


AI IN FIRE ENGINEERING

Strategy and Position Paper

Professional quality, organisational readiness and responsible exploration



BUILDING CAPABILITY. CONNECTING THE PROFESSION. SHAPING THE FUTURE.

A positioning paper on professional quality, organisational readiness and responsible exploration. Not a technical standard, implementation guide or endorsement of any product, supplier or political position.

Core question

How should a safety-critical profession respond to AI while strengthening professional quality, judgement and public trust?

Purpose

To help fire engineering organisations develop the competence, confidence and conditions required to explore AI responsibly, strengthen professional quality and contribute to a shared profession-wide response.

Authors

Liam Charles Obuobie
Neal Butterworth
Peter Lawrence
Roderic Jones
Reece Walker
Simón Santamaria Garcia
Nimra Akhtar
Leslie Marshall
Sophia Ifezue

Corresponding author: Liam Charles Obuobie, Liam.Obuobie@DFC.co.uk

Publisher: SFPE UK Chapter

Contents

Executive summary.....	3
1. The world is changing	5
1.1 Change is already under way	5
1.2 UK fire engineering is changing too	5
1.3 Access, readiness and assurance are different	6
1.4 The profession is not starting from zero	7
1.5 Why collective action matters	7
1.6 Quality is the test	7
2. Change is not new, but AI is different	7
2.1 Concern is part of responsible exploration	8
3. What must remain human	9
3.1 Output quality and professional quality	9
3.2 What clients and society rely on.....	10
3.3 What must remain visible	11
4. Organisational readiness for responsible exploration	12
4.1 Readiness before reliance	12
4.2 Exploration, pilot and operational reliance	14
4.3 Competence, capability and challenge	14
4.4 A practical starting point	15
4.5 Two bounded examples	15
4.6 Governed information and organisational memory	16
4.7 Leadership controls and proportionality	16
5. Our position and the next step	17
5.1 Our position	17
5.2 Improving fire-engineering standards	17
5.3 Why act together.....	18
5.4 How participation will begin	19
5.5 What SFPE UK will do next	19
5.6 Call to action	19
Definitions	20
References.....	21

Reading note: The Executive Summary provides the complete short version of the paper. The main text then moves from the changing UK fire-engineering context and professional choice, through enduring human value and organisational readiness, to SFPE UK's position, participation routes and next steps. Leaders may wish to focus on Sections 1, 4 and 5; practitioners and reviewers on Sections 2–4; and clients, regulators and partner disciplines on Sections 3 and 5.

Executive summary

AI is already changing how fire-engineering organisations communicate, use information, develop knowledge and make decisions. At the same time, UK fire engineering is moving towards stronger competence, clearer responsibility and greater professional accountability. Change is not new, but AI is different in professional effect: it can enter more parts of work with little friction, shape outputs without being visible and influence how people learn as well as what they produce. The question is no longer whether the profession will encounter AI, but whether organisations are ready to shape its influence without weakening professional standards.

Used well, AI may help competent practitioners find and compare authorised information, identify inconsistencies, communicate clearly and release time for judgement, coordination, mentoring and deliberate learning. Used poorly, it may lend credibility to weak reasoning or produce correct, polished work without developing the person behind it. A deliverable may pass review while the practitioner remains unable to explain, defend or adapt the reasoning independently. AI may not create weak practice, but it can make it faster, more convincing and harder to detect; it may also allow professional progression to outpace professional growth.

Access, organisational readiness and professional assurance are distinct. Access means people can use a system. Readiness means that purpose, rules, information, competence, people leadership and support surround that use. Assurance means that a particular output or decision has received scrutiny proportionate to its consequences. Progress in one does not demonstrate progress in the others, and formally retaining human accountability does not by itself show where AI has already influenced sources, assumptions, methods, framing or wording.

AI is already influencing professional work. Does your organisation know where, can it explain where that influence begins and ends, and can it show that people are becoming more competent rather than only more productive?

For leaders, responsible use requires more than approving a platform or issuing a policy. Organisations need a defined professional problem and quality outcome, clear boundaries, governed information, competent users and reviewers, and assurance that rises with consequence. They also need accountable leadership across the hierarchy and working conditions in which people can disclose AI use, admit uncertainty or dependence, challenge fluent output, ask for human review and recommend stopping without disproportionate blame. Psychological safety is not an absence of accountability. It is the practical ability to question, learn and raise difficult evidence when pressure, hierarchy or apparent success might otherwise silence it. Technical authority does not automatically provide the feedback, coaching and inclusive leadership skills needed to create those conditions. Leaders must seek challenge, respond visibly when it is raised, distinguish bounded investigation from operational reliance and be prepared to continue, revise, delay or reject a use according to the evidence.

If leadership means serving the conditions in which others can do their best work, are leaders willing to make their own decisions open to challenge and show that AI is developing capable professionals rather than only capable-looking outputs?

Fire engineering's enduring value lies in defining the real problem, applying science and first principles, navigating uncertainty, integrating disciplines, explaining consequences, building trust and taking responsibility for a defensible decision. AI can extend capability; it does not create competence. Professional quality must therefore be judged through both the work and the people behind it: the strength of reasoning, coordination, learning, accountability, client advice and public value, and whether practitioners can understand, explain and act responsibly beyond the assisted output.

The profession also faces a capacity challenge. AI cannot substitute for competent fire engineers, but well-governed tools may help scarce expertise reach further, make authorised knowledge easier to navigate, support supervised development and help people from adjacent disciplines build fire-specific understanding.

Used to scaffold questions, comparison, feedback and deliberate practice, AI may deepen learning. Used to replace first-principles reasoning, explanation and formative experience, it may bypass it. Technology must remain a bridge into competence, never a shortcut around it.

Each organisation remains responsible for its own people, systems, clients and decisions, but none should have to answer every professional question alone. If assisted performance is mistaken for developed competence, individual gaps may become organisational assumptions and, when repeated across firms, distort the profession's view of the capability on which future safety decisions depend. Shared language, benchmark questions, anonymised experience and open discussion of limitations can reduce repeated mistakes and help the profession compare which uses deepen learning, which bypass it and how authority remains visible, while protecting proprietary methods, client information and genuine commercial differentiation.

