Synthetic Identity, Avatars and AI-Generated Media Policy

20 Oct 2026, 24 min read

> **Operated by:** Nudge Education Ltd · **Version:** v08.26 (draft for approval) · **Owner:** AI Governance Lead · **Safeguarding owner:** Designated Safeguarding Lead · **Data protection owner:** Data Protection Officer / Data Protection Lead

Policy status: Draft for approval **Approved by:** [Board / Executive Leadership Team] **Effective from:** On approval **Review date:** At least annually, and following any material change in technology, law, regulation, safeguarding guidance or supplier arrangements

1. Purpose

Nudge Education uses artificial intelligence where it can improve accessibility, consistency, efficiency and the quality of our work.

This may include the creation of synthetic representations of members of staff: for example, an AI-generated video avatar or synthetic voice used to deliver induction, training or organisational information.

These technologies can reproduce a real person’s appearance and voice with a high degree of realism. The same capabilities that make them useful also create risks of impersonation, deception, identity misuse, loss of control over personal data and safeguarding harm.

This policy establishes the conditions under which Nudge may create, use, publish, retain and withdraw synthetic representations of real people.

Its central principle is:

A person’s identity remains theirs, even when technology makes it reproducible.

Nudge does not acquire an unrestricted right to a person’s face, voice, mannerisms or identity by creating an AI avatar of them.

2. Scope

This policy applies to:

- employees;
- Flexi Practitioners;
- directors;
- contractors;
- consultants;
- volunteers;
- agency workers;
- former members of the Nudge workforce where previously created synthetic material exists;
- anybody creating synthetic media on behalf of Nudge; and
- any third-party supplier processing a person's image, voice or other identity data for Nudge.

It applies to:

- AI video avatars;
- digital twins;
- voice clones;
- synthetic speech;
- lip-synchronised or digitally altered video;
- AI dubbing that reproduces an individual's voice;
- face replacement or face-generation technology;
- photorealistic synthetic images of identifiable people;
- AI-generated or manipulated audio or video that represents a real individual; and
- future technologies capable of materially reproducing an individual's identity.

This policy uses the term **synthetic representation** as the general term for these outputs.

Some synthetic representations may also fall within the EU AI Act definition of a **deepfake**, depending upon their realism, context and capacity to appear authentic. The EU definition includes AI-generated or manipulated image, audio or video resembling an existing person or event which could falsely appear authentic or truthful.

3. Our governance standard

Nudge applies the Diamond Standard to the use of synthetic identity.

The Standard requires:

Safety — risks are considered before implementation and AI does not make safeguarding decisions.

Sovereignty — people retain control over their identity, authorship and personal information.

Symmetry — where technology can see, interpret or convincingly represent a person, the person affected should be able to understand what the technology is doing.

Stewardship — AI may support human work, but human accountability is not transferable.

Sustainability — the use of AI must be proportionate to the benefit it creates.

Synthetic Intimacy — Nudge will not use synthetic systems to manufacture trust, dependency or false relationships.

These principles form part of the wider Verse-ality/Diamond control set.

Nudge also adopts **Consent as Protocol**: a material change of scope creates a new consent event. Consent to one use does not silently migrate into consent for another.

Three practical questions apply to every deployment:

Can it stop? Can it forget? Can somebody refuse without cost?

PART ONE — SYNTHETIC IDENTITY

4. Identity sovereignty

A synthetic avatar of a member of staff is **not an organisational persona owned by Nudge**.

The represented individual retains a continuing interest in:

- their likeness;
- their voice;
- their manner of representation;
- the contexts in which their identity is used;
- the words presented as coming from them; and
- whether future synthetic content may continue to be generated.

Creation of an avatar does not permit Nudge to make an individual appear to:

- endorse something they have not endorsed;
- express an opinion they have not expressed;
- attend an event they did not attend;
- participate in a conversation they did not participate in;
- make a decision they did not make;
- communicate directly with an individual when they did not do so;
- provide evidence that they were present at a particular time or place; or
- assume responsibility for something they have not personally approved.

Synthetic media must never be used to falsify organisational history.

