HADRIAN

EXTERNAL EXPOSURE MATURITY MODEL

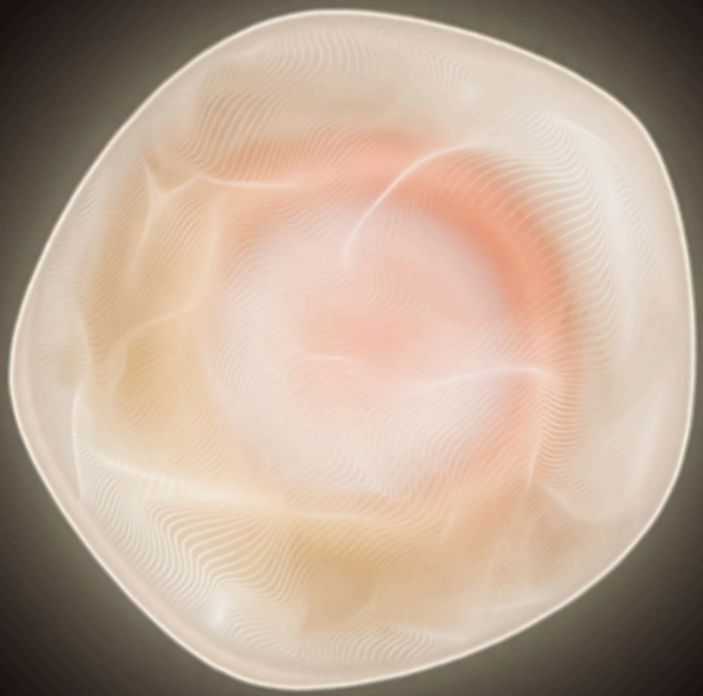

Where is your exposure management programme today, and what does meaningful advancement actually look like?

THE EXPOSURE PROBLEM PROGRAMMES ARE NOT SOLVING

Most organisations have invested heavily in security tooling. They run vulnerability scanners, operate SIEM platforms, conduct annual or quarterly penetration tests, and maintain a vulnerability management programme with defined owners and SLAs. And yet the same pattern repeats: exploitable exposure persists, remediation queues grow faster than they shrink, and when something goes wrong, the finding that enabled it had often been sitting in a report for months.

The issue is not the absence of security activity. It is the absence of a programme structure that connects activity to outcomes. Scanning generates findings. Findings generate tickets. Tickets close, or they do not. But the question that matters, whether the organisation's exploitable exposure is actually declining, often goes unanswered.

This is the gap the External Exposure Maturity Model is designed to address. It describes four recognisable operating postures, each defined not by the tools present but by the structure of how an organisation discovers exposure, confirms it is real, and closes it before an attacker can act. The model gives security leaders a precise language for where their programme is today, what is structurally holding it there, and what has to change to advance.

How this aligns with CTEM

Continuous Threat Exposure Management (CTEM) is a five-phase framework, created by Gartner, that provides structure for moving from reactive vulnerability management to continuous, outcome-driven exposure reduction. The phases are: Scope, Discover, Prioritise, Validate, and Mobilise. This model maps directly to CTEM maturity: Stage 1 and 2 organisations are typically executing the early phases, while Stage 3 and 4 organisations have closed the loop through Validation and Mobilisation. The phase where most CTEM implementations stall is Validation: confirming that findings represent genuinely exploitable risk, rather than theoretical severity. That is the structural inflection point this model is built around.

WHY PROGRESSION MATTERS

The model maps four stages of external exposure management, from reactive and undiscovered to continuous and autonomous.

What advancement actually delivers

The gap between Stage 1 and Stage 4 is not a matter of resources. It is a matter of structure: how an organisation discovers exposure, confirms it is real, and closes it before an attacker can act.

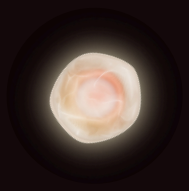

STAGE 1 → STAGE 2

Running Blind to Better Scope, More Hope

Attack surface monitoring coverage rises from below 20% to 40–60%. True-positive rates improve from under 10% to 15–25%. That means fewer than 1 in 10 alerts requiring investigation becomes 1 in 4 to 6. Mean remediation time drops from 90+ days toward 45–90 days. The practical effect: the security team stops spending most of its time on findings that turn out to be irrelevant.

STAGE 2 → STAGE 3

Better Scope, More Hope to Connecting the Dots

Coverage reaches 75–90% of the surface. True-positive rates rise to 40–60%, meaning close to half of all findings are confirmed exploitable before anyone investigates them. MTTR compresses to 15–45 days. SLA compliance reaches 60–80%. The remediation queue gets smaller and more accurate: fewer total findings, higher confidence in each one, and a measurable reduction in exploitable exposure rather than just tickets closed.

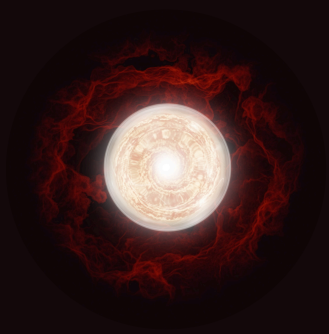

STAGE 3 → STAGE 4

Connecting the Dots to Clear Picture

Coverage exceeds 95%. True-positive rate rises above 70%. MTTR falls below 7 days, with critical exposures closed within 48 hours. SLA compliance above 90%. The benchmark that matters most: when a new exploit is disclosed, the organisation can confirm within hours whether it is affected, not days. At Stage 3, that answer took a team meeting and manual investigation. At Stage 4, the programme answers it automatically.

Use it as a diagnostic, a roadmap, or a shared language with leadership. It does not prescribe a toolset or a timeline. No organisation needs to reach Stage 4 for every asset. The right question is not how advanced are we; it is whether the programme is operating at the level the risk profile requires.

SECTION 2

THE MODEL AT A GLANCE

Knowing where your programme stands is the starting point for meaningful advancement. This section will give you a picture of the specific capability gaps and measurable outcomes that define each stage.

The Lens Table

For each lens, mark the column that most honestly describes your programme today. If your answers land consistently in one column, that's your current stage. If spread across columns, your floor is the leftmost column you marked.

LENS	STAGE 1: RUNNING BLIND	STAGE 2: BETTER SCOPE, MORE HOPE	STAGE 3: CONNECTING THE DOTS	STAGE 4: CLEAR PICTURE
Operating rhythm How frequently does your programme engage with exposure?	REACTIVE	PERIODIC	CONTINUOUS	AUTONOMOUS
Relationship to risk How well do you understand what you're actually exposed to?	EXPOSED	AWARE	VALIDATED	PROACTIVE
Security stance How does your programme position itself relative to attackers?	REACTIVE	STRUCTURED	PROACTIVE	PREDICTIVE
Validation posture How does your programme treat: do our controls actually work?	ASSUMED	TESTED	VALIDATED	CONTINUOUS

The Capability Matrix

Five operational dimensions, four stages. Use this to spot which dimension is your limiting factor.

