Noeva.

CONFIDENTIAL · INTERNAL

VOL. 01 · SAFEGUARDING 2026

The duty *of care.*

Noeva's safeguarding policy. The standard every Spark on this companion is protected by — and the standard every parent has the right to expect of us. Reviewed by the Child Safeguarding & Compliance Lead, the Data Protection & Security Architect, and the Founder.

DOCUMENT	EDITION	STATUS	OWNER
Safeguarding Policy	Vol. 01 · 2026	Adopted · Pre-launch	Safeguarding & Compliance Lead

CONTENTS · TEN SECTIONS

Inside this policy.

A working document. Adopted before any Spark is enrolled, reviewed every twelve months thereafter, and republished whenever the platform, the programme, or the regulatory ground shifts beneath it.

 © Noeva 2026 · Internal

 © Noeva 2026 · Internal

SECTION 01 · STANDING

Purpose, scope *and standing.*

A safeguarding policy is not a marketing document. It is the standard a child can hold us to — and a parent can hold us to on their child's behalf.

> Noeva is an AI learning companion built for children aged *6 to 18*. This policy sets out how we protect every Spark on the companion, the duty of care that runs through every product decision, and the lines we will not cross. It is adopted before any child is enrolled.

Purpose

The purpose of this policy is to make Noeva's duty of care to children operational. It translates the principle — that the best interests of the child are the first consideration in every product decision — into named accountabilities, design defaults, content standards, AI guardrails, data rules, and escalation procedures. It is the policy a regulator, a journalist, a parent, or a future investor can ask to see, and the policy our own people are held to internally.

Scope

This policy applies to every person who builds, runs, supports, sells, or speaks for Noeva, and to every system, model, vendor, contractor and partner involved in delivering the companion to a child. It covers:

- → All interactions between Noeva and a Spark, whether on the web platform, the parent portal, the daily Moves, the three Anchors, or any email, message or notification sent in Noeva's name.
- → All personal data processed about a Spark, a parent, or any other person, from initial enquiry through to deletion.

© Noeva 2026 · Internal

- → All third parties whose work touches a Spark's experience — including Lovable as platform host, Supabase as data backend, model providers, payment processors, email providers, and analytics vendors.
- → All public-facing material — marketing copy, social posts, press, podcast appearances — that depicts, addresses, or could be encountered by a child.

Standing

This is a board-level policy. It is owned by the Child Safeguarding & Compliance Lead, signed off by the Founder, and reviewed in conjunction with the Data Protection & Security Architect. It sits above all other internal policies. Where any other policy, brief, brand guideline, commercial decision, or vendor contract conflicts with it, this policy takes precedence and the conflict is escalated within seventy-two hours to the Founder for resolution.

A NOTE ON LANGUAGE

The vocabulary of safeguarding *is plain.*

Noeva's external voice is editorial and unhurried. Safeguarding language is not. In this document, and in any operational use, we describe risks in the plain words used by regulators, social workers and the police: grooming, abuse, exploitation, self-harm, suicidal ideation. We do not euphemise harm. We do not soften incident descriptions for the sake of brand tone. The brand voice resumes where the harm ends.

© Noeva 2026 · Internal

SECTION 02 · GROUND

The regulatory *ground.*

The five frameworks Noeva operates inside, and how each one applies to a direct-to-consumer AI learning companion serving children in the United Kingdom.

Noeva is a UK-based service. Its primary user base is UK children and their parents. Five frameworks govern the safeguarding posture below. Two are statutory and directly binding; three set the standard of practice we voluntarily meet.

UK GDPR & Data Protection Act 2018	The statutory floor for all processing of personal data. Children are recognised as warranting specific protection. We process the minimum data necessary, on the lawful bases recorded in our DPIA, and we honour all data subject rights — including the right to erasure of a Spark's data on parental request without undue delay.
ICO Age Appropriate Design Code	The "Children's Code." The statutory code of practice, issued by the Information Commissioner, that applies to any online service likely to be accessed by under-18s in the UK. Fifteen standards covering best interests of the child, data protection impact assessments, age-appropriate application, transparency, detrimental use, high-privacy defaults, data minimisation, geolocation, profiling, nudge techniques, parental controls, and online tools to exercise rights. Noeva treats the Code as binding and conforms to all fifteen standards. The Code is currently under review by the ICO following the Data (Use and Access) Act 2025; we will update this policy when revised guidance is published.

