Response to the ICO Call for Evidence: Agentic AI

About CCIA

CCIA is an international, not-for-profit trade association representing a broad cross section of communications and technology firms. For more than 50 years, CCIA has promoted open markets, open systems, and open networks. CCIA members employ more than 1.6 million workers, invest more than £75bn in research and development, and contribute trillions of dollars in productivity to the global economy. Many CCIA members are developing generative and agentic AI services and/or integrating them to improve existing products and services.

Section 2: Data security

Q5. In an agentic AI system, what session isolation controls ensuring data within a session is only accessible to specific parties (eg the user who started the session) are used?

Session isolation in agentic systems is a property of the infrastructure a service runs on as opposed to a feature of the product a person sees, and the specific controls applied are proprietary and vary between different models. Isolation is configured at the infrastructure and deployment layers and therefore the party able to answer for it in any given case is the organisation that deployed the agent and set its configuration rather than the developer who supplied the underlying model. The ICO should recognise that a single prescribed approach to isolation would not be workable with the range of architectures now in widespread use.

Q6. How do organisations evaluate whether to grant an agent access only to internal knowledge systems (eg document repositories, policy libraries, intranet sites) or allow it access to third-party knowledge systems or databases?

This is principally a deployment decision, and the organisation or user making it is weighing the sensitivity of the information involved against what the agent needs in order to do the job. For example, Meta's Agents Rule of Two sets out that an agent should not combine all three of processing untrusted input, reaching sensitive systems or private data, and changing state or communicating externally within a single session, and that where all three are required the agent should not operate autonomously without supervision. [1] Therefore in the context of this question, opening an agent up to third-party sources is the step that introduces untrusted input, which is why it tends to be paired with tighter limits on what the agent can then reach or do.

That judgment is contextual and the ICO should be cautious in seeking overly general rules in guidance without doing harm. What the ICO can usefully expect is that the organisation deploying an agent has made the assessment and should be able to show its reasoning, which is where legal obligations already sit.

Reference

[1] Meta, Agents Rule of Two, October 2025.
https://ai.meta.com/blog/practical-ai-agent-security/

Q7. Do guardrails and permission controls change dynamically in an Agentic system, if so how?

Permissions in agentic systems are generally scoped, not static, and the constraints applied to an agent vary with the context it is operating in. Inputs carrying greater risk, such as untrusted content drawn from the open web, can trigger tighter constraints than the same agent would face on a routine internal task. Particular controls differ between products and for many of our members are proprietary.

This adaptability is a safety feature and should be understood as such. An expectation that an agent's permissions be fixed and declared in advance, so that a person or a regulator can inspect a static list, would remove the capacity to tighten controls in response to risk as it arises. The ICO should frame transparency expectations around the rules by which permissions change, which can be explained to a person in terms they can act on, rather than around a snapshot of the permissions themselves, which cannot always usefully be frozen.

Q8. What are examples of structured logging measures that record AI agent identity alongside agent actions (eg data retrieval, database entry modification), outcome tracking and tamper-evident records?

Industry is developing a range of measures here and is not yet at the stage of converging on a single approach and we anticipate practice continuing to adapt as experience accumulates across a range of use cases.

Structured logging that records which agent acted, under whose authority, and what information it accessed is already recognised as good practice and is being built into many systems. What would help is for the ICO to provide guidance on what such a log needs to achieve, so that an organisation can design to it with confidence, rather than specifying its fields or format in guidance that would date quickly and would then have to be worked around.

Q9. What system setups or functionalities built in by an agentic AI system provider are effective at supporting data security, and where do deployers feel there are challenges to ensuring effective data security?

The call for evidence does not define data security. Such a term can be drawn from the security principle in the GDPR, however the call for evidence states that in agentic workflows data security relies on strict permissions and controls, reaching into data access, tool use and system architecture. Those are broader matters than the security principle addresses, and the ICO should be precise about the scope of what it is asking, because expectations framed around architecture instead of outcomes are harder to apply and date more quickly.

Robust security and governance delivered at the platform or ecosystem level can provide universal guardrails beyond the scope of individual developers, protecting users and supporting adoption by building trust. Operating system and application distribution providers review what runs on a device, withdraw software and credentials that behave maliciously, require applications to use the platform's permission and consent mechanisms before reaching a user's data, and can act across an entire ecosystem at once. Agentic systems inherit those protections. Where regulatory intervention in related areas can constrain that function, by limiting a provider's ability to review what is installed, to require that agents use the platform's permission and consent mechanisms, or to remove software that misbehaves, the effect is to weaken a control the ICO is relying on in this theme. The security outcome the ICO is seeking depends in part on decisions being taken elsewhere. The safeguards the ICO asks about are mostly settings, not fixed features: how one task is kept apart from the next, what an agent is allowed to reach, which actions need a person to approve them, what is written to a log, and what the system remembers afterwards. The end user never sees all of this in its entirety. It is built in by whoever develops the agent and switched on or off by whoever deploys it in a particular setting. Often those are different companies.