CAPABILITY	STAGE 1: RUNNING BLIND	STAGE 2: BETTER SCOPE, MORE HOPE	STAGE 3: CONNECTING THE DOTS	STAGE 4: CLEAR PICTURE
Operating model	Event-driven; exposure found via incidents or audits	Structured programme; defined ownership; CTEM phases introduced	Closed-loop; validation outcomes drive remediation	Offensive by default; adaptive; minimal manual orchestration
Discovery	Incomplete static inventory; periodic scanning of known assets	EASM deployed; cloud, SaaS, APIs, third-party in scope	Continuous monitoring; near-real-time drift detection	Adaptive; correlates identity, infrastructure, SaaS automatically
Prioritisation	CVSS and vendor scores; limited business context	Risk-based with business criticality factored in	Exploitability and attack-path context; fewer, higher-confidence issues	Predictive and context-aware; priority shifts dynamically
Validation	Annual pentests or ad hoc; exploitability assumed, not confirmed	Scenario-based periodic testing within defined scopes	Continuous automated validation; blind windows eliminated	Autonomous adversarial emulation; attack chains adapt to change
Automation	Manual, ticket-driven; siloed tools	Integrated workflows; standardised handoffs	Validation triggers remediation; SOC/IT integrations active	Largely autonomous; humans focus on strategic decisions

Quantitative Benchmarks

These figures are drawn from Hadrian's analysis of exposure management programmes across +300 organisations. They represent observed ranges, not targets, your programme may sit above or below them depending on environment complexity, sector, and starting point.

METRIC	STAGE 1: RUNNING BLIND	STAGE 2: BETTER SCOPE, MORE HOPE	STAGE 3: CONNECTING THE DOTS	STAGE 4: CLEAR PICTURE
Attack surface under continuous monitoring	< 20%	40–60%	75–90%	> 95%
Alert true-positive rate	< 10%	15–25%	40–60%	> 70%
Mean time to remediation (MTTR)	> 90 days	45–90 days	15–45 days	< 7 days (critical: < 48h)
SLA compliance	< 20%	< 40%	60–80%	> 90%
Findings validated as exploitable before escalation	Rarely	Selectively	Consistently	Continuously and automatically

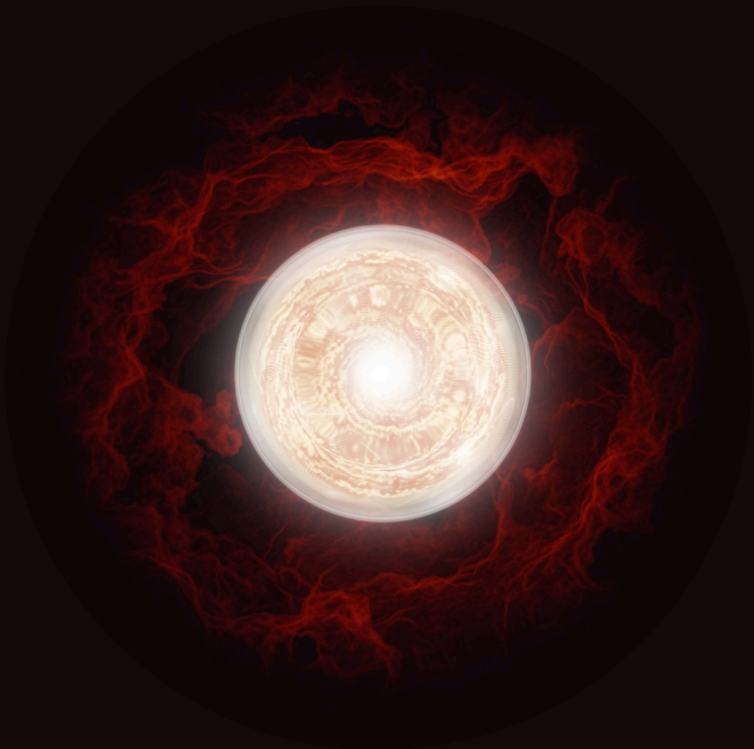

ASSESS YOURSELF:

Take the interactive self-assessment

DISCOVER YOUR MATURITY STAGE

THE FOUR STAGES

Each stage is a recognisable operating posture, not a score, not a grade. Read these as portraits of how organisations actually work. If a description makes you uncomfortable, that's probably the right one.