Public debate ranges from immediate concerns about work, learning, misinformation, environmental demand and human rights to severe or catastrophic future risk. This paper does not attempt to settle those debates. Its practical responsibility is to help a safety-critical profession investigate the uses it can influence now with competence, evidence, proportionality and accountable human judgement. Existing legal, confidentiality, insurance and professional obligations remain, and may be harder to evidence where AI influence is hidden.

SFPE UK's position: AI should augment, not displace, competent professional judgement. Adoption is not the objective; better fire-engineering decisions and outcomes are.

SFPE UK will convene rather than prescribe. It will remain independent of products and providers, present this position, invite wider professional engagement, convene an industry roundtable and begin specialist participation in parallel. Those forums will test priorities, compare evidence of professional benefit and competence, identify limitations and determine where later guidance, research or collaboration can add most value. Evidence and practical implications will be returned to the profession in an accessible form.

If this summary has made you question whether your organisation is ready, do not stop at reflection. Bring your evidence, your experience and your challenge to the next stage. Join the expedition, not to endorse every use of AI, but to investigate possibilities openly and help the profession decide what deserves to continue, what should stop and how AI can strengthen both professional work and the people who deliver it.

PART I | The Professional Case

1. The world is changing

AI is entering fire engineering at the same moment that the profession is being asked to show stronger competence, ethical conduct, clearer responsibility and professional accountability [1]. Change itself is not new: fire engineers have adapted to modelling, digital information and successive regulatory reform before. AI is different in professional effect because it can enter more parts of work with less friction, shape outputs without being visible and influence how people learn as well as what they produce. One development is changing what fire engineers can do; the other is raising expectations of how they justify what they do. Together, they create a sharper professional test: will AI reinforce the standards now being built, or allow unresolved weaknesses to travel further and faster?

1.1 Change is already under way

That test is not theoretical. AI is already used in everyday work for search, transcription, drafting, comparison and analysis. In an SFPE UK Chapter survey of 123 fire-safety practitioners, 70% reported using AI professionally at least occasionally, with large language models the most common type [2]. The aggregated results are intended to be published or made available through the initiative. In some organisations use is visible and encouraged; in others it is cautious, informal or hidden.

Access does not equal control. The same survey found that 27% of respondents worked without a formal organisational AI policy and 68% were unaware of professional-body guidance [2]. Without clear purpose, governed information, suitable learning and permission to challenge, use may outpace the organisation's ability to recognise, understand or assure it.

AI is already influencing professional work. Does your organisation know where?

That question is not an accusation; it is an invitation to lead. Making current use visible allows organisations to set boundaries before informal habits become embedded, support practitioners who are trying to work responsibly and focus on uses most likely to improve knowledge, coordination and professional judgement. The opportunity is to shape stronger practice before habit sets the standard. That matters particularly because UK fire engineering is not standing still.

1.2 UK fire engineering is changing too

Government reform points towards a regulated fire-engineering profession, including future regulation of title and function and a central protected role for the fire safety strategy. The accompanying vision emphasises accredited education, supervised experience, ethics and continuing professional accountability [1]. The statutory regime is not yet complete, but its direction matters now: the profession is being asked to demonstrate not only technical knowledge, but who is competent, who is responsible and why safety-critical advice deserves trust.

Expectations of organisations are changing alongside those of individuals. The building-safety regime increasingly requires duty-holders and employers to demonstrate competence, coordination and reliable information. Draft BS 9994, issued for public comment in 2026, likewise proposes a lifecycle approach to fire strategies, competent professional involvement and guidance on responsible AI use [3]. It remains draft rather than settled guidance, but it shows that AI is already entering discussion about how fire-engineering decisions are recorded, validated and maintained.

Specialist applications are also moving beyond general office support. Peer-reviewed research has demonstrated AI-assisted smoke-flow prediction and performance-based design for complex atrium buildings [4]. One study does not establish routine practice, but it confirms that the profession is already confronting technical as well as administrative uses. The immediate need is not to predict a single future, but to build the judgement, evidence and organisational conditions needed to evaluate different uses well.

Capacity is another part of this context. Government has committed funding to expand fire-engineering education and address shortages in the building-safety workforce [5]. AI cannot substitute for competent fire engineers. Used well, however, digital and AI tools may help scarce expertise reach further by making authorised knowledge easier to navigate, supporting education and supervised development, and helping

people from adjacent engineering disciplines build the fire-specific understanding needed to contribute responsibly. They may accelerate learning when they support questions, comparison, feedback and deliberate practice. They may bypass learning when they replace first-principles reasoning, explanation, interpretation and the formative mistakes through which professional judgement develops. These are opportunities and risks to investigate, not predetermined outcomes. Technology should become a bridge into the profession and a means of extending competent practice, never a shortcut around competence.

1.3 Access, readiness and assurance are different

Three related conditions are often treated as one. **Access** means people can use a system. With AI, that access may begin through familiar everyday interfaces rather than a visible specialist workflow.

Organisational readiness means the organisation has purpose, rules, information, competence and support around that use. **Professional assurance** means a particular output or decision has received scrutiny proportionate to its consequences. In this paper, assurance is an umbrella term: its form may range from competent checking to independent review or specialist validation, depending on how much the output shapes the decision. Progress in one does not prove progress in the others.

This distinction matters because AI combines exceptionally low barriers to use with broad application, fluent output and influence that may not be obvious in the final deliverable. The same system may assist search, drafting, comparison, analysis, communication or learning without a single stable professional remit. A heavily used platform may therefore sit within weak governance, while a cautious organisation may have strong assurance for a small number of uses. AI may not create weak practice, but it can make that practice faster, more convincing and harder to detect, particularly when practitioners rely on fluent outputs beyond their competence. Leaders need to know not only whether AI is used, but for what purpose, with which information, by whom, under what review and with what effect on professional quality and learning.

Figure 1 illustrates the conceptual gap that can open when access grows faster than organisational readiness and professional assurance. The separation represents mismatch between these conditions, not a maturity ranking.




Figure 1. Conceptual gap between access to AI, organisational readiness and assured professional use.

1.4 The profession is not starting from zero

The SFPE Foundation's Grand Challenges white paper already sets out a ten-year research and education agenda for digitalisation, AI and cybersecurity in fire engineering [6]. This UK paper complements that work: the Foundation maps long-term technical priorities and research needs; this initiative focuses on professional quality, organisational readiness and responsible practice.

