POLICY BRIEF

When AI Starts Writing the Record

Why Government Must Lead on Digital Narrative Care

AI is already beginning to reshape the records that drive care, justice, education and public accountability. This brief argues that government must lead in protecting what is human in the record before meaning travels.

Stephen Hall | Digital Narrative Care
stephenhall.uk.consulting@gmail.com
April 2026

The central governance challenge is no longer whether AI can be deployed, but whether it can be trusted, governed, and corrected in the real conditions of work.

When AI tools draft, summarise or reshape institutional records, the central governance question is not only whether the tool is accurate. It is whether the travelling account that later moves through services, decisions, reviews and appeals still preserves who said what, what was inferred, what was changed, and who remains accountable. That travelling account is the real governance object: not only the original record, but the summary, classification, inherited interpretation, and institutional memory built from them. Digital Narrative Care (DNC) is a practical governance framework for that problem. Its framework integrates record integrity, meaning integrity and temporal integrity, supported by evidential care, interpretive care and aftercare.

DNC governs three linked layers that current AI governance language often leaves blurred. First, what entered the record. Second, what the institution made of it. Third, what hardens over time as that account is inherited, reused and acted on. The issue is not the tool in isolation. It is what happens when AI systems begin to shape the accounts through which institutions understand people, justify action and retain memory over time.
The first task is to name the threat. The threat is not AI in the abstract. It is AI-shaped narrative entering an institutional record and then travelling as if it were settled fact, without a sufficient trail of provenance, review, correction or challenge. That risk will not be solved by transparency about tools alone. It requires attention to what the tools produce, how those outputs are checked, and how the resulting account is allowed to move through systems.

Why government must lead

Government is not simply another adopter of AI. It sets standards, writes procurement rules, defines the assurance environment, and remains answerable to Parliament and the public when that environment fails. The UK already has important pieces of this architecture in place. The Government Digital Service's AI Playbook for the UK Government offers guidance on using AI safely, effectively and securely across government. The Algorithmic Transparency Recording Standard provides the framework through which public bodies publish structured

information about how and why algorithmic tools are used, and government has made it mandatory for all central government departments.

That matters because DNC is not a standalone ethics add-on. It is a missing operational layer within governance that already exists. In healthcare, justice and other public services, the UK is moving from abstract AI policy to real implementation. The question is no longer whether AI will be used in record-intensive settings. It already is. The question is whether current guidance and procurement are specific enough about the integrity of the travelling accounts these systems create, edit and carry forward. Government must therefore lead not only as a user, but as the authority that sets the conditions under which everyone else uses these tools.

The next governance challenge is not procurement alone but implementation. AI systems do not become trustworthy because a contract is signed or a framework exists on paper. They become trustworthy when the people using them understand what they are seeing, what must still be checked, when human judgement is required, and how errors can be corrected once the record begins to travel. In that sense, the real test is adoption, not procurement: the policy problem is whether governance survives contact with practice. Productivity claims are not enough if they are purchased at the cost of contestability, accountability, or the quiet transfer of checking burdens and liability onto frontline staff.

Pilot activity, transparency language and orderly rollout plans are not readiness in themselves. Real readiness means role clarity, record-level checking, workable routes for correction, and the ability to see when an interpretation has begun to harden into institutional memory.

Policy considerations

What follows should be read as evidence of one shared governance object, not as a tour of separate sector problems. Different sectors reveal different faces of the same issue: the travelling account and the institutional memory built from it.

1. Record integrity: the live governance gap is the travelling account, not only the model

Assurance frameworks must address not just whether AI tools are accurate, but whether the account they produce remains explainable, attributable and contestable as it travels through systems.

The UK's existing frameworks are beginning to recognise this. NHS England's guidance on AI-enabled ambient scribing products makes clear that these tools are now being adopted for clinical and patient documentation and workflow support. The same guidance warns that clinical hazards may include missing critical information, incorrect information or context, and delayed outputs, and it requires organisations to identify a Clinical Safety Officer and key risks. NHS England's implementation guidance also makes clear that notes, summaries and letters created by ambient scribes may contain errors, and that clinicians remain responsible for checking them before anything is added to the record.

That is already close to the heart of the DNC argument. The issue is not whether an AI system exists. It is whether there is a visible, reviewable chain between conversation, draft, revision, approval, storage and later use. In other words, assurance is a dial, not a switch. The intensity of checking should vary by risk, but the record must remain intelligible and contestable throughout.

2. Meaning integrity: healthcare, social care, policing and justice are the most visibly activated sectors

Where AI reshapes what someone said into what the record says they meant, the integrity of meaning must be protected. These are the sectors where that risk is already live and consequential.

The most extensively evaluated current use of AI-assisted clinical documentation offers an instructive evidential window — not because the problem is confined to healthcare, but because this is where deployment is furthest advanced and independent scrutiny has begun to accumulate.

