ORTENT WHITEPAPER

Due diligence by design

What a SaaS or AI company should build from day one so that diligence lifts its price instead of cutting it.

By Andrew Wyatt, Founder, Ortent Advisory Ltd

July 2026 · For founders, boards and investors in SaaS and AI

ORTENT.CO · SAAS AND AI GROWTH · BOARD ADVISORY · NED

Contents

Executive summary

A buyer does not pay for revenue and technology. It pays for confidence that the revenue will hold, the product will scale, the assets will transfer and nothing ugly will surface after completion. Diligence is where that confidence is won or lost.

Most companies treat diligence as a data-room exercise that starts when a sale starts. By then it is too late to fix the things that cost the most. The founder who wins runs it the other way round. The evidence a buyer will ask for is built while the business is built, so that when the process comes the answers are already true, already proved and already consistent.

This paper measures nine domains that decide how diligence goes for a SaaS or AI company: chain of title, data provenance, revenue quality, commercial and change of control, privacy, security and resilience, AI governance and claims, sector and regulatory perimeter, and insurability and disclosure. For each one it sets out what a buyer tests, what to build from day one, and what a weakness costs in a transaction. The costs are specific: a lower multiple, a bigger escrow, a new indemnity, a deferred payment, a delay, or a buyer who walks.

Three ideas run through all nine. First, a story that stays consistent through diligence is worth more than a better story that shifts, because change reads as either concealment or not knowing your own business. Second, the speed and accuracy of your answers count as much as the answers themselves, because in a competitive process response time tells a buyer whether the evidence was ready or is being built on demand. Third, this is daily work, not a project. The discipline that makes you diligence-ready is the same discipline that makes you better run.

The companion tool at ortent.co/tools/diligence-radar scores your company across the nine domains and shows you the widest gaps. The self-scoring prompt does the same from your own evidence, one domain at a time, and will not lift a score without proof. This paper is the argument behind both. It is written for founders and their boards, and for the investors and acquirers who will one day test them.

What the radar measures

The radar scores a company against nine domains. Each domain gets two scores out of ten.

Exposure is how much the domain can move a deal for your company in particular. A company that sells AI into hospitals carries high exposure on sector regulation and data provenance. A horizontal developer tool carries less. Exposure is not a judgement on the company. It is a reading of where the money and the risk sit for this business.

Readiness is how well you can prove your position today, with evidence that already exists, not evidence you would assemble if asked. A written policy is not readiness. A signed assignment, a reconciled metric, a dated test result, a licence that covers the use: that is readiness.

The gap is exposure minus readiness. It is the part of your transaction risk that is real and unmanaged. A domain can carry high exposure and still score a small gap if the evidence is in order. The point of the radar is to find the domains where exposure is high and readiness is low, and to close that distance while it is cheap to close.

Exposure usually sits above readiness, the same way importance sits above delivery in any honest self-assessment. That is normal. The question is not whether a gap exists. It is which gaps are wide enough to cost you at the table.

How each domain is scored

Readiness follows a simple ladder. Nothing behind a claim scores near zero. A policy with no evidence of practice scores low. A control that runs but is not tested or logged scores in the middle. A control that runs, is tested, is logged and can be traced end to end scores high. So a five is a policy that is really used, not a policy that exists. An eight is a claim you can prove in an afternoon from records you already keep.

Exposure follows the company. It is set by what you do, who you sell to, what data you hold, which markets you touch and what role you play in any AI system. Two companies with identical controls can carry very different exposure, because one sells to regulated buyers and one does not.

Insurability is scored a little differently, because it depends on the others. A domain that a buyer cannot cover with warranty and indemnity insurance carries higher exposure by default, since the risk has nowhere to go but the price or the escrow.

The rules that keep the scoring honest

Four rules keep the scores tied to evidence rather than to confidence.

First, every readiness score names its proof. Not a policy title. The actual artefact: the signed document, the reconciled export, the test result, the approval trail. No artefact, no score above the middle of the ladder.

Second, silence is not readiness and it is not a low score. If a domain has no evidence behind it, it is marked not ready, and the paper says the evidence is missing rather than guessing at a number. Missing evidence is a finding, not a gap to be smoothed over.

Third, a policy is not proof. A company that has written an AI policy, a security policy or a data-retention schedule has not, by that act, done the thing. Where a control appears only as a document with nothing showing it runs, readiness is capped low. Writing it down is not doing it.

Fourth, nothing is invented. A score reflects only what the evidence supports. A plan dated in the future is a plan, not a result. A certificate in progress is not a certificate. A claim in a pitch deck is not a fact until a record stands behind it.

These rules cost comfort. They will show gaps a founder would rather not see. That is the point. The buyer will find them either way. Better to find them first, while the people who can fix them are still in the building and still take your call.

The three ideas the paper rests on

Consistency beats polish. Diligence is not one document review. It is a series of readings, formal and informal, in which a buyer compares your pitch, your board papers, your accounts, your contracts, your product and your answers in the room. Conviction rises when those readings agree and falls when they do not. A story that shifts under questioning reads as concealment or as a team that does not know its own business, and the damage lands whether or not the underlying fact was harmless. One caution, because a sharp founder will push back on this: the point is not that your strategy must never change. Strategy should change, and buyers want to see a company that adapts. What must stay stable and provable is the evidence base, the metric definitions, the ownership, the contract terms. Unexplained change is the killer, not change.

Speed is a signal. In a competitive process the buyer is reading your operation through the diligence itself. Fast, accurate, consistent answers say the evidence was ready and the company is run well. Slow answers that arrive in pieces and do not quite reconcile say the opposite, and they say it about how you will run after completion. Response time is not administration. It is the buyer watching you work.

Readiness is built in, not bolted on. The same discipline that makes diligence fast is the discipline that makes a business good: metrics that reconcile, contracts you can find, data you have the right to use, controls that run. You do not build this for the buyer. You build it because it is how a serious company operates, and the transaction value is a return on work you should be doing anyway.

