AI-QA LEADERSHIP IN EDUCATION

Operational AI Quality Assurance Handbook

For QA Teams, Registry, Programme Leads & Internal Auditors
ASIC Standards for Growth Series: Webinar 1 - Post-Webinar Resource

Audience	Quality assurance officers, academic registrars, programme leaders, internal auditors, compliance teams
Purpose	Policy mapping, evidence collection, compliance assurance, and audit readiness for AI governance
Series	ASIC Standards for Growth - Webinar 1 | 24th March 2026 Operational AI Quality Assurance Handbook

Contents

Introduction: What This Handbook Is For

This handbook is written for the people who make quality assurance work in practice: QA officers, academic registrars, programme leaders, compliance teams, and internal auditors. Your role goes far beyond check-box compliance; you are central to building and sustaining the evidence infrastructure that demonstrates policies and procedures are not only in place but genuinely effective.

AI has entered that evidence infrastructure whether or not your institution has formally acknowledged it. It is shaping how students produce work, how staff give feedback, how decisions are made about student progress, and how your platforms and services operate.

This handbook helps you to map those touchpoints, collect appropriate evidence, and be ready to demonstrate sound governance where it matters.

> **How to use this handbook**
>
> **Section 1 - The AI QA landscape:** Where AI intersects your existing QA processes
> **Section 2 - The full AI QA Evidence Matrix:** What to collect, from whom, and how often
> **Section 3 - Audit readiness:** What reviewers look for and how to prepare
> **Section 4 - The AI Incident Log:** Building and running it in practice
> **Section 5 - RAG self-assessment tool:** Where are the gaps in your QA evidence?
> **Section 6 - Policy templates and working documents you can adapt immediately**

Section 1: The AI QA Landscape

1.1 Where AI Intersects Your QA Processes

AI does not sit alongside your existing QA processes — it cuts across all of them. The table below maps the key intersections. This is the starting point for understanding what your QA function now needs to govern.

QA Area	Policy & Process Touchpoints	Required Evidence	Responsibility	Review Frequency
Assessment Design & Validity	Assessment regulations; programme approval; periodic review; external examiner briefings	Updated regulations citing AI; 'AI-aware' assessment design rationale; moderation commentary	Academic Quality / Registry	Annual minimum
Marking & Feedback	Feedback protocols; staff guidance; tool-use policy; moderation process	Rubrics updated with AI rules; reviewer notes confirming human oversight; sample marked work	Programme Leaders / QA	Each assessment cycle
Academic Integrity	Academic conduct regulations; student declaration processes; appeals procedure	AI-specific conduct policy; student disclosure form; case log with outcomes; appeals trend data	Registry / Academic Registrar	Annual review; live case log
Learning Analytics	Data governance policy; student comms; ethics review; DPIA	DPIA for each analytics tool; fairness/bias audit reports; student information provided	DPO / Data Team / QA	Annual; at tool change
Programme Documentation	Programme approval; annual monitoring; spec templates; employer consultation	Revised ILOs noting AI; AI-permitted-use statements in specs; committee minutes	Programme Leaders / Registry	Annual monitoring cycle
Accessibility & Inclusion	VLE standards; disability support procedures; alternative format processes	Accessibility audits; captioning and alt-text records; alternative format samples	Learning Technology / DSA	Annual; at content change
Staff CPD & Capability	CPD framework; appraisal process; induction; peer observation	CPD log with AI-specific training; appraisal record extracts; observation notes	HR / Academic Dev	Annual CPD cycle
Marketing & Recruitment	Marketing approval workflow; QA sign-off procedures; consumer protection	Signed-off copy; ethics checklist for AI capability claims; complaints log	Marketing / QA Lead	At point of publication
AI Procurement	Risk register; procurement policy; data protection; legal review	Vendor due diligence checklist; DPIA; contract clauses; risk mitigation plan	DPO / Procurement / QA	At procurement; annual review

QA Area	Policy & Process Touchpoints	Required Evidence	Responsibility	Review Frequency
Incident Management	Complaints policy; appeals procedure; risk committee; board reporting	AI Incident Log; trend analysis reports; escalation routes; committee responses	QA Lead / Risk Committee	Quarterly minimum

1.2 What Is New and What Is Not

A common concern is that AI in quality assurance represents an entirely new and uncharted territory. In practice, most of what is needed builds on existing QA infrastructure. The table below outlines areas where you may be extending existing practice and where genuinely new processes may be required.