5. Consent is a process, not a checkbox

5.1 Separate consent events

Nudge distinguishes between:

1. permission to capture source material;
2. permission to create the synthetic representation;
3. any biometric consent required by the technology provider;
4. permission for Nudge to use the representation;
5. the purposes for which it may be used;
6. the audiences to whom it may be shown;
7. approval of what the representation may say;
8. permission to create materially different versions, languages or formats; and
9. permission for any continued use after the person’s relationship with Nudge ends.

These permissions must not be collapsed into one indefinite acceptance.

5.2 No presumed consent

Consent will not be inferred from:

- employment;
- participation in filming;
- acceptance of another Nudge policy;
- previous appearance in conventional photographs or videos;
- use of an avatar on an earlier occasion;
- silence;
- failure to object;
- a general intellectual-property clause in a contract; or
- the fact that another member of staff has agreed.

5.3 Refusal without detriment

Nobody is required to provide their likeness or voice for creation of a synthetic avatar unless an exceptional requirement has been expressly agreed before appointment and lawfully incorporated into the role.

For existing workers, refusal must not adversely affect:

- recruitment or continued employment;
- allocation of work;
- remuneration;
- promotion;
- performance assessment;
- access to opportunities;
- professional standing; or
- relationships with managers.

Where an avatar is proposed for work that could reasonably be delivered another way, an appropriate alternative must be available.

5.4 The employment power imbalance

Permission under this policy is an ethical and organisational authorisation. It must not automatically be treated as Nudge's UK GDPR lawful basis for

processing personal data.

The ICO advises that consent will not usually be an appropriate lawful basis where there is a clear imbalance of power, particularly between employer and worker, unless genuine freedom to refuse or withdraw can be demonstrated.

Nudge's Data Protection Officer or Data Protection Lead must therefore separately determine and record:

- the relevant Article 6 lawful basis;
- whether special-category processing occurs;
- any applicable Article 9 condition;
- international transfer arrangements; and
- the required privacy information.

A signed avatar permission form does **not**, by itself, establish data-protection compliance.

6. Bounded Synthetic Identity Mandate

Each synthetic representation must have a written **Synthetic Identity Mandate**.

The mandate defines what Nudge may and may not do with that representation.

It must state:

- the person represented;
- the technology provider;
- the approved purpose;
- the approved audience;
- whether use is internal, external or both;
- whether young people may see the content;
- authorised subject areas;
- prohibited subject areas;
- who may prepare scripts;
- who may approve scripts;
- whether translation or dubbing is permitted;
- the approved languages;
- whether personalised content is permitted;
- the expiry date;

- the withdrawal route; and
- whether any post-employment use is authorised.

The default position is:

Anything not expressly authorised is not authorised.

A mandate must be renewed at least annually and whenever the scope materially changes.

7. Script and editorial control

An avatar must not be treated as an autonomous spokesperson.

Every published output must have a named human accountable for its substance.

A materially new script must either:

- be personally approved by the represented individual; or
- fall within an explicitly documented, bounded script mandate previously approved by that person.

A bounded script mandate may, for example, permit an avatar to deliver updated induction material where:

- the subject remains within an agreed policy area;
- the factual content has been approved by the responsible Nudge policy owner;
- the tone and type of communication remain within the mandate;
- the avatar is not made to express a new personal opinion; and
- the represented individual retains the ability to stop future use.

Synthetic media must never be used to manufacture retrospective approval.

PART TWO — TRANSPARENCY

8. People must know when media is synthetic

Nudge does not trade on ambiguity.

Where somebody could reasonably believe that synthetic image, audio or video is a conventional recording of a real person, Nudge will disclose its synthetic nature.

As a minimum, video content must include a clear statement such as:

AI-GENERATED AVATAR · HUMAN-APPROVED CONTENT

or:

This video uses an AI-generated representation of [name]. The content has been reviewed and approved by Nudge Education.

For audio-only content, an audible disclosure must be made.

The EU AI Act's Article 50 transparency requirements apply from 2 August 2026 and require deployers to clearly disclose qualifying deepfake content at first exposure. The European Commission says this disclosure must be understandable and perceivable without special technical tools.