Regulators, standards bodies, professional institutions, academia, technology organisations and businesses are already active. Fire engineers are experimenting with administrative support, document comparison, modelling and other specialist applications. The problem is not a lack of activity, but fragmented learning, uneven capability and inconsistent professional expectations.

Visible use is easier to govern than hidden use. Where approved systems, information boundaries and review routes are unclear, practitioners may avoid useful work or proceed informally. Open discussion and proportionate boundaries help leaders understand what is happening before unreliable practice becomes routine.

In this paper, **AI** is an umbrella term. Current uses include large language models and generative AI; specialist work may involve machine learning, BIM-connected systems, computer vision or surrogate models. These approaches are not interchangeable. Where an established engineering method or transparent automation solves the problem reliably, it should remain the preferred approach.

Fire engineering also differs from many general business settings because a small contribution can influence a much higher-consequence judgement. Drafting, retrieval, comparison or modelling support may sit inside a safety-related decision. Assurance should follow how much the output shapes the professional decision, not simply the type of system being used.

1.5 Why collective action matters

Each organisation must decide how AI fits its own responsibilities, clients and systems, but none should have to answer every professional question alone. Shared language, benchmark questions, anonymised experience and open discussion of limitations can reduce repeated mistakes, improve consistency and give smaller organisations access to wider evidence. Practice, academia, regulators and professional bodies see different parts of the problem; bringing those perspectives together is more valuable than parallel experimentation that never becomes collective learning.

This does not require organisations to surrender proprietary methods, client information or commercial advantage. It means strengthening the shared professional baseline while organisations continue to compete through their people, implementation and relationships. If every organisation learns alone, the profession will repeat avoidable mistakes at unnecessary cost.

1.6 Quality is the test

AI should not be judged by licences issued, prompts submitted or documents produced. The relevant test is whether it improves professional quality: reasoning, consistency, learning, coordination, client advice and safety. Does it help competent practitioners work better, or make weak practice easier to scale?

Fragmented use can produce uneven service, duplicated effort and conflicting expectations between practitioners, reviewers and clients. A shared professional approach cannot remove variation, but it can make expectations clearer and help organisations compare evidence before changing practice. Change is already happening. The profession must now decide whether it will shape that change deliberately or allow custom, convenience and commercial pressure to shape it by default.

2. Change is not new, but AI is different

Fire engineering has adapted to major technologies before. CFD, evacuation modelling, BIM, spreadsheets and automation can all be misunderstood, misapplied or over-trusted, and the same duties of competence, checking and accountability still apply. AI is different in professional effect because it can enter more tasks through familiar interfaces, move between technical and non-technical work, shape outputs without being obvious and influence how people learn as well as what they produce. The principles remain familiar; the conditions in which they must operate have changed.

The choice is therefore not between enthusiastic adoption and wholesale rejection. Both are too simple and reductive for a safety-critical profession. It is whether AI use develops through disconnected, partly hidden decisions or through visible, accountable practice shaped by professional competence and shared safeguards. Used well, AI may improve access to knowledge, reduce avoidable administration, support performance-based methods and release scarce professional time for judgement, mentoring and relationships. Used poorly, it may conceal uncertainty, accelerate weak work and bypass the experiences through which competence develops. Its value depends on the task, the user, the evidence and the organisation around it.

Quality is more than technical correctness alone. It depends on whether the right problem was framed, suitable evidence and uncertainty were understood, other disciplines were integrated and the advice can be defended in context. A plausible answer may still be poor professional work if it addresses the wrong question or leaves ownership unclear.

If AI can help competent practitioners interrogate more evidence and identify problems earlier, are your QA processes ready to capture that benefit without letting confidence outrun competence?

That risk is exactly why competent investigation matters. Used within clear boundaries, AI may allow practitioners to interrogate more authorised evidence, compare more options and identify inconsistencies earlier, while preserving human ownership of the judgement. Refusing to examine that possibility can carry its own professional cost; investigation is not the same as adoption, and evidence may still show that a use should not proceed.

Leaders cannot remove uncertainty, but they can decide how their organisations respond to it. They can set purpose, separate experimentation from reliance, protect learning, require proportionate review and give people authority to question or stop unsuitable use. Responsible practice means neither standing still nor moving faster than the organisation can govern.

This is also a commercial choice. Organisations can use AI to increase familiar output or to improve the value of professional time. The stronger case is more capacity for investigation, coordination, client dialogue, supervision and review. That benefit disappears if every saving is absorbed into higher throughput.

2.1 Concern is part of responsible exploration

People may reasonably worry about job security, lost learning, opaque decisions, environmental cost or the changing value of experience. Leaders should not dismiss these concerns as resistance; they may reveal real risk, weak communication or missing safeguards. Equally, broad claims about AI should not replace evidence about a defined use in a real professional setting.

Public debate also ranges from immediate harms to the possibility of severe or catastrophic future risk. The 2026 International AI Safety Report reviews evidence across misuse, malfunction and systemic risk while emphasising substantial uncertainty about future capabilities and outcomes [7]. These concerns should not be dismissed, and this paper does not claim to resolve them. Its narrower responsibility is to help a safety-critical profession make competent, proportionate and accountable decisions about the uses it can influence now. Uncertainty calls for stronger evidence and scrutiny, not paralysis or casual reassurance.

This is consistent with the Health and Safety Executive's 2026 approach to AI: existing goal-setting duties still apply, those who create risk are expected to manage it, and regulation should protect people while enabling innovation through proportionate controls [8]. Fire engineering should apply the same discipline to the professional risks and opportunities within its influence.

A constructive response is to explain the purpose, affected roles, approved boundaries and what will not change; invite challenge before decisions are fixed; provide safe routes to report hidden use or failure; and review workforce, learning and environmental effects as evidence develops. This reduces both fear-driven avoidance and pressure to appear uncritically enthusiastic [9, 10]. That still leaves the profession with a deeper question. If technology changes how work is produced, which parts of good fire engineering must remain visibly human and professionally owned?

3. What must remain human

Fire engineering protects people by bringing together science, human behaviour, design, systems, operations and regulation under uncertainty. Its value lies not only in reports or calculations, but in defining the real problem, understanding context, balancing competing needs, explaining consequences, earning trust, challenging where necessary and taking responsibility for a defensible decision.

