# Technical Note

## Virtual Identity (VI): A Technical Implementation Model for Proportional Identifiability — Achieving Legal Certainty for AI Innovation and Consistent Privacy Protection Without Amending Article 4(1) GDPR

## Context and Purpose of This Note

In November 2025, the European Commission presented the Digital Omnibus package, including a proposal to clarify how pseudonymised data should be treated when assessing whether information is "personal data" under Article 4(1) GDPR, drawing on the Court of Justice's reasoning in EDPS v Single Resolution Board (C-413/23 P) regarding context-dependent identifiability.

In a Council compromise text dated 20 February 2026, EU Member States are set to remove the Commission's proposed new wording on the definition of personal data and also to delete the enabling approach that would have allowed the Commission to specify criteria via secondary implementing measures for when pseudonymised data might fall outside GDPR scope.

This development follows the EDPB–EDPS Joint Opinion 2/2026, which supported simplification goals in principle but raised significant concerns about changing the definition of personal data through a negative formulation and using implementing acts to determine what is no longer personal data after pseudonymisation — noting that such choices directly affect the GDPR's scope.

**Practical implication:** The legislative direction currently indicates low appetite to reopen the core definition of personal data. The policy focus shifts toward how pseudonymisation and identifiability are implemented and evidenced in practice — including through EDPB guidance.

**Purpose of this note:** In that context, this technical note does not advocate a definitional amendment. It offers a cryptographically verifiable implementation model for proportional identifiability — intended to support legal certainty for innovation while preserving consistent protection across processing chains — without changing Article 4(1) GDPR.

# Understanding the Policy Objective Behind the Debate

Before presenting the technical model, it is worth identifying precisely what each party sought to achieve — because the VI architecture addresses all objectives simultaneously without requiring the instrument that created institutional disagreement.

| Party | Core Objective | Position on Definitional Change |
|---|---|---|
| European Commission | Legal certainty for pseudonymised data use — reduce AI compliance burden | Proposed definitional clarification drawing on SRB reasoning |
| Member States — Council | Maintain GDPR's protective scope — prevent fragmentation of rights | Removing proposed amendment — February 2026 compromise text |
| EDPB / EDPS | Consistent protection across processing chains | Joint Opinion 2/2026 — opposed negative formulation and implementing acts |
| AI and technology industry | Access to pseudonymised data without full GDPR compliance burden | Welcomed Commission proposal |
| Data subjects | Protection that follows data regardless of who holds it | Definitional change risks rights disappearing at recipient level |

**The observation this table reveals:**

The Commission's underlying objective — legal certainty for entities that genuinely cannot re-identify pseudonymised data — is shared across institutions. The disagreement is about the instrument, not the objective. A cryptographically verifiable implementation model achieves the objective without the instrument that was rejected.

# The Technical Problem the Debate Has Not Yet Resolved

The SRB ruling established a sound legal principle — that identifiability should be assessed in context, taking into account the means reasonably likely to be used by the specific entity holding the data. The Commission's proposal attempted to generalise this into a workable legal standard. The institutional response has identified why a definitional approach creates problems — but has not yet provided a technical answer to the underlying identifiability question.

The gap is this:

"*Recital 26 already requires account to be taken of means reasonably likely to be used to identify the person. No technical mechanism currently exists to make that assessment verifiable, stable, and consistent across processing chains — rather than entity-relative, technology-dependent, and contested.*"

The VI architecture provides that technical mechanism.

# What the Current Data Classification Landscape Looks Like

| | Personal Data | Pseudonymised Data | Anonymous Data |
|---|---|---|---|
| GDPR definition | Art. 4(1) — directly or indirectly identifiable | Art. 4(5) — identifiable with additional information | Recital 26 — not identifiable by reasonably likely means |
| GDPR scope | Fully in scope | In scope — some relaxed provisions | Outside scope |
| Re-identification | Direct | Possible — key or AI inference | Not feasible |
| Lawful interception | Straightforward | Possible | Not possible |
| AI training use | Full compliance required | Compliance required — disputed extent | Unrestricted |
| Verifiability of classification | Straightforward | Judgment-based — contested | Technically difficult to prove |
| Stability as AI advances | Stable | Weakening — inference makes re-identification easier | Stable |
| Consistent across entities | Yes | No — entity-relative assessment | Yes |

**The gap this table reveals:**

No existing category simultaneously satisfies:

- Genuine non-re-identifiability for unauthorised parties — including AI inference
- Lawful re-identifiability for judicial process
- Cryptographic verifiability of the non-re-identifiability claim
- Consistency across all entities in the processing chain
- Stability as AI capabilities advance

Pseudonymised data fails the first, third, fourth, and fifth. Anonymous data fails the second. The Commission's proposed amendment attempted to address the gap through definitional change. The VI architecture addresses it through a verifiable implementation model — without any definitional change.