DOMAIN 01

Chain of title

do you own what you are selling

A buyer's first question is whether everything required to run the product can transfer with the business. That covers two things: the company itself, and the assets it runs on.

The company side is the cap table and the corporate record. A clean, fully diluted cap table that reconciles to every share issue, transfer, option grant, SAFE, convertible and warrant. Statutory registers that are current. Board and shareholder approvals for the decisions that needed them. A discrepancy here is title risk in the very thing being sold, and it is far easier to fix while the former founder, the early employee and the first investor still answer your email.

The asset side is intellectual property. Every founder, employee and contractor should have signed a present-tense assignment of the work they did, with confidentiality and, where it bites, a waiver of moral rights. This has to cover the pre-incorporation work, the inventions, the designs, the databases, the domains and the model artefacts, not the code alone. Background IP and third-party materials should be scheduled, not assumed. The repositories, the signing keys, the production access and the cloud accounts should sit under company control, not a founder's personal login. The single most common defect I see is a contractor who wrote something important and never signed anything. It is cheap to fix in year one and expensive to fix in a process.

The transaction effect of a gap here is severe and direct. Unassigned IP becomes a closing condition, a price cut, a specific indemnity or a dead deal. A buyer cannot pay full value for an asset it is not sure it will own.

DOMAIN 02

Data provenance

where your data came from, and what you may do with it

For an AI company this is now the domain that moves fastest. A buyer will want to know, for every material dataset, where it came from, who owns it, what you are allowed to do with it, and whether that includes training, fine-tuning, evaluation and commercialising the output. Data found online is not free to train on or sell. Customer data is not yours to reuse unless the contract, the notice and the design all say so.

The reason this has moved is that the risk now carries a price. Anthropic settled a training-data copyright claim in 2026 for around 1.5 billion dollars, roughly 3,000 dollars a book, with final approval in July. That turns a theoretical worry into a number a buyer can put in a model, and it means a licensing market now exists as the clean alternative to argument. A buyer testing your data will not accept a fair-use theory where a licence was available and you did not take it.

The day-one move is a data bill of materials. For each dataset: the source, the owner, the licence, the permitted purposes, the data types inside it, the lawful basis where personal data is involved, and any restriction that a change of control would trigger. Keep the lineage, so you can show how a training set was built and filtered. The company that can hand a buyer a documented, licensed data trail is selling a defensible asset. The company that cannot is asking the buyer to price a lawsuit it cannot see the edges of.

DOMAIN 03

Revenue quality

numbers a buyer can rebuild

A buyer will recalculate your revenue from the ground up, and the multiple it pays turns on whether the numbers survive that. The work is to make every metric reproducible from source before anyone asks.

That starts with definitions. ARR is not a standard measure. Once usage, consumption, services and discounts enter the picture, two companies can mean very different things by the same number. Set a board-approved definition of ARR and MRR that states what is in, what is out, how you annualise, how you treat trials, usage, services and foreign exchange, and how it differs from recognised revenue. Keep customer-level monthly cohorts so retention, churn, expansion and concentration can be rebuilt, not asserted. Preserve the trail from contract to invoice to cash, and reconcile the management metrics to the ledger every month. A metric a buyer can reproduce is worth more than a flattering one it has to unwind, and unwinding one is how a quality-of-earnings adjustment starts.

For AI products there is a newer test: is the revenue durable, or just fashionable. AI features are easy to adopt and easy to cancel, and usage-based revenue is more volatile than seats. Buyers have started separating AI revenue that is embedded in a workflow with real switching costs from AI revenue that could churn out inside a year. Show adoption, repeat use and retention by cohort, not a single headline growth rate. Durability is worth more than speed, and it is something you build in from the first release, not something you find later.

