C4 2: total 8, medium. Owns the shift: no.

Class: adapt. The index is below the severe band and capacity and time are adequate, so the rules give adapt; the pricing alert sets the priority. Actions: (1) put AI into the drafting workflow within 60 days and measure hours per deliverable, owner head of delivery; (2) convert the top 20 clients to fixed-fee pricing by month 12, measured as share of revenue on fixed fees, owner CEO; (3) build a productised regulatory change service on the firm's own templates, owner commercial director. The mitigated revenue at risk, if the top 20 move to fixed fees, roughly halves.

Illustrative example B: Brightline Contact Services Ltd, a fictional contact centre (exit early)

Brightline runs 650 agents handling customer contact and back-office processing for retailers and utilities. Revenue £22.0m, EBITDA £2.6m. Priced per FTE and per transaction. Eighteen months to the planned exit. Net debt 5 times EBITDA with tight covenants.

#	FACTOR	SCORE	EVIDENCE
1	Revenue on replicable tasks	3	70% of contact volume is tier-one queries by chat, email and phone that AI agents now handle
2	Pricing basis	3	86% priced per FTE, the rest per transaction
3	Labour arbitrage	3	Agent cost 64% of revenue
4	Data and workflow	2	Runs on clients' CRM; holds QA data
5	Regulation and accountability	2	A regulated complaints line for two utilities clients; the rest unregulated
6	Switching costs	2	Three-year contracts, six months' termination for convenience
7	AI-native competition	3	Two clients piloting AI agents directly with software vendors; one renewal lost
8	Distribution	1	Long-standing relationships and a framework place

Index 85 (severe). Triage score 18. Alerts on factors 1, 2 and 7.

Revenue at risk (three-year horizon, because exit is under three years away).

LINE	REVENUE	EXPOSED	COMPRESSION	PASS-THROUGH	AT RISK
Tier-one contact handling (per FTE)	£13.0m	70%	45%	80%	£3.28m
Complaints and regulated lines (per FTE)	£6.0m	30%	30%	60%	£0.32m
Back-office processing (per transaction)	£3.0m	80%	50%	80%	£0.96m
Total	**£22.0m**				**£4.56m (20.7%)**

At 35% flow-through, EBITDA at risk is £1.60m, 61% of current EBITDA.

Capacity. C1 0, C2 1, C3 1, C4 1: total 3, low. Owns the shift: no.

Class: exit early. Severe index with low capacity and 18 months left. The rule gives exit early; the judgement still has to be made. The market already prices generic BPO at around 5 to 7 times EBITDA, so a sale doesn't escape the repricing. The case for bringing it forward rests on the regulated complaints line, which a specialist buyer may value

above the sector multiple, and on selling before the lost renewal and the two pilots show up in the numbers. Actions: (1) prepare a sale of the business or of the regulated line within six months, owner deal partner; (2) open talks with lenders on covenant headroom before the next test date, owner CFO; (3) agree with the two piloting clients a managed AI service priced per resolved contact, to protect the renewals that matter to a buyer, owner CEO. If the third action wins both renewals on new terms, re-score: capacity would rise and the case for adapt would reopen.

Illustrative example C: Ravensmoor Claims Data Ltd, a fictional vertical software company (accelerate)

Ravensmoor sells claims-handling software to UK motor and property insurers and loss adjusters. Annual recurring revenue £14.0m, EBITDA £3.5m. 60% priced per handler seat, 40% per claim processed. Three years left in the hold.

#	FACTOR	SCORE	EVIDENCE
1	Revenue on replicable tasks	2	Replicability test: one engineer rebuilt the core triage and reporting workflow in four days; it failed on multi-party claims and insurer integrations
2	Pricing basis	2	60% per seat
3	Labour arbitrage	0	Services under 10% of revenue
4	Data and workflow	0	12 years of structured settled-claims data, with contract rights to aggregated, anonymised use
5	Regulation and accountability	1	Customers are FCA-regulated; the product sits inside audited claims processes
6	Switching costs	0	Gross revenue retention 96%; integrated with policy administration systems
7	AI-native competition	2	Funded AI claims start-ups piloting with two customers
8	Distribution	1	Insurer panels and loss-adjuster networks

Index 37 (moderate). Triage score 7. No alert. Moat strength 0.33.

Revenue at risk.

LINE	REVENUE	EXPOSED	COMPRESSION	PASS-THROUGH	AT RISK
Handler seats (seat compression)	£8.4m	50%	30%	70%	£0.88m
Per-claim revenue (share loss)	£5.6m	15%	30%	100%	£0.25m
Total	**£14.0m**				**£1.13m (8.1%)**

At 85% flow-through, EBITDA at risk is £0.96m, 28% of current EBITDA.

Capacity. C1 2 (an ML fraud-scoring model in production), C2 3, C3 3, C4 2: total 10, high. Owns the shift: yes, the claims data and the insurer integrations.