# What VI Introduces — The Implementation Model

The Virtual Identity architecture proposes a **cryptographically verifiable implementation model** that gives technical precision to the Recital 26 proportionality standard — making it verifiable, consistent, and stable rather than judgment-based, entity-relative, and technology-dependent.

**The core principle:**

*"Recital 26 requires account to be taken of means reasonably likely to be used to identify the person. VI makes this assessment a cryptographic predicate rather than a probabilistic judgment. If defeating re-identification would require compromising the hardware isolation boundary under*

*defined security assumptions — those means are not reasonably likely by any technically defensible standard. The assessment becomes verifiable, not contested."*

# How VI Works — The Split Knowledge Architecture

VI operates on two technically and legally distinct layers:

**Layer 1 — The External Layer** *(What every party outside the platform observes — including AI training pipelines, data recipients, and third-party processors)*

- A cryptographic alias generated independently — not derived from any real identifier
- Session-scoped — expires when the session ends by architectural construction
- Unlinkable across sessions — no stable pattern exists to correlate across contexts or time
- Re-identification by any external party — including AI inference and data brokers — is cryptographically infeasible under defined security assumptions
- This layer contains no re-identifiable material — a data breach of this layer exposes nothing that allows identification

**Layer 2 — The Controlled Internal Layer** *(What exists exclusively within the platform's hardware security boundary)*

- The binding between VI alias and real identity
- Held within a TEE/HSM — hardware-isolated, never transmitted, never accessible via API
- Accessible only through a cryptographically enforced judicial process
- The platform operator cannot access it unilaterally — hardware enforcement prevents silent access within the defined architecture

**What this means for the GDPR processing chain:**

| Processing Chain Position | Current Pseudonymisation | VI Architecture |
|---|---|---|
| Original controller | Holds re-identification key — full controller liability | VI binding in TEE/HSM — no exposed key outside hardware boundary |
| Data processor | May inadvertently hold re-identification capability — accidental controller risk | Processor lacks hardware-gated keys — cannot become accidental controller |
| Third-party AI recipient | Claims cannot "reasonably re-identify" — contested, unstable standard | Cryptographically cannot re-identify under defined security assumptions — verifiable |
| Supervisory authority | Cannot verify recipient re-identification capability | Cryptographic audit trail verifiable on demand |
| Data subject | Rights may disappear when data classified as non-personal for recipient | Re-identification pathway remains judicial — rights preserved across chain |

# How This Achieves the Commission's Objective Without the Definitional Change

The Commission's objective was legal certainty for entities that genuinely cannot re-identify pseudonymised data.

VI achieves this objective more precisely and more durably:

| Commission Objective | Definitional Approach | VI Implementation Model |
|---|---|---|
| Legal certainty for recipients who cannot re-identify | Subjective — entity claims it cannot re-identify | Objective — entity cryptographically cannot re-identify under defined assumptions |
| Reduce AI compliance burden | Reduce scope of personal data for certain entities | Data received is verifiably non-re-identifiable — compliance proportionate to actual risk |
| Reflect SRB context-dependent identifiability | Generalise ruling into definitional text | Implement SRB principle as cryptographic predicate — verifiable not interpreted |
| Enable innovation without fragmenting rights | Creates entity-relative fragmentation | Same verifiable standard for all entities — no fragmentation |
| Legal stability as technology advances | Standard shifts as "reasonable means" evolve | Cryptographic assumption-based — stable standard independent of technology state |

"*AI companies and data processors receive the legal certainty they need — not because the definition of personal data has changed, but because the data they receive is genuinely, verifiably, cryptographically non-re-identifiable under defined security assumptions. The definition does not need to change. The architecture of the data changes instead.*"

