techUK Response to the Department for Education Call for Evidence - AI and Digital Technology in Children's Social Care

techUK is the UK's leading technology trade association, representing more than 1,200 companies from across the digital economy. Our membership is broad and diverse, including large global technology businesses, UK-based SMEs, scale-ups, consultancies, and specialist providers committed to using innovation to drive economic growth, improve public services, and address major societal challenges.

techUK works closely with government, regulators, and industry to help shape policy, build markets, and support the adoption of technology in ways that are safe, inclusive, and effective.

In the context of children's social care, techUK's members include organisations that design, deliver, and support technologies used across local government and frontline services. This provides a broad perspective on both the opportunities and practical barriers associated with adopting AI and digital tools in complex public service environments. This response draws on that collective expertise and on insights gathered through engagement with industry participants active in this space.

The evidence shows that AI's use in children's social care is still in its early, exploratory phase. While there is growing awareness of AI's potential to assist frontline workers, especially by easing administrative tasks and improving access to information, its implementation remains fragmented. Challenges involve data access, governance, procurement, and organisational capacity. Most obstacles are structural, cultural, or regulatory rather than technological.

This response emphasises that AI should support, not replace, professional judgement. Its primary benefit is enhancing practitioner capabilities and facilitating more effective data utilisation.

Externally Developed AI and Digital Technology Tools

AI tools in children's social care are mostly used for low-risk, supportive roles. They are not yet broadly implemented or widely integrated, with most applications aimed at boosting productivity and improving interaction with current data systems.

An important distinction should be drawn between AI capabilities built into existing case management systems and AI tools added on top of or alongside those systems. Where AI is embedded within a case management platform, it benefits from the data governance, access controls, audit trails, and information assurance arrangements already

established for that system. This allows practitioners to adopt new capabilities more quickly and more safely, because the underlying governance questions, who can access what data, under what conditions, and with what oversight, have already been resolved.

By contrast, bolt-on or standalone AI tools typically require separate governance assessments, data-sharing arrangements, and integration work before they can be used with confidence. Both models have a role to play, and standalone tools can offer valuable specialist capabilities that embedded solutions do not. However, recognising this distinction is important when considering the speed, safety, and scalability of adoption, and when prioritising where early progress is most achievable.

At present, the primary application in children's social care involves natural-language query tools that enable practitioners to explore case records and associated datasets via conversational interfaces. This approach greatly diminishes the need for formal system training and speeds up access to pertinent information. Consequently, AI is primarily utilised to extract value from existing data within systems, rather than creating new datasets or workflows.

Additional use cases involve summarising case records, transcribing meetings and interactions, and supporting administrative tasks. These functions are appreciated because they reduce documentation time and enable practitioners to devote more attention to direct engagement with children and families.

Nonetheless, adoption remains cautious, with little evidence of AI being employed for automated decision-making in frontline practice. It is widely agreed that such applications would be inappropriate, considering the sensitivity and complexity of children’s social care.

There is growing evidence of informal or “shadow” AI use, in which practitioners independently use publicly available tools, such as generative AI, to aid their work. This tendency is fuelled by workload demands and the lack of approved options, but poses considerable risks to data security and confidentiality. Additionally, the growing use of AI-generated tools developed outside formal governance increases risk exposure.

It is acknowledged that publicly available translation tools are used to facilitate communication with families for whom English is not the first language. In practice, these tools are often employed in urgent situations where formal interpretation services are unavailable. While this demonstrates an operational need and exposes gaps in current resources, it also raises serious concerns about data protection, confidentiality, and the management of sensitive information. These tools are not designed to securely handle personal data in children’s social care settings, and their use could result in data being processed outside approved environments. This underscores the need for approved, secure alternatives that meet practitioners’ needs while ensuring robust safeguards.

Data Availability and Interoperability

Data access and interoperability are among the main obstacles to the effective use of AI in children's social care, alongside wider challenges relating to digital capability and organisational capacity. Data is distributed across multiple organisations, including social care, health, education, and policing, each operating under different governance structures, technical environments, and interpretations of data-sharing policies. This fragmentation limits the ability to build a holistic view of children and families, which is essential for effective safeguarding and early intervention.

A key challenge is the lack of consistent understanding and agreement on what data can be shared, how it can be shared, and under what conditions. Differing interpretations of legislation and risk often result in overly cautious approaches, preventing the combination of datasets and limiting the ability to derive meaningful insights. At the same time, technical barriers such as legacy infrastructure, inconsistent data standards, and variable system maturity across organisations can further constrain progress. In practice, these issues create barriers to both operational decision-making and the use of AI to identify risk patterns across systems.