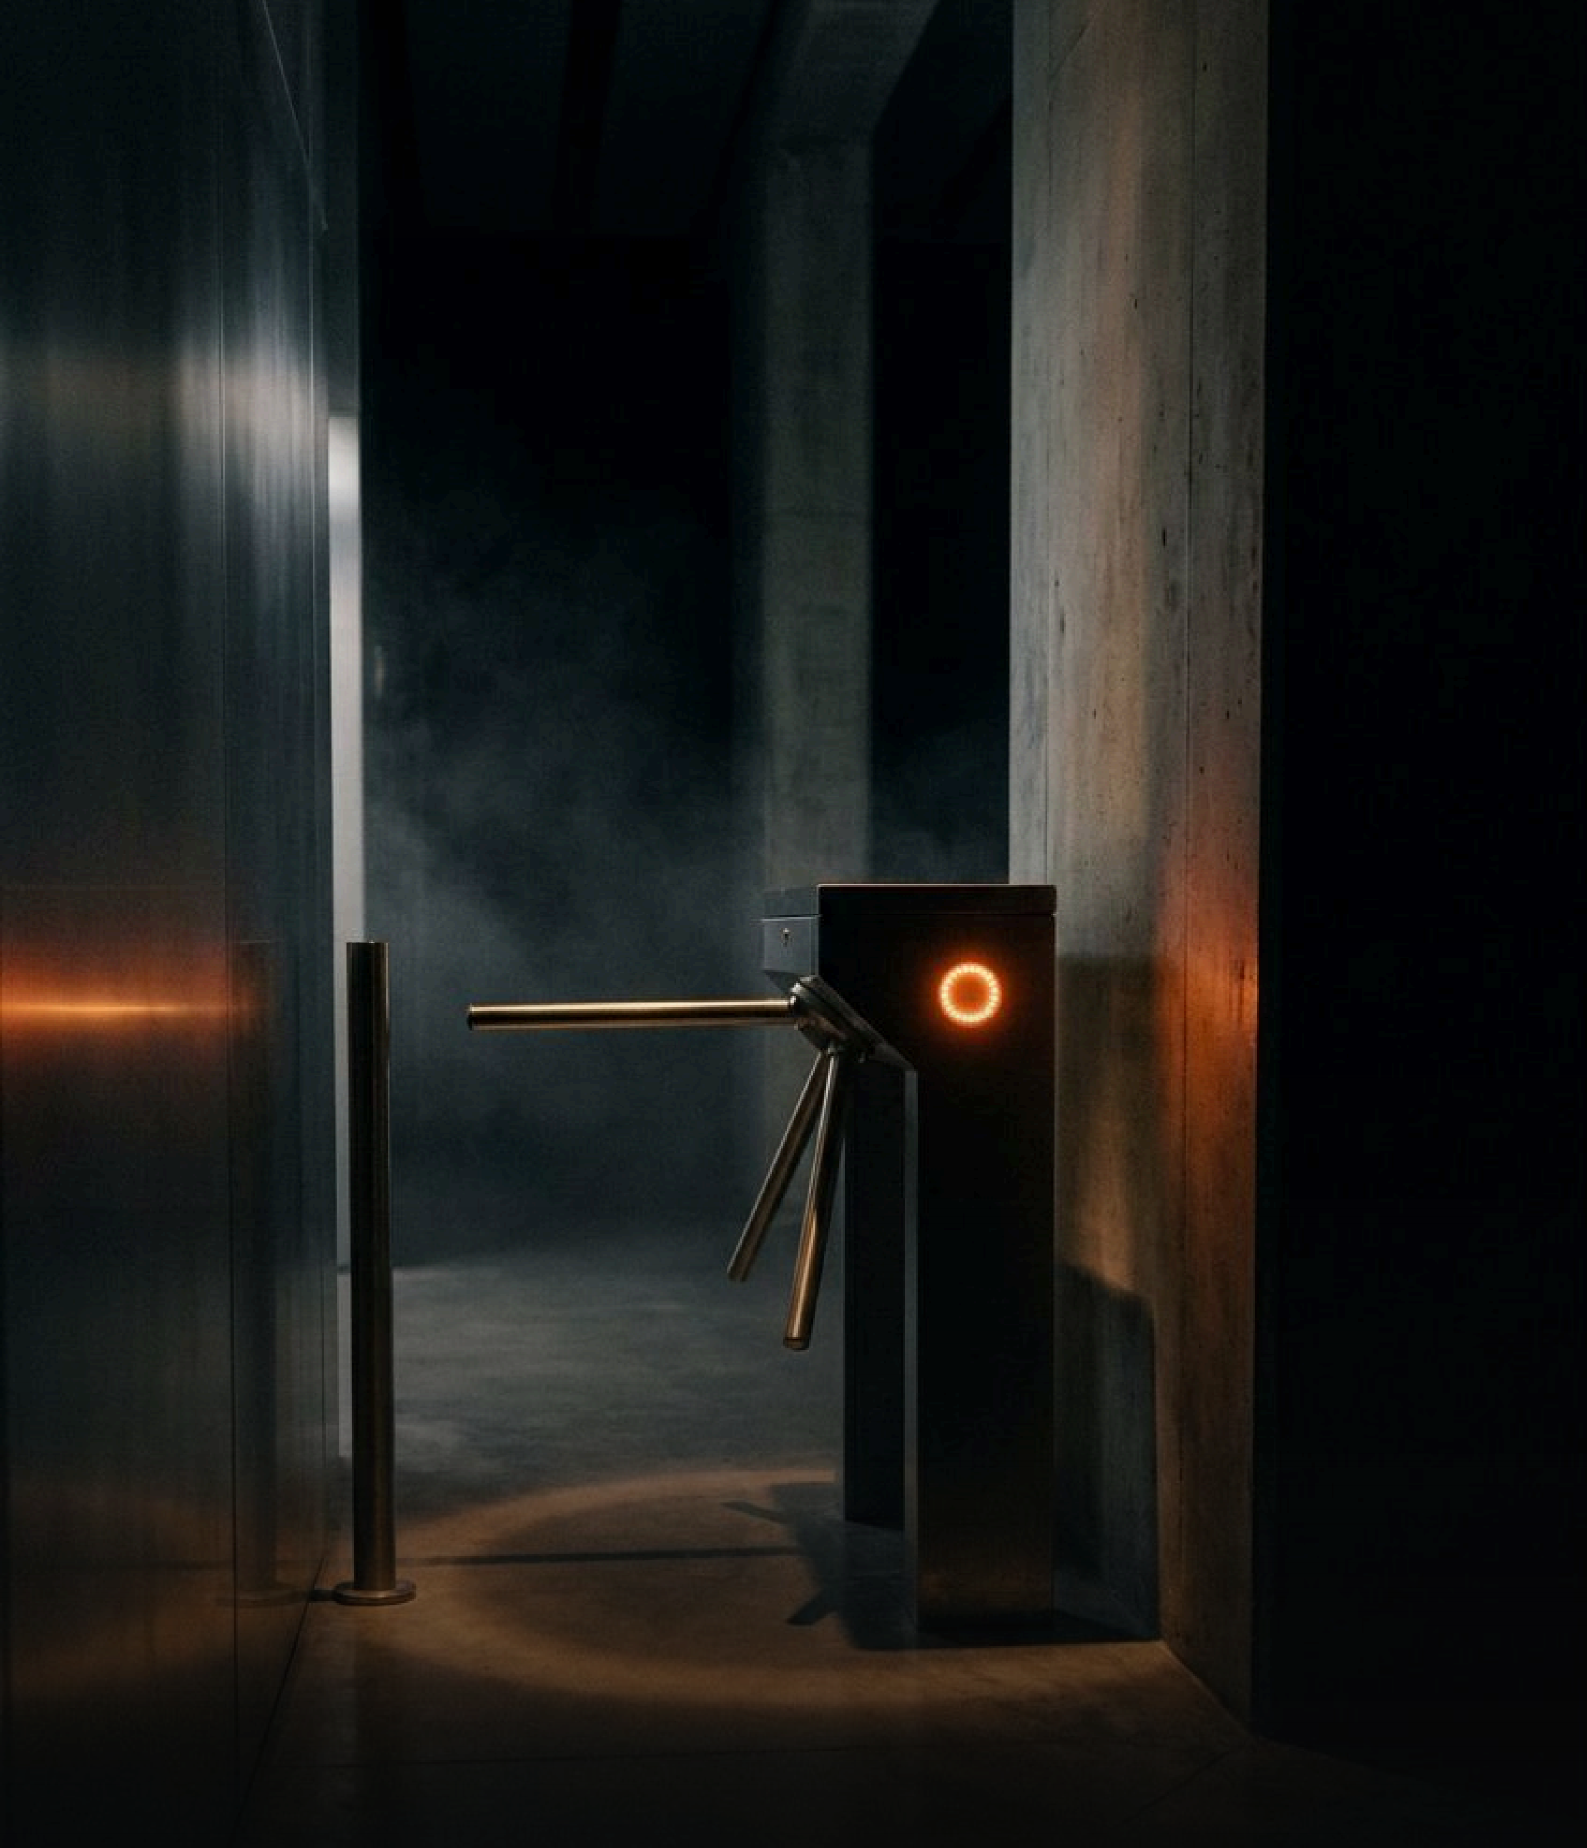

DOMAIN 04

Commercial and change of control

the contracts that either move or block

Your customer and supplier contracts either support a sale or obstruct it, and most founders do not know which until a buyer's lawyer tells them. The fix is to know already.

Track every signed contract, amendment, order form and side letter, and pull from each the terms that matter in a deal: the change-of-control and assignment provisions, the term and renewal, the termination rights, the liability caps, the IP and customer-data rights, the exclusivity or most-favoured terms, and any promise on roadmap, escrow or source code. A single clause requiring customer consent on a change of control, sitting in your largest account, can force a revenue haircut or a closing condition. It is a different conversation if you know about it two years out.

The harder discipline is non-standard obligations. A clause your sales team negotiated to win a deal, that never reached product, security or finance, is a future diligence exception waiting to be found. Record it, own it, and make sure the business is actually delivering it. Verbal concessions should become written amendments. Revenue quality and contract quality are the same asset seen from two sides.

DOMAIN 05

Privacy

mapped to the product, not to a template

Privacy diligence fails most often not because a company has no policy, but because the polished policy does not match what the product actually does. The evidence a buyer trusts is a data map, not a document.

Build a record of what personal data you hold, where it lives, who you share it with, which countries it touches and how long you keep it. Document your role as controller or processor and your lawful basis for each use. Keep current processing agreements and a real subprocessor list. Make sure the data-subject rights you promise, deletion, access, objection, actually work when tested, including through backups and any derived AI artefacts. Where you process at scale or in ways likely to be high risk, do the impact assessment before the processing starts and treat it as a live document, not a retrospective form.

One area deserves its own line: training on customer data. Default to not doing it unless the contract, the notice, the lawful basis and the technical design all clearly allow it, and keep those four things aligned. Possession of data is not permission to reuse it. A buyer that finds service data quietly feeding model improvement, with no basis for it, is looking at a retraining cost and a regulatory exposure at once.

DOMAIN 06

Security and resilience

controls that match the claim

Security diligence is a test of whether your controls match what you have told customers you do. The gap between the two is where the risk sits.