Extend existing QA processes	New processes required
Assessment regulations → add AI-specific clauses	AI Incident Log (a dedicated, governed record)
Moderation notes → add AI-check confirmation	Vendor AI ethics due diligence checklist
External examiner briefing → include AI expectations	AI authorship declaration process for students
Academic conduct policy → update definitions and examples	DPIA/equivalent completed for every AI tool (not just new ones)
Programme spec templates → add AI permitted-use field	Demographic bias audit of learning analytics outputs
CPD records → tag AI-related training	AI Literacy Framework mapped to staff roles
Student induction materials → add AI guidance section	Standing AI governance committee or working group
Accessibility audits → include AI-generated content review	AI procurement ethics checklist

Section 2: Audit Readiness and What Reviewers Expect

Quality reviewers and inspectors are not expecting perfection. The sector is moving quickly and everyone is learning. What they will be looking for is evidence of intentional, documented, and improving governance and oversight. This section tells you what that means in practice for each evidence area.

Note: This document is not exhaustive and exists as a starting point to provide guidance and support. It does not constitute an ASIC inspection requirement.

2.1 The Evidence Standard

Broadly consistent standards will be expected:

- Is this evidence current, i.e. has it been produced or reviewed recently?
- Is it specific: Does it relate to AI, or is it generic policy that has not been updated?
- Is there a human decision behind it or was it simply generated and filed?
- Does it show learning and does the institution respond to what it finds?
- Is it accessible? Can staff and students access the guidance they need?

2.2 Evidence Area by Area

Assessment Validity & Moderation

Reviewers will look at a sample of moderation records. They will expect to see explicit reference to AI considerations not just general commentary. A moderation record that was written before AI was a realistic concern, and has not been updated, would not fulfil the evidence requirement.

> **What strong evidence looks like**
>
> - Moderation notes that include a specific AI-check field or comment (e.g. 'Reviewed for signs of AI generation and no concerns identified; assessment design discussed')
> - External examiner reports that reference AI (as a concern *or* as a feature of good practice)
> - Programme review minutes that document a structured discussion of assessment validity in light of AI capability
> - Assessment briefs that include a clear statement of what AI use is and is not permitted and how this is enforced

Academic Conduct & Integrity

Reviewers will want to see that your academic conduct framework has been actively updated (not just that a legacy policy technically covers AI as a form of misconduct). They will also look at how cases are managed.

> **What strong evidence looks like**
>
> - A conduct policy that explicitly defines AI-generated content, addresses the spectrum of use (from legitimate support to full substitution), and sets proportionate consequences
> - A student-facing AI declaration that is submitted with assessed work and retained on file

- A case log that records AI-related conduct concerns, the process followed, and outcomes
- Evidence that panels and decision-makers have been briefed on how to handle AI cases fairly and consistently

Data Protection & Procurement

This is an area where documentation rigour is non-negotiable. In the UK a DPIA (Data Protection Impact Assessment) is a legal requirement for high-risk processing, not a best-practice aspiration. Internationally, various in-country equivalents exist, and best-practice aligns with having equivalent documentation in place. Reviewers who find AI tools in use for high-risk processing without the appropriate documentation will note this as a significant concern.

What strong evidence looks like

- A DPIA register/ equivalent that lists every AI tool in use, the date of the DPIA/equivalent, the DPO (Data Protection Officer)/equivalent sign-off, and the review date
- Vendor contracts that include explicit data processing clauses
- A procurement checklist that includes AI-specific questions (e.g. model transparency, bias testing, data residency)
- Evidence that the DPO/equivalent is involved in procurement decisions at the point of evaluation, not retrospectively

Staff Capability & CPD

Reviewers will increasingly look beyond whether CPD was offered to whether it was effective and whether it reached the right people. Patchy uptake, particularly among part-time, associate, or professional services staff, will be scrutinised.