AI can extend capability by helping people find information, compare evidence, test options and communicate. It does not create competence: the knowledge, experience and professional behaviour required to exercise sound judgement in context. Figure 2 summarises the boundary established in this section: AI may support professional activity, but the framing of the problem, integration of evidence, exercise of judgement and ownership of accountability must remain visibly human.




Figure 2. AI may support professional activity, while problem framing, competence, judgement and accountability remain human-owned.

3.1 Output quality and professional quality

A convincing output can look complete long before it has earned professional trust. Table 1 shows the questions that separate apparent quality from defensible fire-engineering work.

Table 1. Apparent output quality and professional quality.

Apparent strength	Professional question	Evidence of quality
Fluent, polished deliverable	Does it address the right problem, and who owns the judgement?	A clear brief and approval route, linked to the actual decision and named professional responsibility.
Confident technical language	Are the sources, assumptions and limits suitable and understood?	Authoritative evidence, visible uncertainty and an explanation that a competent reviewer can challenge.
Fast calculation or model output	Was the method appropriate and used within its valid range?	Method selection, input checks, sensitivity or benchmarking, and independent review where

		consequence requires it.
Consistent recommendation	Has it integrated other disciplines, constraints and affected people?	Recorded coordination, considered alternatives and a proportionate account of residual risk.

AI can improve presentation faster than it improves understanding. That creates two distinct competence risks. First, weak or incorrect work may appear credible. Second, the work may be correct while the practitioner behind it cannot explain, defend or adapt the reasoning independently. Traditional review may detect the first risk by checking the deliverable; it may miss the second if repeated output quality is mistaken for professional growth. Quality assurance must therefore test not only whether the work appears complete, but whether the person responsible understands the evidence, assumptions, method and judgement behind it.

Illustrative scenario | A polished report and apparent competence

An early-career practitioner uses an approved generative-AI system to prepare report drafts from authorised project information, previous examples and selected guidance. The reports are clear, technically acceptable and consistently pass competent review. Over time, that repeated output quality could be treated as evidence that the practitioner is ready for greater autonomy. Yet, when asked to explain why a fire scenario was excluded, defend a key assumption or adapt the strategy under client challenge, the practitioner cannot reconstruct the reasoning without AI support. The deliverables were acceptable; the development evidence was not.

A stronger approach would preserve the drafting benefit while deliberately reinvesting released time in learning. The practitioner would expose assumptions, verify sources, compare alternatives and explain the final judgement to a competent reviewer. AI could generate questions, surface comparable cases and help rehearse challenge, but progression would depend on demonstrated understanding rather than polished output alone. Competent review is therefore more than proofreading: it tests the problem definition, evidence, method, uncertainty, coordination and ownership behind the deliverable, and whether the person responsible can reason beyond it.

This does not mean turning every AI-assisted deliverable into an oral examination. It means creating proportionate evidence that learning is keeping pace with assisted performance: through explanation, supervision, client discussion, varied tasks, reflective review or other suitable demonstrations of independent judgement. Emerging research involving university students and early-career knowledge workers suggests that the manner of AI use matters: using AI to scaffold thinking while retaining cognitive agency was associated with more favourable perceived learning outcomes than delegating core thinking to the system, even though both approaches offered similar immediate benefits [11]. AI could help develop fire engineers faster and more deeply than before by making grounded knowledge, challenge and feedback easier to reach. It could also allow professional progression to outpace professional growth if polished delivery is mistaken for independent capability. The relevant workforce question is therefore not whether people used AI, but whether its use increased what they can understand, explain and do responsibly without leaning on it.

That places a direct responsibility on leaders. AI increases the importance of knowing the people behind the work: what they understand, where they need challenge, how they learn and whether their confidence reflects genuine competence. High-performing teams do not leave those conditions to personality or goodwill. Leaders make time for questions, candour, feedback and reflection; they create visible routes for people to admit uncertainty, expose dependence and test their reasoning without lowering accountability. In an AI-supported workplace, those leadership behaviours are part of professional assurance because they help distinguish capability being developed from capability merely being borrowed.

3.2 What clients and society rely on

Clients may buy reports, drawings, calculations and meetings, but they rely on engineers for integrated judgement, clear advice, trusted relationships and accountable ownership. High-quality fire engineering makes assumptions and uncertainty visible, coordinates decisions across disciplines and produces advice that people can explain, challenge and act on. Efficiency matters when released time improves analysis, mentoring, affordability or safety, not when it merely increases output.

Commercial strength matters. Healthy organisations employ people, invest in learning, improve methods and remain capable of serving clients over time. Competition can raise standards and reward organisations that deliver real value. But commercial success is a means of sustaining professional service, not the purpose for which fire engineering exists.

Why did you become a fire safety engineer: to outperform the next organisation, or to shape a built environment where a human life does not need rescuing from fire in the first place?

The answer need not set public purpose against commercial reality. Strong organisations can protect proprietary work while helping to raise the baseline on which every client and member of the public depends. That balance reflects the profession's ethical duties of responsibility to society, accuracy and rigour, leadership and communication, and responsibility for the future consequences of technology [12]. If AI can make authorised knowledge easier to reach, reveal inconsistencies earlier and release experienced engineers to mentor, coordinate and challenge more effectively, investigating that opportunity is part of serving the profession well. The test is whether the benefit reaches professional quality and public safety, not simply the balance sheet.

Trust is built by listening, explaining limits, following through and being willing to deliver unwelcome advice. Relationships matter because fire-engineering decisions span clients, designers, contractors, regulators, operators and occupants. AI may help organise information or adapt an explanation, but it cannot carry the duty of care or professional courage needed to challenge a preferred design, commercial pressure or an established assumption.

Professional quality also depends on multidisciplinary influence. A recommendation that is sound in isolation may fail if it cannot be integrated with architecture, structures, building services, construction, operations or emergency response. The fire engineer helps others understand the consequences and find a proportionate route that the wider team can implement.

A thought experiment from professional sport makes the hidden cost of communication easier to see: **what if every word spoken or typed cost £1, including words later deleted?** The point is not to say less, but to be deliberate about purpose, audience, timing and what another person needs in order to act. Investing early in shared language can reduce repeated explanation when decisions must be made quickly. The same discipline can strengthen fire-engineering teams, client relationships and multidisciplinary work.