Nudge applies that standard **organisation-wide regardless of whether a particular audience or use falls within the territorial scope of the EU AI Act**.

8.1 Persistent disclosure

The Diamond Standard goes further than disclosure only at first exposure.

Where reasonably practicable, Nudge synthetic video will retain a visible disclosure marker throughout playback.

This reflects the principle of **identity non-capture**: people should not gradually lose awareness that they are encountering a synthetic representation.

8.2 Metadata

Where the provider supplies machine-readable provenance, content credentials, watermarking or equivalent indicators that content is AI-generated, Nudge must not deliberately remove or obscure them.

PART THREE — SAFEGUARDING

9. Children and young people

Nudge works with children and young people, including young people who may be vulnerable because of their circumstances, needs or previous educational experiences.

Synthetic identity therefore carries safeguarding implications beyond ordinary corporate communications.

Keeping Children Safe in Education 2026 comes into force on 1 September 2026. It explicitly recognises harmful interactions with generative AI applications as part of online **contact risk**, and recognises that images may be digitally altered or wholly generated by AI, including deepfakes and deep nudes.

Nudge adopts these relevant protections in advance of the statutory commencement date and wherever they represent sensible safeguarding practice.

Where Nudge is operating for, with or within an organisation directly subject to KCSIE, the applicable statutory safeguarding arrangements remain paramount.

9.1 No synthetic replicas of children

Nudge will not create a digital twin, voice clone or photorealistic synthetic replica of a child or young person.

This is a safeguarding invariant under this policy.

Any future proposal to change this position would require a separate safeguarding, legal, ethical and governance process and cannot be authorised merely as an exception to this policy.

9.2 Child-facing staff avatars

Pre-recorded staff avatars may be used for generic educational or informational content seen by young people only where:

- the use has a genuine educational or accessibility purpose;
- the content is clearly labelled as synthetic;
- the represented person has authorised child-facing use;

- the content contains no personal information about individual young people;
- a safeguarding assessment has been completed where appropriate;
- the content cannot reasonably be mistaken for a live or private interaction; and
- a responsible human remains available where the content could generate questions or concerns.

9.3 Prohibited child-facing uses

A synthetic representation of a member of staff must not:

- conduct a private one-to-one relationship with a young person;
- impersonate live contact from a practitioner;
- provide unsupervised pastoral or therapeutic support;
- solicit disclosures;
- conduct safeguarding interviews;
- assess whether abuse has occurred;
- determine safeguarding risk;
- undertake behaviour management;
- make disciplinary decisions;
- create emotional dependency;
- represent itself as personally caring about or remembering a particular child;
- simulate friendship, attachment or exclusivity;
- replace the responsible adult in a safeguarding interaction; or
- be used as the sole source of important information about a child's welfare.

There are no AI-only safeguarding decisions.

This reflects the Diamond controls of Safety, Stewardship and Synthetic Intimacy.

10. Synthetic media as safeguarding evidence

Modern synthetic media means that apparent realism is no longer proof of authenticity.

If a child or adult presents a concerning image, voice recording or video:

- staff must take the concern seriously;
- staff must not dismiss it simply because it may have been AI-generated;
- staff must not assume it is authentic merely because it appears realistic;
- staff must follow Nudge safeguarding procedures;
- concerns involving a child must be escalated through the appropriate safeguarding route; and
- staff must not conduct their own technical or forensic investigation unless specifically authorised to do so.

A disputed or synthetic image may still be evidence of harassment, coercion, sexual abuse, bullying, reputational harm or another safeguarding concern even if the event apparently depicted did not occur.

PART FOUR — DATA PROTECTION AND SECURITY

11. Personal and biometric data

Images and recordings of identifiable staff are personal data.

Some avatar systems also process facial geometry and voice characteristics.

HeyGen currently states that custom-avatar creation involves face geometry and may involve voiceprint data, and that its identity-verification process compares biometric information to confirm that the person granting permission is the person represented.

Where biometric data is technically processed for the purpose of uniquely identifying an individual, it may constitute special-category personal data.

The Data Protection Officer or Data Protection Lead must be involved before Nudge creates staff avatars at organisational scale.