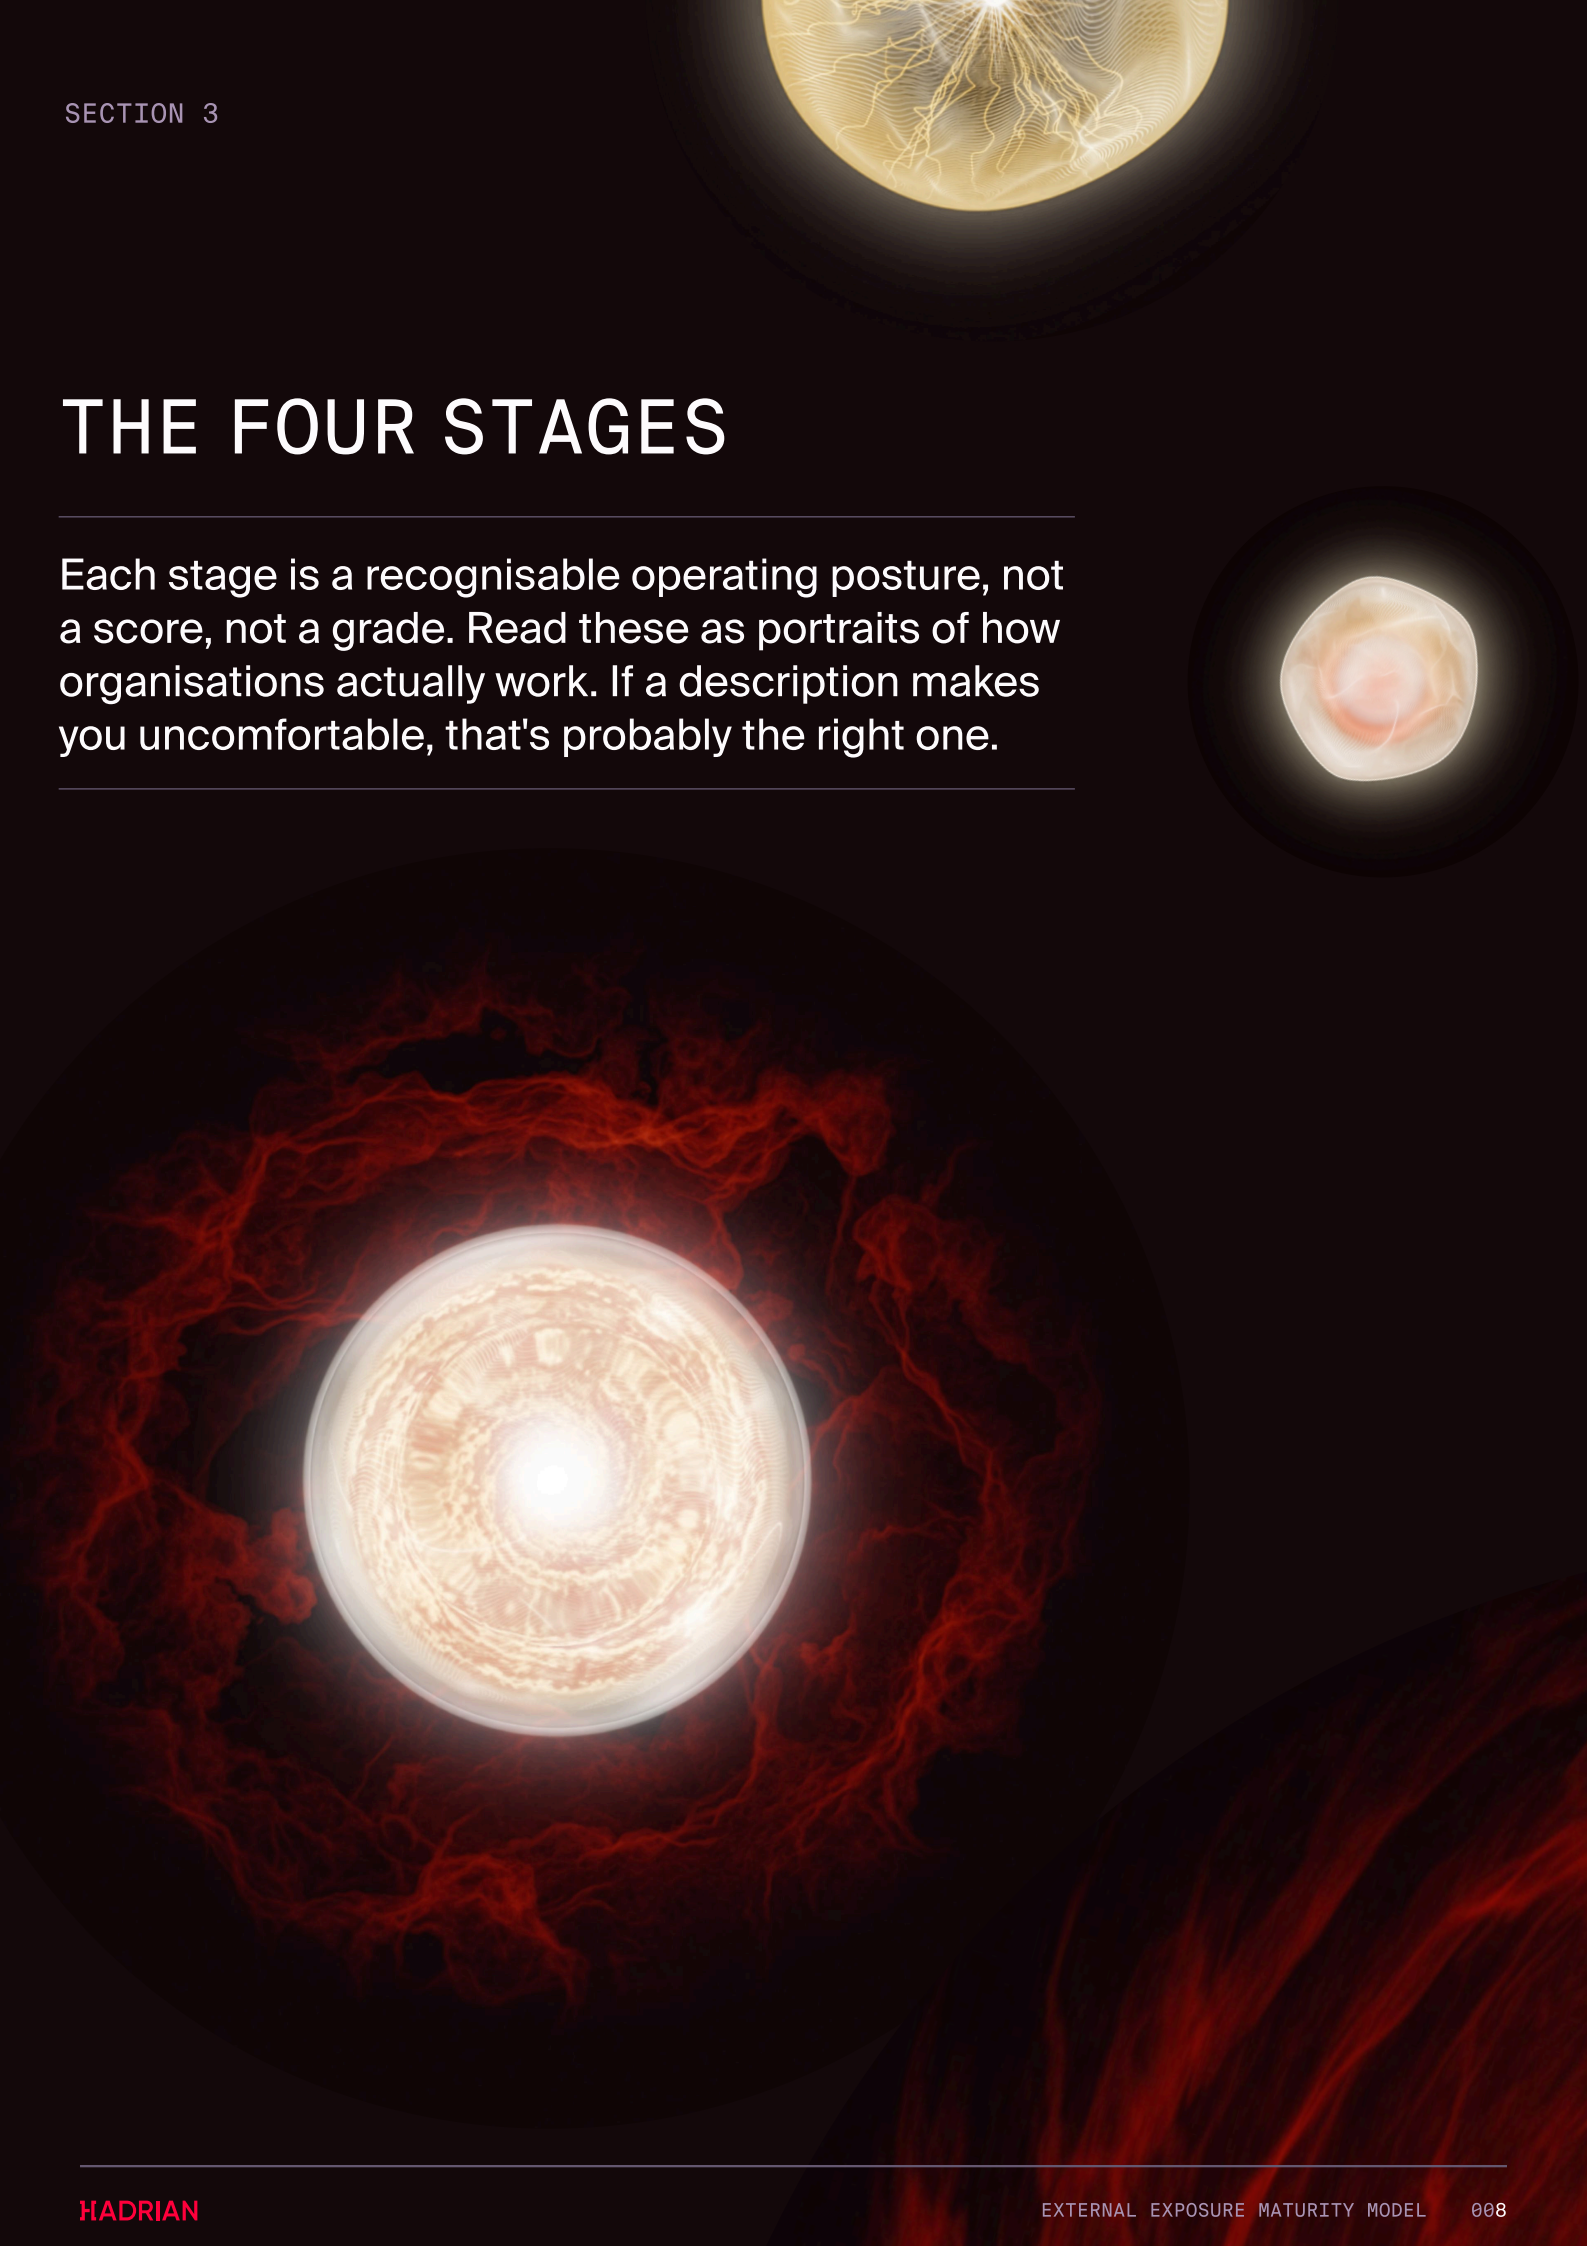

STAGE 1

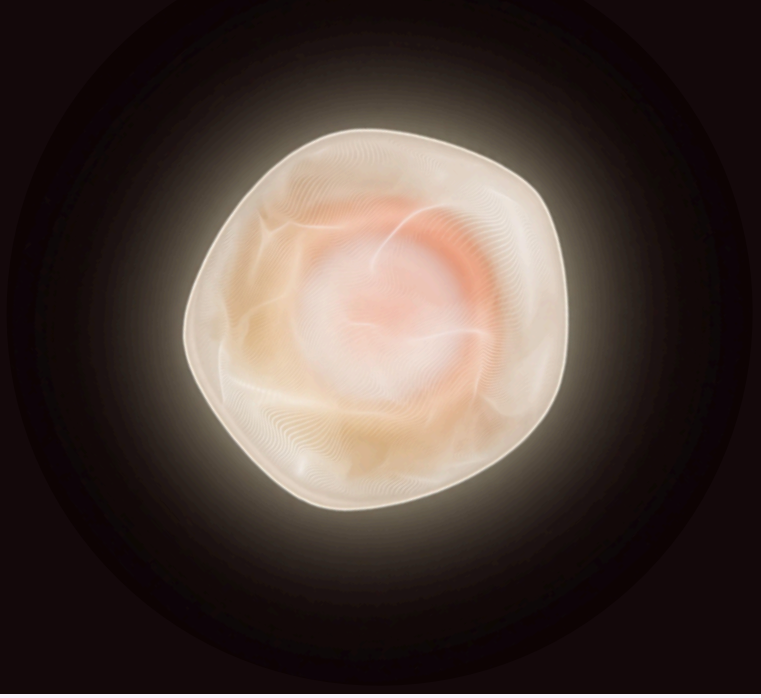

RUNNING BLIND

"We find out about exposure when something goes wrong."

The name is blunt because the experience is blunt. The organisation has no reliable picture of its external attack surface. Assets appear in conversations after incidents, not before them. The security team knows it's missing things; it just doesn't know what, or where, or how significant. Scanning happens, but the output is a list of severities that no one fully trusts and that grows faster than anyone can work through.

What separates Stage 1 from Stage 2 isn't tooling. Many organisations at this stage have scanners, SIEMs, and ticketing workflows. What's absent is structure: no systematic discovery cadence, no defined ownership model, no consistent method for confirming whether findings represent real risk. The activity exists. The programme doesn't.

How teams describe it

"We're always firefighting. We never get ahead of things."

"The pentest report landed and we still have half of it open from last year."

"I don't actually know how many internet-facing assets we have. No one does."

What this looks like organisationally

Security meetings are dominated by the most recent incident or audit finding. The CISO's board update is built around compliance status, not exposure. When a new critical CVE drops, the first question is 'do we have that software?' Nobody is entirely sure of the answer. Remediation conversations start with 'whose team is that?' more often than 'how exploitable is this?'. The pentest report from eight months ago is still open in a spreadsheet that has never been fully actioned.

The relationship between security and engineering is transactional at best. Security raises findings; engineering disputes severity or deprioritises against product work. Without a consistent method for confirming exploitability, security cannot make a compelling case for urgency. The team is reactive and spends a disproportionate amount of time responding to whatever made the news this week rather than to the organisation's actual exposure.

What keeps organisations here

The most common reason isn't budget. It's that the organisation has defined security as a compliance function rather than a risk management function, and compliance doesn't require knowing what's exploitable, only what's documented. This creates a structural ceiling: the programme produces evidence of activity rather than evidence of reduced risk, and leadership doesn't yet know to ask for the difference.

The operational bottleneck at Stage 1 is signal-to-noise collapse. Scanners generate thousands of findings with no mechanism to distinguish real risk from background noise. The result: the team's entire capacity goes to triage, none to remediation. Volume creates the impression of work; it does not create security outcomes.

Ownership is the third blocker. Exposure lives at the intersection of security, IT, and engineering. When nobody owns the full picture, nothing gets treated as urgent until something breaks.