# AI Innovation and Privacy — Simultaneous Achievement

The underlying tension the Commission sought to resolve:

| AI Innovation Requirement | Privacy Requirement | VI Resolution |
|---|---|---|
| Access to behavioral data for training | Purpose limitation — data used only for declared purpose | CJT cryptographically binds data to declared purpose — freely usable within scope, non-executable outside it |
| Legal certainty about data classification | Consistent protection regardless of who holds data | Same cryptographic standard applies to all entities — no entity-relative fragmentation |
| Reduced compliance burden for pseudonymised data | All pseudonymised data is personal data under current law | Verifiably non-re-identifiable data carries compliance burden proportionate to actual risk |
| Cross-border data flows for global AI development | Jurisdiction-scoped data protection | CJT encodes jurisdiction — data remains non-executable outside permitted jurisdictional scope |

| | | |
|---|---|---|
| Processing without persistent identifier exposure | Data minimisation under Article 5(1)(c) | No stable identifier exists in external layer — minimisation by architectural default |

# Controller and Processor De-Risking Under GDPR

A direct practical benefit for the compliance simplification the Commission sought:

| GDPR Role | Current Risk | VI Resolution |
|---|---|---|
| Data Controller | Responsible for re-identification key security | VI binding in TEE/HSM — no exposed key outside hardware boundary |
| Data Processor | May inadvertently hold re-identification capability — Art. 28 liability gap | Processor lacks hardware-gated keys entirely — cannot become accidental controller |
| Data Subject | Rights may disappear when data classified as non-personal for recipient | Judicial unmasking process preserves rights across entire chain |
| Supervisory Authority | Cannot verify processor re-identification capability | Cryptographic audit trail verifiable on demand against defined parameters |

This has direct practical relevance for:

- Data Processing Agreements under Article 28
- Sub-processor chain liability under Article 28(4)
- Cross-border transfer risk assessment under Chapter V

# Lawful Court-Order Unmasking — Democratic Accountability Preserved

**The governing principle:**

*"A court order can unmask one person for one purpose. The architecture preserves judicial responsibility for any scaling decision — proportionality is enforced by the architecture, not aspirationally required by law alone."*

**What a court order can do:**

- Unmask one specific VI session
- For one specific declared investigation purpose
- Under one specific legal authority and jurisdiction
- Within one specific declared time window
- Quota of one unmasking event per order — enforced by sealed monotonic counter

## The unmasking process:

| Step | What Happens | Technical Mechanism |
|---|---|---|
| 1 — Order issued | Court specifies exact session, investigation reference, legal authority, jurisdiction | Human judicial process — architectural bypass would require compromising hardware isolation boundary under defined assumptions |
| 2 — Unmasking CJT generated | Token encodes: legal authority, jurisdiction, session scope, quota of one, temporal validity, expected counter state | Cryptographically signed by judicial authority — modification invalidates token |
| 3 — TEE/HSM evaluates | Hardware evaluates CJT against cryptographic predicate including counter state | Pass or fail — no discretionary override within defined architecture |
| 4 — Binding released | Only specified session identity revealed | Hardware-enforced — exceeding scope would require compromising isolation boundary |
| 5 — Counter incremented | Monotonic counter advances — replay deterministically rejected | Enclave-sealed — reset would require defeating hardware boundary under defined assumptions |
| 6 — Audit receipt | Tamper-evident signed record generated automatically | Court-admissible, immutable, cryptographically verifiable |

**What each stakeholder receives:**

| Stakeholder | What VI Delivers |
|---|---|
| Law enforcement | Specific, targeted, lawful access — operationally effective |
| Citizens | Proportionality by architecture — one order, one session, one purpose |
| Courts | Full accountability — every unmasking cryptographically logged and auditable |
| EU institutions | Sovereignty — foreign authorities require EU judicial process subject to applicable EU and Member State law |
| Platform operators | Legal clarity and processor de-risking — hardware enforcement removes discretionary access liability |
| Regulators | Verifiable compliance — cryptographic audit trail on demand |

**On judicial scaling:**

*"The architecture does not prevent lawful access at scale. It ensures that lawful access at scale requires lawful process at scale — each order judicially authorised, each access cryptographically enforced, each event individually audited. Proportionality is an architectural property, not solely a legal aspiration."*