© Noeva 2026 · Internal

Online Safety Act 2023	Statutory regulation administered by Ofcom. The Protection of Children Codes of Practice came into force on 25 July 2025 and impose specific duties on regulated user-to-user and search services likely to be accessed by children. Noeva is structured as a companion: a Spark does not encounter content posted by other users, and there are no public-facing chat, forum, comment, or sharing functions. Our current legal assessment is that Noeva is not an in-scope user-to-user service. We adopt the Codes' relevant principles voluntarily — particularly on governance, content moderation of AI outputs, complaints handling, and the protection of children from content harmful to them — and document our reasoning in the children's access assessment held with this policy.
Keeping Children Safe in Education 2025	Statutory safeguarding guidance for schools and colleges in England. Noeva is not a school and is not bound by KCSIE. We adopt its principles by reference: the four categories of online risk (content, contact, conduct, commerce), the 2025 inclusion of misinformation, disinformation and conspiracy theories under "content," the safer-recruitment standards we apply to any staff hire, and the designated-lead structure for safeguarding decisions. Where Noeva enters any future partnership with a school, KCSIE compliance becomes contractually binding.
UN Convention on the Rights of the Child	The ethical ground beneath everything above. Article 3 — the best interests of the child as a primary consideration. Article 12 — the right of children to be heard. Article 16 — the right to privacy. Article 17 — access to information that promotes their well-being. Article 19 — protection from all forms of violence. The Convention shapes how we resolve any case where commercial interest and child interest are in tension. The child wins.

© Noeva 2026 · Internal

SECTION 03 · WHAT WE ARE

What Noeva is *and isn't.*

A safeguarding policy that doesn't describe the actual product can't actually protect anyone. Here is what Noeva is, what it isn't, and why the scope matters.

> A Spark using Noeva is doing four daily Moves and three daily Anchors, generated by an AI tuned to their age, phase and pace. They are not posting, not sharing, not chatting with other children, and not meeting strangers. The companion is the surface, and the parent holds the keys.

© Noeva 2026 · Internal

What Noeva is

- An AI learning companion delivering daily, age-calibrated tasks across the eight Pillars.
- A direct-to-consumer subscription service. The parent is the contracting party; the Spark is the beneficiary.
- A Parent Portal that gives the parent visibility of what the Spark is doing without requiring them to sit in every session.
- Three Anchors — Reflect (a private journal), Tune (a curated world-news feed), Track (a markets feed) — all of which the parent can adjust or disable.
- A model-mediated experience: every output a Spark sees is generated by a language model under defined guardrails described in Section 07.

What Noeva is not

- A social network. There is no peer-to-peer messaging, no friend mechanics, no follower mechanics, no public profiles, no comments.
- A user-generated-content platform. Sparks do not publish to anyone outside their own family.
- A school. We do not replace a school, do not award qualifications, and do not assume safeguarding responsibility for a Spark's wider life.
- A general-purpose chatbot. The model serves the programme. It does not roleplay, it does not pretend to be a friend, and it does not present itself as human.
- An advertising business. Noeva does not run ads to children, does not profile children for ad targeting, and does not sell data.

WHY THIS MATTERS FOR SCOPE

The companion architecture *reduces — but does not eliminate —* categories of online risk.

The absence of peer-to-peer features means the dominant pathways to grooming, bullying, and child-on-child harm that exist on social platforms are not present on Noeva. The dominant residual risks are different and are addressed in Section 06: AI-generated content harm, content-related distress in the Reflect Anchor, the Tune feed exposing a Spark to upsetting world news, third-party links a child might follow, parental-account compromise, and the data-handling risks inherent in any subscription service. Scoping the policy honestly is how we make sure each of those is actually covered, rather than buried inside generic platform language.

© Noeva 2026 · Internal

SECTION 04 · PRINCIPLES

Ten principles. *Tested daily.*

Every product decision, every model output, every brief, every vendor contract, every campaign — measured against these. If a decision fails the test, the decision changes.

01

Best interests first.

Where commercial interest and the interest of a Spark conflict, the Spark wins. This is recorded in writing and named the decision-maker.

02

Parental authority is the keystone.