The practical question is: what communication will achieve the intended result with the least avoidable repetition? Clear context, acronyms, constraints, uncertainty and decision rights improve briefs, relationships and organisational memory. Information captured in usable form is also easier to govern and retrieve, turning scattered enterprise data into professional intelligence rather than an inaccessible archive.

AI makes this discipline harder to avoid. Useful support depends on clear purpose, audience, authorised information, assumptions and limits, followed by supervision and checking. Treating AI as a new team member can be a useful briefing metaphor, provided it does not imply personhood, reliability or transferred responsibility. The same practices strengthen shared language, data literacy and accessible organisational knowledge for graduates, clients, collaborators and experienced practitioners.

If communication had a visible price, would your organisation invest earlier in shared understanding and usable knowledge, or continue paying repeatedly through ambiguity, rework and inaccessible proprietary intelligence?

If AI-supported search, retrieval or synthesis can make authorised knowledge easier to find, experienced engineers may spend less time reconstructing context and more time mentoring, coordinating, testing assumptions and speaking with clients. That is not a guaranteed outcome; it depends on disciplined capture, reliable sources and leaders deliberately returning released time to higher-value professional work.

3.3 What must remain visible

Responsible AI use should strengthen the visibility already expected of professional work. Even where the model is difficult to explain, the decision must not be. Organisations should be able to show who defined the problem, which evidence was trusted, what assumptions and limits applied, where human judgement intervened and who approved the outcome. Human accountability may remain formally unchanged while

practical influence has shifted earlier in the work: AI may help select sources, frame alternatives, propose methods or shape wording before a competent reviewer sees the result. That influence may be appropriate, but its remit and material effect must be understood.

Are you giving authority to AI without realising it, and could you explain where that influence begins and ends?

The answer is not to treat every AI-supported choice as an improper transfer of authority. It is to define the system's role, keep material sources, assumptions and interventions visible, and ensure that competent people retain the information, time and authority needed to challenge or stop the work. "The system produced it" is not a professional explanation; the organisation must account for its own reasoning and actions. Existing accountability, confidentiality, insurance and legal obligations do not disappear when AI is used, and may be harder to evidence where influence is hidden or crosses organisational or jurisdictional boundaries. Existing duties of honesty, integrity, respect for life, competence, leadership and public interest still apply. The technology changes the setting in which those duties are exercised; the organisational challenge is to keep them effective when use is fast-moving, distributed or partly hidden. Values only shape practice when organisations give people the conditions, information and authority to act on them. That is why professional purpose must now become organisational readiness.

PART II | Organisational Readiness and Professional Position

4. Organisational readiness for responsible exploration

The competence and authority risks described in Section 3 become leadership responsibilities in the organisation. Many barriers to responsible AI use are organisational before they are technical. Cross-industry research links stronger outcomes to senior leadership, redesigned workflows, governance and workforce development, rather than access to tools alone [13, 14]. Practitioners may already be experimenting while leaders are still setting policy, approved systems, information boundaries and acceptable uses. Elsewhere, leadership uncertainty may suppress useful work and drive it out of sight. For a leader, this is the point at which abstract debate becomes operational responsibility. Organisational readiness enables people to work visibly, remain within competence and show how safety-critical activity is governed. It also helps direct effort towards uses that may improve knowledge access, coordination, mentoring and professional challenge.

If accountability remains human, can your organisation explain where AI influences its work today, and whether that influence is improving the standard of professional decision-making?

4.1 Readiness before reliance

Readiness does not mean knowing every technology, eliminating uncertainty or moving every team at the same pace. An organisation may be ready to test one bounded use and unready to rely on another. The task is to create conditions in which people can learn responsibly, recognise limits and justify when reliance is appropriate. Table 2 sets out six organisational conditions that leaders should establish.

Table 2. Organisational conditions for responsible exploration.

Organisational condition	What leaders should establish	Evidence that the condition exists
Purpose and quality	Define the professional problem, the people affected, the intended benefit, the quality outcome and the effect on professional learning before selecting a system.	A short use-case statement, baseline, named beneficiaries, intended development effect and stop conditions.
Leadership and governance	Set approved boundaries, ownership, escalation routes and assurance proportionate to consequence, including which decisions may be assisted, which authority must remain human, and how leaders will seek and respond to challenge from across the hierarchy.	An accountable owner for the use and for people development, an approved-use register, decision route, structured upward-feedback route, visible leader response and periodic review.

Capability and literacy	Build sufficient AI, data and professional understanding for people to use, question and supervise systems responsibly and explain the work without hiding behind the output.	Role-based learning, competent reviewers and evidence that users can explain assumptions, defend decisions, recognise limits and reject outputs.
Information and knowledge	Provide authorised, current and governed sources, make material AI influence visible, and set clear rules for access, confidentiality, retention and correction.	Source ownership, access controls, retention rules, version management, correction routes and records of material sources, assumptions or system influence.
Culture and psychological safety	Make it safe to disclose AI use, admit uncertainty or dependence, ask basic questions, challenge seniority or claimed benefit, report failure and stop unsuitable use. Leaders should actively invite difficult feedback and create conditions in which people with less positional authority can answer honestly.	Visible disclosure and reporting routes, protected escalation, recorded challenge, reflective discussion, evidence that junior staff speak candidly, and leadership responses that change behaviour or decisions without penalising good-faith questioning.
Learning and feedback	Turn experiments, corrections, near misses and outcomes into shared learning; protect mentoring and challenge; and stop uses that improve output while weakening professional development.	Review records, protected mentoring time, evidence of what people have learned, progression decisions not based on output alone, lessons shared and documented stop, revise or continue decisions.

These conditions form a quality system, not a maturity badge or compliance checklist. Figure 3 shows how an organisation can bound a use, test and observe its effect, make the evidence visible and then decide whether to stop, revise, continue or progress. Learning should feed back into updated controls rather than make progression the assumed outcome. Time saved is not a quality gain if judgement, learning, traceability or trust become weaker.



Figure 3. A responsible learning cycle for bounded AI exploration.

4.2 Exploration, pilot and operational reliance

Organisations should distinguish three states of use. **Exploration** is limited learning with no material reliance on the output. A **pilot** tests a defined use in real work under named controls and evaluation. **Operational reliance** means the output routinely influences professional or organisational decisions. Each transition should be deliberate; repeated use must not become reliance by accident.

