NIST Privacy Framework
Data Minimization

PRIVACY IN PRACTICE | HIGHER TRUST TOMORROW

PEOPLE
PROCESS
PRIVACY
PROGRESS

Goal: make minimization a risk-based outcome, not a compliance slogan

Audience: privacy, data governance, product, and security each own one slice

Roadmap: framework mechanics, Control-P outcomes, operating model, playbook, metrics, workshop

Baseline: PF 1.0 (January 2020); PF 1.1 final expected 2026, aligned to CSF 2.0

Running scenario: a customer analytics feature proposed with broad collection



MINIMIZE TODAY
ENABLE TOMORROW

Why Minimization Is a Strategic Imperative



- Shrinks breach impact, storage cost, and downstream copy exposure
- Over-collection drives purpose spillover, inference risk, orphaned data stores
- GDPR Art. 5(1)(c): adequate, relevant, limited to what is necessary
- CCPA/CPRA: reasonably necessary and proportionate to each purpose
- Business wins: faster reviews, cleaner AI data, quicker rights response
- Law sets a floor; the framework manages risk above it

Regulatory Landscape and Interoperability



US State Laws

- US states (Colorado, Connecticut, Virginia, Texas, Oregon) broadly require adequate, relevant, necessary collection



California (CPPA)

- California treats minimization as foundational across every purpose, including consumer rights requests



GDPR

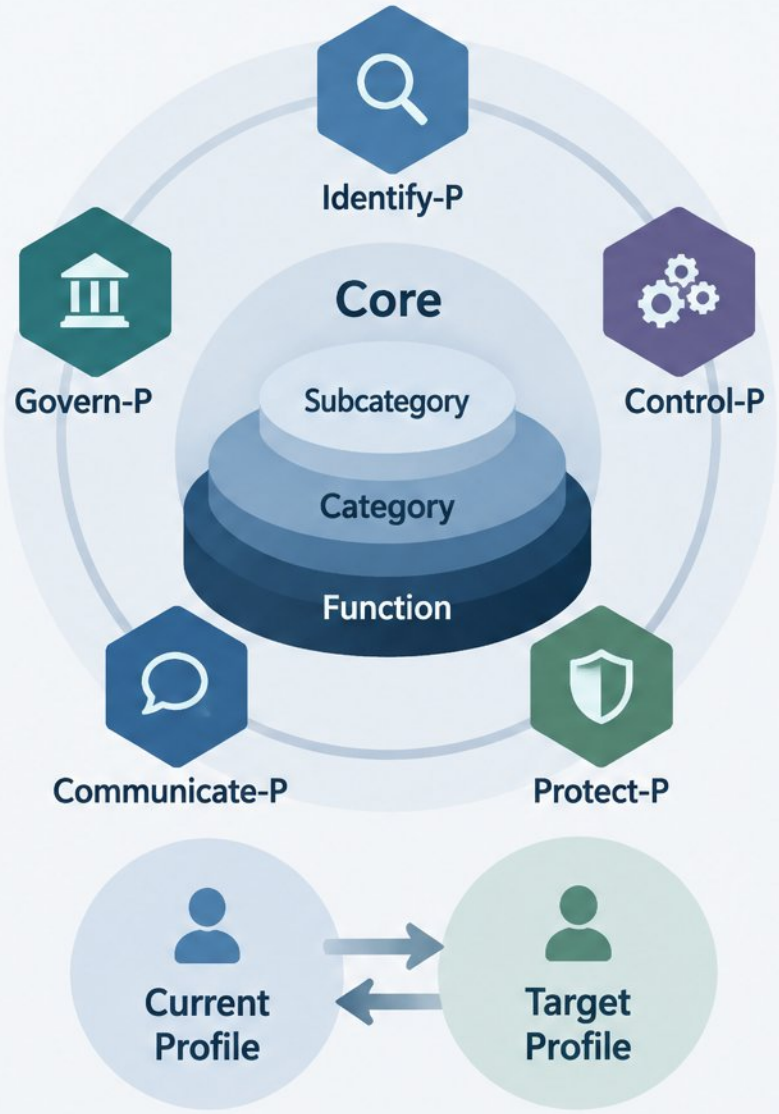
- GDPR needs an Article 6 legal basis; purpose and storage limitation are separate principles
- CPPA advisories flag excessive data requests—e.g., data beyond what's needed to send an opt-out signal



Build one minimization control set for overlapping regimes, not separate per-law projects

Framework Foundations: Core, Profiles, and Tiers

- Core: Functions, Categories, and Subcategories
- Five Functions: Identify-P, Govern-P, Control-P, Communicate-P, Protect-P
- Core is not a checklist — select outcomes by risk and role
- Profiles: Current vs. Target state, driving gap analysis
- Tiers: how risk-informed and repeatable practices are
- PF 1.1 IPD adds Community Profiles for shared interests



Current vs. Target state, driving gap analysis

Tiers

how risk-informed and repeatable practices are



Where Minimization Lives in the Framework

A COORDINATED APPROACH

Policy, inventory, and technical controls work together to make minimization real.