The parent contracts, consents, configures, and holds the right to read, pause, export, and delete. We do not work around the parent.

03

High privacy by default.

The default setting on every privacy and visibility control is the most protective. Loosening any setting requires an explicit parental action.

04

Minimum data, shortest retention.

We collect only what is necessary to deliver the companion. We retain it only as long as it is necessary. We delete on request and on schedule.

© Noeva 2026 · Internal

05

No profiling that harms.

We do not profile Sparks for advertising. We do not use children's data to build engagement-maximising recommender behaviours.

06

No dark patterns.

No nudge techniques designed to extend session time, weaken privacy settings, encourage data sharing, or pressure a parent into a less protective configuration.

07

Honest AI.

The model is never presented as human, never pretends to a relationship, and never claims certainty it doesn't have. It teaches Sparks to use AI as a thinking partner, not a substitute for thinking.

08

Reportable, reachable, responsive.

A Spark or parent can flag a concern in two clicks. We acknowledge within twenty-four hours. We resolve or escalate within defined windows.

09

Transparent to the child.

Privacy information is rewritten for the age group reading it. A nine-year-old sees a nine-year-old's version. A sixteen-year-old sees the full version. Both are honest.

10

Reviewed, recorded, revised.

This policy is reviewed every twelve months and on every material change to the programme, platform, vendor stack, or regulatory ground. The audit trail is kept indefinitely.

© Noeva 2026 · Internal

SECTION 05 · IDENTITY

Age, identity and *parental authority.*

Who is on Noeva, how we know, and the structural commitment that the parent — not the child — holds the keys to the account.

The parent is the contracting party

Noeva is sold to parents. The parent creates the account, authenticates with their own email and a password meeting the security standard set in the Data Protection & Security Architect's controls, and accepts the terms on the Spark's behalf. The Spark does not contract with Noeva and does not pay Noeva. This is the deliberate architecture: it makes parental authority the structural default, not an optional layer.

Age range and the cutoffs

Noeva serves Sparks aged **6 to 18**. A Spark turning 18 retains access for the remainder of the parent's current subscription period, after which the relationship moves to direct adult contracting. We do not enrol Sparks under the age of 6: the programme is not appropriate for them and we do not wish to design for that age group within this companion.

Age assurance

Noeva uses a layered age-assurance approach that is proportionate to the risk profile of a learning companion under parental oversight. Self-declaration alone is insufficient — the ICO has been explicit on this point — and so we combine it with three further mechanisms:

- → **Parent-asserted Spark age** entered at enrolment, with the parent confirming under the terms of service that the asserted age is accurate.
- → **Adult-verified payment** as a proxy signal: subscription is taken via a payment method requiring the account-holder to be 18 or older.

© Noeva 2026 · Internal

→ **Programme calibration**: the model adjusts Move complexity to the asserted age. A material mismatch between asserted age and observed engagement triggers a Parent Portal prompt to confirm or correct.

Where Noeva enters higher-risk feature territory in future — for example, expanding the Reflect Anchor or introducing any feature with peer-visibility — the Child Safeguarding & Compliance Lead will assess whether highly effective age assurance methods (as defined by Ofcom in its 2025 guidance) are required and will update this section accordingly.

Parental controls and the keys

The Parent Portal gives the parent the operational keys to the account. At any time, without contacting support, the parent can:

- Read everything the Spark has written in Reflect — and choose, separately, whether the Spark knows they can.
- Adjust the Tune feed: subject filters, intensity, frequency, off-completely.
- Pause the companion for any duration.
- Change the Spark's asserted age, name, or settings.
- Export the Spark's full data record in a portable format.
- Delete the Spark's account and request erasure of all associated data.
- Update consent on any optional processing.

A NOTE ON REFLECT

The Reflect journal sits at the most sensitive point *in the companion.*

Reflect is a private daily journal. A Spark may write about anything: a difficult day, a row at home, a worry, a friendship, a feeling they don't know what to do with. Three operational rules apply. First, the parent has the right to read — they are the parent — and we make this right clearly visible to the Spark at the point of writing, in age-appropriate language. Second, if a Reflect entry contains language meeting the disclosure thresholds in Section 09 (active self-harm, suicidal ideation, abuse disclosure), the safeguarding escalation in that section is triggered. Third, Reflect content is never used to train any model.

