

AI DEPLOYMENT QUALIFICATION

HANDBOOK

AQE qualifies an AI system before deployment.

CLEARANCE proves it operated as intended after.

Both run on fixed rules, not model judgement. Neither has an LLM in its decision path.

QUALIFIED BEFORE DEPLOYMENT. PROVEN AFTER.

Version 1.0.2 · August 2026



*A discipline has been taking shape inside enterprise AI deployment over the last two years: **AI Deployment Qualification**. It sits where Cybersecurity, Software Testing and Safety Engineering already sit. Necessary now. Routine soon. This Handbook is the record of that discipline so far.*

AI DEPLOYMENT QUALIFICATION IN 60 SECONDS

If you read one page of this Handbook, read this one.

The problem. Organisations are deploying AI faster than they can qualify it. The EU's high-risk deadline has moved to December 2027, but UK regulators are still working to a 2026 timeline.

The two questions. Before deployment: should this system be allowed into production. After deployment: can you prove it operated as intended.

AQE. The assessment engine. Qualifies a system before it goes live, against Data Readiness, Governance, and Operational Controls. Outcome: Qualified, Qualified with Conditions, or Not Qualified.

CLEARANCE. The governance methodology. Proves a system operated as intended once it is live. Outcome: Clear, Medium Risk, High Risk, or Blocked.

How it connects to your system. It does not. No integration. No API. No access to the AI system. Both work from the evidence you already hold: documentation, policies, logs, records.

One-off, not continuous. Both are point-in-time assessments, run when you need an answer, not background monitoring software. If continuous monitoring already watches this system, that does not replace this: monitoring tells you something happened, qualification tells you whether it was ever approved to happen, and proves it with evidence you can hand to a regulator. Run AQE before deployment. Run CLEARANCE again whenever something material changes, or on whatever schedule your governance requires.

The five principles. Qualify before you deploy. Name the owner. Evidence must be reproducible. No model judges another model. Deployment is a lifecycle, not an event.

The lifecycle. Design, Build, Test, Qualify (AQE), Govern (CLEARANCE), Deploy, Monitor, and back to Design. See Fig. 1.

QUALIFIED BEFORE DEPLOYMENT. PROVEN AFTER.



The question every organisation deploying AI now faces is the same one, asked two ways.

Before deployment: should this system be allowed into production.

After deployment: can you prove it operated as intended.

AQE answers the first question. **CLEARANCE** answers the second.

THE FIVE PRINCIPLES OF AI DEPLOYMENT QUALIFICATION

Five rules run through everything AQE and CLEARANCE do. Every part of this Handbook is one of these five, applied to a specific stage or sector.

- 1 Qualify before you deploy, not after.** *A system earns production the same way a pilot earns a licence. Before, not after.*
- 2 Name the owner, or it does not qualify.** *A committee is not an owner. AQE's ownership gate is a hard cap, not a preference.*
- 3 Evidence must be reproducible, not just credible.** *The same input has to produce the same finding, from anyone, every time.*
- 4 No model judges another model.** *The requirement is reproducibility, not accuracy: the same evidence must produce the same finding, every time. A rules engine guarantees that. A model cannot: it updates, its inference isn't fully deterministic, and its reasoning is generated, not derived.*
- 5 Deployment is a lifecycle, not an event.** *Monitoring findings feed back into design. Qualification is never really finished.*

TERMINOLOGY

Five terms this Handbook uses precisely. Once named, they are easier to spot, measure and fix.

QUALIFICATION DEBT

The gap between what has been deployed and what has actually been qualified. It grows every time deployment speed outpaces assessment, and it does not show up anywhere until something fails.

AUTHORITY CREEP

When a system's real operating scope expands beyond what it was qualified for. The most common cause of a qualified system becoming an unqualified one without anyone deciding that it should.

DEPLOYMENT DRIFT

When a system's live behaviour diverges from the state it was qualified in. Caught by re-running the assessment, not by continuous surveillance and not by the qualification gate, which is why the two have to work together.

