


ORTENT WHITEPAPER

Scale-Readiness is a State, Not a Motion

The five questions before you add anything on top.

By Andrew Wyatt, Founder, Ortent Advisory Ltd

August 2026 • A diagnostic for chairs, PE operating partners, and CEOs



Contents

Executive summary	3
Why the problem exists	4
The Five-Foundation Traffic Light	6
The five foundations in depth	11
Applied cross-industry examples	17
Before / after: the November board meeting	21
Where the diagnostic does not apply	23
Next steps	25

Executive summary

Growth-stage SaaS CEOs answer "are we ready to scale?" from feel. The board asks. The CEO looks at the last two quarters. Bookings are up. Pipeline looks healthier than it did in Q1. The team is motivated. The CEO says yes.

Nine months later the sales team has doubled, the new geography has opened, the partner programme has launched, and two of those three moves have quietly stalled. Sandbagged, in the polite language of the November board meeting.

The pattern is not a strategy problem. Every one of those moves was strategically sound. The pattern is that scale-readiness was measured as a forward motion. It is a state. States get evaluated as states.

This paper defines the state. Five foundations. Fifteen questions. A traffic light per foundation. A decision rule that treats amber and red as blocking, not as directional. If any single foundation scores amber, the honest answer to "can we scale?" is not yet. If any single foundation scores red, the answer is fix that first, and stop pretending the others matter.

The paper is not about slowing down. It is about knowing when adding is amplification and when adding is exposure. Growth-stage businesses that skip this diagnostic before adding on top of amber foundations do not fail. They erode. Six quarters of erosion is harder to unwind than one quarter of honest triage.

The instrument sits in Section 3. Read it. Score your business. Come back to Section 4 for what each foundation actually means when it is in the amber or red state and how you rebuild it.

Why the problem exists



Every board of a PE-backed SaaS business between £5M and £50M ARR is having a version of the same conversation this quarter. The CEO has three or four scale moves on the deck. New geography. Doubled sales headcount. First partner programme. First acquisition. First segment expansion. The chair asks the reasonable question, which is: are we ready?

The CEO answers from feel because the CEO has nothing else to answer from. There is no shared instrument. There is a board pack with green traffic lights next to the initiatives that leadership picked to report on, which is not the same thing as a state assessment. There is a set of KPIs that are usually about bookings and pipeline, which are outputs of past decisions, not readiness for future ones.

The chair, if the chair is honest, then does one of two things. They accept the CEO's answer, because the alternative is a hard conversation about what specifically is not ready, and the chair does not have the instrument to have that conversation either. Or they push back on the basis of intuition, which reads as scepticism about the CEO's judgment and starts a dynamic no board wants.

The gap is not competence. The gap is language. Scale-readiness is a state. States require assessment. Assessment requires an instrument. This paper is that instrument.

The other thing worth naming. The scale question is almost never asked when it should

be asked. It is asked at the annual planning meeting, when the deck is already written and the CFO has already modelled the new hires. By that point the honest answer is politically costly, so it does not get asked in the honest form. It gets asked as “any concerns?” which is not the same question.

The Five-Foundation Traffic Light is designed to be run early, quarterly, and by leadership rather than by the CEO in isolation. That mechanic matters as much as the questions. A state assessment run alone by the person who has to defend the state is not a state assessment. It is a defence.

Ninety minutes of leadership time, once a quarter. That is the ask. Everything else in this paper explains what to do with the ninety minutes.

The Five-Foundation Traffic Light



Read this section before anything else. Score your business here. Then read the rest of the paper to understand what the score means.

Five foundations. Three questions per foundation. Fifteen questions total. Each question is a **RED** / **AMBER** / **GREEN** mark against a specific state, not against a target.

Score each question honestly. **Green** means the state is intact and evidenced. **Amber** means the state has drifted but is recoverable inside a quarter. **Red** means the state is not in place, and adding on top of it will make matters worse rather than better.

The overall verdict is not a total. It is a decision rule:

- **All five foundations green:** Scale-ready. Add.
- **Any single foundation amber:** Not yet. Fix the amber foundation. Revisit in ninety days.
- **Any single foundation red:** Not yet, and the red foundation is the only thing that matters. Fix it. Everything else waits.