The moment of readiness

Organisations become ready to advance when something makes the gap undeniable, typically an incident that revealed an asset nobody knew existed, or a board question that exposed the distance between reported posture and actual posture. The shift from compliance to risk orientation usually requires a specific internal advocate: someone who can frame the conversation not as 'we need more tools' but as 'we need to know what we're actually exposed to'.

What has to change structurally

The prerequisite for Stage 2 isn't a tool purchase. It's an ownership decision. Someone needs to be accountable for the completeness of the external asset inventory, not just the security of known assets. Once that accountability exists, the supporting structures follow: continuous discovery of assets beyond predefined scope, risk-based prioritisation that incorporates business context alongside severity scores, and defined SLAs that are actually enforced rather than aspirational.

Capabilities that enable this transition

- External attack surface management that discovers assets outside predefined inventory: cloud, SaaS, subsidiaries, third parties
- Risk-based prioritisation incorporating business criticality and reachability, not just CVSS
- Asset ownership workflows that connect discovered assets to accountable teams

DOWNLOAD THE STAGE 1 TO 2 ADVANCEMENT PLAN

STAGE 2

BETTER SCOPE, MORE HOPE

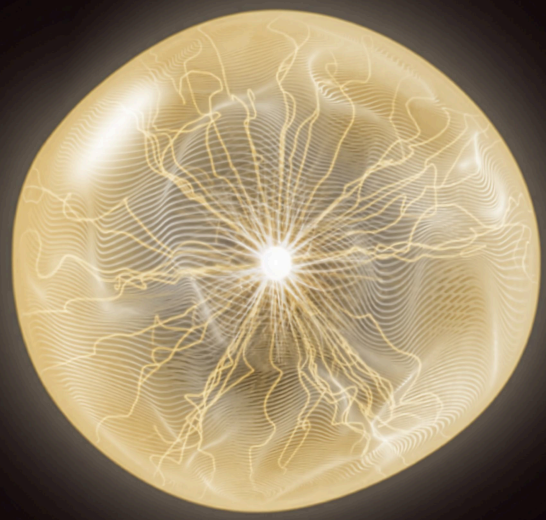

"We know more than we did. We're just not sure it's enough."

Stage 2 organisations have done real work. There is a formal vulnerability management programme. Major assets are inventoried. Testing happens on a schedule rather than only after incidents. Security has a seat at the table and a defined mandate. The CISO can produce a posture report without scrambling.

The name captures both the progress and the limitation. Scope has expanded: more assets, more testing coverage, more structured ownership. And there is genuine hope: hope that the assets you know about are the important ones, hope that the quarterly pentest caught what matters, hope that exploitability roughly tracks severity. That hope is not unreasonable. But it is still hope rather than evidence. The programme is structured; it is not yet validated.

How teams describe it

"Findings from the last pentest are already going stale by the time the next cycle runs."

"We have good coverage of the assets we know about. The problem is the ones we don't."

"We're risk-based in theory. In practice we're still working through CVSS 9s and 10s."

What this looks like organisationally

The security team has a programme rhythm: quarterly pentests, regular vulnerability scans, a ticket queue with defined owners. Board reporting has improved from compliance status to current vulnerability exposure. When a new CVE drops, there is a process for checking, but it is manual and slower than anyone would like.