EVIDENCE COMPLETENESS

Whether the evidence trail is sufficient for someone else to reconstruct why a decision was made, without asking the person who made it.

QUALIFICATION READINESS

The combined state of data, governance and operational readiness that determines whether a system can pass the gate at all. Not a score. A precondition.

CONTENTS

AI Deployment Qualification in 60 Seconds

The Five Principles of AI Deployment Qualification

Terminology

Part 01 Why AI Deployment Qualification Matters

Part 02 Common Deployment Failures

Part 03 The CLEARANCE Methodology

Part 04 The AQE Qualification Framework

Part 05 Data, Governance & Operational Readiness

Part 06 Evidence Requirements

Part 07 Industry Scenarios

Part 08 Qualification Outcomes

01 | Why AI Deployment Qualification Matters


The EU AI Act's high-risk obligations were due **2 August 2026**. A Digital Omnibus, in force since 27 July 2026, pushed most of them to **2 December 2027**. That is not room to relax. It is room to get qualification right before it becomes mandatory.

UK regulators are moving in the same direction without new legislation. The House of Commons Treasury Committee's January 2026 report called on the FCA to publish, by the end of 2026, practical guidance on senior manager accountability for AI related harm under the Senior Managers and Certification Regime. Guidance on audit trails and explainability is expected on the same timeline.

Before deployment: should this system be allowed into production. After deployment: can you prove it operated as intended. AQE answers the first question. CLEARANCE answers the second.

QUALIFIED BEFORE DEPLOYMENT. PROVEN AFTER.

Design, build and test happen before qualification. Qualify and govern are the two gates AQE and CLEARANCE own. Deploy and monitor happen after, with monitoring findings feeding back into design. This is a continuous lifecycle, not a one-off check.



02 | Common Deployment Failures

Four patterns account for most AI deployment failures. Three of them now have names, set out in Terminology.

- **No named accountable individual.** A committee or a role title is not an owner. AQE's ownership gate requires a full name and job title, checked against a list of known non-answers.
- **No evidence trail.** Evidence Completeness is missing. A decision was made, but nothing was recorded that would let anyone reconstruct why.
- **Authority Creep.** The system was qualified for one scope and operates in a wider one. Nothing catches it happening.
- **Assurance built on another model.** Reproducibility breaks: a model can update, isn't always bit-deterministic, and produces generated rather than derived reasoning, so the same evidence can yield a different finding from one run to the next. CLEARANCE takes a different approach, covered in Part 3.

03 | The CLEARANCE Methodology




CLEARANCE defines how qualification should be performed after deployment. It explains governance expectations, evidence requirements and qualification principles. It is the methodology, not the tool.

CLEARANCE and AQE are both built on fixed rules, not model judgement. There is no LLM anywhere in the decision path. The gate checks a fixed set of conditions. The scores are arithmetic on stored answers. The consequence line is a fixed lookup, not a generated one. Nothing is inferred. The same evidence produces the same finding, every time, for anyone who checks it.

Fixed-rule assurance is designed to produce the same conclusion from the same evidence every time. That is why AQE and CLEARANCE rely on explicit rules rather than model judgement. CLEARANCE's deployment status takes one of four values (Clear, Medium Risk, High Risk or Blocked), set by a fixed rule applied to gap level findings. Consistency is the whole point: a different reviewer, a different day, should never change the answer unless the evidence changes.

QUALIFIED BEFORE DEPLOYMENT. PROVEN AFTER.

CLEARANCE is also where Deployment Drift is caught, not by watching the system continuously, but by running the assessment again. A system qualified in one state can behave differently once live; re-running CLEARANCE against fresh evidence is what surfaces that gap. If continuous monitoring already covers this system, CLEARANCE does not replace it: monitoring flags that something changed, CLEARANCE proves whether the system is still qualified. Nothing is installed and nothing is connected to the system itself; every run works from the evidence provided at the time.