Class: accelerate. The seat revenue is exposed, but the company owns exactly what an AI-native entrant lacks: years of structured claims outcomes and live integrations. Actions: (1) launch AI triage and reserving built on the claims data, priced per claim, owner CPO; (2) migrate seat customers to per-claim pricing at renewal, measured as

share of revenue on per-claim pricing, owner CRO; (3) assess one of the AI-native entrants as an acquisition, owner deal partner. This is the class funds forget; a company like this deserves more capital, not a risk flag.

The portfolio view

Adding three more fictional holdings shows how the Portfolio tab presents a whole book. Figures are illustrative.

COMPANY (ALL FICTIONAL)	MODEL	INDEX	BAND	CAPACITY	YEARS TO EXIT	CLASS
Fellside Fire and Security Ltd	Field services	22	Low	Medium	4	Defend
Ravensmoor Claims Data Ltd	Vertical software	37	Moderate	High	3	Accelerate
Northgate HR Software Ltd	Seat-priced HR software	52	Moderate	Medium	3	Adapt
Calderbrook Compliance Ltd	Compliance consultancy	66	High	Medium	4	Adapt
Kestrel Digital Engineering Ltd	Offshore development	87	Severe	Medium	3	Adapt (severe)
Brightline Contact Services Ltd	Contact centre	85	Severe	Low	1.5	Exit early

Kestrel shows the rule at work: severe exposure, but with medium capacity and three years left, it is an adapt case with a large, urgent plan, not an exit. Fellside shows the defend case: core work that is physical and certified, and an upside score of 3 in scheduling, quoting and engineer reports that the fund should capture.

18. Common mistakes

1. **Letting each deal team score its own holdings.** The people who championed an investment at IC are the last to score it as exposed. One small team, one standard.
2. **Scoring the sector instead of the company.** "IT services" is not a score. A firm with 70% of revenue on multi-year managed services contracts and its own monitoring platform is a different asset from one placing contractors by the day.
3. **Debating factor 1 instead of testing it.** Ten deliverables, current tools, a blind grade from a practitioner. One day settles an argument that otherwise eats a management meeting.
4. **Taking management's word on pricing.** "Value-based" usually means an hours estimate with a fixed price on top. The sales ledger says what the pricing is.
5. **Counting data the company can't use.** Check the contracts. Years of client data with no reuse rights is not a moat.
6. **Netting upside against threat.** A large back-office saving doesn't offset a revenue line that is about to shrink. Score them separately and show both.
7. **Changing weights to rescue a company.** Weights change for a class of business, with a written reason, before scoring starts.
8. **Applying a percentage haircut to the whole revenue.** Revenue at risk is built line by line from exposed share, compression and pass-through. A flat 10% haircut tells the committee nothing about what to fix.
9. **Ignoring flow-through.** A 10% revenue loss in a services firm that keeps its team can be a 40% EBITDA loss. Show both.

10. **Stopping at the score.** The value is in the class and three owned actions. A heatmap without decisions is something to file.
11. **Putting everything in adapt.** If almost every company lands in the middle, the scoring has been softened. Expect defend, accelerate and at least one hard conversation in a typical portfolio.
12. **Treating it as a one-off.** Factor 7 can move in a quarter. Refresh on the cadence above, and before any exit process.

> **Red flag.** A management team that can't name a single customer who has asked about AI, or a single competitor using it, hasn't been asking. Score factor 7 from the lost-deal records and your own competitor scan, and lower C1 and C4.

19. Glossary

TERM	MEANING
Exposure index	The weighted sum of the eight factor scores, scaled 0 to 100. Higher is more exposed
Triage score	The unweighted sum of factors 1 to 7, 0 to 21, as used in the portfolio triage
Band	Low, moderate, high or severe, from the exposure index
Single-factor alert	A score of 3 on factor 1, 2 or 7, which blocks a defend classification
Moat strength	The average of factors 4, 5 and 6. Lower is stronger
Revenue at risk	Revenue × exposed share × compression × pass-through, summed across lines, over the longer of the hold and three years
Pass-through	The share of an AI-driven saving or volume reduction that customers capture
Flow-through	The share of lost revenue that falls to EBITDA if costs don't change
Capacity to respond	Four sub-factors scored 0 to 3: AI delivery track record, funding headroom, board backing, change capacity
Owns the shift	Evidence that the company owns data, distribution or the automation that AI makes more valuable
Replicability test	A timeboxed attempt to reproduce the core deliverable or product with current AI tools, from public material
Labour arbitrage	Earning margin from the gap between labour cost and bill rate

20. Sources and further reading

Market evidence, 2026

- EY, Private Equity Pulse, Q1 2026, April 2026.
- Bain & Company, Private Equity Midyear Report 2026.

- Livingstone Partners, 2026 IT Services Market Update, August 2026.
- Bloomberg, Apollo is screening all software investments for AI threat risk, June 2026.
- Apollo, Torsten Slok, Not all software is bad, July 2026.
- J.P. Morgan Asset Management, Disruption decoded, May 2026. The most concrete public framework for private credit.
- Protiviti, Due diligence reimagined: AI's impact on valuations, March 2026. Useful on the AI maturity side of diligence.
- Coller Capital, Global Private Capital Barometer, Summer 2026.
- Skills England, Annual skills report 2026.
- Office for National Statistics, Artificial intelligence in UK businesses, 2023 to 2026.
- Private Equity Wire, Bain uses AI vibecoding to build software replicas in PE due diligence, June 2026, summarising Financial Times reporting.