© Noeva 2026 · Internal

SECTION 06 · THE FOUR C'S

Content, contact, conduct, *commerce.*

The four-category framework adopted from KCSIE 2025. Below we set out where Noeva's exposure actually sits in each category, and the specific controls in place for each.

Content. What a Spark might encounter.

This is Noeva's principal risk surface, because the model generates content and the Tune feed surfaces world news. Within "content," KCSIE 2025 now explicitly includes misinformation, disinformation and conspiracy theories alongside the longer-standing categories of violence, sexual content, hate, and material promoting self-harm or eating disorders.

Controls

- → Model outputs are filtered against an explicit prohibited-content list — Section 07 — with both system-prompt and post-generation moderation layers.
- → The Tune feed is curated. Stories are scored for age-appropriateness against the Spark's asserted age before being surfaced. Stories that meet but exceed the threshold are de-prioritised, not removed, on the principle that Sparks deserve to understand the world they are about to inherit.
- → The parent can adjust Tune intensity by topic and by overall sensitivity.
- → External links generated by the model are blocked from rendering as clickable. Where a citation is offered, the source name is shown, not a live hyperlink.

Contact. Who a Spark might be reached by.

Sharply scoped by the companion architecture: there is no peer-to-peer messaging, no chat with strangers, no public profile, and no friend mechanic. The contact surface is therefore narrow.

© Noeva 2026 · Internal

Controls

- → No feature on Noeva allows a Spark to communicate with anyone outside their own family account.
- → Noeva-originated email and notifications are sent to the parent, not to the Spark, by default. Direct Spark notifications are opt-in via the Parent Portal.
- → Customer support is provided only to the verified parent account-holder. We will not discuss a Spark's account with anyone else, including the Spark themselves except through Parent-authorised channels.

Conduct. How a Spark might behave or be encouraged to behave.

This includes the risk that an AI companion subtly shapes a child's behaviour in ways that are not in their interest — for example, encouraging dependency, validating unhealthy patterns, or modelling unhealthy communication.

Controls

- → The model is explicitly tuned away from sycophancy. It does not flatter, does not foster emotional dependency, and does not present itself as a friend or substitute for human relationship.
- → Move design is built around action — making, doing, deciding — not screen time. The companion is not optimised for session length.
- → Streak mechanics, if used, are designed around weekly cadence rather than daily punishment. A missed day does not break the relationship.
- → The model does not give medical, legal, financial, or specialist mental-health advice. It directs the Spark — and the parent — to appropriate human help where the topic warrants it.

Commerce. How a Spark might be marketed to, profiled, or pressured.

Noeva does not run ads to children, does not profile children for ad targeting, and does not sell data. Commercial messaging is directed to the parent through the verified parent account.

Controls

- → No third-party advertising in any Spark-facing surface.
- → No in-app upsells inside the Spark's daily Moves or Anchors.

© Noeva 2026 · Internal

- → No tracking pixels, no advertising SDKs, and no analytics that profile a Spark for any purpose beyond product quality. Analytics on Spark interactions are aggregated and pseudonymous.
- → No partnerships in which a Spark's data is shared with a third party for marketing.
- → Subscription, billing, retention and renewal communications are sent to the parent, not the Spark.

 © Noeva 2026 · Internal

SECTION 07 · MODEL

AI safety *for Sparks.*

Noeva is built around a language model. This section sets out what the model is permitted to do, what it is not permitted to do, and how we verify both.

> Children should use AI to think *better,* not to think for them. That is the editorial position of the Intelligence Pillar and the operational position of the safety layer that sits underneath it.

Prohibited outputs

The model is configured to refuse to produce, to a Spark, any of the following — irrespective of how the request is framed:

- Sexual content of any kind, or content that sexualises, grooms, or otherwise harms a child.
- Content depicting, instructing in, or normalising self-harm, suicide, eating disorders, or other behaviours harmful to physical or mental health.
- Content promoting violence against any person or group; hateful or discriminatory content against any protected characteristic.
- Information that would meaningfully aid the construction of weapons (chemical, biological, nuclear, radiological, or conventional improvised) or the commission of serious crime.
- Personalised medical, legal, financial, or specialist mental-health advice. The model directs to qualified human help.
- Content that misrepresents the model as human, presents the model as the Spark's friend, or fosters dependency on the model in place of human relationships.
- Misinformation, disinformation, or conspiracy content presented as fact.