12. Data Protection Impact Assessment

A DPIA screening must be completed before adopting a new synthetic-identity platform or materially changing how an existing platform is used.

A full DPIA must be undertaken where required by law or where the screening identifies likely high risk.

The ICO identifies innovative AI technology and certain biometric processing as circumstances capable of triggering DPIA requirements, and advises organisations to use DPIAs as substantive risk-management tools rather than compliance tick boxes.

The DPIA or screening record must consider:

- purpose and proportionality;
- identity and biometric processing;
- employee power imbalance;
- supplier access;
- international transfers;
- data retention;
- model training;
- subprocessors;
- security;
- impersonation risks;
- withdrawal;
- residual copies;
- children who may see the resulting material; and
- less intrusive alternatives.

13. No child data in avatar-generation tools

Personal information concerning a child or young person must not be entered into a synthetic-avatar or video-generation system unless that specific processing has received prior approval under Nudge's data protection and safeguarding procedures.

As the default:

- scripts must use generic examples;
- names must be removed;
- case details must be fictionalised or anonymised; and
- confidential safeguarding information must never be uploaded.

DfE guidance says schools should understand how approved AI tools process personal data and warns that pupil personal data should not be entered into generative AI systems without appropriate data-protection assessment and assurance that it is protected and not used for further model training.

14. Supplier requirements

Before an avatar platform is approved for Nudge use, the responsible owner, Information Security/Data Protection function and AI Governance Lead must establish, as appropriate:

- the supplier's contractual role;
- a Data Processing Agreement;
- relevant international-transfer safeguards;
- where data is hosted;
- the categories of personal and biometric data processed;
- subprocessor arrangements;
- retention periods;
- deletion mechanisms;
- security controls;
- access controls;
- model-training arrangements;
- incident-notification arrangements; and
- the process for deleting or disabling an avatar.

Where a supplier offers a mechanism to prevent staff identity data or customer content being used to train or improve general AI models, **Nudge's default is to opt out unless a separate, explicit governance decision authorises otherwise**.

HeyGen's current privacy information states that personal data may be used for research and development of AI models but provides an opt-out; this must therefore be checked and configured appropriately for Nudge's account and contractual tier.

HeyGen's current Data Processing Addendum specifically recognises image and voice recordings and biometric information associated with facial identity verification.

Supplier terms and settings must be rechecked periodically because platform capabilities and contractual arrangements change.

15. Access and security

Synthetic identities are privileged organisational assets, but they remain representations of real people.

Access must therefore follow least-privilege principles.

Nudge must:

- use organisational rather than personal accounts;
- enable multi-factor authentication where available;
- restrict avatar-generation permissions to approved users;
- keep a record of authorised operators;
- remove access promptly when responsibilities change;
- prohibit credential sharing;
- protect source footage;
- restrict downloading of reusable avatar assets where possible; and
- maintain a record of consequential synthetic outputs.

Nobody may create an unofficial Nudge avatar of a colleague using a personal AI account.

PART FIVE — USE CONTROLS

16. Traffic-light use framework

GREEN — normally acceptable with standard approval

Examples include:

- internal induction;
- generic staff training;
- policy explainers;
- procedural guidance;
- software walkthroughs;
- accessibility versions of approved material;
- factual organisational updates within an agreed Synthetic Identity Mandate.

These still require transparency and appropriate script approval.

AMBER — additional review required

Examples include:

- public website content;
- social media;
- external marketing;
- commissioner-facing material;
- recruitment campaigns;
- child-facing pre-recorded material;
- multilingual voice cloning;
- AI dubbing;
- personalised video;
- public statements on sensitive topics.

Amber uses require specific approval by the represented person plus the relevant Nudge accountable role.

Child-facing Amber uses require safeguarding review.