The difficulty deployers might see, therefore, is less the availability of controls than knowing which configuration the ICO would regard as adequate. Good security here depends not on what any single control looks like, but on whether the deploying business has the tools to use it properly and is held responsible for doing so.

Q10. Do you have any other comments about data security?

We have used this box to set out a number of points that apply across the call for evidence as a whole, as well as to data security.

Our central point is that legal clarity on agentic AI should be achieved by applying established data protection principles, not by constructing a new framework specific to agentic systems. The questions agentic AI raises about accountability, lawfulness and oversight are not new in kind. Legal clarity should avoid artificial distinctions between fast-evolving technologies, because such distinctions date quickly and create uncertainty about which rules apply to which products and services. This is the position CCIA set out in its response to the Digital Regulation Cooperation Forum's call for views on agentic AI, where we stated that the best starting point for regulators is to understand the current legal position, review whether it is creating artificial barriers to innovation, and only then implement changes that are as risk-based and technology-neutral as possible. [2]

Relatedly, the call for evidence describes an AI agent as a system based on large language models which has access to tools, memory and an environment, and which has some level of autonomy in its approach to achieving a task. As a description of the technology that is reasonable. The practical difficulty comes when we treat an agentic AI system as any computing system that includes one or more AI agents. On that reading a large system with a single agentic feature becomes an agentic AI system in its entirety, and guidance addressed to agentic AI systems would apply across all of it.

The potential breadth is easy to understate here: the ICO's own illustration of an environment is access to a computer's calendar and clock in order to know the date and

time, and a great many systems with no meaningful autonomy would satisfy that. The ICO should direct its expectations at the agentic processing itself rather than at every system that contains an agent, and to calibrate them by what an agent can actually do, meaning whether it plans and selects its own data sources and tools, whether it initiates multi-step action without the user intervening, and how consequential the actions available to it are. We welcome the ICO's reference to its secondary duties to support responsible innovation and economic growth. Those duties now have a statutory footing. Section 120B of the Data Protection Act 2018, inserted by the Data (Use and Access) Act 2025, requires regard to be given to the desirability of promoting innovation and promoting competition in the exercise of data protection functions. We also support the Government's Regulating for Growth objectives. A regulator whose decisions bear as directly as the ICO's do on whether AI is adopted across the economy should be able to show how that duty has shaped its guidance.

Survey research conducted for CCIA by Freshwater Strategy, published in October 2026, provides UK evidence on the effect of data protection uncertainty on business decisions. Data protection compliance is estimated to cost UK technology businesses approximately £5.29 billion each year, with almost half of that cost (48%) falling on large firms. A majority of businesses (52%) have been prevented from undertaking a business activity because of data protection regulation, most commonly by avoiding beneficial uses of data where it was unclear whether they were permitted (28%), scaling back data use (23%) and not adopting new technology (18%). 68% consider that the compliance risk they carry is disproportionate to the risk to individuals, and 62% consider that uncertainty about enforcement is a bigger issue than the rules themselves. [3]

References

[2] CCIA, response to the DRCF call for views on agentic AI, November 2025. https://ccianet.org/wp-content/uploads/2025/11/Agentic-AI-CCIA-Submission.pdf
[3] Freshwater Strategy for CCIA, UK Data Regulation Survey, October 2026. https://ccianet.org/research/stats/freshwater-strategy-uk-data-regulation-survey/

Section 3: Transparency

Q11. How do developers or deployers of an agentic AI system explain whether they handle personal information on a person's device, in a remote environment or through a combination of both?

On-device and hybrid processing is already a substantial part of how these systems are built, and the published explanations are reasonably detailed. Some models run entirely on a device. Heavier requests are passed to a hardware-secured remote environment, with the link between the two established through remote attestation and encryption, and the provider states that it does not retain the data processed there or make it accessible to itself. Some implementations add private information retrieval, so that a device can fetch what it needs without revealing what it asked for, and some allow a person to export a log of their own requests.

What a person needs to understand is where the processing happens and what the provider can see, and that it is being explained. An expectation that it be restated in a prescribed

form or location would not help, because the architectures differ enough that a single formula would describe some of them inaccurately.

Q12. How do organisations provide timely information to people when an agent: accesses a new data source; uses an external tool; or shares personal information with another agent or organisation?

This question covers three different things, and industry practice is at a different stage on each.

Where an agent reaches a new data source or uses an external tool, the arrangements are well established. Sharing personal information with another agent is a newer problem and we deal with it separately at question 18.

On sharing with another agent, the position is moving quickly and we set out what is and is not available at question 18. If the ICO expects a person to be told at the moment their information passes from one agent to another, it would be requiring something that no product or published specification currently provides, rather than recording something that already exists.