# Proportionality Against Disproportionate Surveillance

**The structural condition in current systems:**

*"Whoever holds the key holds the capacity to re-identify the entire dataset — simultaneously, through a single access event, without per-individual judicial process."*

**How VI addresses this condition:**

| Surveillance Capability | Technical Status Under VI | Basis |
|---|---|---|
| Stable identifier to locate and track | Not present in external layer | VI alias changes every session by construction |
| Persistent cross-session pattern to correlate | Not present in external layer | No linkage exists between sessions |
| Single access point to re-identify bulk population | Not present | No key re-identifies more than one judicially specified session |
| API surface to reach identity binding | Not present within defined architecture | TEE/HSM presents no accessible API layer |
| Foreign authority compelled bulk access | Requires EU judicial process per session — subject to applicable EU and Member State law | Sovereign by architectural construction within applicable legal framework |
| Platform operator silent unilateral access | Prevented within defined architecture | Hardware enforcement gates every access with mandatory audit |
| Replay of unmasking token | Prevented by sealed monotonic counter | Counter state enclave-sealed — reset would require defeating hardware boundary |

# Threat Model and Security Scope

All security claims are relative to a defined threat model — not absolute unconditional guarantees:

| Threat Category | VI Security Posture | Basis |
|---|---|---|
| External re-identification by data brokers | Cryptographically infeasible under defined assumptions | No stable cross-session identifier in external layer |
| AI inference re-identification | Cryptographically infeasible under defined assumptions | No persistent correlatable pattern across sessions |
| Database breach of external layer | No identity exposed within defined architecture | External layer contains no re-identifiable material |
| API-based access to identity binding | Not possible within defined architecture | TEE/HSM presents no accessible API surface |
| Replay attacks on unmasking tokens | Prevented by sealed monotonic counter | Enclave-sealed — defeating requires compromising hardware boundary |

| | | |
|---|---|---|
| Side-channel attacks on TEE/HSM | Subject to hardware vendor security assumptions | FIPS 140-2/3 Level 3 or equivalent EU certification schemes |
| Post-quantum cryptographic threat | Requires planned migration — see Technical Note | CJT structure is algorithm-agnostic |

**Security assumption statement:**

*"Security properties are asserted under standard cryptographic assumptions applicable to deployed TEE/HSM infrastructure consistent with FIPS 140-2/3 Level 3 or equivalent EU certification schemes. Defeating these properties would require compromise of the cryptographic or hardware isolation boundary under defined security assumptions — not merely computational effort at currently feasible scales."*

# Technical Note — Algorithm Agnosticism and Post-Quantum Readiness

The CJT structure and VI architecture are **cryptographic algorithm-agnostic:**

- The enforcement predicate evaluates binding, purpose, jurisdiction, quota, and temporal constraints as structured inputs — independent of any specific cryptographic primitive
- Underlying signature schemes, key derivation functions, and hash functions are modular — replaceable without architectural change
- Migration to post-quantum primitives requires algorithm substitution, not architectural redesign
- Directly relevant to Digital Euro and eIDAS 2.0 timelines extending to 2030 and beyond

*"The VI+CJT architecture will not become cryptographically obsolete as post-quantum standards mature. Institutions planning for 2030+ deployment horizons can treat the architecture as cryptographically forward-compatible."*

# Complete Comparison

| Dimension | Anonymous Data | Pseudonymised Data | VI — Proportional Identifiability Model |
|---|---|---|---|
| Re-identification for recipient | Not feasible | Feasible — key or inference | Cryptographically infeasible under defined assumptions |
| Lawful interception | Not possible | Possible | Yes — cryptographically gated judicial process |
| Standard stability as AI advances | Stable | Weakening | Stable — cryptographic assumptions not probabilistic |
| Verifiability by regulator | Not required | Judgment-based | Cryptographic audit trail — on demand |
| Consistency across processing chain | N/A | Entity-relative — fragmentation risk | Same verifiable standard for all entities |
| GDPR definition change required | No | No | No |
| Purpose limitation | Not applicable | Policy-declared | Cryptographically enforced by CJT |
| Bulk unmasking risk | Not applicable | Present — key access | Absent — quota-bound per judicial order |
| Processor accidental controller risk | Not applicable | Present | Absent — hardware-gated keys inaccessible to processor |
| Post-quantum readiness | Not applicable | Algorithm-dependent | Algorithm-agnostic — migration by substitution |
| Democratic accountability | Lost | Maintained | Maintained and architecturally enforced |
| AI innovation access | Unrestricted | Compliance-burdened | Cryptographically clear — stable legal basis |
| Replay attack resistance | Not applicable | Varies | Sealed monotonic counter — deterministic prevention |