The reason the rule is not a total is that these foundations do not average. Amplification depends on the weakest, not the mean. A business with four green foundations and one red foundation is not four-fifths ready. It is a business about to spend a quarter building on top of a foundation that will crack.

FOUNDATION One

Commercial Model

Whether your pricing, packaging, and rev-share structure can hold the additional weight scale will place on them.

Q1 Can your CRO explain, in one sentence, the primary reason a customer chooses you over the closest alternative? And can two other members of your leadership team give the same one-sentence answer without prompting?

Q2 Do you have one pricing structure, not three or four running in parallel? If you have more than one, is the reason for each documented and defended, or is it because different sales cycles produced different structures and nobody has consol-

dated?

Q3 If you doubled the sales team tomorrow, would the new hires close deals at the same discount profile as the current team, or would discounting drift as green reps chased quota? Can you point to the mechanism that prevents that drift?

FOUNDATION Two

Go-to-Market

Whether your motion to identify, qualify, and win customers is defensible under load.

Q4 Would a commercial hire understand your ideal customer profile from your documented ICP in fifteen minutes, or would they need three coaching sessions with the CRO to piece it together from lived context?

Q5 When a deal loses, do you know within a week whether it was product, pricing, positioning, or timing? Or does it get logged as "competitive loss" and never revisited?

Q6 Is your CRO's forecast built bottom-up from named opportunities with dated close plans, rather than top-down from an ARR growth target the board asked for?

FOUNDATION Three

Deployment

Whether your ability to turn a signed customer into a live customer is repeatable and predictable.

Q7 If a customer signs today, can you tell them the exact date they will be live, or is the honest answer "it depends on your team's availability"?

Q8 Do your implementations follow a defined playbook, or does each customer implementation reinvent the shape based on who is running it?

Q9 Can a new implementation lead take a customer through go-live end-to-end without the customer feeling the handover, or does the founder or head of Delivery still need to be in the room for the harder calls?

FOUNDATION Four

Retention and Support Discipline

Whether your ability to keep customers, expand them, and support them can carry additional volume without breaking.

Q10 Is your net revenue retention driven by upsell into your existing base, rather than by not-churning?

Q11 Can your customer success team name, right now, the three customers most at risk this quarter, why they are at risk, and what is being done about it? And do those three names match what the CRO would say if asked separately?

Q12 When something breaks in production, does the person who found it own the fix, without a three-team triage before someone takes it on? And does that fix loop back into the product roadmap without requiring executive escalation?

FOUNDATION Five

Organisation

Whether the operating model itself has the ownership, accountability, and succession discipline to hold under scale pressure.

Q13 Does every operating area of your business have one named person, who is accountable to the board, not just to the CEO?

Q14 In the last twelve months, has the board changed a decision because of what an operating owner said, rather than what the CEO said? If the answer is no, the question is whether operating owners have real voice at the board.

Q15 Have you identified and documented your top three single-person operating dependencies, with either succession plans or knowledge-transfer commitments against each? If any of those three people were unavailable for six weeks, would the business hold?

How to score honestly


The instrument only works if the scoring is honest. Two ways to make it honest.

First, do not score alone. Have your CFO, CRO, and head of Delivery run the same fifteen questions independently, without conferring. Compare. The questions where you scored green and they scored amber are the interesting ones. That is where your view of the state and the state itself have diverged.

Second, if you cannot cite specific evidence for a green mark, downgrade it to amber. The instrument treats “we probably do this” as amber. If you have to think about whether the state is in place, it is not fully in place.

The instrument is unforgiving on purpose. Growth-stage businesses drift into amber quietly. The instrument names the drift before it becomes a scaling failure.

The five foundations in depth



Section 3 gave you the scoring instrument. This section explains what each foundation actually is, why it matters at scale, and what rebuilding an amber or red foundation actually looks like.

Read the foundation your business scored weakest first. Fix that. Revisit the diagnostic in ninety days.

Foundation One. Commercial Model.

The commercial model is the answer to why anyone pays you, how much they pay, and how the payment structure holds up when you increase volume.