Q13. How do organisations use pre-authorisation, prompts, approval steps, live status indicators or other techniques / approaches to improve transparency for users?

Techniques of this kind are already in widespread use, and the pattern is that they are applied where the consequences of an action are highest, not uniformly. Confirmation is commonly sought before an agent completes a financial transaction, changes sensitive data, sends a communication, creates an account, manages consent or agrees to legal terms, and in some implementations it is sought for those categories even where a person has set a standing permission to allow them.

Requiring confirmation at every step would train people to click through prompts without reading them, which is the outcome the transparency principle exists to prevent, and an expectation framed around the number of approval points rather than their placement would make these services harder to use without making anyone better informed. Expectations in this area therefore need to be risk-proportionate, and a prescribed list of action types requiring confirmation should be avoided in particular. Any such list would be overtaken quickly by the range of things agents are asked to do, and it would displace the judgement about consequence that currently drives these decisions and that produces better outcomes than a fixed schedule could.

Q14. How do organisations allow end users to set limits on the websites, tools, data sources and actions available to an agent?

Limits can be set at two levels, and the distinction is important for where the ICO directs its attention. At the consumer level, people are typically given a choice over which applications

and data sources an agent may connect to, with the categories of data involved set out for each and the ability to disconnect at any time. At the deployment level, the organisation configuring an agent can restrict the sites it may visit, register only the tools a particular workflow requires, and limit access to local files. The specific controls are proprietary and vary by product, so we have not described them here.

The two layers are complementary, and the more granular limits sit at the deployment layer because that is where the purpose of the processing is actually known. Guidance addressed to deploying organisations on how to configure those limits for their own context will therefore do more for individuals than guidance addressed to the design of consumer interfaces, where the choices available have to be expressed broadly enough to make sense to everyone.

Q15. How do post-task execution summaries and audit trails explain: what an agent did; what personal information it used; which tools and recipients were involved, and what happened as a result?

A substantial range of agentic tools has come to market in a short period, and providers are navigating requirements in different ways and have not converged on a single form of summary or audit trail. That variation is to be expected at this stage and it is useful, because the level of detail a person needs depends heavily on what the agent was asked to do. Approaches in use include live interfaces showing the steps an agent has taken and plans to take, work logs that can be reviewed after the event, and notification with a summary on completion where a task has been carried out outside the interface.
No one of these should be treated as the model. Prescribing the form a post-task summary must take would arrest a line of development that is moving quickly. The more valuable contribution the ICO can make is to include in guidance what a summary needs to achieve as an outcome, and to leave the form to the service and the context.

Q16. How can people challenge, correct or reverse an agent's actions?

Reversal is inherently complex. Permissions in these systems are scoped so that an agent can take a step forward as an action but not a step back. An agent may be permitted to draft a message but not to send it, or to create a record but not to delete one, precisely because the destructive action carries the greater risk or is not practicable. However, once an agent has booked, sent or paid for something in a third party's system, the organisation that ran the agent is generally not in a position to reverse it, and the route to correction runs through the counterparty on the same terms as it would if the person had acted themselves. An expectation framed as a right to reverse an agent's action would therefore be undeliverable in a large proportion of cases, and would sit awkwardly with the scoped permissions that make these systems safer in the first place.

An expectation framed around telling people promptly what an agent did, keeping a clear record of it, and giving them a route to challenge it and seek correction is both achievable and more useful to the individual, and the ICO should take that approach.

Q17. How can agent identification, authentication, permissions and tracing support transparency to people?

Identification is developing along two lines, and it is worth distinguishing them because they do different work. In machine-to-machine interactions, system identities rely heavily on cryptographic validation. These credentials are temporary and cycle automatically. They are strictly partitioned to prevent cross-workload sharing and engineered to resist impersonation.

The signing of agent requests is being standardised in parallel, so that a site can verify that traffic comes from a declared agent. [4] Between an agent and a person, disclosure is delivered in the interaction itself, as where an agent placing a call on someone's behalf opens by identifying itself as an AI assistant on a recorded line and naming the person it is acting for.

Agent identification is a control that supports oversight and accountability, and it should not be assumed to carry transparency to the individual with it. The two are related but they are not the same thing, and conflating them would push the industry towards an architecture that is worse for privacy. Identification is what allows a log to record which agent acted under whose authority, which in turn is what makes oversight possible.

Reference
[4] IETF, Web Bot Auth HTTP message signatures, October 2026. https://www.ietf.org/ietf-ftp/internet-drafts/draft-ietf-webbotauth-httpsig-protocol-01.html

Q18. How do organisations provide meaningful transparency when personal information moves between agents in a multi-agent system?