# Relevant Implementation Contexts

| EU Policy Context | Specific Relevance |
|---|---|
| EDPB updated pseudonymization guidance — in progress | Cryptographically verifiable standard for Recital 26 proportionality — reference implementation model |
| Digital Omnibus — post-Council compromise | Technical implementation model achieving Commission objective without definitional change |
| AI Act — high-risk system implementation guidance | Purpose-bound execution enforcement for AI training and inference pipelines |
| Digital Euro privacy design — 2027+ | Non-persistent payment identifiers with proportionate lawful access — algorithm-agnostic for 2030 horizon |
| eIDAS 2.0 wallet privacy architecture | Session-scoped identity presentation without cross-context linkage — post-quantum forward-compatible |
| ENISA architecture security standards | Hardware-enforced enforcement gate as reference security architecture — FIPS 140-2/3 aligned |
| Article 28 — Data Processing Agreements | Processor de-risking — architectural prevention of accidental controller status |
| NIS2 implementation | Hardware-isolated identity binding as security-of-processing reference model |

# Closing Statement

This note does not take a position on the Commission's proposed definitional amendment or the Council's response to it. It observes that the underlying policy objective — legal certainty for entities that genuinely cannot re-identify pseudonymised data, enabling innovation while preserving consistent protection — remains unresolved by the current legislative trajectory and is achievable through a technical implementation model that requires no change to Article 4(1) GDPR.

*"The VI architecture offers a cryptographically verifiable implementation model for the proportionality principle that GDPR has required since 2018. Proportional identifiability — accessible to judicial process on a per-individual basis with replay-resistant quota enforcement, inaccessible to disproportionate bulk access, consistent across all entities in the processing chain, enforceable under defined security assumptions, and forward-compatible with post-quantum cryptographic standards — becomes a verifiable property of the architecture rather than a declaration of intent. AI innovation and privacy protection are not competing objectives requiring a legal trade-off. They are simultaneously achievable through an architecture that makes identifiability proportional by construction. This note is offered as a technical complement to the EDPB's ongoing work on pseudonymisation guidance and to the Commission's reform objectives — not as a position in the institutional debate, but as a contribution to its technical resolution."*

# Comparison Tables — AI Innovation, Privacy, and VI+CJT

## Table 1 — The Current Tension: What AI Innovation Has vs What Privacy Requires

*The gap that the GDPR reform debate is attempting to close*

| Dimension | What AI Innovation Currently Has | What Privacy Currently Requires | The Gap |
|---|---|---|---|
| Data access | Broad access to pseudonymised datasets — re-identification risk present | Purpose limitation — data used only for declared purpose | Same data cannot be simultaneously freely accessible and purpose-limited |
| Identifier persistence | Stable persistent identifiers across sessions — enables behavioral profiling | Data minimisation — no more data than necessary | Persistent identifiers exceed minimisation requirement by design |
| Cross-context data use | Data collected for one purpose reused for AI training | Purpose binding — collected purpose = used purpose | Repurposing is the norm — purpose limitation is the exception |
| Legal classification certainty | Contested — same pseudonymised data may or may not be personal data depending on entity | Consistent protection across all processing chain entities | Entity-relative classification creates compliance uncertainty and rights gaps |
| Cross-border processing | Data flows globally for AI training — jurisdiction of processing varies | Jurisdiction-scoped protection — EU law applies to EU citizen data | Data exits EU jurisdiction — protection cannot follow |
| Processor re-identification capability | Processors often hold or can access re-identification keys | Controllers accountable — processors should not hold controller-level capability | Processors carry accidental controller risk — liability unclear |
| Compliance burden | Full GDPR compliance required even where re-identification is genuinely infeasible | Proportionate compliance — burden should match actual risk | Burden does not scale to actual risk — innovation friction created |
| Lawful interception compatibility | Pseudonymisation sometimes prevents lawful access — friction with law enforcement | Democratic accountability — authorities must access data for legitimate investigations | Tension between privacy and lawful access unresolved technically |
| Audit and verifiability | Compliance declared by controller — not independently verifiable | Demonstrable compliance — controllers must be able to show compliance | Declaration is not demonstration — supervisory authorities cannot verify |
| AI training data provenance | Training data purpose often undeclared or loosely declared | Every processing operation requires lawful basis and declared purpose | AI training pipelines operate with loose or retroactive purpose declarations |