The baseline is unglamorous and expected: an asset and data inventory, identity control with multi-factor authentication on privileged access, a secure development process with review and testing, vulnerability management with real remediation, encryption, central logging, tested backups with stated recovery targets, and an incident response plan you have actually rehearsed. Vendor risk is part of this now, not a footnote, because your suppliers are your exposure too.

Certifications such as ISO 27001, SOC 2 or Cyber Essentials can speed enterprise sales and shorten diligence, but only if the scope covers the real product, people and cloud, and only if the exceptions are tracked, not buried. Two traps to avoid. Do not claim your cloud provider's certificate as your own. And do not let a certificate stand in for a control that is not really running underneath it. A buyer that finds a certified control failing in practice trusts nothing else on the list.

DOMAIN 07

AI governance and claims

governed through its life, and provable

For an AI company this domain carries the premium and the risk in the same place. A buyer will not pay a premium for having AI. It pays for AI that owns a workflow, runs on data you are allowed to use, and produces a measured outcome you can prove. Everything in this domain is about making that provable and keeping it governed.

Govern the system across its life, not at launch. Keep a register of every model, its purpose, its data, its version, its owner and its limits. Keep model or system cards, evaluation results against set thresholds, and testing for the failure modes that matter for your use: bias, hallucination, privacy leakage, prompt injection and abuse. Keep evidence of human oversight where a decision affects a person, and monitoring for drift once it is live. For any third-party foundation model you depend on, record the terms that matter: whether your inputs can be used to train the provider's model, the retention, the deprecation notice, the indemnity, and your plan if the model is withdrawn or repriced. A buyer will tell the difference between a real moat and a thin dependency on someone else's API.

Two developments make this sharper in 2026. First, agents. A system that takes actions, rather than only generating text, raises questions a model card does not answer: what can it do without a human, who is liable when it acts wrongly, how are its actions logged and reversed, and what access does it hold. If you run agents, govern them as actors, not as features. Second, claims. AI-washing is now an enforcement risk, not a marketing choice. The SEC and FTC have brought cases through 2025 and 2026, including business-to-business ones, and a buyer inherits that exposure. Every categorical claim, autonomous, AI-powered, an accuracy or bias figure, needs a substantiation file behind it. A claim you cannot evidence is the fastest way to break the consistency your whole story depends on.