Two different things could be meant by this question. Where agents exchange information with one another or with an external service, the arrangements for identifying who is on the other end, verifying what they are permitted to do, and recording what passed between them are a solved engineering problem. That work is now quite advanced. It also operates between systems and is largely invisible to the person.

Disclosure to the individual at the moment their personal information passes from one agent to another is a different matter. The ICO has itself identified privacy-preserving communication between agents as an area for research rather than as an available control, [5] which is the right characterisation. Work is moving in that direction at the level of standing permissions, where some providers now give a person a single place to review and withdraw the access an agent holds over their account. [6]

The ICO should recognise that an expectation framed around disclosure at the point of transfer would be creating practice rather than recording it, and to develop it with industry on the basis of what is achievable.

The ICO should not treat any particular technical specification as the benchmark for what transparency requires. The protocols in this area are young, most are governed by foundations whose membership extends well beyond any one firm, and they are revised frequently. An expectation founded on a specification as it reads today would be out of date before guidance built on it was published, and it would commit the market to one approach at precisely the point when competing approaches are still being tested. That is a poor

outcome for the people the rules exist to protect, who may benefit from a better approach still to bed in.

References
[5] ICO, Tech futures: agentic AI, January 2026.https://ico.org.uk/about-the-ico/research-reports-impact-and-evaluation/research-and-reports/technology-and-innovation/tech-horizons-and-ico-tech-futures/ico-tech-futures-agentic-ai/
[6] Google, third-party connections. https://support.google.com/accounts/answer/13533235

Q19. Do you have any other comments about Transparency?

Experience from other disclosure regimes, notably cookie consent banners under the Privacy and Electronic Communications Regulations (PECR), demonstrates that proliferating consent and disclosure requirements can degrade a person's understanding instead of improving it, producing meaningless click-through in place of informed choice. Industry experience with AI content labelling suggests that contextual, in-product transparency is more effective than static, consent-based disclosure. [7]

Disclosures should not be standardised while agentic services remain diverse and fast-changing. The ICO should focus on whether a person can understand and control what an agent does, rather than on the volume of information presented before they begin.

This theme brings together a number of related but distinct issues, and it may be worth separating them. Agentic systems generate two quite different kinds of information. There is what an organisation keeps in order to supervise its own systems and to demonstrate compliance, which means agent identity records, permission configurations, routing records and audit logs. Separately, there is what Articles 13 and 14 of the UK GDPR require an organisation to tell a person about the processing of their data. Both are commonly described as transparency, although they serve different purposes and different audiences. It would help organisations considerably if the guidance were clear that records of the first kind support oversight, accountability and investigation without themselves becoming transparency obligations. The reason is practical, not legal. If detailed logging becomes something that must be surfaced to individuals, organisations acquire a reason to log less, because every additional field widens the disclosure surface. That is unlikely to be the outcome the ICO is looking to achieve.

Reference
[7] CCIA, response to the DRCF call for input on consumer interest and AI, July 2026. https://ccianet.org/wp-content/uploads/securepdfs/2026/07/DRCF-consumer-attitudes.pdf

Section 4: Accountability

Q20. How do organisations allocate controller, joint controller and processor accountabilities in agentic AI systems?

Where a model is made available openly or via an API and a third party deploys it in a specific context, accountability should follow the party with control over the deployment

decision and the end-user relationship, rather than the upstream developer, who has limited visibility of downstream use. Controllership should be determined by who in fact decides the purposes for which personal information is processed and the essential means by which it is processed, which is the test the UK GDPR already applies.

The business that decides what an agent is for, what personal information it can reach and what it is allowed to do will usually be the one deploying it, and that is where controllership should sit. A developer who supplies the model, the framework or the protocol, but has no say in what a particular deployment is used for, should not become a legal controller of that processing. The same principle applies for an organisation deploying an agent. It should not be expected to account for upstream processing it does not control and cannot observe, and requiring it to explain the development of a model it did not build would be as misplaced as treating a developer as responsible for a deployment it did not determine.

The courts have already drawn this line. In Fashion ID the Court of Justice held that a party is a controller only of the operations for which it actually determines the purposes and means, and does not become a controller of operations beyond its control (Case C-40/17, 29 July 2019, retained EU case law). A developer's defaults, including the tools an agent ships with and what it retains and logs, are product decisions. They do not decide what a deployer's agent is then used for.

In its guidance for businesses using AI agents, the Competition and Markets Authority placed responsibility on the organisation deploying the agent, and stated that this holds "even if someone else designed or provides the AI agent on your behalf". [8] This does not seem consistent with the ICO's approach.

Reference
[8] Competition and Markets Authority, Complying with consumer law when using AI agents, March 2026.
https://www.gov.uk/government/publications/complying-with-consumer-law-when-using-ai-agents

Q21. How do organisations show their allocation of accountability in practice and in their contracts?