The frustration at Stage 2 is subtler than at Stage 1 because the programme looks healthy from the outside: coverage metrics are improving, SLA compliance is tracked, and posture reports can be built without scrambling. But the team knows something is off: findings from the last pentest are going stale before the next cycle runs, the remediation queue contains items nobody has confirmed are genuinely exploitable, and the gap between what the programme reports and what an attacker could actually do remains uncomfortably wide.

What keeps organisations here

The structural blocker at Stage 2 is the gap between testing cadence and environment velocity. Quarterly penetration tests were designed for environments that changed quarterly. Most enterprise environments today change daily. The programme is not broken; it is mismatched to the pace of the problem.

The operational bottleneck at Stage 2 is manual validation. Confirming that a finding is genuinely exploitable requires human investigation. When that step does not scale, remediation queues fill with unvalidated findings, and teams resort to informal heuristics: anything over CVSS 8 gets treated as exploitable. That is faster. It is also far less accurate, and it neither reduces backlog nor improves confidence in what remains.

Measurement reinforces both problems. Most Stage 2 programmes track coverage and activity: findings discovered, findings closed, SLA compliance. These metrics feel like progress but don't answer the question leadership actually cares about: is our exploitable exposure going down?

The moment of readiness

The trigger is usually a finding that should have been caught before it was. An exploitable exposure discovered between test cycles, or during an unscheduled assessment, or by an external researcher. The question that follows, how long was that open, is often the catalyst for recognising that periodic validation can no longer keep pace. The readiness question is whether the organisation is prepared to move from scheduled testing to continuous monitoring, and whether leadership will fund the structural change that requires.

What has to change structurally

The shift from Stage 2 to Stage 3 requires closing the gap between testing cadence and environment velocity. That means continuous external monitoring that detects new exposure as it emerges, not just at the next scheduled scan. It also requires treating exploitability validation as a workflow step, not a manual investigation. When confirmation that a finding is real is automated rather than human-dependent, the triage bottleneck breaks. Measurement needs to change too: the programme should track exploitable risk removed, not just findings closed.

Capabilities that enable this transition

- Continuous external attack surface monitoring that flags new assets and changes between pentest cycles
- Automated exploitability validation that confirms which findings are genuinely reachable and testable
- Integrated remediation routing so validated findings reach the right owner without manual handoff

DOWNLOAD THE STAGE 2 TO 3 ADVANCEMENT PLAN

STAGE 3

CONNECTING THE DOTS

"We know what's exploitable. We're starting to see how it fits together."

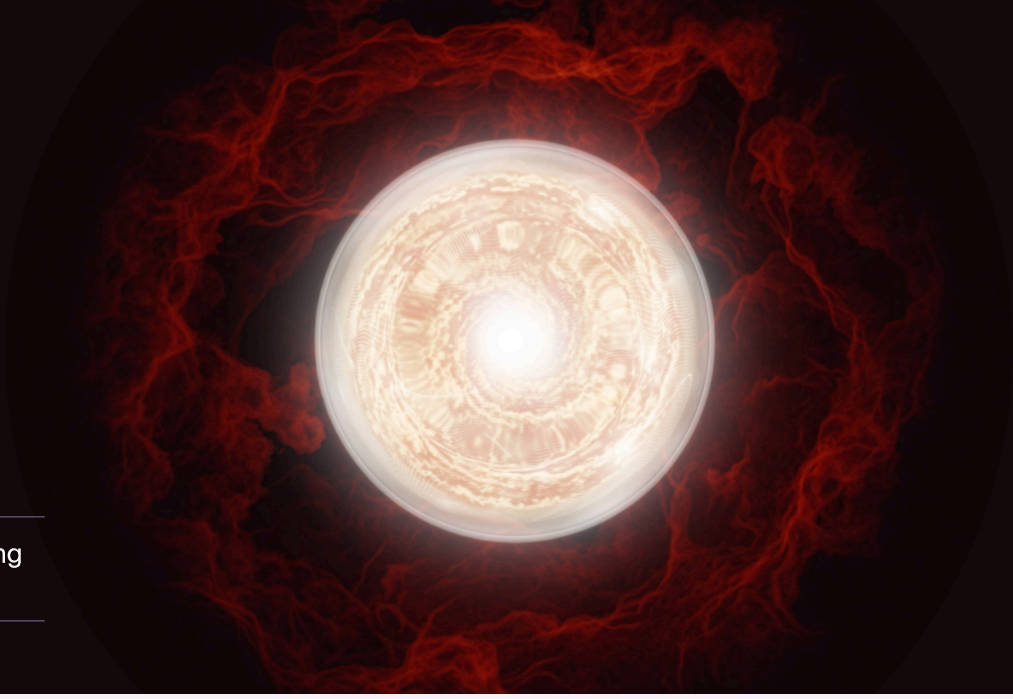

Stage 3 is a genuine operational inflection point. Validation is no longer periodic; it is continuous. Blind windows between test cycles have largely closed. Exploitability is confirmed before findings are escalated, not assumed from severity scores. Remediation workflows are integrated with security tooling, so validated findings route automatically rather than landing in a shared inbox.

The name marks what's new: for the first time, the programme isn't just cataloguing weaknesses in isolation; it is assembling them into a picture of how an attacker might actually move. Individual validated findings get mapped against attack paths. Prioritisation shifts from 'what's the highest CVSS?' to 'what can be chained to cause real damage?'. This is a fundamentally different view. The ceiling is that building and maintaining this picture still requires meaningful human orchestration. The programme sees the dots; connecting them at speed and at scale is still manual.

How teams describe it

"We know what's exploitable. The challenge is keeping pace with how fast the environment changes."

"Our findings are validated before they hit the queue. The problem is connecting them into a full attack picture."

"We can tell the board what's confirmed exploitable. We can't yet tell them how an attacker would chain it."

What this looks like organisationally

The security team's daily rhythm looks different. Findings arrive pre-validated. Tickets are routed to owners automatically. The conversation in triage has shifted from 'is this real?' to 'how does this fit into the broader exposure picture?'. MTTR is actively measured, trending down, and defined by exploitable risk closed rather than ticket count.

Reporting, board metrics, and KPIs are often still built around periodic snapshots because that is what leadership has always asked for. The board still wants a number between one and ten. KPIs that made sense at Stage 2 are now measuring the wrong things, but changing them requires a conversation nobody has prioritised. Meanwhile the security team is operating with a level of precision that the governance structure was not designed to support. The CISO can tell you what is exposed. They cannot yet tell you, fast enough to matter, what an attacker could exploit.

What keeps organisations here

The constraint is orchestration throughput. Running a complex adversarial scenario that follows attack chains across systems, tests realistic intrusion paths, and adapts to environment changes requires significant expert time. Individual findings can be validated continuously; simulating what a sophisticated attacker would actually do with them is still manual and resource-intensive. The programme sees the dots. Connecting them at speed is still a human job.

The operational bottleneck at Stage 3 is emulation scale. The programme can validate individual findings continuously, but translating them into a complete, real-time picture of how those findings combine into exploitable attack paths still requires human orchestration per scenario. As environments grow in complexity, the gap between what can be validated and what can be emulated becomes the rate-limiting constraint.

There is also an organisational lag. Operations are running continuously, but governance has not caught up. Leadership reporting, board metrics, and KPIs are often still built around periodic snapshots rather than live posture data. The programme produces more signal than the organisation has learned to consume.