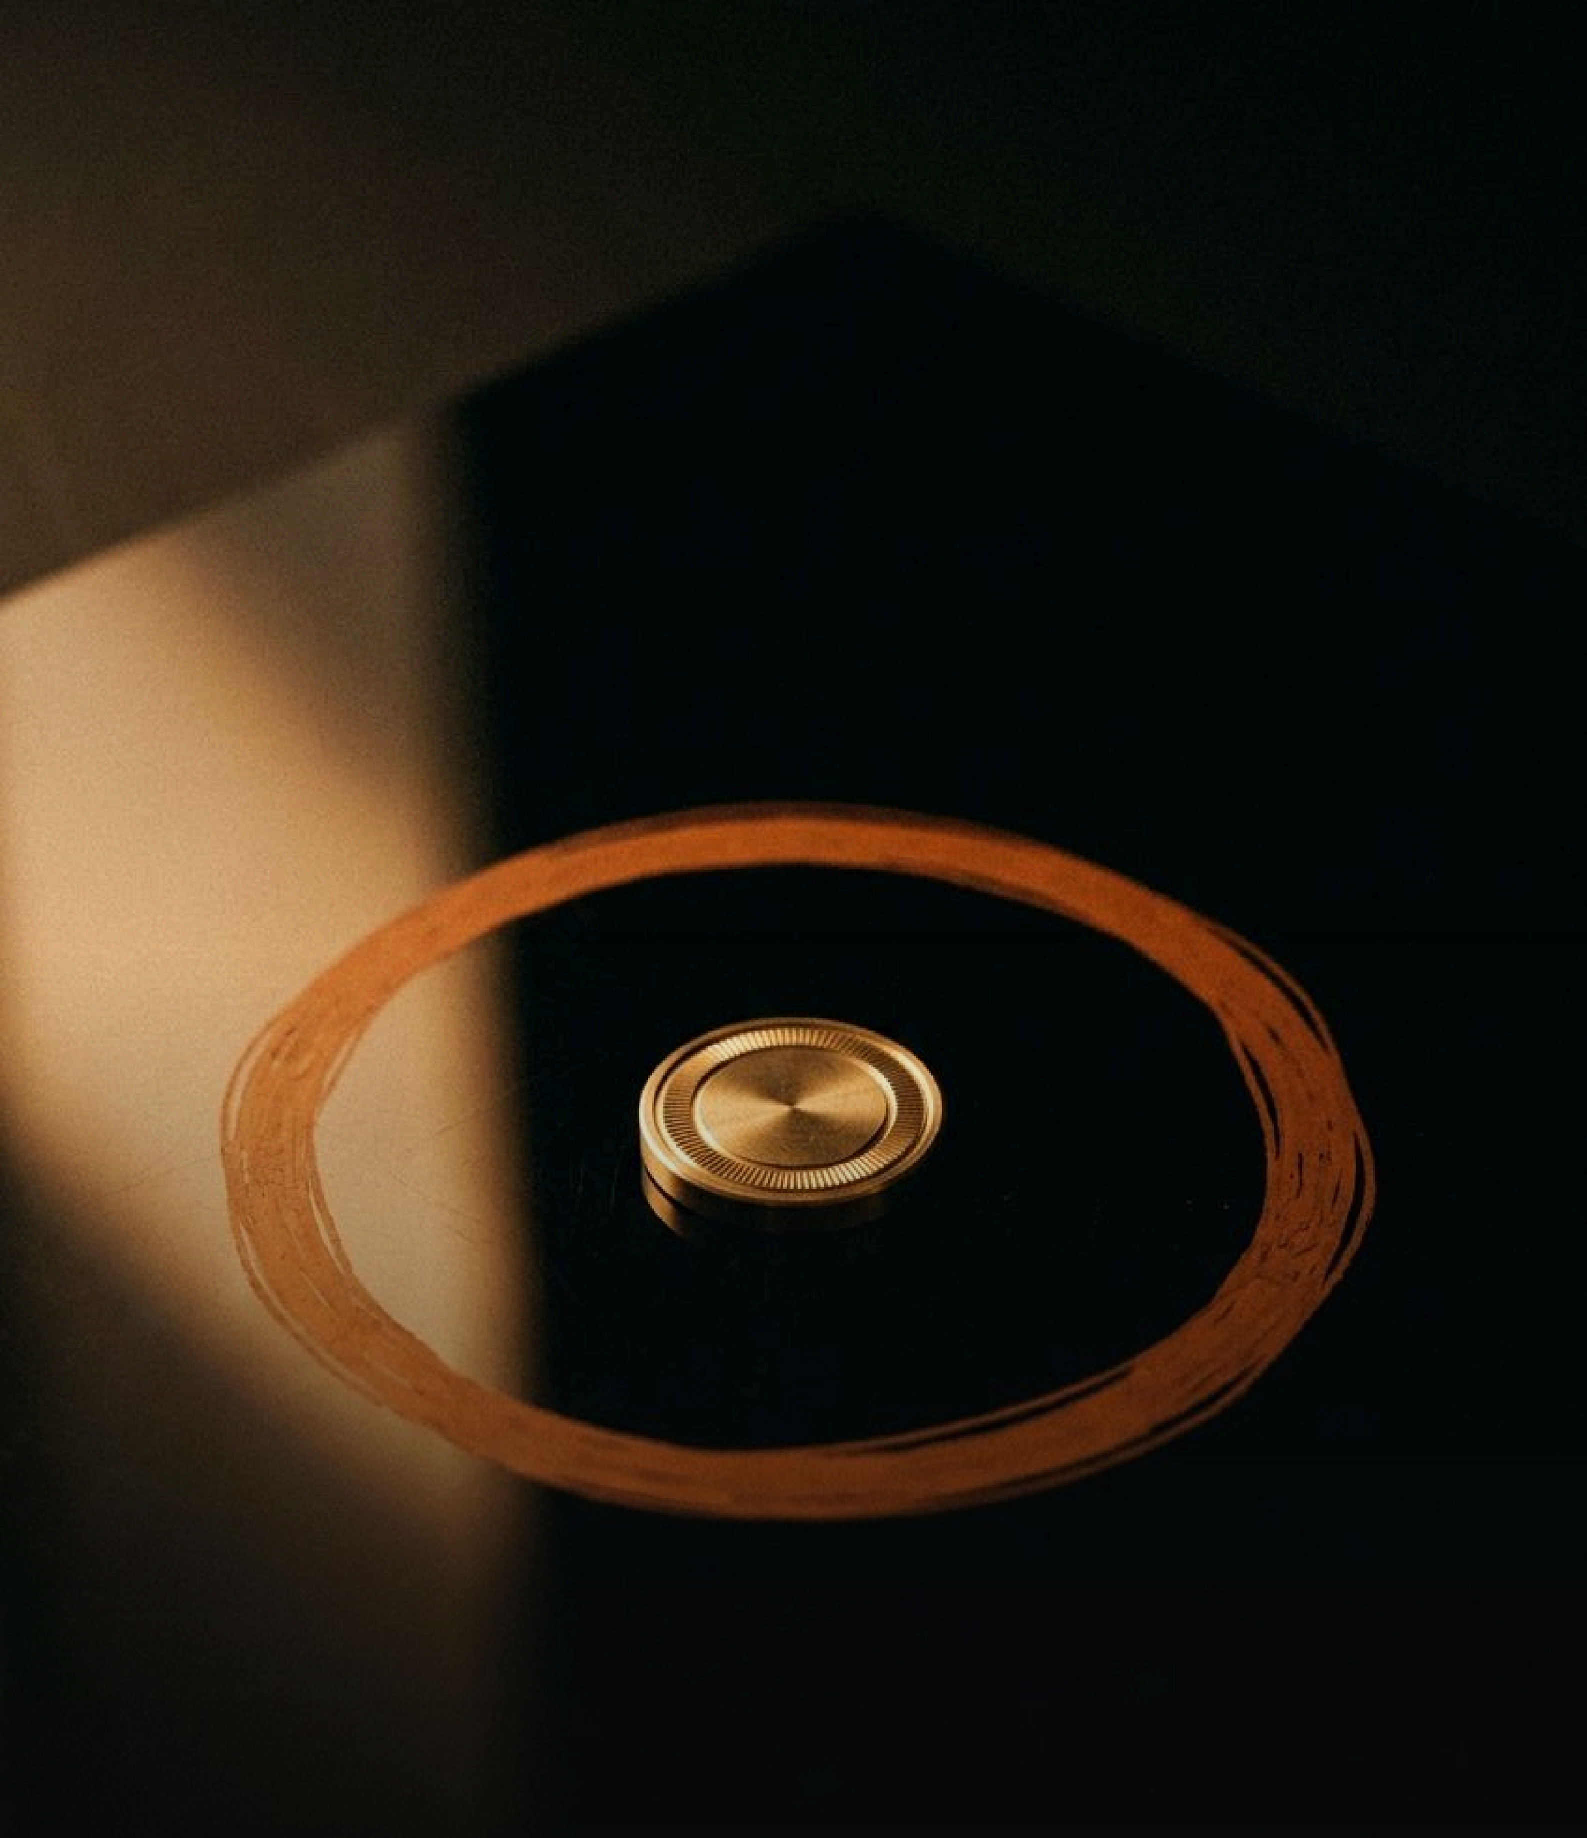

DOMAIN 08

Sector and regulatory perimeter

classify before you scale

The most expensive regulatory mistake is discovering after product-market fit that your product is a class of thing you did not build it to be. Classification comes before compliance, and it comes before scale.

Work out, in writing, what your product is in each market you sell to. In health, whether its purpose and claims make it software as a medical device, and in the UK the MHRA and in the US the FDA both treat much clinical software as regulated. In financial services, whether you perform or support a regulated activity such as payments, lending, advice or insurance distribution, or whether you are an unregulated supplier to a regulated firm, which still inherits audit, resilience and exit obligations through the customer contract. Employment, education, children's data and biometrics each carry their own overlay, and the trigger is the real-world decision your product affects, not whether marketing calls it AI.

The EU AI Act deserves a specific and current note, because its dates have just moved. The Digital Omnibus agreed in late 2025 pushed the high-risk obligations back: standalone high-risk systems now fall due on 2 December 2027 and high-risk systems embedded in regulated products on 2 August 2028. The transparency duties under Article 50 stayed at 2 August 2026, with a short grace period for watermarking, and new prohibitions were added. Formal adoption was expected around the middle of 2026. Diligence has to establish your role for each system, provider, deployer, importer or distributor, because the duties differ, and a change of purpose or a rebrand can move you into the provider role without you noticing. Do not wait for the final date. A buyer will price the distance you still have to travel.

The lesson of those shifting dates is itself a discipline. Regulation moves, so hard dates in your own documents rot. Keep a dated regulatory map and refresh it before a launch or a process, rather than trusting a memo written a year ago.