# Table 2 — What AI Gets Today vs What It Expects vs What VI+CJT Delivers

*The three-column answer to the reform debate*

| AI Requirement | What AI Gets Today | What AI Expects From GDPR Reform | What VI+CJT Actually Delivers |
|---|---|---|---|
| **Legal certainty on data classification** | Contested — entity must assess whether data is personal on case-by-case basis | Definitional clarification — pseudonymised data non-personal for recipient who cannot reasonably re-identify | Cryptographic certainty — recipient verifiably cannot re-identify under defined security assumptions — no assessment required |
| **Access to training data without full GDPR burden** | Full compliance required for all pseudonymised data regardless of actual re-identification risk | Reduced scope — data outside GDPR for entities without reasonable re-identification means | Proportionate compliance — data that is cryptographically non-re-identifiable carries burden proportionate to actual risk |
| **Cross-platform behavioral signals** | Persistent identifiers — available but privacy-invasive and legally contested | Definitional relaxation enabling freer use of behavioral data | Session-scoped behavioral signals across millions of contexts — richer and more diverse than persistent profiling — without identifier persistence |
| **Purpose flexibility for AI training** | Legitimate interests proposed as catch-all basis — contested | New Article explicitly recognising AI training as legitimate interest | CJT purpose binding — training within declared scope is cryptographically clear — outside declared scope is cryptographically prevented |
| **Cross-border data processing** | Legal uncertainty — transfer mechanisms fragile and politically dependent | Simplified transfer rules reducing compliance friction | CJT jurisdiction binding — EU citizen data non-executable outside permitted jurisdictional scope regardless of where processing occurs |
| **Processor legal clarity** | Processors risk accidental controller status — liability gap | Simplified processor obligations | Processor architecturally cannot hold re-identification capability — accidental controller risk eliminated by hardware design |
| **Inference and derived data use** | Derived data inherits personal data classification — compliance required | Derived data from non-personal pseudonymised data also non-personal | Derived data inherits CJT purpose and jurisdiction binding — non-executable outside declared scope regardless of derivation chain |
| **Re-identification risk management** | Organisational controls — policy-based, audit-dependent | Legal safe harbour for entities without reasonable re-identification means | Cryptographic safe harbour — re-identification infeasible under defined assumptions — verifiable by regulator on demand |

| | | | |
|---|---|---|---|
| **Lawful interception compatibility** | Pseudonymisation creates friction — some systems impede lawful access | No specific proposal — tension unaddressed | Judicial unmasking process — one session, one order, one purpose — lawful access preserved proportionately |
| **Audit and supervisory verifiability** | Self-declaration — supervisory authority cannot independently verify | No specific proposal — gap unaddressed | Cryptographic audit trail — supervisory authority can verify compliance on demand against defined parameters |
| **Long-term architectural stability** | Technology-dependent — AI advances erode pseudonymisation protection over time | Implementing acts to update criteria as technology evolves — rejected by Council | Algorithm-agnostic architecture — stable as AI advances — post-quantum forward-compatible to 2030 and beyond |
| **Data subject rights preservation** | Rights may disappear when data classified as non-personal for recipient | Rights gap created — data subjects lose GDPR rights at recipient level | Judicial unmasking process preserves rights across entire processing chain — rights do not disappear at any entity level |
| **Foundation model training** | Broad data access — legally contested | Freer access to pseudonymised datasets | Session-scoped diverse behavioral signals across millions of users — technically sufficient for foundation model capability without persistent profiling |
| **Personalization and recommendation** | Persistent cross-session profiling — legally contested | Continued access to longitudinal behavioral data | Domain-scoped VI continuity within declared purpose — personalization within context without cross-context surveillance |
| **Fraud detection and AML** | Cross-session behavioral continuity — legally justified but architecturally identical to surveillance | No specific proposal | Purpose-bound VI with scoped longitudinal continuity within regulated context — fraud detection enabled, cross-context surveillance prevented |