The moment of readiness

Stage 3 organisations become ready to advance when they confront a specific question they can't yet answer quickly: if a new exploit class was disclosed today, how long would it take us to confirm whether our full surface is affected? If the honest answer is 'days' rather than 'hours', the gap between detection capability and emulation capability has become the limiting factor. The readiness conversation is about whether leadership will invest in closing that gap, and whether the programme governance can accommodate autonomous operation.

What has to change structurally

The shift to Stage 4 requires replacing manual orchestration of adversarial scenarios with automated emulation that adapts to environment changes without human intervention per scenario. It also requires closing the loop between live operational data and leadership visibility: real-time posture dashboards derived directly from programme output, not compiled periodically. And it requires a governance decision: pre-authorising rapid response protocols so that when a new exploit surfaces, the programme responds in hours rather than waiting for a triage meeting.

Capabilities that enable this transition

- Automated adversarial emulation that adapts to environment changes without manual orchestration per scenario
- Real-time posture scoring and board-ready metrics derived directly from live programme data
- Continuous attack chain analysis that maps how validated findings combine into exploitable paths

DOWNLOAD THE STAGE 3 TO 4 ADVANCEMENT PLAN

STAGE 4

CLEAR PICTURE

"For the first time, we know what an attacker sees at any moment."

Clear Picture is not an aspirational tier. It is what exposure management looks like when the structural constraints of the previous stages have all been resolved: reactive discovery, periodic validation, manual orchestration and lagging leadership visibility. The programme is continuous, adaptive, and largely autonomous. The security team's attention is on strategic decisions and exception handling, not on triaging noise or manually assembling attack paths.

The name is exact. For the first time, the organisation has a complete, real-time picture of its external attack surface from an attacker's perspective: not just what assets exist, but which are reachable, which are exploitable, how they chain together, and what the likely intrusion paths look like right now. This picture updates continuously as the environment changes. It doesn't wait for the next scan cycle or the next pentest.

What makes Clear Picture distinctive is not any single capability. It is the absence of the structural gaps that define every earlier stage. There are no blind windows. There is no triage backlog built from unvalidated findings. There is no gap between what the programme knows and what leadership sees. The organisation is no longer reacting to exposure; it is anticipating it.

What resolved to get here

The organisation replaced manual adversarial scenario orchestration with automated emulation that adapts continuously to environment change. It closed the loop between live operational data and leadership reporting. And it made a governance decision: pre-authorised rapid response protocols and defined escalation paths for autonomous operation, which allowed the programme to respond at machine speed without requiring a human decision at each step.

The team is smaller relative to surface coverage than at any earlier stage — not because headcount was cut, but because automation has absorbed the work that previously required human triage. Expertise has concentrated upward. The security team works on problems that require judgement: threat modelling, adversarial scenario design, the strategic questions that cannot be automated. The CISO's relationship with the board has changed too. The conversation is no longer about what happened last quarter, it is about what the programme sees right now.

What this looks like organisationally

The security team's posture conversation has changed register. The question is no longer 'what are we exposed to?' The programme answers that continuously. The question is 'how fast can we close it?' and 'what are the highest-leverage things our people should be working on that automation can't yet handle?'. Board reporting is derived directly from live programme data. Engineering teams receive validated, routed findings without security having to hand-deliver them. When a new exploit is disclosed, the programme confirms within hours whether the surface is affected, not days.

The security team at this stage is smaller relative to surface coverage than at any earlier stage, because automation has absorbed the work that previously required human triage. Expertise has shifted upward: the team works on the problems that require judgement, not the problems that require volume.

How teams describe it

- "We know about new exposure before attackers can probe it."
- "The question we ask now is not whether we are exposed. It is how fast we can close it."
- "I can tell the board our current posture from data generated in the last 24 hours. That wasn't possible before."

What this stage requires to sustain

Clear Picture is not a destination that, once reached, maintains itself. Adversarial techniques evolve, and the emulation quality of the programme must evolve with them. New categories of infrastructure, including AI systems, agentic workflows and novel SaaS integrations, create surface area that requires the discovery and validation models to be updated continuously.

The organisational requirements are as important as the technical ones. Governance structures need to support autonomous operation without becoming an obstacle to it: pre-authorised response protocols, clearly defined escalation thresholds, and leadership alignment on what 'real-time posture visibility' actually means in practice. Programmes that reach Clear Picture and then introduce bureaucratic gates around automated response will find their capability degrading toward Stage 3.

The most common path back to Stage 3 is leadership change: a new CISO who reintroduces manual review gates, or a board that loses confidence in autonomous operation after a false positive. Sustaining Clear Picture requires active maintenance of the organisational mandate, not just the technical capability.

Capabilities that define this stage

- Autonomous adversarial emulation that adapts to environment changes continuously, without manual orchestration per scenario
- Predictive, dynamic prioritisation that shifts automatically as new exposures emerge and attack chain context changes
- Real-time attack surface model correlating identity, infrastructure, SaaS, and third-party changes into a unified, continuously updated picture
- Full programme integration from exposure detection through validated finding to remediation completion, with minimal manual handoffs

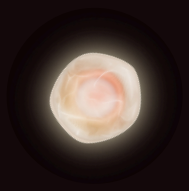

STAGE 1 → STAGE 2

Moving from Stage 1 to Stage 2 can be achieved in a few months with focused effort and clear ownership.

STAGE 2 → STAGE 3

Stage 2 to Stage 3 typically takes six to eighteen months, depending on environment complexity and integration depth.

STAGE 3 → STAGE 4

The transition to Stage 4 is ongoing; it requires sustained operational investment rather than a defined project. What accelerates advancement is not more tooling; it is clarity: clear ownership of exposure, clear definitions of what done looks like, and measurement of exploitable risk removed rather than activity generated.

SECTION 4

WHERE DO YOU STAND?