RED — prohibited

Nudge will not use synthetic representations to:

- conceal the fact that content is AI-generated;
- impersonate somebody without permission;
- manufacture evidence;
- fabricate attendance at a meeting or event;
- issue disciplinary or performance-management messages;
- dismiss or recruit somebody;
- communicate a safeguarding decision;
- make a decision about an individual child;
- simulate a private conversation that did not occur;
- manufacture a testimonial;
- make somebody appear to endorse a commercial, political or personal position they did not approve;
- create sexualised or intimate synthetic material;
- bully, humiliate or harass;

- exploit vulnerability;
- manufacture emotional dependency;
- make an avatar appear sentient or personally attached to a user;
- recreate a child;
- create a synthetic representation of somebody who has refused permission; or
- continue creating new material after permission has been withdrawn.

17. Live and interactive avatars

A pre-recorded avatar and an interactive synthetic agent are not the same risk.

This policy permits approved **pre-recorded** synthetic media.

It does **not** automatically authorise an avatar to become:

- conversational;
- autonomous;
- memory-enabled;
- personalised;
- agentic;
- capable of sending messages;
- capable of responding to young people; or
- capable of taking action.

Turning a staff avatar into an interactive system is a **material scope change** and requires a separate AI risk assessment, data-protection assessment, safeguarding review and explicit renewed permission.

PART SIX — WITHDRAWAL AND END OF RELATIONSHIP

18. Withdrawing permission

A represented person may withdraw permission for future synthetic use through the route stated in their Synthetic Identity Mandate.

They do not need to justify the decision.

On withdrawal, Nudge must:

1. stop generating new content using that identity;
2. suspend unnecessary operator access to the avatar;
3. identify content currently in circulation;
4. stop future scheduled distribution where reasonably practicable;
5. remove or replace Nudge-controlled content as soon as reasonably practicable;
6. request relevant deletion or restriction from suppliers where appropriate; and
7. record the withdrawal and actions taken.

Unless there is a documented reason why this cannot reasonably be achieved, Nudge should aim to remove affected material from active internal channels within **10 working days**.

Public or third-party copies outside Nudge's direct control may not always be capable of complete recall. This limitation must be explained honestly rather than promising impossible deletion.

A restricted record may be retained where necessary for legal, safeguarding, contractual, audit or accountability purposes, subject to the appropriate lawful basis and retention requirements.

Withdrawal stops future permission; it does not rewrite history.

19. Leaving Nudge

Permission to generate new synthetic content ends automatically when an individual leaves Nudge unless a separate post-employment agreement expressly provides otherwise.

On or before departure:

- generation access must be suspended;
- active synthetic content must be reviewed;
- continuing business need must be assessed;
- the individual's preferences must be recorded;
- internal materials should be replaced or updated where practical; and
- suppliers must be instructed appropriately.

Nudge will not rely upon a general employment intellectual-property clause as permission to keep generating new speech or video in a former colleague's likeness.

A post-employment licence, where genuinely required, must be separately agreed and specifically scoped.

PART SEVEN — INCIDENTS AND MISUSE

20. Suspected misuse of a synthetic identity

Misuse includes:

- unauthorised creation;
- unauthorised publication;
- an avatar saying something outside its mandate;
- compromised avatar credentials;
- fraudulent imitation;
- removal of AI disclosure;
- synthetic harassment;
- creation of sexualised imagery;
- unauthorised voice cloning;
- publication following withdrawal; or
- use affecting a child or safeguarding concern.

Staff must report suspected misuse promptly through:

AI / data protection route: [insert route] **Information security route:** [insert route] **Safeguarding route:** [insert DSL procedure]

Where a child may be at risk, safeguarding procedures take priority.

Nudge should preserve sufficient information to understand the incident while avoiding unnecessary replication or circulation of harmful synthetic content.

Serious incidents must trigger consideration of:

- safeguarding escalation;

- personal-data-breach procedures;
- supplier notification;
- legal advice;
- law-enforcement referral where appropriate;
- affected-person support;
- removal or containment;
- credential reset;
- DPIA reassessment; and
- policy or system change.

PART EIGHT — HUMAN ACCOUNTABILITY

21. AI does not inherit authority

A synthetic avatar may communicate an approved decision.

It may not **become the decision-maker**.

For every synthetic output there must be a human or accountable Nudge function capable of answering:

- Who commissioned this?
- Who approved the substance?
- Whose identity is represented?
- What permission existed?
- What technology produced it?
- When was it generated?
- Who can withdraw or correct it?
- Where is the current authoritative information?