The productivity case is real and should be conceded clearly. A London multi-site study across nine sites — 165 clinicians, 16,470 patient encounters — found a 23.5 per cent increase in direct care time and an 8.2 per cent reduction in total appointment time. A Dutch general practice study recorded 42.7 seconds less documentation time per consultation. These are meaningful gains in clinical attention and workforce burden.

But the evidence on what AI-assisted clinical documentation produces — the account itself — tells a more complex story. An independent evaluation of AI documentation tools in acute NHS services found that 37.3 per cent of logged uses required editing for clinical accuracy, that 44.4 per cent of participating clinicians had encountered invented information in at least one encounter, and that 24 hallucination-related instances were formally logged. A separate study reported drawbacks including inaccurate summaries and possible interference with clinician reasoning. That is not an argument against deployment. It is an argument for governance that matches the pace of it.

The independent evidence base makes that case more sharply still. A rapid review identified 1,450 studies on AI-assisted clinical documentation — and found only six that met inclusion criteria for real-world implementation evidence. The NIHR Rapid Service Evaluation Team and Nuffield Trust described the existing evidence as fragmented and often industry-led, and identified record quality, patient outcomes, fairness and value for money as areas requiring robust independent evaluation. That gap — between deployment speed and governance depth — is not unique to healthcare. It is the pattern this brief is designed to address.

Healthcare is currently the clearest UK proof point for that pattern in action. Oxford University Hospitals described its AI-assisted documentation pilot as one of the most comprehensive evaluations undertaken in an acute NHS setting to date. OUH reported that nearly 90 per cent of clinicians reduced documentation time, more than half saved 30 minutes or more per day, 73 per cent reported improved wellbeing, and consent rates were 99.7 per cent.

Royal Devon provides a second, strongly UK-facing anchor. In March 2026, the trust said AVT had already supported more than 2,700 outpatient consultations, that it had become the first trust to integrate AVT fully into its Epic electronic patient record, and that its NHS England-backed emergency department pilot was intended to inform wider NHS learning and adoption.

These examples matter not because they prove the technology is settled, but because they show the opposite: deployment, evaluation and assurance are now happening at speed inside the NHS. That is exactly where DNC belongs. The UK's AI challenge is no longer only invention or procurement. It is whether adoption is outpacing the governance, training and professional practice needed to keep meaning human once accounts start to travel.

Social care is similarly active, but with weaker and less consistent safeguards. The Ada Lovelace Institute says councils in England and Wales are rapidly piloting AI-powered transcription and summarisation tools in social work. Its 2026 research found that social workers often carry the accountability burden for checking outputs, that oversight is inconsistent, and that inaccuracies, omissions and hallucinations can enter official

documentation. Unlike healthcare, social care does not yet benefit from an equally developed assurance infrastructure. That makes the risk sharper, not smaller.

Policing and the wider justice system are also active and should be treated as such. The Alan Turing Institute has argued that algorithmic transparency is achievable in policing and could bring significant rewards. The Crown Prosecution Service states that AI outputs must be verified by an appropriately qualified and trained person and that individuals remain responsible for each use and output. The Ministry of Justice's AI Action Plan for Justice sets out a responsible and proportionate approach to AI adoption across courts, tribunals, prisons, probation and supporting services. Judicial guidance in England and Wales adds that any use of AI by or on behalf of the judiciary must protect the integrity of the administration of justice and uphold the rule of law.

The most significant recent development in this space is the creation of Police.AI, a national centre backed by over £115 million to accelerate AI adoption across all 43 forces in England and Wales. Its named early priorities include call transcription, case file production, crime recording and disclosure processes. These are not peripheral administrative functions. They are the mechanisms through which spoken encounters become institutional records — the precise point at which meaning integrity, record provenance and temporal accuracy are most at risk. The governance frameworks accompanying Police.AI emphasise transparency, bias testing and ethical oversight. What they do not yet address with sufficient specificity is what happens inside the record: whether the account that enters the file still reflects who said what, what was inferred, what was changed, and who remains accountable as that record travels. That is the upstream gap this brief is designed to name.

If these protections are not built into the governance of AI-assisted records, the consequences will not be abstract. A flawed clinical summary may shape treatment. A distorted care note may influence statutory intervention in a family's life. A police or justice record that cannot be traced or challenged may affect liberty, evidence and redress. A student record, misconduct file or appeals file shaped by AI may influence progression and sanction. In each case, the same risk appears: once a machine-shaped account enters the file, it can begin to govern the person before the person has any meaningful chance to answer back.

3. Temporal integrity: higher education is structurally important even where named deployments are thinner

AI systems treat contingent, time-sensitive captures as permanent institutional facts. In higher education, that risk is compounded by a second temporal problem: deployment is already outpacing the governance, formation and legitimacy structures needed to contain it.