What strong evidence looks like

- A CPD log that tags AI-related training, broken down by staff group and role
- An AI Literacy Framework that defines what different roles need to know and be able to do
- Appraisal documentation that includes AI literacy as a development area where relevant
- Qualitative evidence (such as observation notes or staff feedback) that training has changed practice

Student Experience & Communication

Reviewers may speak directly with students. They are likely to ask whether students understand the AI guidance they have been given, whether it was clear, and whether it felt fair. Inconsistency across modules or programmes is a key concern.

What strong evidence looks like

- Module handbooks or briefs that contain a specific, consistent AI guidance section
- Induction records showing AI guidance was covered, with attendance or confirmation records
- Student survey data or focus group notes that include AI-related questions
- A student-facing AI declaration process that is consistently applied across the institution

Section 3: The AI Incident Log: Building It in Practice

The AI Incident Log is one of the most important pieces of governance infrastructure your institution can put in place. This section gives you everything you need to set one up, run it, and use it to drive improvement.

Note: The log should not be used to find fault and blame, its purpose is to serve as a learning instrument and facilitate the building and maintenance of a strong quality culture approach to AI governance.

3.1 What Counts as an AI Incident?

Institutions sometimes interpret 'incident' narrowly limiting it to dramatic failures or formal complaints. The log should capture a much broader range of events:

- AI tools producing incorrect, misleading, or harmful outputs (including hallucinations)
- Student or staff complaints about AI, including concerns about fairness or transparency
- Cases where AI was used in a way inconsistent with policy by staff or students
- Assessment concerns where AI may have affected the validity of a submission
- Accessibility failures caused by AI-generated content
- Chatbot or virtual assistant responses that gave students incorrect information
- Vendor outages or failures that affected teaching or assessment
- Bias concerns arising from learning analytics or algorithmic decision-making
- Appeals or complaints linked to AI-related conduct decisions

Important note on culture
For the log to work, staff must feel safe to report incidents without fear of blame. Logging an incident should be framed as a positive, professional act with the log serving as evidence of responsible practice.

Consider naming this the 'AI Learning Log' in student and staff-facing communications if 'incident' creates anxiety.

3.2 The AI Incident Log: Recommended Fields

Field	What to record	Why it matters
Reference number	Unique ID for each entry (e.g. AI-2024-001)	Enables tracking, cross-referencing, and trend analysis
Date reported	Date the incident was logged (not necessarily when it occurred)	Allows you to track reporting lag and timeliness
Date occurred	When the incident actually happened (if known)	Reveals patterns - e.g. clusters around assessment deadlines
Type of incident	Category: output failure / conduct / complaint / access / analytics / other	Enables trend analysis by category over time
AI tool or system involved	Name of the tool, platform or model	Reveals which tools carry most risk; informs procurement review

Field	What to record	Why it matters
Who was affected	Staff / students / both — and approximate scale	Enables proportionate response and reporting to committees
Description	Plain-English account of what happened	Straightforward, factual and non-judgmental records lead to better learning and outcomes
Immediate action taken	What was done in the short term to address the incident	Demonstrates responsiveness and safeguarding intent
Escalation route	Who was notified and when (DPO/equivalent, department head, board, etc.)	Demonstrates governance awareness and appropriate escalation
Resolution or status	Resolved / ongoing / under review (with notes)	Tracks closure and flags items requiring ongoing attention
Lessons identified	What this incident tells us about policy, practice, or tools	Where you learn where issues exist and progress stems from
Action taken	Changes made to policy, practice, or procurement as a result	Provides evidence that learning is applied
Reviewed by	Name and role of person who reviewed the entry	Demonstrates oversight and prevents entries being filed and forgotten

3.3 Governance and Review Cadence

The log is only as useful as the governance structure around it. Without regular review, entries accumulate without generating learning or action.

