Reframing AI Governance & Education

Beyond User Certification: Governing AI as Embedded Infrastructure

A policy response to the BCS AI Skills and Adoption Report 2026 *(by Michael Gbadebo)*

Executive proposition

The BCS report identifies a real and consequential problem: AI adoption is running ahead of organisational governance and skills. Its evidence shows widespread unapproved use, shallow training, weak confidence in responsible use, and strong demand among BCS members for recognised standards. This briefing accepts that diagnosis while challenging the proposed locus of intervention. As AI becomes an embedded capability within operating systems, search, development environments, productivity suites and enterprise platforms, responsible adoption cannot be governed primarily through credentials attached to individual users. A durable framework should combine proportionate human competence with vendor accountability, architectural controls, procurement standards, organisational governance and a national education strategy centred on durable cognitive and epistemic capabilities.

1. Engage the BCS evidence on its own terms

BCS surveyed 844 members between 25 June and 30 July 2026. The respondents were predominantly UK-based and included IT professionals and leaders across sectors and organisation sizes. The findings therefore provide valuable evidence about the technology profession, but they should not be treated as a nationally representative survey of the entire UK workforce.

The survey identifies AI as the leading technology skills gap: 35% named it their single biggest gap and 70% placed it in their top three. More than 42% reported using AI tools that their organisation had not formally provided or approved; 59% identified both accuracy/trust and security/data protection as leading adoption challenges; and 49% cited lack of skills. BCS members also strongly supported recognised ethical training or certification for workplace AI users, with an average response of 4.06/5 and 52% selecting the strongest agreement.

These findings establish a demand for stronger standards and better judgement. They do not, by themselves, establish that universal end-user certification is the most effective mechanism for achieving those outcomes. That distinction is central to this response.

2. Where the BCS diagnosis and this framework agree

There is substantial common ground. AI adoption is accelerating; governance is uneven; training is often too shallow; users need to question and verify AI output; leadership needs greater technical understanding; and high-impact technical work requires accountability. BCS itself explicitly argues that training should build workers' judgement rather than merely teach operation of a particular tool.

The disagreement is therefore not about whether people need competence. They do. It is about where competence sits within the governance architecture, and how responsibility should be distributed between users, employers, executives, software vendors and platform architects.

3. AI is becoming an embedded capability

The policy category of an 'AI user' is becoming less stable. Generative assistance, summarisation, recommendation, search augmentation and code generation are increasingly delivered as features inside software that organisations already procure, and employees already use. This changes the governance problem.

First, not all AI use is equivalent. There is a material difference between AI functionality embedded in an approved enterprise product, AI functionality deliberately enabled by an organisation, and an employee independently introducing an external service that creates new data-transfer or security risks.

Second, 'shadow AI' is therefore an incomplete governance category. BCS is right that more than two in five respondents report using unapproved AI tools. But the same phenomenon can contain both individual experimentation and structural failure: organisations may not know what AI capabilities are already present in their approved software, while employees may adopt external tools because formal procurement has not kept pace with practical demand. BCS itself notes that the finding can be read both as enterprising workers and as governance running behind practice.

Third, tool-specific training depreciates quickly. BCS reports that training is often focused on operating a particular tool and that course material can become dated as the technology changes. **The more durable objective is therefore not mastery of a product interface but competence in evaluation, verification, data handling, risk recognition and appropriate delegation to automated systems.**

4. From user certification to layered governance

BCS recommends that anyone using AI at work should complete recognised ethical training or certification, while recommending professional registration for technologists in high-stakes IT roles. The first recommendation should be treated cautiously, not because competence is unimportant, but because a universal credential risks confusing a baseline of awareness with responsibility for risks that are principally controlled elsewhere.

A person can be trained to identify hallucinations and still be unable to control an enterprise platform's retention settings, an API's data boundary, a vendor's model update, an integration's permissions, or an organisation's procurement decision. Training is therefore one layer of control, not the control plane itself.

A more robust model is layered:

1. **Vendor accountability.** Vendors should provide clear data-flow documentation, privacy and retention controls, provenance information where applicable, model/version change information, security assurances, documented limitations and meaningful auditability.

2. **Architectural governance.** Organisations should map where AI capabilities exist across approved software, integrations and infrastructure; define access and data boundaries; log material AI interactions where appropriate; and assess risk when capabilities or models change.

3. **Procurement governance.** AI-enabled products should be assessed as part of ordinary technology procurement rather than treated as exceptional applications. Contracts should address data use, retention, security, model changes, service continuity, audit rights and exit arrangements.

4. **Organisational leadership.** Boards and executive teams need clear ownership of AI risk, value and deployment. BCS itself calls for stronger technical leadership at board and executive level.

5. **Role-proportionate human competence.** Users should receive training appropriate to the risks of their work. High-impact technical and professional roles should face stronger competence and accountability requirements than low-risk routine use.

6. **Professional standards for high-stakes roles.** BCS's case for independent standards for technologists working in high-stakes AI environments is materially different from universal certification of everyone who encounters an AI feature. Those two policy instruments should not be conflated.

This distributes responsibility according to control. The person using an AI assistant remains accountable for appropriate judgement but is not made the sole governance mechanism for risks they cannot technically control.