At sub-£5M ARR, commercial model is often loose by necessity. The business is finding product-market fit. Pricing is being tested. Rev-share structures with early partners are custom. That is fine at that stage.

By £10M ARR, commercial model should be one structure with documented exceptions, not three structures running in parallel. By £20M, the CRO should be able to explain the reason a customer chooses you over the closest alternative in one sentence, and that sentence should be the same one every member of the leadership team gives.

Commercial model becomes the amber foundation for growth-stage businesses in three specific patterns.

First, the business grew by acquisition and never consolidated the pricing structures in-

herited from acquired products. Two products, two pricing philosophies, one CRO trying to sell both. New reps close whichever product feels easier to explain. Discounting drifts because there is no defensible primary product.

Second, the business grew by adding segments and priced each segment differently for defensible reasons at the time, then never revisited whether those reasons still held. Enterprise pricing, mid-market pricing, self-serve pricing, all live. A sales rep talking to a mid-sized enterprise buyer does not know which structure to use.

Third, the business has been reactive to competitive pricing pressure for eighteen months and has quietly given up on defending a primary price point. Discounts show up on every deal above a threshold. The finance team stops flagging them because everyone flags them.

Rebuilding an amber commercial model is a ninety-day exercise, not a six-month one. The move is not to redesign the pricing. It is to consolidate to one primary structure with documented, defensible exceptions, and to give the CRO the language to hold price on the primary structure.

A **RED** commercial model is different. Red means the primary reason a customer chooses you cannot be articulated in one sentence, or the articulation the CRO gives is not the articulation the customer would give. That is a positioning problem, not a pricing problem. Positioning problems take two quarters to fix, not one. Adding sales headcount on top of a red commercial model amplifies the positioning gap. New reps do not have the CEO's context and default to price when the story does not land. Discounting accelerates.

If the commercial model is red, everything else in the paper waits. Fix positioning first.

Foundation Two. Go-to-Market.

Go-to-market is not the sales team. Go-to-market is the motion the sales team runs, and whether that motion is defensible without the founder or CRO in the room.

The scale test on go-to-market is not "can you close deals?" It is "can a new hire close deals at the same profile as the current team, and can the CRO forecast the deals bottom-up without wishful thinking?"

The three questions in Section 3 name the three most common places go-to-market slips into amber.

ICP that lives in the CRO's head rather than in a document new hires can read is the first slip. It looks fine while the sales team is small enough that the CRO onboards each hire personally. It reveals itself the moment you double the team. Sales cycles get long. Deals fall out of the pipeline late. The CRO starts saying the new hires need more time.

Loss analysis that logs deals as "competitive loss" without diagnosis is the second slip. This one hides in plain sight. Every business has some deals labelled competitive loss. The question is whether that label represents an actual finding or a shorthand for "we do not know why." If it is the second, the business is not learning from losses, and the losses will repeat when scale amplifies the pipeline volume.

CRO forecast built top-down from board targets rather than bottom-up from named opportunities is the third slip, and the most dangerous. Top-down forecasts get met by discounting, sandbagging, or, most commonly, by quietly moving deals across quarters until the number lands. That is not a forecast, that is a narrative. Under scale, the narrative fractures within one quarter and the CRO loses credibility with the board.

Rebuilding an amber go-to-market takes one quarter of disciplined loss analysis, a written ICP that new hires can run without the CRO in the room, and a rebuilt bottom-up forecast the CRO can defend in front of the CFO. The rebuild is not glamorous. It is discipline more than design.

Red go-to-market is when the motion itself is broken. Deals close because of a specific relationship rather than because the motion works. Losing that relationship or the person who holds it loses the deal. If your go-to-market is red, no scale move works because the motion cannot be replicated. Fix the motion before adding anything on top.

Foundation Three. Deployment.

Deployment is the boring foundation. It is also the one growth-stage businesses most consistently underrate.

At £10M ARR, most businesses can implement customers. By £20M ARR, most cannot promise the customer a live date, and the honest answer if the sales team pushed for one would be "it depends on your team's availability, which I cannot predict." That is amber deployment.