Structural challenges are compounded by organisational and geographic boundaries. Data sharing between local authorities and across devolved administrations remains inconsistent. This leads to fragmented perspectives, particularly when children and families engage with multiple services across regions. As a result, no single organisation has a complete view of the information required to support timely, informed decision-making.

There are two broad models for addressing this challenge: centralised data systems and federated approaches. A centralised national data repository could offer some advantages. It has the potential to create a single, unified view of data, simplify access for practitioners, and enable large-scale analytics across the system. A central system could also support more consistent standards and reduce duplication of effort across local authorities. However, this approach introduces significant challenges. Centralisation would require substantial investment, long delivery timelines, and complex coordination across multiple organisations. It also raises concerns around data ownership, accountability, and public trust, particularly given the sensitivity of children's social care data. There is also a risk that large, centralised programmes become inflexible and struggle to adapt to local needs or evolving technologies.

In contrast, a federated approach keeps data within existing systems while enabling authorised access and analysis across organisational boundaries through secure, controlled mechanisms. Rather than duplicating or relocating data, federated models use shared standards, secure connections, and robust access controls to provide visibility where needed. This approach supports local autonomy while enabling cross-

system insight and allows incremental implementation rather than large-scale transformation. It is particularly well-suited to children's social care, where safeguarding depends on connecting information from multiple services, including family and household context, without requiring wholesale restructuring of existing systems.

Federated technologies and delivery models already exist, but implementation maturity varies significantly across UK public services. In many cases, the main barriers are governance-related, including legal interpretation, risk appetite, organisational alignment, and the practical challenge of integrating with legacy systems. Variability in how data-sharing rules are understood and applied continues to slow progress, even where suitable technical options are available. While federated models also require coordination, investment, and common standards, they are generally more flexible, scalable, and aligned with the current structure of public services.

Managing Risks and Ethics

Ethical considerations are key when implementing AI in children's social care. Due to the sensitive nature of the data and the possible effects on vulnerable people, a careful and principled approach is essential.

A key principle is that AI should not replace professional judgement. Practitioners must retain decision-making authority, with AI providing insights and assistance rather than making decisions. This aligns with the broader public-sector view that AI should enhance, not replace, decision-making in high-stakes contexts.

A further consideration is that AI models trained on incomplete or partial data will incorporate bias. This is particularly relevant in children's social care, where historic data reflects the structures, thresholds, and decision-making patterns of the care system as it has operated over time, rather than an objective record of need or risk. There is a real risk that, without active attention, these biases become unwittingly baked into the tools being developed and deployed. While some degree of bias is unavoidable - in part a consequence of the care system itself and the historical methods by which services have been delivered - this does not absolve developers or commissioners of responsibility. Bias should be actively identified, documented, and managed throughout the lifecycle of an AI tool, including careful attention to training data, ongoing monitoring of outputs, and transparency with practitioners about the limitations of the models they use. It reinforces the importance of human oversight, professional judgement, and clear accountability in the use of AI to support decisions about children and families.

Data security remains a crucial concern. AI tools should operate in secure, controlled environments, and there is rightly strong opposition to using open or consumer-grade tools to handle sensitive data. The distinction between enterprise-level, governed AI solutions and publicly accessible applications is particularly significant here. Using unapproved tools to process sensitive information poses serious risks to confidentiality

and data protection. It is therefore vital to deploy AI in secure environments, with appropriate licensing and safeguards, to uphold trust and ensure compliance.

A major obstacle to adoption is risk aversion rooted in concerns about reputation and uncertainty. When organisations have defined governance frameworks and approved tools, they tend to be more confident in adopting AI and more willing to explore its possibilities. On the other hand, without such safeguards, organisations often shy away from adoption due to perceived dangers. This hesitation is often exacerbated by misunderstandings about AI's role, such as fears of job loss, which further limit engagement.

Tackling these challenges demands a focus on education, capacity building, and clear national policies. There should be straightforward, practical standards that specify what safe and appropriate AI use in children's social care entails. Complex frameworks may hinder adoption, while clear and balanced guidance can help organisations proceed confidently.

It is important to find a balance between safety and innovation. Overly strict regulations, such as those for medical devices, would raise costs, slow progress, and restrict testing and the improvement of new methods. Conversely, the absence of clear boundaries fosters uncertainty, risk-averse behaviours, and unregulated use. A proportionate framework is needed, one that offers clear, practical standards on data security, oversight, and accountability without unnecessary complexity. Achieving this balance will promote confident adoption while safeguarding children and families.