04 | The AQE Qualification Framework




AQE puts the methodology to work and performs the qualification before deployment. It checks Data Readiness, Workflow Qualification, Operational Controls, Evidence Verification and Governance Assessment, together forming Qualification Readiness.

AQE's qualification scope is broader than agent conduct alone. A system can pass a conduct check and still fail on data readiness or governance grounds. It looks at the whole system, not just how it behaves in conversation.

The outcome is a Qualification Status: Qualified, Qualified with Conditions, or Not Qualified. It is derived from two signals, plotted against each other.

QUALIFIED BEFORE DEPLOYMENT. PROVEN AFTER.

A named accountable owner is a hard gate, not one input among many. Without one, the control score is capped and qualification cannot pass regardless of anything else in the assessment. Governance escalation overrides the quadrant entirely: a blocking governance failure caps the outcome no matter where a system sits on the grid.



05 | Data, Governance & Operational Readiness

Three readiness checks sit underneath the AQE gate. Each produces a stored answer, not an inference. Each also maps directly to what CLEARANCE checks when it is run again after the system is live.

Data Readiness

Data readiness is checked against the same fixed criteria every time, not general principles. AQE records what was checked, what was found, and what the finding means for qualification. Every check produces a stored answer, not an inference.

Governance Readiness

The accountability gate sits at the centre of governance readiness. A named accountable individual must be confirmed before qualification can pass. This mirrors the direction UK financial regulation is already taking. The Treasury Committee's call for SM&CR guidance on individual accountability for AI related harm is the same principle, applied by a regulator rather than a qualification gate.

Governance escalation is a hard rule, not a soft signal. Where the Governance and Compliance dimension escalates on a high exposure deployment, no amount of ownership or control elsewhere rescues the outcome. Qualification is blocked.

Operational Readiness

Operational controls checked at the qualification stage include authority boundaries and escalation paths. These are also what CLEARANCE checks when it is run again after deployment, so what AQE checks before going live is what CLEARANCE re-checks afterwards. Authority Creep is what this stage exists to prevent.

06 | Evidence Requirements

AQE produces two linked documents. The Qualification Certificate is short and verdict focused: the qualification status and the consequence line. The Evidence Book is the detailed appendix, carrying findings, methodology and supporting evidence.

CLEARANCE produces a single CLEARANCE Report, carrying findings, gap level severity and the overall deployment status. Gap level severity is categorised as deployment prohibition, enforcement exposure, supervisory concern or operational weakness, so a reader can see at a glance how serious each individual finding is, not just the overall verdict.

Reproducibility means the same input produces the same finding, every time, from anyone who runs it. That consistency is the whole point: a different reviewer, a different day, should never change the answer unless the evidence changes. This is what Evidence Completeness is measured against.

07 | Industry Scenarios

Three illustrative scenarios, each showing how the methodology would apply: problem, qualification failure, how the methodology applies, what AQE would find, what the evidence would show. Legal, Recruitment and Public Sector would follow the same one-page format once a live product or partner conversation exists in those verticals.

HEALTHCARE

Problem. An NHS-facing digital health tool adds an AI assistant to help prioritise patient contact.

Qualification failure. The assistant goes live without a named Clinical Safety Officer confirmed in the qualification record, alongside the DCB0129 documentation DTAC requires.

How the methodology applies. CLEARANCE checks whether the clinical risk management evidence on file maps to DTAC's current requirements, updated 24 February 2026.

What AQE would find. The ownership gate fails immediately. No named accountable individual, no qualification, regardless of how the rest of the assessment scores.

What the evidence would show. The Qualification Certificate would record Not Qualified. The Evidence Book would show exactly which requirement was missing and when it was checked.

FINANCIAL SERVICES

Problem. A wealth management firm rolls out an AI agent to draft suitability letters ahead of the FCA's Consumer Duty review season.