Higher education belongs in this brief for a reason that goes beyond its deployment map. Universities are part of the governance pipeline itself — they are where future professionals form their understanding of responsibility, review, record-making and judgement in workplaces where AI systems are increasingly present. AI is entering practice before the professionals using it have been shaped to interrogate what it does to records, to meaning, and to where responsibility sits. That makes the formation argument the primary justification for including higher education here, not a footnote to it.

The visible deployment story is less mature than in health or justice — but that is not the same as lower risk. The Higher Education Policy Institute's 2026 survey found that 95 per cent of students report using AI in at least one way, 94 per cent use generative AI to help with assessed work, only 36 per cent feel encouraged by their institution to do so, and only 38 per cent say they are provided with AI tools. AI use is already deeply embedded in student practice while institutional rules, support and auditability are still catching up. Universities are record-making institutions: assessment, progression, misconduct, appeals, disability support, student wellbeing and research governance all depend on records that must remain reviewable and contestable.

The pace question sharpens this further. As HEPI argued on 9 April 2026, universities are "slow" for reasons bound up with fairness, standards, accountability and public trust, while AI is "fast" in ways that quickly reset expectations and practice. Too slow, and AI use becomes unofficial, uneven and privately optimised; too fast, and policy becomes reactive, implementation inconsistent, and legitimacy fragile. In higher education, governing pace without losing legitimacy is now part of the policy problem itself.

Higher education is therefore not only a governance priority in its own right. It is part of the governance pipeline. The brief's final recommendation follows from that directly.

4. Financial services and workforce systems: important emerging lanes, not the primary evidence base here

Financial services and workforce systems are real future lanes for DNC, but they should not displace the stronger public-interest spine of this brief. They raise analogous questions about provenance, review, contestability and the downstream force of records. The strongest current public proof points sit elsewhere. The proportionate conclusion, for now, is that these are important emerging sectors, not the primary evidence base. In those sectors too, efficiency gains should not be treated as sufficient justification if the practical burdens of checking, explanation and liability are simply pushed downward onto frontline staff or individuals affected by the record. That principle travels.

Recommendations

1. Embed DNC criteria into central AI governance and procurement guidance

Government should update core AI governance and procurement materials so that any AI system that creates, edits or materially summarises an official record must meet explicit DNC criteria. At minimum, those criteria should require clear output logging, a named human reviewer, version visibility, and a route for correction or challenge. Existing frameworks already provide part of the scaffolding. The task now is to name the travelling account explicitly as the governance object — and to build assurance around how it is interpreted, acted on, and carried forward.

2. Require auditable human review for high-consequence records

Where AI affects records used in clinical care, social work, policing, justice, education, or other high-consequence decisions, no AI-produced summary should enter the official record without documented human review and approval. The record should show when AI was used, what it produced, who reviewed it, what changed, and how later corrections can be made visible over time.

3. Name accountability for the institutional-memory layer and require periodic audit

Government should require public bodies deploying AI in high-consequence record systems to name a lead responsible for the institutional-memory layer: for monitoring whether AI-generated summaries, classifications and inherited interpretations are accumulating in ways that distort how people are represented and judged. That brief should include periodic audit of persistent record effects, especially where classifications, summaries or risk flags travel across time, teams or agencies. This is not a new bureaucracy. It is a specific governance responsibility that must be owned.

4. Commission sectoral pilots focused on integrity in travel, not only productivity

Government should commission and evaluate a small number of implementation pilots in sectors where the pressure is already real: NHS settings, local authority social care, policing

and justice workflows, and higher education casework or assessment processes. These pilots should test not only time savings, but whether record integrity, meaning integrity and temporal integrity are preserved in practice, and whether the resulting accounts remain reviewable, contestable and correctable once they begin to travel.

5. Use a proportionate risk model: stronger checks where accounts travel furthest

Assurance is a dial, not a switch. Low-risk administrative uses may justify lighter-touch controls. Accounts that shape care, liberty, sanctions, progression, eligibility or formal redress should trigger stronger requirements. This is where DNC's care triad becomes practical: evidential care in capture, interpretive care in summarisation and review, and aftercare in correction, explanation and follow-through when something goes wrong.

6. Build workforce capability into implementation, not around it

Government should require high-consequence AI deployments to include role-based training, supervision and review responsibilities as part of implementation, not as optional follow-on support. Auditability without workforce capability is not enough. Where AI systems shape records used in care, justice, education, regulation or formal redress, implementation plans should show who checks outputs, how judgements are reviewed, how corrections are made visible, and how responsibility is held in practice.

7. Treat professional formation as part of the governance pipeline

Government should work with professional regulators and higher education quality bodies to establish that initial professional education for roles in high-consequence settings must address AI-mediated record-making, interpretive judgement, and professional accountability as conditions of practice readiness — not as optional digital literacy add-ons, but as core formation requirements.