Leadership in local authorities is crucial for shaping the organisational culture and strategy towards AI. When leaders set clear expectations, foster safe experimentation, and effectively communicate AI's role, it leads to better alignment among teams and increased workforce engagement. On the other hand, uncertain or risk-averse leadership can hinder progress across the organisation. Therefore, it's vital that leaders are well-informed and confident in their ability to support AI adoption, facilitating safe and successful large-scale implementation.

Implementation Challenges

Implementation of AI in children's social care is constrained by a range of structural and operational barriers. It is important to recognise that the core technologies required to support children's social care are, in most cases, already available and proven in other public-sector and commercial settings. The critical gap is not technological readiness but the commissioning capability and governance clarity needed to deploy these technologies safely and at scale. Addressing this gap should be the primary focus of central government action.

Recent procurement reform provides an opportunity to improve transparency, flexibility, and engagement with the market. However, evidence from broader public-sector technology procurement suggests that uncertainty remains about the extent to which the Procurement Act alone will resolve longstanding barriers. While the new regime may support better publication of opportunities and more flexible approaches to awarding contracts, legislative change in itself is unlikely to be sufficient without accompanying improvements in commissioning capability, commercial confidence, pre-market engagement, governance clarity, and organisational willingness to adopt more innovative approaches. Without this, many of the structural barriers that limit participation and innovation are likely to persist, not because suitable technology is unavailable, but because the conditions required to commission and deploy it with confidence are not consistently in place.

Funding and budget cycles constrain investment in innovation. Local authorities are often required to demonstrate clear financial benefits before adopting new technologies, which is particularly difficult in the early stages of AI adoption, where benefits may emerge over time or through service redesign rather than immediate cashable savings. This can lead to delayed procurement, limited experimentation, and a preference for lower-risk continuation of existing arrangements. Suppliers may also be expected to invest in discovery, proof-of-concept activity, or solution development at their own risk before any formal commitment is made. This dynamic can reduce market participation, particularly among smaller or more innovative organisations, and slow the pace at which established technologies are adapted, tested, and scaled for use in children's social care contexts.

Procurement processes are often ill-suited to emerging technologies. Traditional approaches tend to focus on procuring predefined tools or detailed technical specifications rather than addressing clearly defined service challenges or desired outcomes. This can inhibit innovation by requiring suppliers to meet fixed requirements rather than proposing adaptable, iterative solutions that evolve with operational needs. It also makes it harder for organisations to respond to the rapid pace of change in AI, where capabilities develop quickly, and implementation often requires continuous improvement rather than a single one-off purchase. The result is that mature, deployable technologies can remain out of reach, not because they do not exist, but because procurement processes are not structured to adopt them effectively.

A move towards outcome-focused procurement is essential. Rather than specifying technologies, procurement should begin with clearly defined service problems, user needs, and intended outcomes. This would enable a broader range of existing solutions to be considered and encourage collaboration between commissioners and suppliers to develop approaches better aligned with frontline realities. It would also support a shift in organisational risk appetite, moving from a transactional model based on compliance

towards one centred on shared problem-solving, continuous improvement, and measurable impact.

Existing procurement mechanisms that support innovation are frequently underused. Frameworks such as pre-commercial procurement and innovation partnerships provide more flexible and collaborative routes for applying and adapting emerging technologies, but are rarely adopted at scale. These approaches can enable earlier engagement with the market, co-design with practitioners, phased testing in live environments, and independent evaluation of impact. They have historically supported effective collaboration across the public, private, and academic sectors. Greater use of these models could reduce implementation risk, improve the quality of solutions, and support wider learning and evidence sharing across the system.

Procurement should also place greater emphasis on value rather than excessive weighting towards the lowest upfront cost. In wider digital procurement, price can dominate evaluation models, sometimes encouraging under-scoped bids, unrealistic delivery commitments, or short-term decisions that fail to deliver sustainable value. In children's social care, where safeguarding, implementation quality, workforce adoption, interoperability, and long-term service outcomes are critical, procurement should assess whole-life value rather than headline price alone.

Capability within commissioning and buying teams is a particularly important factor, and one of the most significant determinants of whether available technologies can be deployed effectively. Procuring AI solutions requires commercial, digital, data, and service-design expertise that is not always consistently available across the sector. Limited in-house capacity can make it harder to assess competing offers, scrutinise technical claims, evaluate implementation plans, or distinguish genuine value from marketing claims. Where this capability is weak, even well-established technologies can be difficult to adopt with confidence. Strengthening commercial and digital capability, alongside access to specialist advice where needed, would materially improve decision-making and procurement outcomes.