The three questions in Section 3 test three specific weaknesses.

Q7 tests whether implementation timeline is a knowable number or an emergent property. Emergent property is amber. Knowable number, with a defensible variance range, is green.

Q8 tests whether implementation shape is a playbook or a personality. Playbook is green. Personality is amber, and it degrades to red as soon as the personality leaves.

Q9 tests whether the founder or head of Delivery is still in the room for the harder calls. If they are, the business has not decoupled the executive from delivery. That is amber. Scaling the sales team on top of amber deployment produces a specific failure mode: signed deals that cannot be delivered, customer success ARR that grows slower than bookings ARR, and a slow-motion churn spike two years later when the delivery gap catches up with the customer base.

Rebuilding amber deployment is playbook work. Document the implementation motion. Test it with a new implementation lead running an actual customer. Revise. Do that twice. That is a quarter of work and it is high-leverage.

Red deployment is when signed customers do not go live at all, or go live at a rate materially below what the sales pipeline assumes. Red deployment blocks every scale move. There is no point adding sales headcount, geographies, or partners on top of a business that cannot deliver what it sells. Fix delivery first, and be honest that the scale question is off the table for at least one quarter.

Foundation Four. Retention and Support Discipline.

Retention and support discipline is the foundation most PE operating partners fail to diligence properly in the deal, and the one that most consistently determines whether the eventual exit is a good one.

The three questions in Section 3 test three states.

Q10 tests whether net revenue retention growth is driven by upsell or by simply not churning. NRR at 110% because customers grow their usage inside your product is a different signal than NRR at 110% because you kept everyone at the same level and none of them left. The first is scale-ready. The second is fragile.

Q11 tests whether the customer success team knows which customers are actually at risk, and whether that view matches the CRO's view. Divergence between customer success and the CRO on at-risk customers means one of the two functions has bad information. Under scale, that gap widens.

Q12 tests whether production issues get resolved without a triage that requires three teams and executive escalation. If they do, support is playing catch-up. Under scale, the catch-up compounds and the product roadmap gets hijacked by production firefighting.

Rebuilding amber retention and support discipline is a two-quarter exercise, not a one-quarter one. The move is to align customer success and the CRO on a shared at-risk view (usually via a weekly review with a joint list), to reintroduce upsell motion into customer success (which most businesses lose as they scale), and to give production support a mandate to own fixes without triage. None of this is glamorous. All of it holds under scale.

Red retention discipline is when churn is running above 15% annually in a business that should be under 8%. Red retention is a growth question in disguise. You cannot outrun a leaky bucket by pouring more in. Fix retention. Everything else follows.

Foundation Five. Organisation.

Organisation is the foundation the other four sit on. It is also the one CEOs most consistently score themselves higher on than the leadership team does.

The three questions in Section 3 test three things.

Q13 tests whether ownership is distributed. Every operating area with one named owner accountable to the board, not just to the CEO. The test is not the org chart. The test is whether operating owners have voice at the board.

Q14 tests whether that voice is real. Has the board changed a decision because of what an operating owner said in the last twelve months? If the answer is no, the operating owners are advisors to the CEO, not accountable to the board. That is amber organisation. It is remediable, and the remediation is procedural, not personnel: give operating owners actual board time and let them present their own areas rather than have the CEO present on their behalf.

Q15 tests succession. Not the formal succession plan for the CEO. The three single-person dependencies that would break the business if unavailable for six weeks. Every business at £20M ARR has at least three of these. The question is whether the business has named them, documented what each person knows, and either built succession or accepted the risk knowingly.

Rebuilding amber organisation is a quarter of governance work and a quarter of documentation work. It is not glamorous. It is the foundation the other four foundations stand on. If organisation is amber, the fixes to the other four foundations do not compound. If organisation is green, every fix to the other four compounds.

Red organisation is rare and it is fatal for scale. Red means the operating model itself lives in the CEO's head. Under scale, the CEO becomes the bottleneck for every decision that used to route through them by choice. The CEO burns out inside two years and the business

either sells at a discount or resets. If your organisation is red, the answer is not to scale. It is to build the operating model before scale is even discussed.