Movement between these states should depend on evidence, not enthusiasm or elapsed time. Before a pilot, leaders should confirm scope, users, information, review, duration and stop conditions. Before operational reliance, they should confirm accountable ownership, repeatable controls, change management, monitoring and a fallback if the system becomes unavailable or unsuitable. Figure 4 shows that movement from exploration to pilot and operational reliance should occur through explicit evidence-based decisions, rather than through repeated use gradually becoming accepted practice.




Figure 4. Evidence-based transitions from exploration to pilot and operational reliance.

4.3 Competence, capability and challenge

Organisations can buy capability through licences, platforms and specialist support. They cannot buy professional competence in the same way. Competence brings together skills, knowledge, experience and behaviours (SKEB). It is demonstrated when people can define the problem, apply judgement, recognise limits, challenge outputs and remain accountable. Digital familiarity does not confer authority to approve engineering work. As UK fire engineering moves towards stronger individual accountability and greater scrutiny of organisational capability, an organisation's approach to AI may become part of how it demonstrates responsible professional practice.

Challenge must be designed into the organisation. Review is not meaningful if people lack the time, information, competence or permission to disagree. In this paper, psychological safety means that people can disclose AI use, say that they do not understand an output, admit dependence, question a claimed benefit, report an error or near miss, ask for human review and recommend stopping without fear of disproportionate blame or reputational harm. It does not mean friendliness, unrestricted experimentation, absence of accountability or protection from legitimate scrutiny. Leaders should make these behaviours visible in normal work, state who may pause or redirect a use, explain how concerns are escalated and show how good-faith challenge changes decisions. Psychological safety is therefore an assurance and learning condition, not an optional cultural benefit [10]. Existing professional ethics, legal duties, confidentiality and data-protection obligations apply whether or not AI is involved.

Psychological safety cannot depend on whether a leader has time, is in the right mood or happens to invite challenge. It must be part of how teams review work, make room for different voices, respond to mistakes, discuss uncertainty and make decisions under pressure. Culture is not defined by the best behaviour an organisation can point to. It is set by the lowest standard of behaviour and performance that people repeatedly accept, excuse or leave unchallenged. One mistake does not define a culture, but behaviour that is allowed to become normal does. Without a consistent structure, the people who most need to challenge a decision, often those with the least positional authority, remain dependent on the goodwill of those they may need to question.

Leaders cannot judge psychological safety from their own intentions. The test is whether people with less positional authority can offer difficult feedback, whether leaders seek it before formal appraisal cycles and whether visible change follows. Saying that a door is open still leaves the least powerful person carrying the risk of walking through it. Leadership becomes accountable when feedback can travel upwards as well as downwards, and when those who speak can see that their contribution changed behaviour, decisions or support.

Outside appraisal season, when did you last ask someone several grades below you what your leadership makes harder for them, and change your own behaviour because of the answer?

Technical authority does not automatically create people-leadership capability. Many technical leaders take on responsibility for people because of their professional expertise, but effective feedback, coaching, inclusive decision-making, constructive responses to challenge, conflict management and the development of future leaders all require deliberate preparation and practice. Organisations that expect rigorous development of engineering competence should apply comparable discipline before promotion into substantial people responsibility and throughout leadership practice, particularly where AI can conceal whether genuine learning is taking place. Future practitioners should encounter workplaces where equity, inclusion and voice are normal professional expectations. Current leaders should aim to leave better conditions than those they inherited.

The same standard should apply to development. Leaders need evidence that people are becoming more capable, not simply that assisted work is becoming easier to approve.

If your assurance system checks the deliverable but not what the person understands or has learned, what exactly are you developing: engineers or outputs?

4.4 A practical starting point

Responsible practice should begin with visible, bounded uses close to real work and low in consequence. Authorised retrieval, meeting support, comparison, organisation or drafting may provide useful starting points where a competent person can readily verify the result. A checked meeting summary, for example, may reduce administration without determining a fire-engineering decision, provided the source, confidentiality, retention and human review are clear. Each trial should state the quality outcome sought and the conditions for stopping. The aim is not to maximise activity, but to expose barriers and learn what governance and support are needed.

Starting with accessible uses also widens participation. Early-career practitioners may know current interfaces but lack engineering experience; established practitioners may bring deep judgement but less digital confidence; smaller organisations may face different resource and infrastructure constraints. Leaders should use plain language, provide equitable access and offer more than one route to participate, while keeping competence boundaries explicit.

4.5 Two bounded examples

The following examples are illustrative. They neither prescribe a workflow nor endorse a system. They show how responsible investigation may lead either to controlled professional decision support or to a justified decision not to use AI.

Illustrative scenario | Professional decision support

A competent engineer uses an approved AI-supported process to compare authorised project documents and flag apparent inconsistencies in fire-resistance requirements. The output focuses review but is not treated as a

resolved design decision. The engineer checks each material finding against the source, confirms versions and jurisdiction, coordinates with affected disciplines and records how discrepancies were resolved. A separate technical reviewer challenges the scope, sources and conclusions before client advice is issued. Because missed or invented inconsistencies could affect design, named engineering ownership, authoritative evidence and independent challenge are essential.

Illustrative scenario | When not to use AI

A practitioner needs to complete an established calculation with defined inputs, transparent equations and an accepted checking method. A conventional worksheet already performs the task reliably and allows another engineer to follow each step. An AI approach would add uncertainty without improving the decision, so the organisation retains the established method and limits automation to controlled data entry or presentation. This is proportionate method selection, not reluctance to innovate. Choosing not to use AI should be as visible and evidence-based as choosing to use it.

4.6 Governed information and organisational memory

AI-supported work is only as reliable as the information, context and controls around it. Organisations should identify authorised sources, make ownership and decision status clear, and control access, confidentiality, retention, correction and reuse. Better capture and retrieval can improve continuity and accountability, but indiscriminate recording or use of public systems can create surveillance, privacy, confidentiality and information-quality risks.

Good organisational memory depends on purposeful capture of decisions, assumptions, responsibilities and professional reasoning. Recording should remain proportionate, governed and limited to a clear need. It should not become indiscriminate surveillance. The same standard should apply across the hierarchy. Leaders should not require detailed traceability from practitioners while allowing their own material decisions to remain inaccessible, undocumented or dependent on recollection. Accountability becomes more credible, and psychological safety stronger, when people can see how decisions were reached and can question them without carrying all of the organisational risk.