Governance clarity is equally critical. In many cases, the technical means to share data, apply AI tooling, or connect systems already exist, but adoption is held back by uncertainty about what is permissible, who is accountable, and how risk should be managed. Variability in how data-sharing rules, information governance requirements, and AI-specific standards are interpreted across organisations leads to inconsistent decisions and risk-averse behaviour. Clearer national guidance on acceptable use, accountability, and oversight would enable commissioners and practitioners to deploy existing technologies with confidence, rather than defaulting to inaction in the face of ambiguity.

Capacity constraints within children's services further limit adoption. High workloads, workforce pressures, and day-to-day operational demands reduce the time available for training, experimentation, and change management. Without dedicated support, protected capacity, and leadership commitment, even well-designed and readily available solutions may struggle to gain traction in practice.

Finally, AI implementation should be understood as an ongoing transformation process rather than a one-off procurement exercise. Many solutions require refinement after deployment as operational needs evolve, users provide feedback, and evidence of impact develops. Contracting and delivery models that support iteration, review, and responsible scaling are therefore likely to be more effective than static procurements based on assumptions that all requirements can be fully defined in advance. The focus should be less on waiting for new technology to emerge and more on building the commissioning and governance conditions that allow existing and maturing technologies to be applied effectively to the challenges children's services face today.

Evaluation of AI and Digital Tools

The current evaluation of AI tools in children's social care is largely informal and limited. Most are in pilot phases, primarily assessing operational metrics such as time savings and administrative efficiency. While these metrics are helpful, they do not fully capture the value of AI in this setting.

There is limited evidence on AI's impact on children and families. This is due to its early adoption and the challenges of measuring outcomes in social care. Evaluation should go beyond checking whether tools work as intended and focus instead on whether they enable better, faster decisions. In safeguarding, decision speed and quality are vital. AI should be judged on its capacity to improve access to information, enable earlier action, and help professionals make more informed judgements.

A key point is that, in many cases where harm has occurred, relevant information was available but not effectively shared or integrated across systems. This underscores the need to evaluate AI within the broader context of information exchange and system-level decision-making, rather than as an isolated tool. A more effective approach is to assess a clear "before and after" scenario to determine whether AI-enabled processes enhance risk visibility, enable quicker escalation when necessary, and lead to better outcomes for children and families.

Evaluation should also reflect the breadth of children's social care objectives. While preventing harm is a central priority, there are also wider goals focused on helping children and young people achieve positive, long-term outcomes. AI has the potential to contribute across this spectrum, and evaluation frameworks should capture its impact on both safeguarding and broader developmental outcomes.

There is a clear need for more structured and consistent evaluation approaches. Establishing common frameworks and metrics would enable comparisons across projects and support the development of an evidence base. Evaluation should focus on a defined set of high-impact use cases, identified in collaboration with practitioners, rather than attempting to assess all potential applications at once. Measuring impact from multiple perspectives, including practitioner experience, outcomes for children and families, and system-level efficiency, will provide a more comprehensive understanding of value.

There is also value in independent evaluation and collaboration with academic partners to ensure robustness and credibility. This would support the generation of high-quality evidence and enable continuous learning across the sector.

A pragmatic approach to scaling is also required. Large, single-nation solutions are unlikely to be effective given the diversity of local contexts and needs. Instead, a model based on interoperable, “good enough” solutions, supported by common standards and governance frameworks, is more likely to enable rapid and effective adoption. This approach allows local flexibility while maintaining consistency in safety and quality.

To support this, central government has an opportunity to invest in regional demonstrators and collaborative innovation programmes. These should be practice-led and focused on real-world challenges, enabling local authorities, technology providers, and research partners to test, evaluate, and refine solutions. Such an approach would accelerate learning, build the evidence base, and support scaling based on proven impact rather than theoretical potential.

Children’s Single View Systems

The concept of a single view of the child is strongly supported as a means of improving safeguarding and decision-making. Access to integrated data across services is seen as critical for identifying risk patterns and enabling timely intervention.

As set out earlier in this response, there are important considerations regarding the relative merits of centralised and federated data models. Implementation should therefore adopt a proportionate approach that recognises the strengths and limitations of both. While a fully centralised data repository may offer a simplified, unified view, it introduces significant complexity, cost, and governance challenges. A federated approach is therefore likely to provide a more practical and scalable foundation, enabling data to remain within existing systems while allowing secure access across organisational boundaries. In this context, a single view should enable practitioners to access relevant information across systems through interoperable, connected services, without requiring the creation of a single, centralised database.