Exposure research

- Eloundou, Manning, Mishkin and Rock, GPTs are GPTs: labor market impact potential of LLMs, Science, 2024.
- Cazzaniga and colleagues, Gen-AI: Artificial Intelligence and the Future of Work, IMF Staff Discussion Note, 2024.
- Gmyrek and colleagues, Generative AI and Jobs: a refined global index of occupational exposure, ILO working paper 140, 2025.
- Felten, Raj and Seamans, How will language modelers like ChatGPT affect occupations and industries?, 2023.
- Tomlinson, Jaffe, Wang, Counts and Suri, Working with AI: measuring the applicability of generative AI to occupations, Microsoft Research, 2025.
- McKinsey Global Institute, The economic potential of generative AI, 2023.
- Department for Education, The impact of AI on UK jobs and training, 2023.
- Bain & Company, Will agentic AI disrupt SaaS?, Technology Report 2025. The best public framework for software workflows.

Productivity evidence

- Noy and Zhang, Experimental evidence on the productivity effects of generative AI, Science, 2023.
- Brynjolfsson, Li and Raymond, Generative AI at work, NBER, 2023.
- Dell'Acqua and colleagues, the jagged technological frontier study of BCG consultants, Harvard Business School, 2023.
- METR, early-2025 AI and experienced developer productivity, 2025, and February 2026 update.

Related reading on this site

- AI disruption risk assessment: will AI hollow out the target?
- AI disruption risk across a portfolio: a triage framework
- Could a small team with AI coding tools rebuild this product?
- AI washing in due diligence
- Will buyers pay for your AI?

Anyone wanting this run across a portfolio, or on a live deal, can find how we do it on the AI disruption risk review page.

21. About this resource

Written by Anyway Ventures, the York technology and AI practice led by Brannan Coady, who has been on both sides of technical and AI due diligence as an operator, acquirer and seller. Version 1.0, October 2026. Free to use, adapt and share with attribution. Not legal, tax or investment advice.

ABOUT ANYWAY VENTURES

Senior hands for technology and AI decisions

Technology and AI work for private equity funds, venture investors and UK mid-market management teams: before a deal, during the hold and at exit. Every engagement is led by Brannan Coady, an exited founder and former CEO. Fixed price, built to be handed over.

FOR INVESTORS

Technical due diligence

Code, architecture, security, IP and one-to-one team interviews, every material finding costed. Red-flag review in five to seven days.

ANYWAYVENTURES.COM/TECHNICAL-DUE-DILIGENCE

AI due diligence

Whether a target's AI claims are real, defensible and theirs, and a costed AI value-creation plan for the IC.

ANYWAYVENTURES.COM/AI-DUE-DILIGENCE

AI disruption risk review

Is AI a threat to the business or its moat? One deal, or every holding in a portfolio ranked by exposure.

ANYWAYVENTURES.COM/AI-DISRUPTION-RISK

Technical due diligence for VCs

A two to three day check before the term sheet: the product tested on evidence, not the demo.

ANYWAYVENTURES.COM/VC-TECHNICAL-DUE-DILIGENCE

Retained AI operating partner

Senior AI capacity held by the fund and deployed into the company with the strongest AI value case.

ANYWAYVENTURES.COM/AI-OPERATING-PARTNER

Exit readiness and vendor technical due diligence

Find what a buyer's diligence will find, 12 to 24 months out, fix it, then publish a report buyers can rely on.

ANYWAYVENTURES.COM/EXIT-READINESS

FOR OPERATORS

AI audit and readiness assessment

Two to three weeks to a ranked, costed list of AI opportunities, the gaps to close and a 90-day plan.

ANYWAYVENTURES.COM/AI-AUDIT

Embedded AI delivery

Internal tooling, automation and AI agents in production inside your business, built with your team.

ANYWAYVENTURES.COM/EMBEDDED-AI-DELIVERY

Fractional Chief AI Officer

One senior person who owns AI on your leadership team, two to four days a month.

ANYWAYVENTURES.COM/FRACTIONAL-CHIEF-AI-OFFICER

Hands-on AI training

In-person training where every attendee builds a working tool on their own job.

ANYWAYVENTURES.COM/AI-TRAINING

Board and NED roles

Non-executive seats where technology, product and AI decide the outcome.

ANYWAYVENTURES.COM/BOARD-AND-NED

Every enquiry starts with the same free 30-minute call.

A straight answer on whether we can help, what it would cost and how long it would take.

EMAIL
brannan@anywayventures.com

BOOK A CALL
anywayventures.com/contact

MORE RESOURCES
anywayventures.com/resources

Anyway Ventures Limited · Registered in England and Wales, company number 16318119 · Registered office: Newbury House, Main Street, Little Ouseburn, York, YO26 9TD. This resource is general guidance, not legal, tax or investment advice. Free to use, adapt and share with attribution to Anyway Ventures. Version 1.0, October 2026.