© Noeva 2026 · Internal

Designed behaviours

The model is configured, in addition to the prohibitions above, to behave in specific ways aligned with the programme:

- → To teach, in age-appropriate terms, that AI is a tool — useful, fallible, and not a substitute for the Spark's own thinking.
- → To show working: when the model offers an answer, it explains how it got there in language the Spark can follow.
- → To redirect to a parent or trusted adult when a Spark raises a topic outside the model's competence: a friend in trouble, a worry about home, a question that needs a real person.
- → To recognise emotional distress and to respond with care, without simulating clinical assessment, without minimising the Spark's experience, and without providing means-related information of any kind.
- → To respect the Spark's age: a Move generated for an eight-year-old reads differently from a Move generated for a sixteen-year-old, both in vocabulary and in conceptual depth.

How we verify

The safety layer is not a one-time configuration. It is a living set of controls verified continuously.

01 **Red-team testing before any model change.**

No update to the system prompt, the model provider, the model version, or the moderation layer reaches a Spark before red-team testing has been run against the prohibited-output list and the results reviewed by the Child Safeguarding & Compliance Lead.

02 **Output sampling.**

A randomised sample of generated Move and Anchor outputs is sampled and human-reviewed weekly against the safety standard. Failure rates are tracked. A single Severity-1 finding pauses the affected feature.

03 **Standing report mechanism.**

A parent or Spark can flag any single Move, Anchor entry, or response in two clicks. Flagged content is reviewed within twenty-four hours and the safety standard updated if the flag reveals a gap.

© Noeva 2026 · Internal

04 **Quarterly safety review.**

The Child Safeguarding & Compliance Lead reviews the safety layer end-to-end every quarter and reports findings to the Founder. Failure patterns drive prompt changes, moderation tuning, or — in extremis — a change of model provider.

Training data and Spark data

Spark data is not used to train any model — neither the model Noeva uses, nor any third-party model accessed via Noeva. This is set out contractually with each model provider and is verified at vendor onboarding. Where a model provider's standard terms permit training on customer data, we explicitly opt out in writing as a precondition of use.

© Noeva 2026 · Internal

SECTION 08 · DATA

Data, privacy and *the DPIA.*

The data discipline that sits under the safeguarding policy. Owned operationally by the Data Protection & Security Architect; reviewed jointly with the Child Safeguarding & Compliance Lead.

Lawful basis

Noeva relies on two primary lawful bases for processing personal data. Performance of contract is the basis for processing parent data necessary to deliver the subscription. The Spark's data is processed on the basis of parental consent, recorded in the Parent Portal and revocable at any time. Special category data — including any health, ethnicity, religious belief, or sexual-orientation information disclosed by a Spark in Reflect or elsewhere — is processed only where necessary and on explicit consent, with additional protections set out below.

Data minimisation

We collect: parent identifiers (name, email, payment), Spark first name and age, the programme's record of which Moves have been completed, Reflect entries, configuration preferences, and standard service-operation logs. We do not collect: a Spark's last name, address, school, photographs, voice recordings, biometric data, behavioural-profiling tags for advertising, or any data we have not been able to articulate a clear product reason for collecting.

High-privacy defaults

Every default setting is the most protective option. Geolocation is off and not requested. Discoverability mechanics do not exist on the companion, so the question does not arise. Profile fields beyond the necessary first name and age are off by default. The Spark's daily activity is visible to the parent and to no-one else.

© Noeva 2026 · Internal

Retention

Retention periods are defined for each data category in the Data Retention Schedule held with this policy. Reflect entries are retained for the duration of active subscription unless the parent shortens this — for example, to a rolling 90-day window. On account closure, all Spark data is deleted within 30 days, save for the minimum financial and audit records required by law. The parent receives written confirmation of deletion.

Third parties and vendor risk

Noeva is built on Lovable and uses Supabase as data backend. Both are bound by data-processing agreements that meet UK GDPR Article 28 standards. Specific controls in place include:

- → Row Level Security enabled on every Supabase table containing Spark or parent data, configured and tested by the Data Protection & Security Architect.
- → API keys rotated on a published schedule; any keys created before November 2025 audited and replaced.
- → Vendor security postures reviewed at onboarding and annually thereafter, with the right to terminate written into every agreement.
- → A vendor register maintained and reviewed quarterly. Any sub-processor not listed in the register may not process Spark data.