Contractual arrangements between our members and their customers vary by product and by terms. The more useful observation is described at question 23 on the work the contract alone cannot achieve. Terms of service set out the allocation of responsibility in legal language, and at a level of generality that has to cover every customer. The technical and compliance documentation published alongside them tells a particular customer what the system can actually reach, what it can do, what is logged, and what that customer must configure for itself. That is what allows the customer to understand the terms as they apply to its own deployment. The ICO should read the two together and not treat the contract as the whole of the picture.

Q22. What mechanisms should organisations use for logging, tracing and monitoring AI agent actions (including data access) to ensure controllers retain control and accountability for the data processing that serves the purposes they set?

The relevant mechanisms are logs that make material events traceable: which agent acted, under what authority, and what information it accessed. Amazon's AgentCore Observability structures telemetry as sessions, traces and spans. The aim should be traceability of material events and key decision points rather than a complete record of everything an agent does. Agent logs will frequently contain personal information, so an expectation of exhaustive logging would sit awkwardly with data minimisation and storage limitation.

For a controller to retain control in the sense this question describes, three things are needed: a record of which agent acted and under whose authority, a defined point at which a significant action requires human approval, and the ability to reconstruct afterwards what information was used. All three are available in products on the market today, and the documentation cited at questions 23 and 27 covers each of them. That approval point should not be assumed to require review before the action in every case. In fraud prevention and cybersecurity, intervention and contestability after the event can protect people better than mandatory prior review, which risks delaying action needed to prevent harm.

Q23. What information do developers of agentic AI systems provide to organisations deploying them about how the systems work, how they process personal information and what risks could emerge through use, to help them demonstrate their data protection compliance?

There is established practice here. For example, Amazon's AI service card for Nova Act states that safety is a shared responsibility between the provider and the customer, sets out intended uses and limitations, publishes evaluation results, confirms that the service is not intended to support practices prohibited under applicable law, and gives deploying organisations specific mitigations including URL allow and block lists, registering only the tools a workflow requires and restrictions on local file access. It tells customers they should consider disclosing their use of machine learning to end users, and that they may reference the service card in their own documentation. [9] That is a provider setting a safe default while leaving the deploying organisation accountable for any departure from it.

Reference
[9] AWS, Nova Act AI Service Card.
https://docs.aws.amazon.com/ai/responsible-ai/nova-act/overview.html

Q24. What information should developers of agentic AI systems provide to set out clear controller and processor accountabilities?

The objective should be that a deploying organisation can understand, in practical terms, how an agentic system meets its obligations under the UK GDPR and the Data (Use and Access) Act, and can explain that to its own customers and to the ICO if asked. That is a

question about whether what developers already publish is workable and intelligible for the people who have to rely on it, rather than one calling for a prescribed list. The market has strong commercial reasons to get this right, because a customer that cannot satisfy itself on compliance does not deploy.

In practice what a deploying organisation most needs to understand is the categories of personal information the system can access, the actions it can take and which of those are gated by default, the logging and audit facilities available and what they capture, the data flows associated with tools, plugins and protocols including any transfer to the developer, what the system retains between tasks and how that can be scoped or deleted, evaluation results and known limitations, and the configuration choices that are the deployer's to make together with the provider's recommended defaults.

Q25. Do you have any other comments about Accountability?

On the statement in the call for evidence that organisations should interpret the concept of a controller broadly, the ICO should consider the consequences of this more fully. Allocating controllership to a party that does not determine the purpose, and cannot see how a system is used downstream, does not give that party the ability to protect anyone. It risks giving individuals additional words on a privacy notice and a diffusion of responsibility. The UK GDPR's existing parameters, particularly who decides the purposes and the essential means, already reaches developers where they genuinely make those decisions.
The duty at section 120B, noted above, bears directly on the choices in this call for evidence. Agentic AI is at an early stage of deployment in the UK, and guidance that resolves contested questions of scope before the evidence this exercise is gathering has been assessed would be difficult to reconcile with it.

There is recent precedent for the approach CCIA is urging. Addressing the same technology in March 2026, the CMA published research alongside the guidance cited at question 20, and framed what it expected of businesses as the application of existing law rather than as a new set of obligations specific to AI. Its research recorded that UK consumer law "applies whether decisions are made by people or by AI", and that the CMA did not favour heavy-handed regulation at this stage of the market's development. [10]

Where a single agentic deployment engages the ICO alongside the Competition and Markets Authority, Ofcom or the Financial Conduct Authority, there is a strong case for lead-regulator designation, the avoidance of duplicative reporting, and a clear hierarchy where multiple duties apply to the same system. This is the position CCIA set out in its response to the DRCF's call for input on consumer interest and AI. [11] Note that the CMA has already issued guidance to organisations deploying the same systems, as covered at question 20.