Human oversight must be exercisable in practice, not merely stated in policy. This mirrors the wider Diamond architecture in which accountability is routed to named humans and cannot be transferred to AI.

22. The avatar is not the source of truth

Synthetic induction or training material may become outdated.

Every avatar-produced policy, process or training video must therefore signpost the current authoritative source.

For example:

“This video summarises the policy. The current policy in the Nudge Practitioner Portal is the authoritative version.”

Where a video conflicts with a current policy or safeguarding procedure, the current approved policy or procedure prevails.

Outdated synthetic media must be corrected, withdrawn or clearly superseded.

PART NINE — AI LITERACY

23. Training

Anyone authorised to create, approve or publish synthetic media for Nudge must receive proportionate training covering:

- what synthetic media is;
- the limits of AI-generated representations;
- consent and withdrawal;
- disclosure;
- data protection;
- safeguarding;
- impersonation risks;
- script approval;
- hallucination and factual accuracy;
- platform security;
- prohibited uses; and
- incident reporting.

Article 4 of the EU AI Act requires providers and deployers within its scope to take measures supporting AI literacy among staff and others operating AI systems on their behalf. The European Commission recommends that such measures reflect the system, its risks and the context in which it is used.

Nudge treats appropriate AI literacy as good organisational practice regardless of whether a particular use is legally within EU jurisdiction.

PART TEN — ACCOUNTABILITY AND REVIEW

24. Responsibilities

Executive Leadership

Responsible for:

- approving organisational use of synthetic identity;
- ensuring appropriate resources and oversight;
- ensuring technology does not bypass safeguarding or employment responsibilities.

AI Governance Lead

Responsible for:

- maintaining the approved-use register;
- risk classification;
- Synthetic Identity Mandates;
- supplier governance alongside Data Protection/IT;
- monitoring emerging AI regulation;
- periodic policy review.

Data Protection Officer / Data Protection Lead

Responsible for:

- lawful-basis assessment;
- DPIAs;
- privacy information;
- processor arrangements;
- biometric-data assessment;
- international transfers;

- data-subject rights.

Designated Safeguarding Lead

Responsible for:

- child-facing risk assessment;
- safeguarding incidents involving synthetic media;
- ensuring AI-generated material does not replace human safeguarding judgement.

People Team

Responsible for:

- ensuring staff permission is freely sought;
- providing alternatives;
- recording withdrawal;
- integrating synthetic-identity review into leaver processes;
- preventing detriment following refusal.

Content or Policy Owner

Responsible for:

- factual accuracy;
- currency;
- human editorial approval;
- superseding outdated material.

Person Represented

Has the right to:

- understand what is proposed;
- set boundaries;
- refuse;
- ask questions;
- receive a copy of their mandate;
- request review;
- withdraw future permission.

25. Audit record

Nudge will maintain an auditable record of synthetic identity deployments including, where appropriate:

- represented individual;
- approval date;
- platform;
- permitted scope;
- relevant DPIA;
- relevant privacy information;
- approval/consent events;
- material scope changes;
- published outputs;
- withdrawal;
- incidents;
- review dates; and
- deletion or deactivation actions.

The purpose of the record is accountability, not surveillance.

Where something is removed for lawful privacy reasons, Nudge may retain a minimal record that removal occurred and on what basis, without retaining the removed content itself where this is not necessary.

This follows the Diamond approach of preserving provenance and accountability without treating raw personal data as something that must be kept indefinitely.

Schedule A — Synthetic Identity Authorisation Record

Person represented: [Name]

Role: [Role]

Platform/provider: [Provider]

Representation being created:

- [ ] Video avatar
- [ ] Synthetic voice
- [ ] AI dubbing
- [ ] Still image
- [ ] Other: _____

Purpose

This synthetic identity may be used for:

[Specific purposes]

Audience

- [ ] Internal core team
- [ ] Nudge practitioners
- [ ] Contractors
- [ ] Commissioners/professional partners
- [ ] Parents/carers
- [ ] Children/young people — safeguarding approval required
- [ ] General public
- [ ] Other: _____

Authorised subject matter

[Describe]

Explicitly prohibited subject matter

[Describe]

Default prohibited areas include safeguarding casework, individual children, disciplinary matters, personal endorsements and statements presented as the individual's private opinion unless separately approved.