Cadence	Activity	Responsible
Ongoing	New entries logged by reporting staff or QA officer within 5 working days of incident	Any staff member / QA officer
Monthly	QA officer reviews all open entries; chases outstanding actions; updates status	QA Lead / Registrar
Quarterly	Full review by committee: trend analysis, patterns, actions arising, board summary prepared	QA Committee / Risk Group
Annual	Year-end analysis: categories, frequency, resolution rates, lessons applied; fed into annual quality report	QA Lead + SLT sign-off

Section 4: RAG Self-Assessment Tool: Evidence Gaps

Use the following RAG (Red/Amber/Green) tool to assess the current state of your AI-related QA evidence. Be honest: an accurate picture of current gaps is far more useful than an optimistic one.

Note: This is a working document for internal QA use not a formal inspection instrument.

Evidence Area	RED: Not yet in place	AMBER: Partial / inconsistent	GREEN: Consistent and evidenced
Assessment Regulations	Regulations have not been reviewed or updated to address AI	Some updates made but inconsistent across programmes or not formally approved	Regulations explicitly address AI; reviewed in current cycle; externally scrutinised
Moderation Records	No AI-specific checks recorded in moderation notes	Some moderators note AI considerations; practice is inconsistent	All moderation records include AI-check confirmation; reviewed by QA
Conduct Policy	Policy has not been updated; no AI-specific definition or examples	Policy updated but student-facing guidance is unclear or inconsistent	Policy current; student guidance clear; declaration process in use; case log maintained
DPIA Register	No DPIAs/equivalent completed for AI tools; no register exists	Some tools have DPIAs/equivalent; register is incomplete or not reviewed	AI tools have current DPIAs/equivalent; register maintained by DPO/equivalent; reviewed at procurement
Vendor Contracts	Contracts not reviewed for data processing clauses	Some contracts reviewed; AI-specific clauses absent or weak	All active contracts include explicit AI data clauses; new contracts require ethics checklist
AI Incident Log	No log exists; incidents handled informally	Log exists but entries are sparse, stale, or not reviewed	Active log; quarterly review; trend analysis; actions fed back to committee
CPD Records	No AI-specific training recorded; no framework	Some training offered; uptake patchy; no framework to map against	Framework in place; training logged by role; appraisal linkage; completion tracked
Student Guidance	No AI guidance given to students; or only generic policy reference	Guidance exists but varies by module or programme; not consistently delivered	Clear guidance at induction; module-level specifics; declaration process; evidenced delivery
Programme Specs	No AI-related content in specs; ILOs not reviewed	Some specs updated; AI-use statements absent or vague	All specs include AI-permitted-use statements; ILOs reviewed; committee-approved

Evidence Area	RED: Not yet in place	AMBER: Partial / inconsistent	GREEN: Consistent and evidenced
Accessibility Audit	No audit of AI-generated content for accessibility	Audit of some content; captioning or alt-text records incomplete	Full audit cycle; all AI-generated materials checked; alternative formats documented

Using your RAG results

Step 1: Complete the RAG individually, then compare across your team. Expect differences in opinion and examine these areas.

Step 2: List all RED items. These are your immediate priorities.

Step 3: For each RED item, assign a named owner and a target date.

Step 4: Review AMBER items. Identify which are closest to GREEN and tackle those first.

Step 5: Share your RAG results with your quality committee as a standing agenda item. Repeat every six months.

Section 5: Working Documents & Policy Templates

This section contains working templates you can adapt for your institution.

Each is designed to be functional immediately: edit the bracketed fields and adjust the content to reflect your institutional context and regulatory framework.

Template 1: AI Acceptable Use & Authorship Disclosure Policy - Core Components

The following sets out the minimum components that an institution-wide AI Acceptable Use Policy should contain.

Note: This is a framework, not a full drafted policy. Your legal and academic teams will need to develop this in full.