The ICO should also align with established international frameworks, including ISO/IEC 42001 and the NIST AI Risk Management Framework. Divergence imposes cost on organisations operating across jurisdictions without delivering additional benefit to individuals and businesses across the UK.

References

[10] Competition and Markets Authority, Agentic AI and consumers, March 2026. https://www.gov.uk/government/publications/agentic-ai-and-consumers
[11] CCIA, response to the DRCF call for input on consumer interest and AI, July 2026. https://ccianet.org/wp-content/uploads/securepdfs/2026/07/DRCF-consumer-attitudes.pdf

Section 5: Automated decision making

Q26. How do organisations identify whether ADM provisions apply?

The ICO's expectation that organisations assess the provisions across a complete workflow instead of at individual decision points is understandable in principle. Applied to multi-stage agentic systems it creates a practical difficulty however. The number of points at which an assessment might be required rises with the complexity of the workflow, while the question that matters (whether a decision with significant effects is being taken without meaningful human involvement) is unchanged.

The appropriate level of human supervision varies with context. In high risk applications such as medical diagnostics, human supervision and significant disclosure may be appropriate. In lower risk applications the same requirements would be disproportionate and would deter beneficial innovation. [12] The ICO should anchor the assessment in the consequence for the individual rather than an overly prescriptive architecture.

Reference
[12] CCIA, response to the DRCF call for input on AI policy tools, August 2026. https://ccianet.org/wp-content/uploads/securepdfs/2026/08/DRCF-policy-tools-CCIA-response.pdf

Q27. How does an organisation ensure, where they have judged that the ADM provisions do not apply, that human involvement is meaningful in agentic AI systems?

Meaningful human involvement is the test that determines whether the automated decision making provisions apply, so an organisation which has concluded that they do not apply has concluded precisely that the involvement is meaningful. It is not a separate and continuing obligation running alongside them. Article 22D of the UK GDPR, inserted by section 80 of the Data (Use and Access) Act 2025, gives the Secretary of State power to make regulations defining meaningful human involvement, and no such regulations have yet been made. The ICO should take care not to settle that meaning through agentic AI guidance in advance of them.

Human involvement is meaningful when the person deciding has the information, the authority and the practical opportunity to reach a different conclusion. In agentic systems this is delivered by building the decision point into the system, not by bolting a review onto the end of it. Logging that records which agent acted, under what authority and what information it accessed gives the person reviewing a decision the material they need, and it matters that they can see the action taken and the information it relied on, rather than only the account the agent gives of its own work. Defined escalation points before a significant

decision takes effect give the reviewer the opportunity to intervene. Permission and interface design requiring approval for specified actions gives them the authority, because the action cannot proceed without them. Where those three things are present the human involvement is real, and an organisation is in a position to evidence it. Prior escalation should not be the default in every context. Where a decision protects a person against fraud or a security threat, requiring review before it takes effect can itself cause harm through delay. In those cases meaningful involvement is better delivered through prompt intervention and a real ability to contest the decision after the event.

Each of them is available in products on the market now. Confirmation gates applied by default to consequential actions, per-action approval at the level of an individual action group, and audit logging that attributes an action to a named agent acting under a named user are all documented and in use. [13]

What would assist organisations is confidence that decision points designed this way can be relied upon, rather than a separate assessment obligation layered on top of them. CCIA has previously supported certainty of this kind where the benefit outweighs the inflexibility that comes with prescribed approaches. [14]

References
[13] AWS, Bedrock Agents user confirmation, https://docs.aws.amazon.com/bedrock/latest/userguide/agents-userconfirmation.html
[14] CCIA, response to the DRCF call for input on consumer interest and AI, July 2026. https://ccianet.org/wp-content/uploads/securepdfs/2026/07/DRCF-consumer-attitudes.pdf

Q28. How do organisations using ADM in the context of agentic AI apply the relevant safeguards?

Most agentic activity does not produce the kind of decision these provisions reach. Under Article 22A of the UK GDPR, inserted by section 80 of the Data (Use and Access) Act 2025, a significant decision is one producing a legal effect or a similarly significant effect for the person, and it is based solely on automated processing only where meaningful human involvement is absent. Most agentic tasks are operational and satisfy neither of those conditions, so this question only arises in a lesser proportion of deployments. Where it does have an impact, the Article 22C safeguards are the familiar ones and nothing about the agentic context alters them: information about the decision, the ability to make representations, a route to obtain human intervention, and the ability to contest the outcome.

The harder question in practice is not what the safeguards are but at which point in a multi-step agentic workflow the decision requiring them is actually taken. That depends on how the particular deployment is built and does not admit a general answer, which is a further reason for the ICO to address this through engagement with deploying organisations instead of a uniform expectation.

Q29. How can technical logging and audit mechanisms support compliant multi-stage ADM through improved transparency and accountability?