If everyone below you must make their decisions traceable, what evidence lets them understand, question and hold accountable the decisions made above them?

4.7 Leadership controls and proportionality

Policy should make responsible work easier, not merely prohibit unknown activity. It should define approved systems and information, prohibited uses, disclosure and review expectations, who may authorise exceptions, and how controls increase when work influences technical judgement, client advice or operations.

Procurement should test more than headline capability. Leaders should ask which problem a system is designed to solve, what information and models it depends on, how performance and change are managed, what audit and support are available, and how the organisation can exit. A supplier's assurance cannot replace the organisation's own professional judgement.

Hidden use is less likely when people have a simple route to ask, declare and learn. Organisations should record material incidents and near misses without turning every experiment into a burdensome compliance exercise. The aim is to correct affected work and improve boundaries before weak practice becomes routine.

Smaller organisations need the same professional outcomes but may use simpler controls. A short approved-use statement, named technical owner, controlled information location, peer-review route and periodic discussion of lessons may be more effective than a complex framework. Proportionality means matching the control to the consequence, not lowering the standard of responsibility.

These controls turn curiosity into disciplined work: they allow useful possibilities to be tested without lowering professional standards and give people a safer route to improve practice. But organisational readiness cannot answer every professional question or reveal every repeated pattern. The remaining question is therefore whether learning stays within individual organisations or becomes a shared professional effort.

5. Our position and the next step

5.1 Our position

SFPE UK's position is that AI should augment, not displace, competent professional judgement. Organisations should assess where it can improve professional quality while keeping accountability with identifiable people and preserving public trust. Adoption is not the objective; better fire-engineering decisions and outcomes are. The question is no longer whether AI deserves attention, but what professional response it demands from you and your organisation.

SFPE UK will convene rather than prescribe. It will remain independent of products and providers, invite evidence and challenge, and support detailed guidance only where the profession identifies a clear need and sufficient evidence. Choosing not to use AI is consistent with this position when an established method is clearer, more reliable or easier to assure.

5.2 Improving fire-engineering standards

The initiative will support disciplined professional improvement by applying the paper's arguments to shared evidence and action. AI is a catalyst for asking harder questions about how fire engineering is practised; it is not the objective. Organisations should begin with a defined fire-engineering problem and the quality outcome they want to strengthen, then establish the competence, information, boundaries and assurance needed before relying on an AI-supported approach. Evidence should determine whether a use continues, changes, pauses or stops. This work will be judged by whether it raises standards of reasoning, coordination, accountability, learning and public value, not by how much AI is used.

The commitments in Table 3 provide a common basis for raising those standards. They guide leadership choices across different technologies and levels of consequence; they do not replace professional standards, organisational policy or technical assurance.

Table 3. Principles for improving professional standards.

Principle	Commitment
Human-centred purpose	Begin with the professional problem, people affected and quality outcome sought, not the availability of a system.
Human accountability	Identifiable practitioners and organisations retain responsibility for decisions, approvals and consequences, including where AI has materially influenced the work pathway.
Competence before reliance	People must be able to define the problem, understand limits, explain and defend AI-assisted work, challenge evidence and reject or stop unsuitable use.
Capability before scale	Build literacy, governed information, leadership confidence and learning systems before expanding use, and show that scale will not conceal competence gaps.
Proportionate assurance	Increase evidence, review and control with consequence, uncertainty, autonomy and difficulty of detecting failure.
Traceable decisions	Keep sources, assumptions, limitations, system influence, human interventions and ownership visible enough to explain and challenge.
Quality and evidence over activity	Judge value through reasoning, learning, coordination, trust and safety, not usage, speed or volume alone.
Responsible learning and collaboration	Share transferable evidence, including limitations, failure and effects on competence and professional practice, while protecting confidentiality, intellectual property and development.

These principles work together. Human oversight requires enough competence, information, time and authority to challenge the result; adding a name at the end of a workflow is not enough. Traceability does not

mean recording everything, but preserving the material evidence and decisions needed for proportionate review and accountability.

As illustrated in Figure 5, the more an AI-supported output shapes a professional or operational decision, the stronger the assurance should be. Where AI assists the work, competent checking may be sufficient. Where it supports professional judgement, defined scope, authoritative evidence, understood limits, independent review and named approval are needed. Where it materially influences a high-consequence or operational decision, specialist validation, monitoring, fallback and authority to suspend or withdraw the use may be required. At each transition, leaders should ask what decision could change and be prepared to restrict, pause or stop the use. More checking does not make every use acceptable; the evidence may still show that the competence or controls are insufficient.




Figure 5. Assurance increases as AI-supported output has greater influence on professional or operational decisions.

Existing professional ethics remain the foundation [12]. Environmental proportionality follows the same logic: use the least resource-intensive reliable approach suited to the task [9]. Where an established engineering method or transparent automation is sufficient, a less explainable approach should not replace it. Professional responsibility includes delaying, restricting or rejecting a use when improved quality cannot be demonstrated.

5.3 Why act together

Every organisation remains responsible for its own people, clients, systems and decisions. But no organisation sees the whole professional picture or should repeat avoidable mistakes alone. If assisted performance is repeatedly mistaken for developed competence, individual gaps can become organisational assumptions about capability. If those assumptions are repeated across firms, the profession may overestimate the competence on which future safety decisions depend. Collective work can create shared language, benchmark questions, anonymised experience and stronger evidence while preserving proprietary methods, client information and genuine commercial differentiation. The profession should be honest about the boundary: what knowledge and methods genuinely create commercial advantage, and what learning should be shared because keeping it fragmented would weaken public safety?

If every organisation learns alone, how many avoidable mistakes will the profession repeat... and who ultimately carries the cost?

The answer is not compulsory disclosure or uniform practice. It is a credible place to compare what evidence of competence is meaningful, which uses deepen learning, which uses bypass it and how organisations keep

authority, accountability and limitations visible. Sharing transferable evidence, including unsuccessful approaches, allows the profession to test difficult questions before each organisation has to learn the same lesson for itself. The initiative's participation routes are designed to provide that place.

5.4 How participation will begin

The Core Working Group will safeguard the initiative's coherence and professional position. A Wider Industry Roundtable will bring practical experience, concern and challenge from organisations at different stages of readiness. Specialist AI and fire-engineering groups will investigate defined technical or professional questions. Evidence, limitations and practical implications from both routes will be returned to the profession in an accessible form.