## Table 3 — The Three-State Comparison: Today vs Reform Proposal vs VI+CJT

*For policymakers and regulators — the complete picture in one table*

| Policy Dimension | State Today | Commission Proposal — November 2025 | Council Compromise — February 2026 | VI+CJT Implementation Model |
|---|---|---|---|---|
| **Personal data definition** | Art. 4(1) — identifiable directly or indirectly | New paragraph — data not personal if recipient cannot reasonably re-identify | Removed — definition unchanged | Unchanged — no amendment required |
| **Pseudonymised data classification** | Personal data for all entities in all circumstances | Non-personal for recipient without reasonable re-identification means — entity-relative | Unchanged — personal data | Verifiably non-re-identifiable for external parties under defined assumptions — consistent across all entities |
| **Identifiability standard** | Objective — all reasonably likely means considered | Subjective — entity-relative assessment | Unchanged — objective | Cryptographic predicate — verifiable, stable, technology-independent |
| **Legal certainty for AI** | Low — contested on case-by-case basis | Medium — definitional clarity but legally fragile | Low — unchanged uncertainty | High — cryptographic verifiability removes contestation |
| **Purpose limitation enforcement** | Declaratory — policy-based | Unchanged | Unchanged | Cryptographic — CJT enforces at data level |
| **Compliance burden for AI** | High — full GDPR regardless of actual risk | Reduced — for entities without reasonable re-identification means | Unchanged — high | Proportionate — scales to actual cryptographic re-identification risk |
| **Processing chain consistency** | Inconsistent — entity-relative assessments | Fragmented — same data personal for one entity, not another | Unchanged — consistent but burdensome | Consistent — same verifiable standard for all entities |
| **Processor accidental controller risk** | Present — processors may hold re-identification capability | Unaddressed | Unaddressed | Eliminated — processor architecturally lacks hardware-gated keys |
| **Lawful interception** | Possible — key-based | Unchanged | Unchanged | Preserved — judicial process cryptographically gated, quota-bound, audited |

| | | | | |
|---|---|---|---|---|
| **Bulk surveillance risk** | Present — key breach re-identifies all | Unchanged | Unchanged | Absent — no single access point, quota-bound per judicial order |
| **Supervisory authority verifiability** | Self-declaration | Implementing acts to specify criteria — rejected | EDPB guidance — in progress | Cryptographic audit trail — on demand verification |
| **Data subject rights across chain** | May disappear at recipient level | Rights gap created — data non-personal for recipient | Unchanged | Preserved — judicial unmasking process maintains rights across entire chain |
| **Cross-border jurisdiction** | Fragile — politically dependent transfer mechanisms | Simplified transfer rules | Unchanged | CJT jurisdiction binding — protection follows data regardless of processing location |
| **Post-quantum stability** | Technology-dependent — weakening | Implementing acts to update — rejected | EDPB guidance evolution | Algorithm-agnostic — post-quantum forward-compatible by design |
| **GDPR Article 25 — privacy by design** | Documentation requirement | Unchanged | Unchanged | Hardware-enforced — privacy is architectural baseline |
| **AI training lawful basis** | Contested | New legitimate interests provision for AI training | Partially accepted — with conditions | CJT purpose binding provides cryptographically clear lawful basis scoped to declared training purpose |
| **Innovation vs privacy trade-off** | Perceived as binary — one reduces the other | Attempts to shift balance toward innovation | Balance unchanged | Not a trade-off — both simultaneously achievable through proportional architecture |

# Note -

*"Today, AI innovation and privacy protection are perceived as competing objectives because the architecture forces a choice — persistent identifiers enable innovation and enable surveillance simultaneously, and no technical mechanism exists to separate them. The Commission's reform proposal attempted to resolve this by adjusting the legal boundary — which created fragmentation, rights gaps, and institutional resistance. The Council's response preserves the boundary but leaves the underlying tension unresolved. VI+CJT resolves the tension without moving the boundary — by making identifiability a cryptographic property rather than a legal classification. AI gets what it technically needs: verifiable legal certainty, proportionate compliance burden, diverse behavioral signals, and cross-border processing clarity. Citizens get what GDPR promises: consistent protection across the entire processing chain, purpose limitation that travels with the data, and lawful access through judicial process. The trade-off disappears because the architecture that created it is replaced."*