The mechanisms are those described at questions 8 and 22. Their value for multi-stage decision-making is that they allow the chain to be reconstructed after the event. I.e which agent acted, under what authority, what information it used, and at which point a person approved or declined. Without a record of this kind, an assertion that human involvement was meaningful cannot be tested, either by the organisation or by the ICO.

Q30. How are organisations using permission setting and user interface design to support compliance with the ADM provisions?

Permission setting and interface design are doing a good deal of the work here. The common pattern is a defined list of consequential action types gated by default, a documented mechanism for a deploying organisation to vary that, and an organisation-level policy that can prevent variation.

Section 6: Fairness and purpose limitation

Q31. Can you provide examples of instances where organisations have assessed the compatibility (or non-compatibility) of a new processing purpose?

Purpose limitation does not require that every downstream use of personal information be specified in advance. It requires that further processing be compatible with the purpose for which the information was obtained, judged by the relationship between the two purposes, the context in which it was obtained, its nature, the consequences for the person and the safeguards applied.

Applied to an agent, the purpose is the goal the user has set. Context the agent retains and applies in pursuing that goal is processing for that same purpose, not for a new one, and the compatibility test is capable of accommodating it without strain.

Treating each action an agent takes within a user's instruction as a fresh purpose requiring its own assessment would be a departure from how the principle is applied elsewhere, and it would be unworkable at the frequency with which agents act. The ICO should assess compatibility at the level of the purpose the user has set and the scope of the agent's authority. Processing reasonably necessary to a goal the user has set, and within the agent's permitted scope, should ordinarily be regarded as compatible. Processing outside that scope should not.

Q32. What are the guardrails and control measures that support compliance with the purpose limitation and fairness principles?

Most of the work here is done by how an agent's memory is scoped. The usual approach is to keep long-term memory separated by user and by session, so that what is retained from

one person's task cannot surface in another's, with the degree of separation configurable and with deletion available for individual events as well as for stored records. That is the practical control, and it does more for purpose limitation in agentic systems than anything happening at the level of policy wording. Implementations differ and the detail is proprietary, so it is described here in general terms.

Fairness is harder to articulate. Personalisation across sessions is something people value and ask for, while purpose limitation requires that what an agent learns in one context does not quietly become available in another. Those two have to be reconciled in an interface a general audience can follow, which makes it a design consideration before it is a compliance one. The ICO should work through where the balance should sit with developers before settling an expectation in guidance.

Q33. Do you have examples of where monitoring of agent actions, data access and tool interactions, alongside other safeguards can reduce the risk of AI agents using personal information for purposes that have not been appropriately assessed and authorised?

Section 7: Lawful basis

Q34. Which lawful bases do organisations use for the different processing activities involved in agentic AI development and deployment?

Lawful basis depends on which activity is in question. Development and training of an underlying model, adaptation or fine-tuning of an orchestrator, the processing involved in carrying out a task the user has asked for, retention of information between tasks, and security, abuse prevention and service maintenance are distinct activities with distinct justifications.

Carrying out a task the user has specifically asked for is the clearest of them, and in most agentic interactions it is the performance of what the user has requested that does the work rather than any separate justification. Security, abuse prevention and service maintenance have long been recognised as resting on the controller's legitimate interests, and the ICO's existing guidance treats them that way. Model development raises questions the ICO has already consulted on at length and we do not reopen them here. Retention between tasks is the genuinely new question, and we address it at questions 31 and 32.

Guidance treating agentic AI as a single processing activity with a single lawful basis would not be usable. It would push organisations either towards seeking consent for operations that do not require it, which devalues the consent people give for the things that do, or towards stretching one justification across activities it was never meant to cover.

Q35. What challenges do organisations face in identifying the appropriate lawful basis for their agentic AI processing?

The largest unresolved question for our members is not the identification of a lawful basis under the UK GDPR but the interaction with PECR.

The ICO correctly notes that storing information on, or accessing information stored on, an end user's device engages the specific requirements of PECR, and regulation 6 in particular, in addition to the UK GDPR. Agents that operate on a user's own device, or that interact with the files, applications and operating system of that device in order to carry out a task the user has asked for, may engage those requirements repeatedly in the course of a single instruction. If each such interaction requires separate consent, the practical consequence is a consent interface of the kind that has already been shown not to work.

The ICO should consider how the PECR requirements apply where a user has instructed an agent to perform a task that necessarily entails access to information on their own device, and whether consent to the task can properly be understood to extend to the access required to complete it.

The consequence of leaving this unresolved is perverse. The services most likely to keep personal information on a person's own device instead of transmitting it to a remote operator, are the services most exposed to a repeated consent requirement. That is the opposite of the outcome the framework should produce.

Q36. What analysis or evaluation supports organisations' decisions to identify a given lawful basis?