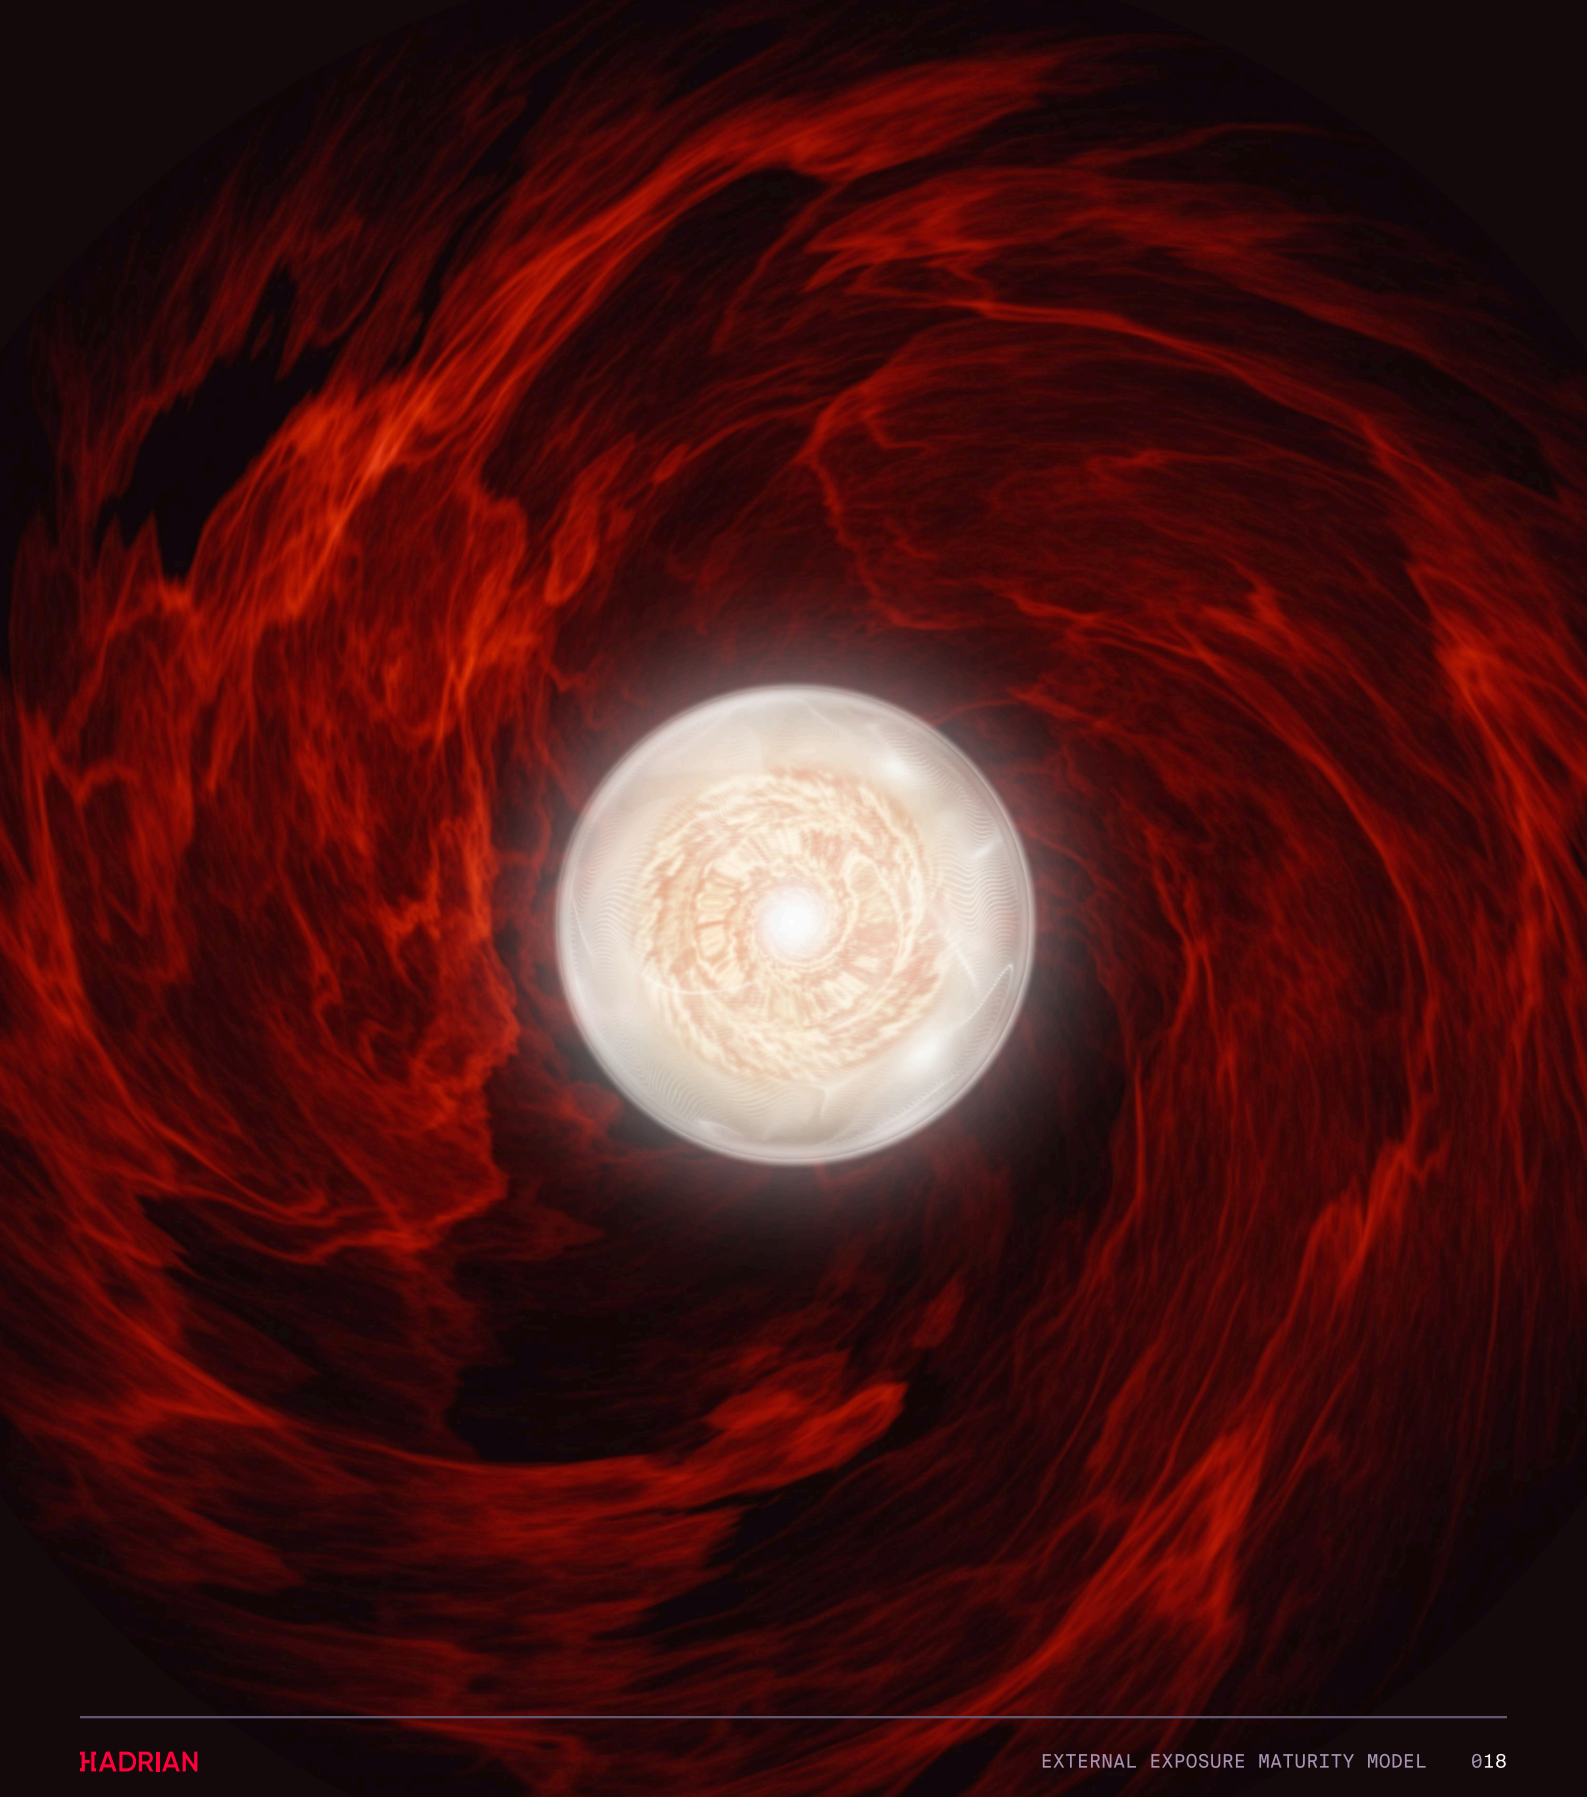

Most organisations that work through the four stages find that their honest answer sits between two of them. Discovery is more mature than validation. Validation is more mature than remediation routing. An overall stage score can tell you roughly where you are. It cannot tell you which specific dimension is holding everything else back.