A meaningful single view must extend beyond the individual child to include family and household context, recognising that risk often stems from wider circumstances.

The primary barrier to implementation is not technology but governance. Achieving alignment across organisations and establishing shared data-sharing agreements remain the key challenge.

What techUK Recommends from Central Government

There is a clear need for central government to provide leadership, clarity, and direction to enable the safe and effective adoption of AI in children's social care. In most cases, the core technologies required to support children's services are already available and proven in other public-sector and commercial settings. The priority for central government should therefore be to create the commissioning capability and governance clarity needed to deploy these technologies safely and at scale. This includes not only setting national policy but also strengthening local leadership capacity. Ensuring that senior leaders within local authorities have the confidence, understanding, and frameworks required to make informed decisions about AI will be critical to driving consistent adoption and shaping organisational culture.

Clear, practical guidance is needed on data sharing, governance, and the acceptable use of AI. This is particularly important because, in many cases, the technical means to deploy AI and share data already exist; what is missing is clarity on what is permissible, who is accountable, and how risk should be managed. Guidance should be proportionate and risk-based, enabling effective data use while maintaining appropriate safeguards. There is also a need to avoid overly complex or restrictive regulatory approaches, such as those modelled on medical device frameworks, which risk increasing costs and slowing innovation. Instead, the government should establish simple, actionable guardrails that provide clarity without stifling progress.

As outlined earlier in this response, a proportionate approach to data architecture is required, recognising the benefits and limitations of both centralised and federated models. In practice, greater emphasis should be placed on enabling federated data approaches, which are better aligned with the current structure of public services and allow for scalable, incremental progress. This should be supported by establishing common data standards, clearer data-sharing frameworks, and mechanisms to facilitate collaboration across organisations, alongside sufficient national coordination to ensure interoperability and consistency.

The government should also support the adoption of secure, enterprise-grade AI environments across children's social care. This would reduce reliance on unapproved or informal tools and ensure that sensitive data is processed within controlled, compliant systems. Where possible, priority should be given to enabling AI capabilities within existing, governed case management environments, where data governance, access

controls, and audit arrangements are already established. This offers the fastest and safest route to adoption, while standalone tools can complement this where they provide specialist capabilities not available within embedded solutions. Providing accessible, approved solutions that meet practitioners' needs is essential to addressing the growing use of "shadow AI" and to maintaining trust in how data is handled.

Procurement approaches should shift towards outcome-based models that focus on solving defined service challenges rather than procuring predefined tools. Greater use of innovation-friendly procurement mechanisms, such as pre-commercial procurement and innovation partnerships, would enable more flexible, collaborative solution development. This would support stronger partnerships across the public, private, and academic sectors, enabling co-design, real-world testing, and continuous improvement. Alongside this, investment is needed in commissioning capability within local authorities, including commercial, digital, data, and service-design expertise. Without this, even well-designed procurement frameworks and readily available technologies will be difficult to deploy effectively.

Dedicated funding for pilots and innovation is essential to enable local authorities to test, iterate, and scale solutions. This should include support for regional demonstrators and collaborative, practice-led innovation programmes that bring together local authorities, technology providers, and research partners to address real-world challenges. A focus on smaller-scale, high-impact use cases will enable faster learning and more effective scaling than large, complex national programmes.

There is also a need for national evaluation frameworks to ensure consistent measurement of impact and to build a robust evidence base for AI adoption. Evaluation should focus on real-world outcomes, such as improvements in decision-making, practitioner efficiency, and outcomes for children and families, rather than on purely technical performance.

A pragmatic approach to delivery should be prioritised. Rather than pursuing large, single national solutions, the government should support the development and adoption of interoperable solutions that can be deployed quickly and improved over time. This approach better reflects the diversity of local contexts and enables more rapid progress.

Investing in training and capability-building is essential to address skills gaps and improve understanding of AI across the workforce. This includes addressing misconceptions, such as concerns about job displacement, and building confidence in the safe and responsible use of AI among practitioners and leaders alike. Government should also support the development of proportionate expectations for bias identification, monitoring, and transparency in AI tools used in children's social care, recognising that while bias cannot be fully eliminated, it should be actively managed throughout a tool's lifecycle to protect the interests of children and families.

Finally, consideration should be given to how existing data-sharing and decision-making frameworks can evolve to support more proactive, AI-enabled approaches while maintaining appropriate human oversight and accountability. Enabling AI to better support early identification and intervention, within clear governance boundaries, offers a significant opportunity to improve outcomes for children and families.