Members apply the usual instruments: data protection impact assessments, legitimate interests assessments where that basis is relied on, and internal legal review, applied to existing ICO guidance on AI and data protection that was not written with agentic systems in view. The cost of that approach falls disproportionately on smaller organisations.

Q37. How do organisations conduct Legitimate Interests Assessments for agentic AI data processing, including training or adapting orchestrator models?

Legitimate interests remains available as a lawful basis under the UK GDPR and in many agentic deployments it will be the appropriate one.

The practical difficulty for an orchestrator is that its processing exists to give effect to an instruction the user has given, so the balancing exercise is conducted against an interest the person has themselves created.

Section 8: Additional questions

Q38. Which best describes your organisation's current and expected use of agentic AI?

Not applicable to CCIA. We are a trade association and do not deploy agentic AI in the sense the question contemplates. Our members' positions range across the full scale offered, from exploration to wide deployment, and they are responding individually.

Q39. Please describe any current or planned uses of agentic AI, including: the business function or activity involved; whether the use is internal, customer-facing or both; and the types of decisions, actions or workflows it supports.

Not applicable to CCIA.

Q40. To what extent have the following issues acted as barriers to developing, obtaining or deploying agentic AI systems?

On lack of regulatory certainty over data protection requirements, recent Freshwater research for CCIA found that 62% of UK technology businesses consider uncertainty about enforcement a bigger issue than the rules themselves, and that a majority (52%) have been prevented from undertaking a business activity because of data protection regulation. On that evidence it is a major barrier. [15]

On cost of compliance safeguards, the same research estimates the total cost of data protection compliance to UK technology businesses at approximately £5.2 billion each year, with almost half (48%) falling on large firms.

Reference
[15] Freshwater Strategy for CCIA, UK Data Regulation Survey, October 2026. https://ccianet.org/research/stats/freshwater-strategy-uk-data-regulation-survey/

Q41. Has your organisation delayed or abandoned plans to deploy agentic AI because of uncertainty relating to data protection requirements?

The Freshwater research mentioned above found that 18% of businesses did not adopt new technology and 23% scaled back their use of data as a consequence of data protection regulation, and that 28% avoided beneficial uses of data where it was unclear whether they were permitted. The pattern for agentic AI is likely to be at least as pronounced, given that no specific guidance has yet been issued. [16]

Reference
[16] Freshwater Strategy for CCIA, UK Data Regulation Survey, October 2026. https://ccianet.org/research/stats/freshwater-strategy-uk-data-regulation-survey/

Q42. In the absence of specific regulatory guidance on agentic AI, how does your organisation obtain assurance that its use complies with data protection requirements?

The answer reported to us by members is a combination of internal legal advice, data protection officer review, data protection impact assessments, existing ICO guidance, external legal advice, vendor assurances and industry standards or frameworks.

The common feature is that each of these applies guidance written before agentic systems existed. That is workable for a large organisation with an in-house team. It is a material cost for a smaller one, and it is a further reason to resolve the questions in this call for evidence clearly.

Q43. To what extent, if at all, does uncertainty about data protection requirements affect your organisation's decisions about whether and how to develop, obtain or deploy agentic AI?

To a great extent, on the sector evidence set out above. 52% of UK technology businesses report having been prevented from undertaking a business activity, and 62% report that enforcement uncertainty is a bigger issue than the rules themselves. [17]

Reference
[17] Freshwater Strategy for CCIA, UK Data Regulation Survey, October 2026. https://ccianet.org/research/stats/freshwater-strategy-uk-data-regulation-survey/

Q44. If applicable, what are the main data protection challenges, uncertainties or areas of difficulty you have encountered in developing, obtaining or deploying agentic AI where greater regulatory certainty can be beneficial?

The allocation of controllership between developers and deployers is the most consequential, because it determines where every other obligation falls. The position we set out at question 20 would help resolve it. The broad interpretation signalled in the call for evidence would not, because it would allocate responsibility to parties unable to discharge it.

How purpose limitation applies to information an agent retains between tasks is the question most likely to constrain product design, and the one on which there is least established guidance.

How PECR, and regulation 6 in particular, applies to an agent acting on a user's own device at the user's instruction is the question on which the current position most clearly produces a result nobody intends.

Q45. What benefits has your organisation experienced, or does it expect to experience, from deploying agentic AI?

The sector evidence indicates what is contingent on the clarity this exercise can provide: if the regulatory burden were eased, 68% of UK technology businesses say they would use data more extensively and 62% say they would adopt new data or AI technologies. [18]

Those are the benefits the ICO's secondary duty on economic growth is concerned with, and they are within the ICO's gift to a greater extent than is sometimes assumed. CCIA would welcome the opportunity to discuss this response with the ICO.

Reference

[18] Freshwater Strategy for CCIA, UK Data Regulation Survey, October 2026. https://ccianet.org/research/stats/freshwater-strategy-uk-data-regulation-survey/