That distinction matters more than it sounds. A programme that averages out at Stage 3 but has Stage 1 validation is not a Stage 3 programme. It is a programme with a Stage 1 bottleneck that is suppressing performance across every other dimension. The remediation SLAs, the board reporting, the level of automation - all of it is constrained by the weakest link, not the strongest. Most internal assessments miss this because they score overall posture rather than breaking it down by operational dimension.

The interactive self-assessment is built differently. It scores your programme across six dimensions independently and weights them to identify where the gap between your current posture and your risk profile is largest. The output is not a number. It is a breakdown of where you are by dimension, a callout of your highest-leverage bottleneck, and a stage-specific action plan your team can act on directly.

If you already know your overall stage, the assessment is still worth completing. The dimension-level profile is what makes it useful in a leadership conversation, not "we are Stage 2" but "we are Stage 3 on discovery and Stage 1 on validation, and here is what closing that gap is worth in remediation time.

ASSESS YOURSELF:

Take the interactive self-assessment for a personalised maturity profile and action plan

DISCOVER YOUR MATURITY STAGE

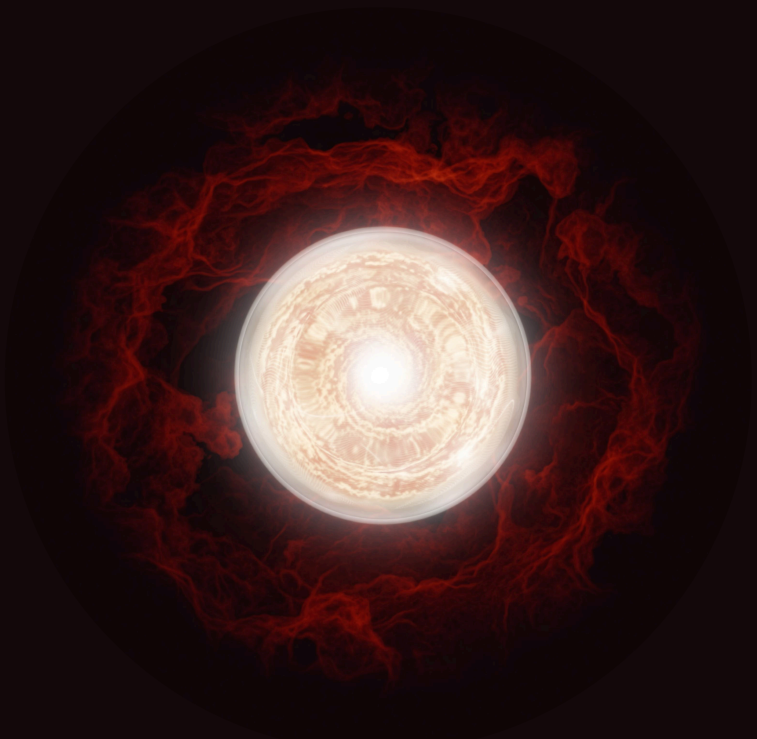

Frequently asked questions

Does every organisation need to reach Clear Picture?

No. Clear Picture is appropriate for organisations where the consequences of compromise justify continuous, autonomous validation. For many organisations, Stage 2 or Stage 3 is the right target for most of their attack surface, with Stage 4 practices reserved for the highest-value assets. The right question is not how to reach Stage 4; it is which assets and risk classes require Stage 4 treatment.

We already use many tools. Why does our maturity remain low?

Tool adoption is not programme maturity. 93% of organisations use vulnerability scanners; only 40% have adopted automated penetration testing. High tool adoption has not translated into faster remediation for most organisations because the limiting factor is not which tools are present; it is whether those tools operate within a structure that validates exploitability, integrates with remediation, and measures outcomes rather than activity.

We know we need to advance but have competing priorities and limited budget. How do we make the case internally?

The most effective argument is not about maturity for its own sake, it is about the cost of staying where you are. At Stage 1, the security team is spending most of its capacity on findings that turn out to be irrelevant. At Stage 2, quarterly testing is generating a false sense of coverage while the environment changes daily. The business case for advancement is not "we need to be more mature." It is "our current structure is consuming resource without reducing exploitable risk, and here is what that costs us in remediation time, SLA compliance, and leadership visibility." The maturity model gives you the language and the data to make that argument precisely rather than generally.

Our CTEM programme is active. Why doesn't it translate into faster remediation?

Most CTEM programmes measure success in ways that do not reflect actual risk reduction. 67% of organisations measure CTEM by coverage gaps identified; only 33% track reductions in exploitable exposure over time. Without exploitability validation as a first-class output, and without ownership alignment connecting validated findings to the people who can fix them, CTEM generates more findings without improving the outcomes that matter.

We have full pipeline visibility. Why is our exposure not going down?

Discovery and validation are separate problems, and most programmes treat them as one. Expanding EASM coverage or increasing scan frequency generates more findings, but if there is no automated mechanism to confirm which of those findings are genuinely exploitable, the result is a larger backlog, not improved posture. More discovery without validation capacity increases the load on analysts without reducing exploitable risk. The constraint is rarely how many findings the programme finds; it is how many of those findings can be confirmed real and routed to remediation before an attacker acts on them.

We haven't had a breach. How do we know our current stage isn't sufficient for our risk profile?

Absence of a known breach is not evidence of adequate posture — it is evidence that you have not yet discovered a breach, or that attackers have not yet acted on access they already have. The median time between initial compromise and detection remains measured in weeks, not hours. Organisations at Stage 1 and Stage 2 typically do not know what is exploitable on their external surface at any given moment, which means they cannot distinguish between being secure and being undetected. The right question is not whether you have been breached. It is whether your programme would know if you had been, and how quickly it could confirm whether a newly disclosed exploit affects you right now.

About Hadrian

Hadrian is an offensive security platform that finds and validates external exposure the way an attacker does — starting from no predefined scope, following attack paths wherever they lead, and confirming every finding as real before it reaches your team. Recognised in the Gartner Hype Cycle for Security Operations, and by GigaOm and Frost & Sullivan.

LEARN MORE: HADRIAN.IO

FROST & SULLIVAN BEST PRACTICES AWARDS

GIGAOM RADAR REPORT LEADER 2025

Gartner Peer Insights 4.9/5 ★

AICPA SOC · ISO 27001

ANALYST-BACKED PERFORMANCE

"Hadrian is classified as an Outperformer due to the deployment of its agentic AI capabilities and the steady expansion of its automated offensive testing modules over the last 12 months"

CHRIS RAY
FIELD CTO, GIGAOM

10x VISIBILITY OF CRITICAL RISKS

80% REDUCTION IN MTTR

10h SAVED PER WEEK ON AVERAGE

TRUSTED BY

NBC · amadeus · McKESSON · BLINQX · Crédit Agricole · Leroy Merlin · SHV · ABN·AMRO · London Business School · RITUALS · SIEMENS ENERGY · LOTTOMATICA · FP Francisco Partners · CHRISTIE'S · BIOLANDES · WeatherTech · BLINQX · DAMEN · jamf · exact · ICT GROUP · nedap · celio* · Inova · XM Cyber · AROMA360

AND 300+ MORE