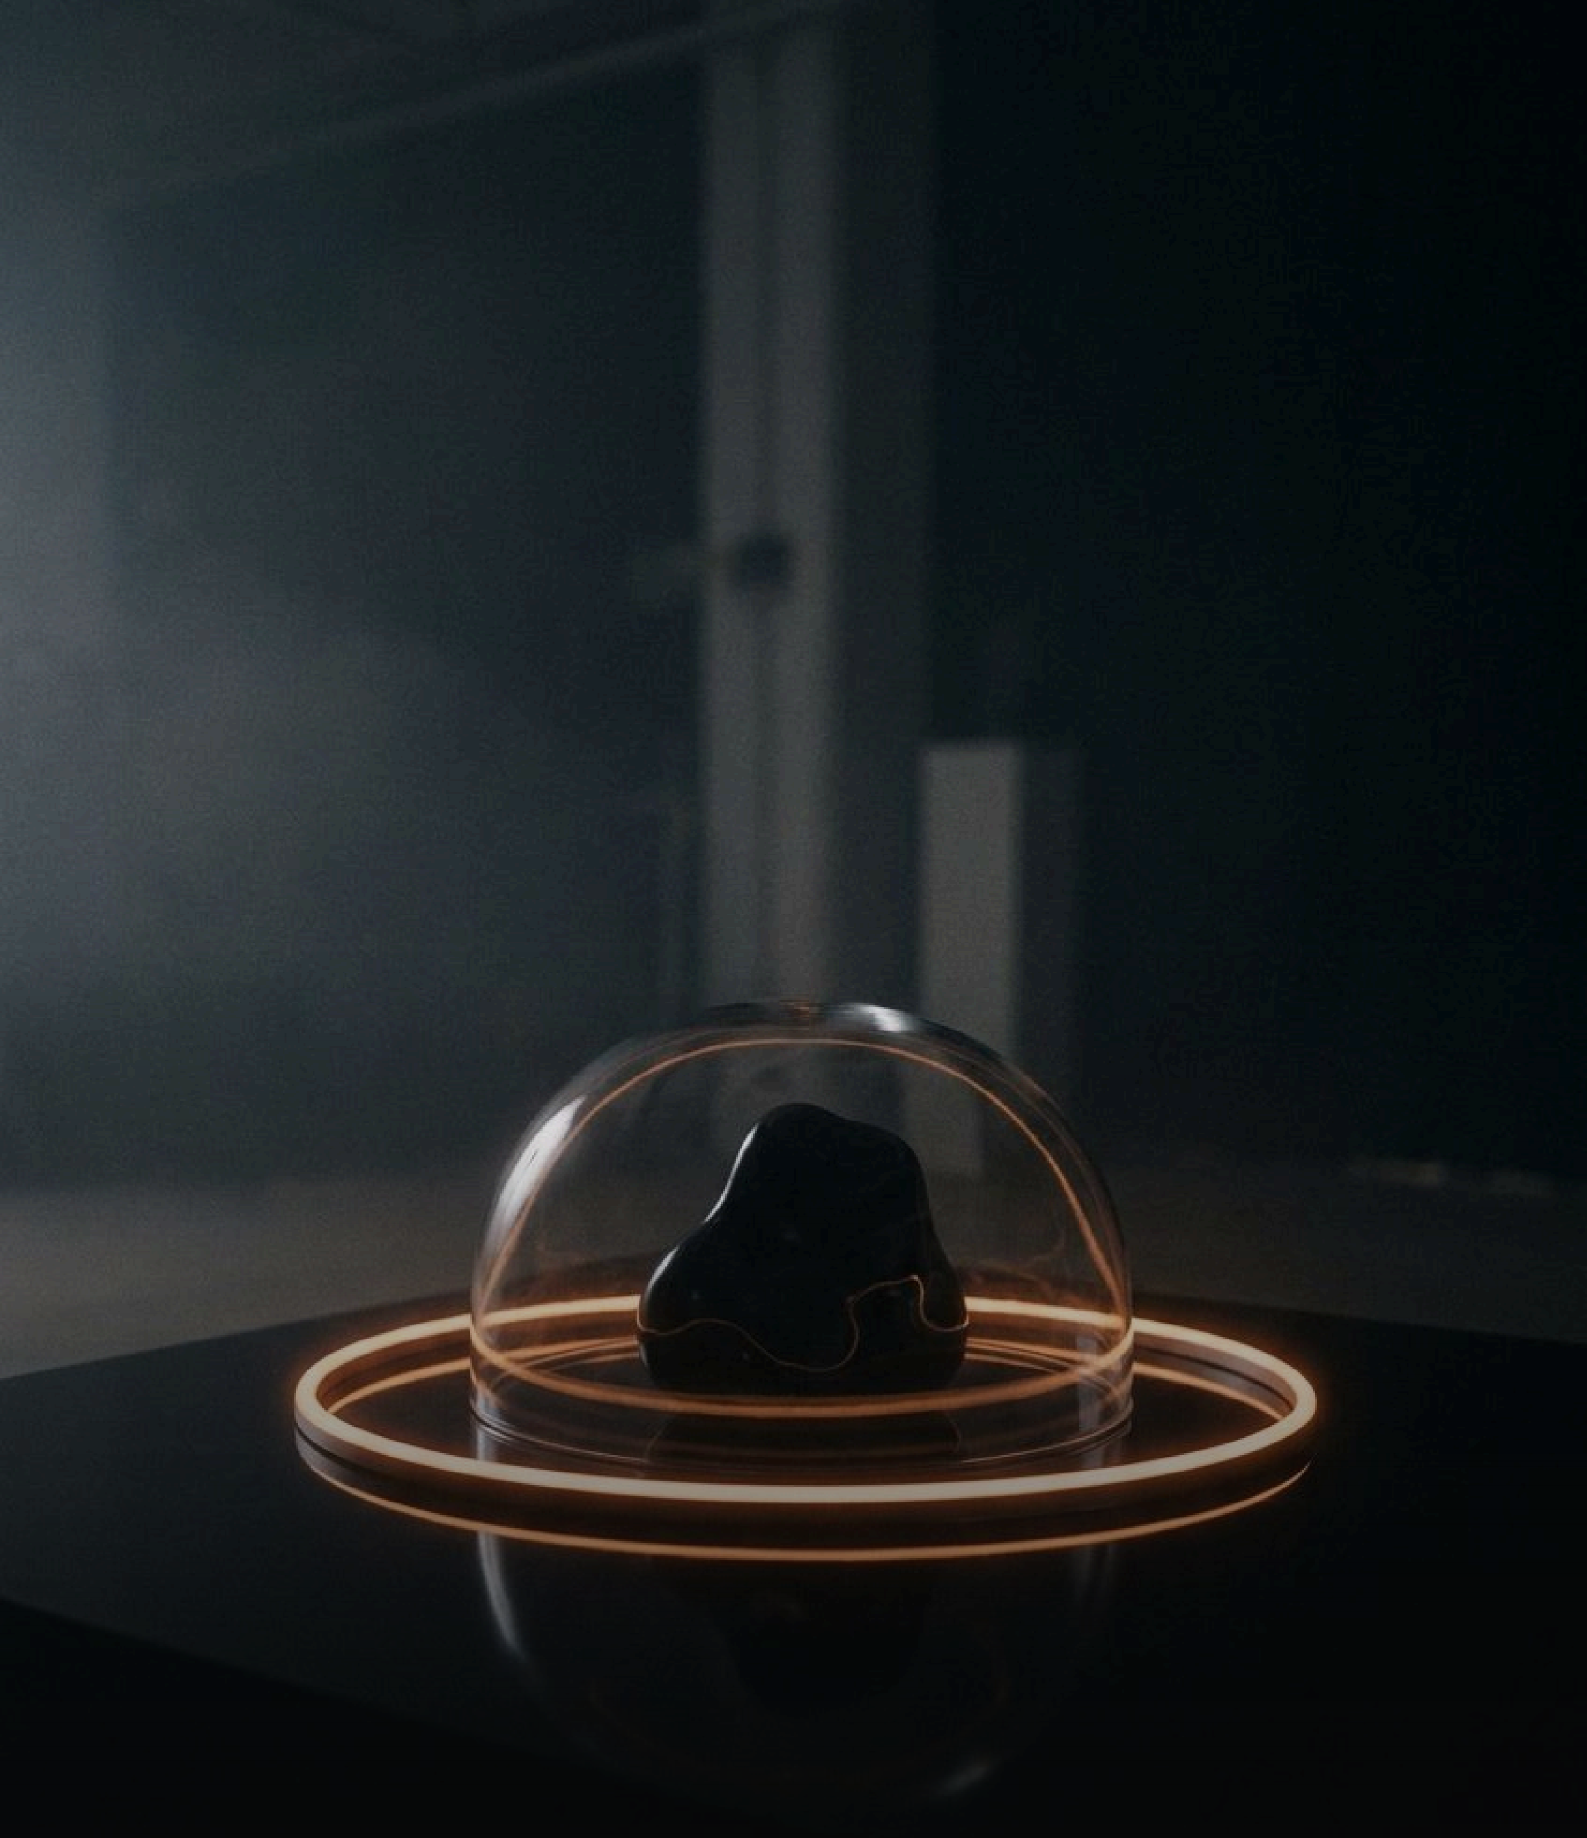

DOMAIN 09

Insurability and disclosure

can a buyer insure this

This is the domain most founders have never thought about, and it is where the other eight get priced. In most deals of any size the buyer transfers risk off its own balance sheet through warranty and indemnity insurance. Whether it can do that for your company depends on the evidence you have built.

Insurers are ending what they call silent AI. They are writing AI-specific exclusions into cyber, technology errors and omissions, directors and officers, and warranty and indemnity cover, and a single AI failure can now fall between all of them. The practical consequence for a sale is direct. If your AI governance, your data rights and your security cannot be evidenced, the insurer carves out those representations, and the risk they would have carried falls back onto you, into a larger escrow, a specific indemnity or a lower headline price. Readiness is the difference between a risk someone else insures and a risk you keep.

The other half of this domain is disclosure discipline. The company that goes into a process with an indexed data room, a known list of issues, and a costed plan for each one is in a far stronger position than the company whose problems the buyer discovers first. A known, disclosed, costed issue is a negotiation. A surprise found in diligence is a loss of trust that colours everything else. Buyers now run their own AI across your data room and surface inconsistencies at a scale a human reviewer would miss, which raises the bar on both the consistency and the speed this whole paper rests on. Your evidence has to be indexed, machine-legible and reconcilable, so it survives that kind of reading.

What this means if you are a founder

Build the evidence while you build the business. The registers, the assignments, the reconciled metrics, the data trail and the tested controls are not a diligence project you run before a sale. They are how a serious company operates, and they compound. A defect caught in year one costs a signature. The same defect caught in a process costs a price chip, a delay, or the deal.

Pick your gaps by exposure, not by comfort. Run the radar, find the domains where exposure is high and readiness is low, and close those first. For most SaaS and AI companies the widest gaps are chain of title early on, then data provenance and AI claims as the product matures, then insurability as a transaction comes into view. You do not have to be a bank. You have to control the risks that could actually cost you at the table.