The routes are connected, but participants need not contribute in the same way. Leaders, practitioners, researchers, clients, regulators, early-career professionals and sceptics can offer relevant experience, evidence, challenge or well-framed questions without presenting themselves as AI specialists.

5.5 What SFPE UK will do next

SFPE UK will use this paper to open the next stage of the initiative. It will present the position, invite wider professional engagement, convene the first industry roundtable and begin specialist participation in parallel. Those forums will test priorities, identify evidence gaps and determine where later guidance, research or collaboration can add most value. The paper is therefore not the end of the argument; it is the point at which the profession begins to answer it together. Further details will be announced through SFPE UK's published channels.

The expedition metaphor is practical rather than promotional. The profession is entering uncertain ground from different starting points, without assuming that every organisation should take the same route or that every use should continue. Its value lies in sharing what is found, testing claims against evidence and changing course when professional quality demands it.

5.6 Call to action

The first action is reflection made visible. **Organisational leaders** should ask where AI already influences work, which professional problems are worth investigating and who has authority to own, review or stop each use. **Practitioners** should bring evidence, experience and challenge, remain within competence and make assumptions, limitations and system influence visible. **Clients, regulators and partner disciplines** should ask how AI shaped material advice, what review supports it and who owns the judgement.

Then turn reflection into participation. If you are a leader, nominate the people who should help shape what comes next. If you are working technically, share learning that could raise the professional baseline. If you are early in your career or joining from another discipline, bring the questions and perspectives the profession needs to hear. If you are sceptical, test the initiative with the concerns it must be able to answer. Register interest through SFPE UK's published channels and join the expedition, not to endorse every use of AI, but to investigate possibilities openly, challenge them rigorously and help the profession decide what deserves to continue.

Bring your evidence, your experience and your challenge. Help decide what the profession should carry forward, together.

Definitions

The terms below are used in a practical sense to support a shared conversation. They do not replace formal technical definitions. More detailed terminology should be developed only where later guidance needs it [15, 16].

Table 4. Essential terms used in this positioning paper.

Term	Plain-language meaning
Artificial intelligence (AI)	An umbrella term for machine-based systems that generate outputs such as predictions, recommendations, decisions or content in response to human-defined objectives.
Generative AI (GenAI)	AI that creates new text, images, audio, code or other content from patterns learned from data.
Machine learning (ML)	Methods that learn patterns from data to classify, estimate, predict or identify relationships.
Automation	A process that performs defined steps or rules with limited human intervention. Automation may use AI, but it does not necessarily do so.
Competence	The skills, knowledge, experience and behaviours needed to perform a professional role responsibly in context.
Capability	The practical ability of a person or organisation to perform an activity, supported by people, systems, information and resources.
SKEB	Skills, knowledge, experience and behaviours, considered together as components of professional competence.

References

- [1] Ministry of Housing, Communities and Local Government (2025), *Fire Engineers Advisory Panel: Authoritative Statement*; and *Next Steps on Fire Engineering Profession Reform*, published 17 December 2025.
- [2] SFPE UK Chapter (2026), *Aggregated Results of the Survey on the Use of AI in Fire Safety Engineering*, 123 respondents; publication or access through the initiative intended.
- [3] British Standards Institution (2026), *Draft BS 9994: Fire Strategies in Buildings — Specification*, Draft for Public Comment 26/30492160 DC, published 17 June 2026; supported by public consultation summaries from the Fire Protection Association and Fire Industry Association.
- [4] Zeng, Y., Zheng, Z., Zhang, T., Huang, X. and Lu, X. (2024), 'AI-powered fire engineering design and smoke flow analysis for complex-shaped buildings', *Journal of Computational Design and Engineering*, 11(3), 359–373, doi:10.1093/jcde/qwae053.
- [5] Ministry of Housing, Communities and Local Government (2026), *£70 million to boost building safety workforce capacity*, March 2026; including £15 million to expand fire-engineering education and increase the number of qualified fire engineers.
- [6] Zhang, T., Ko, Y., Marshall, L., Tarbet, A., Collins, K., Elsagan, N., Flores-Quiroz, N., Hodges, J., Mozaffari, M. H., Pallett, M., Sharma, A., Tam, W. C., Huang, X. and the GCI Digitalization, Artificial Intelligence, & Cybersecurity Working Group (2023), *Grand Challenges in Digitalization, Artificial Intelligence, & Cybersecurity: A 10-Year Plan for Strategic Cooperation in Research and Education to Advance Fire Engineering*, SFPE Foundation, July 2023.
- [7] Bengio, Y., Clare, S., Prunkl, C. et al. (2026), *International AI Safety Report 2026*, February 2026, DSIT 2026/001.
- [8] Health and Safety Executive (2026), *HSE's Regulatory Approach to Artificial Intelligence*, 12 June 2026.
- [9] National Engineering Policy Centre (2025), *Engineering Responsible AI: Foundations for Environmentally Sustainable AI*, February 2025.
- [10] MIT Technology Review Insights (2025), *Creating Psychological Safety in the AI Era*, research report produced in partnership with Infosys Topaz, 16 December 2025.
- [11] Zhu, Q., Li, X., Dong, Y., Chang, P. and Fan, M. (2026), 'Not all cognitive offloading is equal: distinguishing dependent and autonomous offloading to generative AI', *Frontiers in Psychology*, 17, doi:10.3389/fpsyg.2026.1878629.
- [12] Engineering Council and Royal Academy of Engineering (2026), *Statement of Ethical Principles for the Engineering Profession*, February 2026.
- [13] Microsoft (2025), *2025 Work Trend Index Annual Report: The Year the Frontier Firm Is Born*, 23 April 2025.
- [14] McKinsey & Company (2025), *The State of AI: How Organizations Are Rewiring to Capture Value*, 12 March 2025.
- [15] Tabassi, E. (2023), *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*, NIST AI 100-1, National Institute of Standards and Technology, doi:10.6028/NIST.AI.100-1.
- [16] Atherton, D., Schwartz, R., Fontana, P. C. and Hall, P. (2023), *The Language of Trustworthy AI: An In-Depth Glossary of Terms*, NIST AI 100-3, National Institute of Standards and Technology, doi:10.6028/NIST.AI.100-3.