Script authorisation

- [ ] Every script must be approved by me
- [ ] I approve the following bounded script mandate:

[Describe]

Translation / dubbing

Approved languages:

[Languages / none]

AI disclosure

I understand that Nudge will clearly identify resulting material as AI-generated or AI-manipulated and will not deliberately present it as a conventional recording of me.

Refusal and withdrawal

I understand that:

- I may refuse this use without inappropriate detriment;
- I can change the boundaries of future use;
- I may withdraw authorisation for future synthetic generation;
- withdrawal cannot guarantee deletion of copies already legitimately distributed outside Nudge's control;
- Nudge will nevertheless stop new generation and take reasonable steps regarding active content.

Post-employment use

- ☐ No use after I leave Nudge. **Default**
- ☐ Existing approved recordings may remain for _____ after departure but no new content may be generated.
- ☐ Separate post-employment agreement attached.

Expiry

This mandate expires on:

[Date — maximum 12 months from approval unless renewed]

Data protection

UK GDPR lawful basis determined separately by: [Name / role]

DPIA reference: [Reference]

Privacy notice supplied: [Yes / date]

This authorisation does not, by itself, constitute Nudge's UK GDPR lawful basis.

Confirmation

I have had the opportunity to ask questions and understand the proposed use.

Person represented: _____

Date: _____

Nudge approving role: _____

Date: _____

Schedule B — Pre-publication check

Before publishing synthetic media representing a real person, the publisher must be able to answer **yes** to all applicable questions:

- [ ] Is this an approved Nudge use?
- [ ] Is there a current Synthetic Identity Mandate?
- [ ] Is this audience authorised?
- [ ] Is this subject matter authorised?
- [ ] Has the script received the required human approval?
- [ ] Is the content factually accurate?
- [ ] Is the AI nature clearly disclosed?
- [ ] Is available provenance metadata being preserved?
- [ ] Does the content avoid personal information about individual children?
- [ ] Has safeguarding approval been obtained for applicable child-facing material?
- [ ] Is the underlying policy/information still current?
- [ ] Is the authoritative source signposted?
- [ ] Can the content be stopped or withdrawn?
- [ ] Has the represented person not withdrawn permission?
- [ ] Is there a named human accountable for publication?

If the answer to a required question is **no**, the content must not be published.

26. Related policies and documents

This policy should be read alongside:

- Nudge Safeguarding and Child Protection Policy;
- Online Safety Policy;
- Data Protection Policy;
- Privacy Notices;
- Information Security Policy;
- Acceptable Use Policy;
- Staff Code of Conduct;
- AI / Diamond Standard Policy;
- Records Retention Schedule;
- Practitioner Guide;
- relevant supplier Data Processing Agreements and DPIAs.

27. Regulatory context

This policy has been designed with regard to:

- UK GDPR;
- Data Protection Act 2018;
- Data (Use and Access) Act 2025;
- Keeping Children Safe in Education 2026;
- Department for Education guidance on generative AI and data protection;
- the EU AI Act where applicable, particularly transparency and AI-literacy requirements; and
- the Nudge/Verse-ality Diamond Standard and Consent as Protocol framework.

KCSIE 2026 is statutory guidance for relevant schools and colleges in England from **1 September 2026**.

Article 50 AI Act transparency obligations have applied from **2 August 2026**.

Legal and regulatory requirements evolve. This policy must therefore be reviewed when material guidance, legislation, supplier technology or Nudge's use of synthetic identity changes.

Policy commitment

Nudge Education will use artificial intelligence where it allows us to work more effectively in service of children, young people and the people who support them.

But efficiency does not give an organisation ownership of a human identity.

We will not conceal synthetic representation.

We will not manufacture consent.

We will not make people say things they did not authorise.

We will not create synthetic children.

We will not use simulation to replace safeguarding relationships.

And we will not transfer human accountability to a machine.

AI may extend our capacity. It does not reduce our responsibility.