The Data Protection Impact Assessment

A full DPIA, conducted under Article 35 UK GDPR and ICO Children's Code Standard 2, accompanies this policy. The DPIA documents data flows, risks, mitigations, residual risk, and the decision-maker for each residual risk. It is reviewed alongside this policy on the annual cycle and updated on any material change to data processing.

© Noeva 2026 · Internal

A NOTE ON THE SECURITY FLOOR

Safeguarding is not separable from *information security.*

A child's data being exposed in a breach is a safeguarding event. Accordingly, the technical security baseline — RLS on all Supabase tables, 2FA on every administrative account, secure-headers configuration tested against Mozilla Observatory, DMARC on the sending domain, and the incident response procedure in Section 09 — is part of this policy by reference, not a separate document. The Data Protection & Security Architect owns the technical controls; the Child Safeguarding & Compliance Lead owns the assurance that those controls are present and effective.

© Noeva 2026 · Internal

SECTION 09 · RESPONSE

Reporting, escalation *and response.*

What a parent does, what a Spark does, what Noeva does, what happens next, and who gets called when the threshold for external referral is met.

How to raise a concern

A parent or Spark can raise a safeguarding concern through three routes, each of which reaches the Child Safeguarding & Compliance Lead directly:

- → **In-product flag.** Two taps on any Move or Anchor surface a "Report this" action. Reports are routed automatically and tagged by severity.
- → **Direct email.** A dedicated safeguarding inbox — *safeguarding@noeva.live* — published in the Help Centre, the parent welcome email, and the footer of the Parent Portal. Monitored every working day.
- → **Standing escalation.** Any concern raised through customer support that meets the safeguarding threshold is escalated to the Lead within four working hours of receipt.

Severity and timelines

Severity 1 IMMEDIATE RISK	A Spark is in immediate danger, has disclosed active suicidal ideation with intent or plan, has disclosed current abuse, or has been the subject of a credible threat. **Action:** acknowledged within 1 hour; Child Safeguarding & Compliance Lead engaged immediately; external referral to police or local authority children's services made within 4 hours where appropriate; parent informed unless doing so would increase risk to the Spark.

© Noeva 2026 · Internal

Severity 2 SIGNIFICANT CONCERN	Disclosure of historic abuse, indication of an ongoing safeguarding concern at home or school, severe distress, eating disorder language, or a meaningful failure of the AI safety layer affecting a Spark. **Action:** acknowledged within 4 hours; reviewed within 24 hours; parent informed; appropriate external referral considered and made where indicated; reviewed by the Founder within 48 hours.
Severity 3 MATERIAL ISSUE	A model output that crossed a safety boundary without immediate harm; a Tune story surfaced inappropriately; a Reflect entry of concern not meeting Severity 1 or 2 thresholds; a complaint about a single feature. **Action:** acknowledged within 24 hours; resolved or escalated within 7 working days; logged and counted against the relevant safety metric.
Severity 4 SERVICE ISSUE	General complaint, feature feedback, billing question with no safeguarding dimension. **Action:** handled by customer support within published service levels.

External referral

Where a disclosure meets the threshold for external referral, the Child Safeguarding & Compliance Lead, having taken any urgent action required to protect the Spark, will refer to the appropriate authority. In England that is the local authority's children's services in the area where the Spark lives, or, where there is an immediate risk of harm, the police on 999. The NSPCC helpline (0808 800 5000) is used as a consultative resource where the Lead requires advice on whether a referral threshold is met. The Lead records every referral decision, the reasoning, the action taken, and the outcome.

Data breaches

A personal data breach involving a Spark is, by definition, a safeguarding concern. Breaches are managed under the Data Protection & Security Architect's Incident Response Plan, which requires notification to the ICO within 72 hours of becoming aware where the breach is likely to result in a risk to the rights and freedoms of individuals, and direct notification to affected parents without undue delay where the risk is high. The Child Safeguarding & Compliance Lead is informed of every breach involving Spark data, regardless of severity.