Required components: AI Acceptable Use & Authorship Disclosure Policy

1. SCOPE: Who the policy applies to (students, staff, associate/part-time staff, research students)

2. DEFINITIONS: Define 'AI tool', 'generative AI', 'AI-assisted work', and 'AI-generated work' in plain terms

3. PERMITTED USES: What AI tools can be used for, in what contexts, with what level of disclosure required

4. RESTRICTED AND PROHIBITED USES: Specific contexts where AI use is not permitted (e.g. certain assessment types)

5. DISCLOSURE REQUIREMENTS: How students and staff must declare AI use in submitted or published work

6. ACADEMIC CONDUCT: How undeclared or impermissible AI use is treated under your conduct framework

7. STAFF USE: Guidance for staff using AI in teaching, feedback, assessment design, and research

8. DATA AND PRIVACY: What staff and students must not input into AI tools (personal data, confidential information)

9. ACCESSIBILITY: How the policy applies equitably to students using AI as an accessibility tool

10. REVIEW: Who owns the policy, when it will be reviewed, and how updates will be communicated

Template 2: Student AI Authorship Declaration

The following is a ready-to-use student declaration. It should be submitted alongside assessed work. Adapt the wording to match your assessment regulations and submission system.

AI AUTHORSHIP DECLARATION — STUDENT SUBMISSION

Module: _____

Student ID: _____

Assessment title: _____

Please tick the statement that accurately describes your use of AI tools in this submission:

□ I have not used any AI tools in producing this work.

□ I have used AI tools for permitted purposes only (e.g. grammar checking, initial research, accessibility support). I have not used AI to generate any substantial part of the submitted content.
Details of use: _____

□ I have used AI tools in the production of this work as permitted by my [module guidance]. I have declared all such use within the work itself in accordance with [institution] guidance.

I confirm that this declaration is accurate and that I understand the academic conduct implications of a false declaration.

Signature: _____ Date: _____

Template 3: AI Procurement Due Diligence Checklist

Complete this checklist before approving any new AI tool for institutional use. File with the DPIA, vendor contract, and risk register entry.

Question	Yes	No	N/A
Has a DPIA/equivalent been completed for this tool?	□	□	□
Has the DPO/equivalent reviewed and signed off the DPIA/equivalent?	□	□	□
Has the vendor confirmed data residency (where data is stored and processed)?	□	□	□
Does the contract explicitly prohibit use of institutional/student data to train third-party AI models?	□	□	□
Has the vendor provided documentation on how the model was trained and what data was used?	□	□	□
Has the tool been reviewed for accessibility compliance (WCAG 2.1 AA minimum)?	□	□	□
Has a bias or fairness review been completed or requested from the vendor?	□	□	□
Is there a clear process for reporting errors, harmful outputs, or data breaches to the vendor?	□	□	□
Does the contract include provisions for data deletion on contract termination?	□	□	□
Have staff who will use the tool received appropriate training or guidance?	□	□	□

Question	Yes	No	N/A
Has a pilot or trial period been agreed before full deployment?	□	□	□
Has the risk been recorded on the institutional risk register?	□	□	□
Has the tool been reviewed against relevant regulatory guidance (ICO, sector body)?	□	□	□
Is there a named institutional owner for this tool's ongoing governance?	□	□	□

Any 'No' answer should be resolved before sign-off. Escalate to the DPO/equivalent and relevant senior leader. Record the outcome on the risk register.

Template 4: Quarterly AI Governance Committee Report Outline

Recommended structure for quarterly report to QA committee or board:

SECTION A: AI Incident Log Summary

- Number of new incidents logged this quarter (total and by category)
- Status of outstanding incidents from previous quarters
- Notable incidents requiring committee attention
- Trend observations (increasing / stable / decreasing by category)

SECTION B: Policy and Evidence Status

- Any policy updates or reviews completed this quarter
- Any new AI tools approved or under consideration
- DPIA register status update/equivalent
- CPD completion update (headline figures by staff group)