Cross-foundation patterns worth naming

Two patterns show up often enough to name.

Pattern one: all amber, no red. Businesses that score every foundation amber but nothing red are the most common growth-stage pattern. The instinct is to score this as “close to scale-ready, one more quarter of work.” That is wrong. All-amber is not close to green. All-amber compounds. Fixing one foundation while the other four sit amber is treading water, because the four amber foundations continue to erode while you are fixing the fifth. All-amber businesses need to pick the two foundations most load-bearing under their specific scale move and fix both simultaneously.

Pattern two: green everywhere except organisation. This is the founder-led business at £15M–£25M ARR. Commercial model works, go-to-market works, deployment works, retention works. But organisation is amber because everything routes through the founder. Under scale, this business does not fail. It stalls. The founder cannot grow the business past £25M–£30M ARR without breaking. The paper’s specific recommendation here is that the scale move needs to be paired with an organisational rebuild. Adding sales without adding operating owners produces a two-year stall.

Applied cross-industry examples



Three examples, three different scale moves, three different diagnostic outcomes. All names and specifics generalised. The patterns are drawn from actual engagements.

Example one. PE-backed vertical SaaS at £30M ARR. US expansion.

The business was three years post-Series C, growing 40% year-on-year in its European home market, with a PE investor pushing for US expansion. The board was aligned. The CFO had modelled the US operation. The CEO was ready.

The diagnostic scored green on Deployment, green on Retention and Support Discipline, green on Organisation. Go-to-Market scored green with one amber note on ICP documentation. Commercial Model scored amber. Specifically, the business had three parallel pricing structures: an enterprise structure, a mid-market structure inherited from an acquisition, and a self-serve structure. The CRO could not articulate the primary reason a customer chose them in one sentence, and the leadership team gave three different one-sentence answers when asked separately.

The diagnostic's verdict was clear. Do not open the US on an amber commercial model. Fix pricing and positioning first. Six months of consolidation work, one primary structure with documented exceptions, one sentence the leadership team all shared. The US opened on time twelve months later, delivered on plan in the first year, and by the second year was outpacing the European base.

The counter-factual is unknowable but the pattern is well-observed. Opening the US on amber commercial model would have produced US reps discounting to close deals, US customer profile drifting from European customer profile, and a set of US customers whose economics never matched the ones the CFO had modelled. Two years of remediation to unwind rather than six months of preparation to prevent.

Example two. Digital health platform at £18M ARR. First partner programme.

The business was five years old, growing 60% year-on-year, entirely first-party sales. The CEO wanted to launch a partner programme with a global systems integrator to accelerate reach into the largest health systems. The board was cautious but not opposed.

The diagnostic scored green on Commercial Model, green on Go-to-Market, green on Retention, amber on Organisation, and red on Deployment.

Deployment was red because the founder was still personally involved in every major implementation. Not because the founder was clinging. Because the technical complexity of onboarding a health system was genuinely nuanced, and the head of Delivery had joined

only nine months earlier and was still learning. Signed customers went live, but on timelines that varied from four months to fourteen months depending on which health system, which internal champion, and how much the founder personally shepherded the process.

The diagnostic's verdict was that a partner programme built on top of red deployment would fail specifically. The partner would sell deals, the business would sign them, and then the business could not deliver them on any predictable timeline. The partner would lose confidence within twelve months. The customers signed via the partner would churn or renegotiate. The programme would be quietly wound down at eighteen months.

The recommendation was to fix deployment before adding a partner. Nine months of play-book work. Reduce the founder's involvement from every implementation to escalation-only. Document the health-system onboarding pattern such that a new implementation lead could run one end-to-end. Once deployment was green, launch the partner programme.

The business took the recommendation. The partner programme launched twelve months later than originally planned. In the two years after launch, that alliance drove 45% of new-name ARR. In the counter-factual, the programme would have launched on schedule and delivered a fraction of that value while damaging the partner relationship.

Example three. AI-adjacent scale-up at £8M ARR. Doubling sales team.

The business was eighteen months post-Series A, growing 90% year-on-year, and the board had approved a plan to double the sales team from twelve to twenty-four reps in a single year. The CRO was excited. The CFO was already recruiting.