5. A better interpretation of 'shadow AI'

The appropriate response to unapproved use is not simply prohibition. BCS reports that over 42% of respondents use AI tools their organisation has not formally approved and concludes that boards and employers should take this seriously.

The policy question should therefore be: **why did the organisation's governance boundary fail to contain or support useful adoption?** A mature response would combine discovery, risk classification, approved alternatives, rapid procurement pathways, technical controls and proportionate user education. Where an external service creates unacceptable data or security risk, it should be restricted. Where employees are using AI because the organisation has failed to provide a viable approved capability, the remedy may be organisational redesign rather than disciplinary enforcement.

6. Education: from AI operation to epistemic capability

The long-term educational response should not be built around teaching students how to operate today's AI products. Those interfaces will change. The durable educational objective is to develop people who can interrogate automated systems, evaluate evidence and make informed decisions about when to accept, challenge, verify or reject machine-generated outputs.

Five foundational capabilities should therefore be developed across the curriculum:

- **Epistemic judgement and output scrutiny:** verify claims against primary or authoritative sources, recognise uncertainty, distinguish evidence from assertion, and test generated outputs rather than treating fluent language as evidence of truth.

- **Critical reasoning and formal logic:** identify flawed premises, causal errors, unsupported inference and misleading framing in both human and machine-generated material.

- **Independent inquiry:** formulate precise questions, conduct structured research, compare sources and pursue lines of investigation without outsourcing intellectual responsibility to an automated system.
- **Creativity and conceptual synthesis:** develop original problems, hypotheses, designs and interpretations, using automated tools as instruments rather than substitutes for intellectual agency.
- **Cross-disciplinary AI fluency:** embed these practices in English, history, science, mathematics, computing and other subjects so that AI judgement becomes a general intellectual habit rather than a specialist IT topic.

This complements rather than rejects AI literacy. Students still need to understand how AI systems work at an appropriate level, how to use them effectively and how to manage privacy and security. The distinction is between durable capability and product-specific operation.

7. Strategic model comparison

Dimension	BCS emphasis	Embedded-capabilities framework
Problem	Skills deficit and weak governance	Skills deficit plus architectural and organisational governance
Workplace training	Recognised ethical training/certification; deeper judgement	Role-proportionate training focused on judgement, verification and risk
High-stakes technical roles	Independent professional registration	Independent professional standards, combined with organisational/vendor controls
Shadow AI	Clear policies and development plans	Discovery, classification, approved alternatives, technical controls and proportionate enforcement
Accountability	Users, professionals, organisations and leadership	Responsibility distributed according to who controls the relevant risk
Vendor/platform role	Part of procurement and governance	Explicit control layer: data, security, auditability, model/version change and provenance
Education	AI skills and judgement	AI literacy plus durable reasoning, epistemic judgement, inquiry and synthesis

8. Recommendations

For government:

- Treat recognised AI training as a baseline competence mechanism, not as a substitute for organisational and technical controls.

- Develop proportionate standards for AI-enabled products covering data use, privacy, security, auditability, transparency of material model changes and organisational control.
- Support a professional standards regime for high-impact AI and technical roles, while avoiding the assumption that every interaction with an AI-enabled software feature requires the same credential.
- Reform school curricula around durable capabilities: epistemic judgement, critical reasoning, source verification, independent inquiry and creative synthesis, while retaining practical AI literacy.
- Develop public-sector procurement and assurance capability so organisations can evaluate AI-enabled products before and after deployment.

For enterprise boards and executives:

- Map AI capabilities already present in the organisation's software estate, including embedded and automatically enabled functionality.
- Replace blanket 'shadow AI' prohibition with an active governance model: discover, classify, assess, approve, monitor and restrict according to risk.
- Require procurement teams to assess data flows, retention, security, auditability, model changes and vendor dependency.
- Provide role-specific training that teaches employees when to use AI, how to verify outputs and what information must not be exposed.
- Give technical leadership a formal role in AI strategy, procurement and governance.

For education leaders:

- Move assessment away from rewarding unverified generated text and towards evidence, original synthesis, reasoning and source validation.
- Teach students to interrogate AI outputs across subjects rather than confining AI judgement to computing lessons.
- Develop assessment methods that test independent reasoning as well as effective and transparent use of automated tools.

9. Conclusion

The BCS report makes a valuable contribution by documenting a genuine gap between AI adoption, organisational capability and governance. Its survey shows that technology professionals want stronger standards, deeper training and clearer accountability.

The next policy question is where those controls should reside. In an environment where AI is increasingly embedded in ordinary software, the category of the 'AI user' is no longer sufficient as the primary unit of governance. Human judgement remains essential, but it should sit inside a wider control architecture that assigns responsibility to the actors best positioned to manage each risk.

The objective should therefore not be to choose between skills and infrastructure. It is to connect them: competent people, accountable organisations, governable architectures and responsible vendors. Education should develop the durable cognitive capabilities required to supervise automated systems, while regulation and procurement should ensure that those people are not asked to solve through personal vigilance problems that are fundamentally architectural.

Author: Michael Gbadebo | mothcloud.com | 26/09/2026