SECTION C: Assessment and Academic Integrity

- Number of AI-related conduct cases (without personal details)
- Any assessment validity concerns raised by external examiners or moderators
- Any appeals linked to AI conduct decisions

SECTION D: Actions Arising

- Actions from previous meeting: status
- New actions recommended: owner, deadline

SECTION E: Items for Board Awareness

- Any items requiring board-level attention or decision

Closing Note: This Is Ongoing Work

AI capabilities, regulatory expectations, and institutional practice are all moving quickly. QA frameworks that are appropriate today may need updating within twelve months. This should not prevent QA leadership and teams from engaging with AI QA. It is a reason for building systems that are designed to learn and adapt.

The qualities needed for a strong quality culture equipped to deal with the challenges AI presents are honesty (about current gaps) and a willingness to improve.

This handbook is a starting point to help you build the infrastructure to capture what is happening and respond to it.

The quality of your institution's AI governance will be built, incident by incident and decision by decision, in the years ahead.

Appendix A: Plain English Glossary

This glossary explains key terms you may encounter when discussing AI governance. It is deliberately non-technical.

Term	Plain English Explanation
AI (Artificial Intelligence)	Computer systems that can perform tasks which normally require human intelligence (such as writing text, analysing data, or generating images).
Large Language Model (LLM)	The type of AI that powers most text-based AI tools. It has been trained on vast amounts of text and predicts what words or sentences should come next.
DPIA (Data Protection Impact Assessment)	A legal assessment UK institutions must complete before using any system that processes personal data in a high-risk way.
DPO (Data Protection Officer)	The person in a UK institution responsible for data protection compliance. They should be involved in all AI procurement decisions.
AI Incident Log	A record of instances where AI tools produced problematic outputs, caused complaints, or were used inconsistently with policy. A governance essential.
Learning Analytics	The use of data about student behaviour (logins, submissions, attendance) to identify patterns, predict outcomes, or target support. These systems carry fairness risks.
Algorithmic Bias	When an AI system produces systematically unfair outcomes for particular groups, for example, being less accurate for students from certain backgrounds.
Academic Integrity	The expectation that student work is their own, honestly presented. AI changes what this means in practice and requires policy to be updated.
Acceptable Use Policy	A document that defines how staff and/or students may use AI tools. It should cover both permitted and prohibited uses, and consequences for misuse.
Authorship Declaration	A statement that a student submits alongside assessed work, confirming whether AI was used and in what way. Increasingly common in higher and further education.
Hallucination	When an AI tool confidently produces incorrect information e.g., fabricating a reference or a statistic. A significant risk in educational contexts.
Open-source AI	AI models whose underlying code is publicly available. These may be deployed by institutions directly, which creates different governance responsibilities.
Agent / Agentic AI	An AI system that can take sequences of actions autonomously to complete a goal, not just responding to a single prompt, but planning steps, using tools, browsing the web, writing and executing code, or

Term	Plain English Explanation
	interacting with other systems with minimal human input. Unlike standard generative AI, agentic AI can initiate actions, not just produce content. *Note: This raises additional leadership and governance questions around oversight, accountability, and unintended consequences*.
Generative AI	AI systems that can produce text, images, code, or other content in response to a prompt. Examples include ChatGPT, Microsoft Copilot, and Google Gemini.
CPD - Continuing Professional Development	Ongoing professional training and learning activities undertaken by staff to maintain and enhance skills, including AI-specific upskilling.
ILOs - Intended Learning Outcomes	Statements describing what students should know or be able to do after completing a module or programme. Updated in the template to reflect AI capability expectations.
VLE - Virtual Learning Environment	Institutional digital platform (e.g., Moodle, Canvas, Blackboard) used for online teaching, learning resources, and assessment delivery.
WCAG - Web Content Accessibility Guidelines	International accessibility standards (e.g., WCAG 2.1 AA) that ensure digital materials are usable by learners with disabilities.
DSA - Disabled Students' Allowance	UK support scheme providing adjustments, equipment, or assistance to ensure accessibility and inclusion for disabled students.