© Noeva 2026 · Internal

Confidentiality

Where a Spark or parent has disclosed something sensitive, Noeva treats the disclosure with appropriate confidence. Information is shared only on a need-to-know basis: to deliver safeguarding response, to comply with legal obligations, or to protect the Spark or another person from harm. The Spark is, where developmentally appropriate, told what we are doing and why.

© Noeva 2026 · Internal

SECTION 10 · GOVERNANCE

Governance, training *and review.*

Named accountabilities, the training every Noeva person receives, the annual review cycle, and the audit trail. The structural commitments that make the rest of this policy real.

Named accountabilities

Child Safeguarding & Compliance Lead POLICY OWNER	Owns this policy end-to-end. Accountable for ICO Children's Code conformance, the DPIA, AI safety verification, incident response, and the audit trail. The named escalation point for every Severity 1 and 2 event. Reports to the Founder.
Data Protection & Security Architect TECHNICAL CONTROLS	Owns the technical security and data-protection posture. Accountable for Supabase RLS, vendor risk, breach response, and the Article 28 contractual chain. Jointly accountable with the Safeguarding Lead for the DPIA.
Founder FINAL SIGN-OFF	Final sign-off on this policy and every annual revision. Receives every Severity 1 and 2 escalation. Decides where commercial interest and child interest are in tension. Accountable to the parents who entrust Noeva with their children.
Strategic & Commercial Director COMMERCIAL ALIGNMENT	Ensures commercial decisions, partnership choices, and growth mechanics are consistent with this policy. Any campaign mechanic touching Sparks is reviewed against the policy by the Safeguarding Lead before deployment.

© Noeva 2026 · Internal

All Noeva personnel UNIVERSAL	Every person engaged by Noeva, paid or unpaid, reads this policy on day one, attests in writing that they have read and understood it, and refreshes that attestation annually. Every person knows how to escalate a concern.

Training

The Child Safeguarding & Compliance Lead holds a recognised child-protection qualification — at minimum the NSPCC Child Protection in Education advanced training or equivalent — and refreshes that training every two years. All other Noeva personnel complete a Noeva-specific safeguarding induction within their first week, covering this policy, how to recognise disclosure language in a Reflect entry or support email, the escalation routes in Section 09, and the prohibited-output list in Section 07. Refreshers are run annually.

Safer recruitment

Anyone whose role gives them access to Spark data or could put them in direct contact with a Spark is subject to safer-recruitment checks proportionate to the role: identity verification, right-to-work, employment history, and — for any role with substantive Spark-facing responsibilities — an enhanced DBS check. Until Noeva has employees of this kind, the standard applies prospectively and is written into every offer letter.

Review cycle

TRIGGER	CADENCE	OWNER & OUTCOME
Scheduled review	Every 12 months	Lead drafts; Architect reviews; Founder signs off; new version published internally and externally where appropriate.
Regulatory change	Within 60 days	Triggered by an ICO Children's Code revision, an Ofcom Codes update, new statutory guidance, or relevant case law.
Material product change	Before release	Any feature affecting a Spark's experience is reviewed against this policy before it ships. The Lead can hold a release.
Severity 1 event	Within 14 days	

© Noeva 2026 · Internal

		Post-incident review by the Lead, Architect and Founder. Lessons learnt translated into policy or control changes.
Vendor change	Within 30 days	Architect updates the vendor register and the DPIA; Lead confirms safeguarding implications are addressed.

Audit trail

Every revision of this policy, every safeguarding incident, every escalation decision, every external referral, every DPIA review, and every vendor change is recorded and retained. The audit trail is held by the Child Safeguarding & Compliance Lead and is available, on request, to the Founder, to a regulator with proper authority, and — within applicable confidentiality limits — to a parent who has raised a concern about the handling of their own Spark's case.

Public availability

A parent-facing version of this policy is published at noeva.live and linked from the footer of every page. The full internal version is available on request to a parent who reasonably asks for it. Future independent review of this policy by a recognised children's-rights or child-safety organisation is welcomed and will be sought before Noeva reaches a scale that warrants it.

Kate Bell

FOUNDER · FINAL SIGN-OFF

Child Safeguarding & Compliance Lead

POLICY OWNER

Data Protection & Security Architect

JOINT REVIEW · TECHNICAL CONTROLS

Steve Pratt

STRATEGIC & COMMERCIAL DIRECTOR

© Noeva 2026 · Internal