The diagnostic scored amber on Go-to-Market, red on Organisation, and green on everything else.

Organisation was red for a specific reason. The operating model lived in the CEO's head. Every meaningful decision routed through the CEO by default, not by design. There was no head of Operations. The head of Delivery reported to the CEO. The CRO reported to the CEO. The CFO reported to the CEO. Marketing reported to the CEO. In practice, seven people reported to the CEO and one of them was doing the CEO's job on evenings and weekends without title or authority.

The diagnostic's verdict was that doubling the sales team on top of red organisation would break the CEO within eighteen months and, before that broke, would starve the twelve new reps of the coaching, forecast discipline, and deal support that the twelve current reps were absorbing from the CEO informally.

The recommendation was to postpone the doubling by a quarter, hire a head of Operations

to take five of the seven direct reports off the CEO, and let the CRO run the sales function without the CEO in every deal review. Once organisation was green, the CRO could double the team knowingly.

The business took the recommendation. The postponement was uncomfortable. The board asked hard questions about the delay. Six months later, the head of Operations was in place, the CEO's calendar had recovered, the CRO was running sales with real authority, and the sales team doubled through the following year without the coaching gap that the diagnostic had warned about. By the second year, the business had scaled past £20M ARR.

The counter-factual, again, is unknowable but well-observed. Doubling the sales team on red organisation produces a specific failure mode: half the new hires ramp slowly because the CRO cannot coach twenty-four reps the way the CRO could coach twelve, the CEO becomes the bottleneck for every non-standard decision, and the business grows for three quarters and then stalls in the fourth as the operating model catches up with the go-to-market.

Before / after: the November board meeting



The CEO of a PE-backed B2B SaaS business at £22M ARR opened the November 2025 board meeting with the annual plan. Four scale moves. Open a second European market. Launch a partner programme. Add a new product line adjacent to the core. Double the sales team.

The chair asked the reasonable question. Are we ready?

The CEO said yes.

The board approved the plan. The chair, in the back of his mind, was not sure but did not have the language to push.

August 2026, nine months later. Two of the four moves had quietly stalled. The second European market was sandbagged, in the polite language of the quarterly update. The new product line had not launched because engineering had been pulled into remediation on the core product. The partner programme was live but delivering less than a third of what the plan modelled. Only the sales team doubling had held, and even that was ramping slower than plan.

The chair had a different conversation with the CEO in early September 2026. Not about the plan. About the underlying state. The chair asked the CEO to run the Five-Foundation diagnostic before the November board meeting. The CEO agreed, with reservations.

The CEO ran it with the CFO, CRO, and head of Delivery independently, without conferring, as the instrument specifies. The comparison surfaced two things.

First, three of the five foundations scored amber where the CEO had scored green. Specifically, Commercial Model (three parallel pricing structures nobody had consolidated), Deployment (founder still in the room for the harder health-system onboardings), and Organisation (five direct reports to the CEO where two should have gone through a COO layer).

Second, none of the four moves the CEO had originally proposed for the year would have compounded on top of those three amber foundations. Every single one would have either stalled or delivered at a fraction of the plan. In fact, three of the four had already done exactly that.

The CEO took the diagnostic honestly. The November 2026 board meeting opened not with the annual plan but with a state-of-the-business assessment. The CEO showed the traffic lights. The CEO named the three amber foundations. The CEO's proposal was not to add anything for the coming year. It was to fix Commercial Model in ninety days, Deployment in six months, and Organisation across the full year. The scale moves would come back on the agenda when the foundations were green.

The board's response was not scepticism. It was relief. The board had known something was off. The chair had not had the language. The diagnostic gave them the language.

Eighteen months later, in May 2028, the board approved a scale plan that was leaner than the original November 2025 plan. One geography, not two. One product extension, not one plus doubling. A partner programme launched into a green Deployment foundation with a rebuilt Commercial Model. The plan delivered on time and on number in the first year.

The lesson is not that the CEO had been wrong in November 2025. The lesson is that the CEO had answered the wrong question. The chair had asked “are you ready to scale?” and the CEO had heard “are you confident?” Confidence is not scale-readiness. The diagnostic separates the two.