The current governance conversation is largely focused on deployment: procurement, audit, liability, and checking. That is necessary. But it addresses people already in post. It does not address the formation gap — the fact that AI is now entering practice before the professionals using it have been shaped to interrogate what it does to records, to meaning, and to where responsibility sits.

Higher education is part of the governance pipeline, not downstream of it. Where AI systems are already embedded in student practice while institutional rules and professional formation frameworks are still catching up, the pace asymmetry is itself a governance risk. Government should therefore require that professional regulators review whether current standards for initial education and CPD adequately address AI-mediated practice — and that higher education quality bodies treat professional formation for AI-mediated work as a distinct and assessable expectation, not a subset of general AI literacy.

The formation gap is part of the governance gap. Closing it requires action at both ends of the professional pipeline.

Conclusion

The UK does not need to wait for a wholly new legal regime to act. Important elements of the policy infrastructure already exist: government-wide AI guidance, transparency standards, sectoral implementation frameworks, and increasingly explicit public-sector use cases. What is missing is a disciplined way of protecting the travelling accounts these systems now help to create, and the institutional memory that accumulates around them.

Digital Narrative Care names that problem clearly. It offers a way to move beyond generic calls for responsible AI and towards something more operational: accounts that remain explainable, attributable, reviewable and contestable as they move; interpretations that do not quietly harden beyond challenge; and institutions that accept responsibility for what their

records come to do over time. If government leads now through procurement, guidance, auditability, workforce capability, memory-layer accountability and targeted pilots, it can help ensure that AI strengthens public services without weakening the integrity of the accounts on which those services depend.

Guidance, procurement standards and sector pilots can move first, and should. But voluntary frameworks have a ceiling. Where AI systems shape records that determine care, liberty, sanction, progression or redress, the direction of travel must be towards enforceable duties — not as a remote aspiration but as the logical end point of a governance conversation that has already begun. It is a core condition of administrative legitimacy, public trust and accountable decision-making.

Digital Narrative Care provides that operational layer. Its instruments are available for piloting and independent evaluation.

References

Ada Lovelace Institute. Scribe and Prejudice? and Transcribing Trust. 2026.

Crown Prosecution Service. How the CPS Will Use Artificial Intelligence. Accessed April 2026.

Data in Government blog. Making the Algorithmic Transparency Recording Standard (ATRS) Mandatory Across Government. 8 May 2025.

Government Digital Service. AI Playbook for the UK Government. 10 February 2025.

GOV.UK. Algorithmic Transparency Recording Standard: Guidance for Public Sector Bodies. Updated 8 May 2025.

Higher Education Policy Institute (HEPI). Student Generative AI Survey 2026. 12 March 2026.

Higher Education Policy Institute (HEPI). Tempo Conflict at the Heart of AI in Higher Education. 9 April 2026.

Judiciary of England and Wales. Artificial Intelligence (AI) — Judicial Guidance. 31 October 2025.

Ministry of Justice. AI Action Plan for Justice. 31 July 2025.

NHS England. Guidance on the Use of AI-Enabled Ambient Scribing Products in Health and Care Settings. 2025.

NHS England. Implementation Guidance for AI-Enabled Ambient Scribing Products in Health and Care Settings. 2025–26.

Oxford University Hospitals NHS Foundation Trust. Successful AVT Pilot at OUH Shows Promising Results for Staff and Patients. 23 February 2026.

Royal Devon University Healthcare NHS Foundation Trust. Royal Devon Leads NHS England-Backed Pilot to Evaluate Ambient Voice Technology in Emergency Departments. 5 March 2026.

The Alan Turing Institute. New Report Explores Police Perspectives on the Algorithmic Transparency Standard. Accessed April 2026.

Wray, J. et al. Time to Care: The Value of AI-Powered Ambient Voice Technology Across Diverse Clinical Settings. NHS England, January 2026.

van Linschoten, R.C.A. et al. Ambient Scribe in General Practice: A Multi-Perspective Before-After Longitudinal Mixed-Methods Study. npj Digital Medicine, 2026. medRxiv preprint doi: 10.1101/2025.10.30.25339135.

TheHill / Oxford University Hospitals NHS Foundation Trust. Ambient Voice Technology (AVT) in Acute Services: An Evaluation. Oxford University Hospitals NHS Foundation Trust, February 2026.

Kanaparthy, N.S. et al. Real-World Evidence Synthesis of Digital Scribes Using Ambient Listening and Generative Artificial Intelligence for Clinician Documentation Workflows: Rapid Review. JMIR AI, 10 October 2025; 4:e76743. doi: 10.2196/76743.

Shand, J. and Morris, S. Ambient Voice Technology in Health Care: What's the Evidence So Far? NIHR Rapid Service Evaluation Team / Nuffield Trust, 4 February 2026.