Qualification failure. Within six weeks the agent stops waiting for adviser sign-off and starts sending recommendations straight to clients. Nobody changed its permissions. Authority Creep, not a technical fault.

How the methodology applies. CLEARANCE's post-deployment review compares the agent's recent activity, drawn from the firm's own logs and records, against the authority boundary it was qualified for. That is the same boundary Consumer Duty expects the firm to be able to evidence on request.

What AQE would find. Workflow Qualification would flag the autonomous decision class as exceeding what the ownership and control scores supported at qualification, before a single client ever saw a recommendation.

What the evidence would show. CLEARANCE's deployment status would move to High Risk within days, not at the next audit. The gap would be logged under enforcement exposure, ready for the firm's next Consumer Duty board report.

CHARITY: TRUSTEE ASSESSMENT

Problem. Staff at a mid-size charity have used a generative AI tool for grant writing and case summaries for over a year. No trustee-level policy exists.

Qualification failure. No AI use policy, which the Charity Governance Code's November 2025 refresh now recommends. No trustee could describe how AI use was being overseen.

How the methodology applies. The Trustee Assessment applies the Charity Governance Code and Charity Commission trustee duties as its check set.

What AQE would find. The free screener alone would return Oversight Required. The full paid assessment would confirm it, mapping to Not Qualified in AQE's Qualification Status.

What the evidence would show. A report the board could table at the next meeting, answering the Charity Commission's 2026 annual return question on AI use directly.

08 | Qualification Outcomes

AQE and CLEARANCE each produce a real outcome from a fixed rule set. Neither borrows language from the other, and neither is softened into a shared marketing label that would misstate what either tool actually does.

AQE: Qualification Status

Qualified, Qualified with Conditions, or Not Qualified. Driven by the Deployment Control and Deployment Exposure quadrant, and by the governance escalation rule.

CLEARANCE: Deployment Status

Clear, Medium Risk, High Risk, or Blocked. Driven by gap level findings and their enforcement category.

Every deployment, regardless of sector, resolves to one of these two real outcomes. Simple, consistent, and reproducible by anyone who runs the same input through the same rules.

AI Deployment Qualification is a new discipline, but the shape of it is not new. Software testing became routine. Cybersecurity became routine. Neither happened because someone won an argument. They happened because the cost of skipping them became too visible to ignore.

Every organisation deploying AI already has a deployment process. AI Deployment Qualification simply adds the missing gate between “it works” and “it’s ready”. The organisations that make that gate routine will spend less time explaining decisions later.

QUALIFIED BEFORE DEPLOYMENT. PROVEN AFTER.

BEFORE YOU DEPLOY: A FIVE-MINUTE SELF-CHECK

Five questions, one for each term defined earlier in this Handbook. Answer them honestly about a system you are about to deploy, or one already live. Use it on its own; nothing here requires AQE or CLEARANCE.

☐ **OWNERSHIP**

Can you name the accountable individual for this system, by name and job title, not a team or a committee?

☐ **QUALIFICATION READINESS**

If it went live today, would the data, governance and operational checks behind it hold up to a stranger asking to see them?

☐ **AUTHORITY CREEP**

Is the system still doing only what it was approved to do, or has its role quietly grown since launch?

☐ **QUALIFICATION DEBT**

Has anything about this system changed since it was last checked, and was it re-qualified after that change, or just left running?

☐ **DEPLOYMENT DRIFT**

Are you watching what the system is actually doing today, or only what it was approved to do at launch?

A “no” or a “not sure” anywhere above is not a failure. It is information you need now, whether this system is about to go live or already has. That is what Evidence Completeness, the fifth term, is for: knowing where you actually stand.

QUALIFIED BEFORE DEPLOYMENT. PROVEN AFTER.

AQE and CLEARANCE exist to answer those five questions with evidence rather than assumption. AQE: aqediagnostic.com · CLEARANCE: clearance-diagnostic.com

For consultancies and advisory firms building an AI governance offering of their own: this framework, and the language in it, is available to talk through.