Where the diagnostic does not apply

The Five-Foundation Traffic Light is designed for growth-stage SaaS businesses between roughly £5M and £50M ARR that are considering a specific scale move. Four situations where the diagnostic does not apply cleanly.

Turnaround situations

In a turnaround, the foundations get evaluated in a different order. Retention and Support Discipline typically becomes the leading indicator, not the trailing one. The diagnostic, run as written, will produce a lot of red and the recommendation will be “fix everything first.” That is not useful advice in a turnaround. In a turnaround, the instrument to use is a stabilisation plan focused on stopping the bleed. Come back to the Five-Foundation instrument once the business is stable.

Pre-Series-A businesses

Businesses under £5M ARR do not have a Commercial Model in the sense this paper uses the term. They are still finding it. Scoring Commercial Model on a Series-A business will produce misleading results. The instrument to use pre-Series-A is a product-market-fit assessment, not a scale-readiness one.

Deep tech with long R&D cycles

Businesses with R&D cycles measured in years rather than quarters have a different Deployment definition. The instrument's Deployment questions assume a deploy-to-customer motion measured in weeks. For deep tech, adapt those questions to fit the actual cycle. The other four foundations translate cleanly.

Founder-led businesses at sub-£5M ARR

In a business at that stage, the founder is the operating model. Scoring Organisation on a founder-led business under £5M will always come out amber or red, and the recommendation will always be “hire a COO,” which is not always the right move at that stage. The instrument is not designed for that stage. Come back to it at £8M–£10M ARR.

Two additional situations that do not disqualify the instrument but are worth flagging.

Recent acquisitions. Businesses inside the first eighteen months of a significant acquisition should expect amber marks on Commercial Model (integration of pricing structures) and Organisation (integration of leadership). Those amber marks are situationally normal

and time-bounded. Score them, but the ninety-day fix window may need to extend to six months. Do not treat integration-driven amber as a scale-blocker unless it has persisted beyond eighteen months, at which point it is no longer integration and is now structural.

Businesses with a specific single-product bet. If the business is a single-product bet with intentional operational simplicity as a strategic choice, some of the foundations may score amber for reasons that are strategic, not accidental. Retention and Support Discipline in particular. If the bet is deliberate and defensible, score amber, note the reason, and treat it as a known constraint rather than a foundation to fix. The point of the instrument is to name state, not to enforce uniformity.

Next steps

Read the paper. This paper explains why the diagnostic scores what it scores, walks through each foundation in depth, gives three worked cross-industry examples, and names where the instrument does not apply.

Run the diagnostic. Score your business against the fifteen questions in Section 3. Score honestly. If you cannot cite specific evidence for a green mark, downgrade it to amber.

Run it with your leadership team. Have your CFO, CRO, and head of Delivery score the fifteen questions independently. Compare. The questions where you scored green and they scored amber are the interesting ones. That is where your view of the state and the state itself have diverged.

Run the interactive version.

The Ortent Five-Foundation Diagnostic lives at ortent.co/tools/five-questions/prompt. Fifteen questions, traffic-light output, foundation-by-foundation verdict, recommended sequence.

Book a working session. If the diagnostic surfaced something you want a second pair of eyes on, book a session at ortent.co/contact. Forty-five minutes. No pitch, no deck, no slides. Just the paper, your score, and the specific work required to move the amber foundations to green.

Subscribe to The Growth Chair. Weekly notes for founders, boards, and investors trying to make growth-stage SaaS actually scale. One issue a week.

That is the whole services page.

ortent.

Andrew Wyatt is founder of Ortent Advisory Ltd, an operator-led advisory for growth-stage SaaS businesses in Life Sciences, Digital Health, and AI markets. Four exits across three decades, including Lotus to IBM, Paragon to Phone.com, Apertio to Nokia at \$240M, and Clearswift to Lyceum. Board Advisor, Non-Executive Director, and Fractional Chief Growth Officer. London.

ortent.co • linkedin.com/in/andrewcwyatt • © 2026 Ortent Advisory Ltd