And protect your consistency. The most valuable thing you own going into a process is a story that holds up from your pitch to your accounts to your contracts to your product. Guard your metric definitions, your ownership record and your claims the way you guard your cash, because a buyer prices the inconsistency it finds, not the explanation you offer afterwards.

What this means if you are an investor or a board

The gap between exposure and readiness is a useful board instrument, not a compliance score. Three uses stand out.

Watch the quiet domains. Every board tracks revenue. Far fewer track whether the company can prove it owns its IP, has the rights to its data, or can evidence the AI claims in its own marketing. Those are the domains that surface late and cost most, and a board that has never asked to see the evidence is trusting that it exists.

Treat readiness as a value lever, not a cost. The work that closes these gaps is the same work that makes the company better run. A portfolio company that builds diligence-ready from the start carries less risk, sells faster and defends a higher price. That is a return, not an overhead.

Ask for proof, not policy. The failure mode this whole paper guards against is the polished document that nothing stands behind. A board adds the most value here by asking the same question a buyer will ask: show me the evidence, not the policy.

What the radar does not tell you

The radar reads evidence, not intent. A company doing the right things without recording them will score low, and the score is a fair reading of what a buyer could verify, which is the thing that matters in a transaction. The answer is not to argue with the score. It is to build the evidence.

It is also not advice. It is a commercial and governance read that tells you where your exposure and your readiness diverge. It does not replace the legal, tax, accounting and regulatory advice a real transaction needs, and it does not certify that any individual system is compliant or safe. It tells you where to look and what it is likely to cost. The specialist work comes after.

And it is a point in time. Product features, law, data flows and customer expectations all move, and the shifting EU AI Act dates in this paper are a live example. Readiness is not a state you reach once. It is a discipline you keep.

Next steps

Score your company at ortent.co/tools/diligence-radar. It runs the nine domains, sets exposure against readiness, and shows you where the gaps are widest for a business like yours.

Then prove it. The companion prompt at ortent.co/tools/diligence-radar/prompt scores your company from your own evidence, one domain at a time, and will not lift a score without proof. It is the honest version of the exercise a buyer will run on you.

If you want to talk through what the radar shows about your company, or you want a second read before a raise or a sale, ortent.co/contact.