- **CT.DM-P** (Data Processing Management) is the explicit home of minimization language
- **CT.DM-P1–P5:** access for review, transmission, alteration, deletion, and destruction
- **CT.DM-P6–P10:** standardized formats, transmitted permissions, minimized audit/log records
- **CT.DP-P** carries technical intent: limit observability, linkability, identification, and inference
- **Govern-P (GV.PO-P, GV.MT-P)** and Identify-P inventories make it repeatable, not one-off

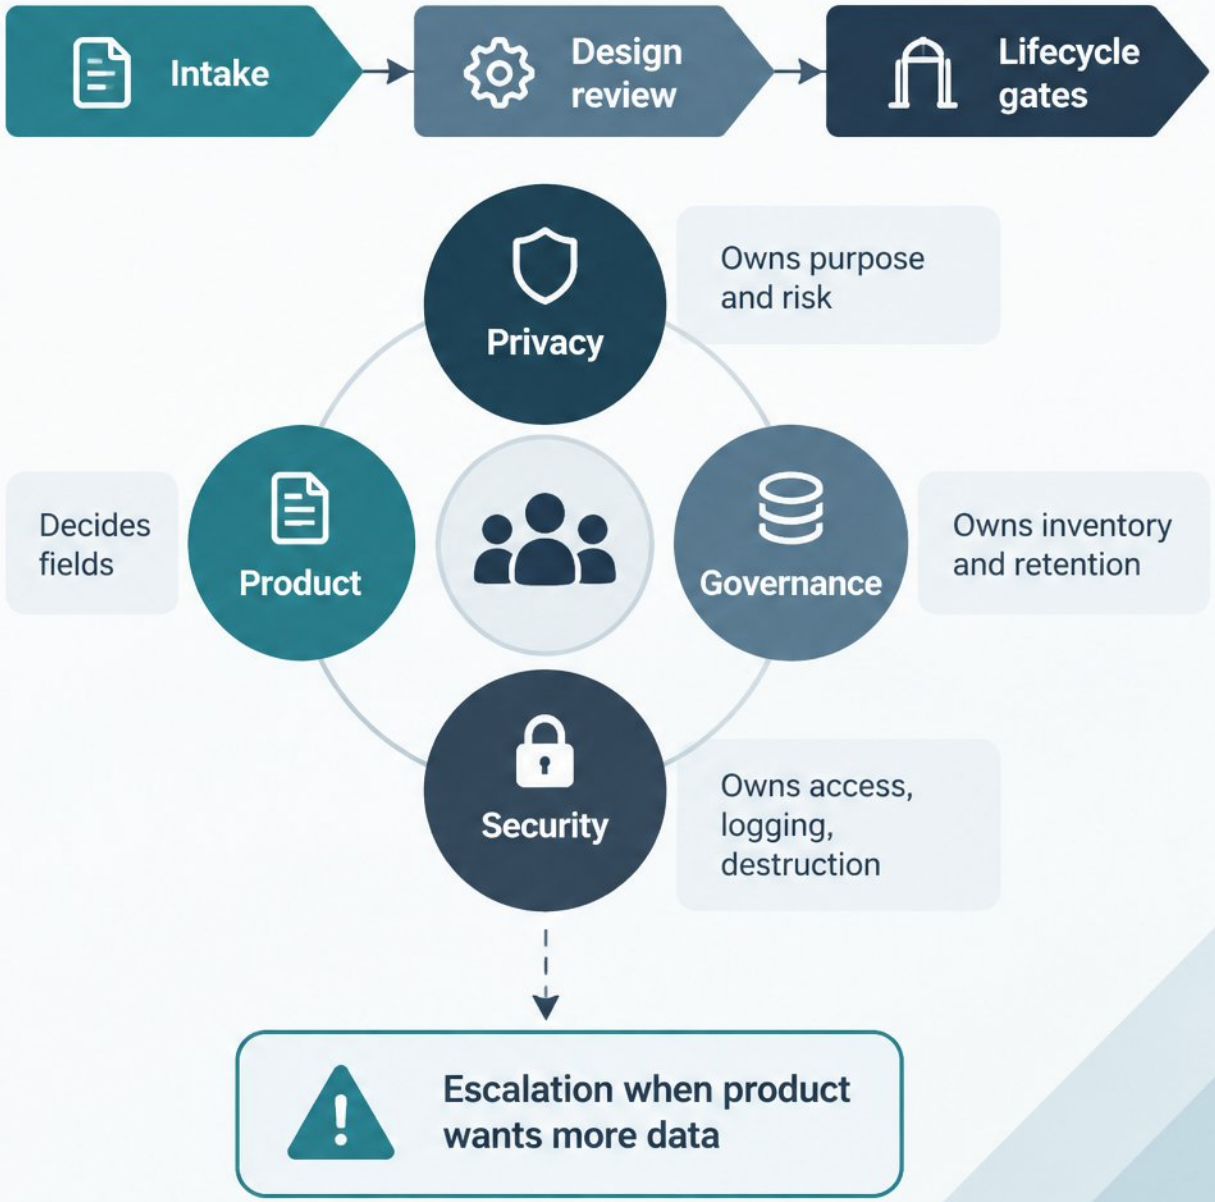
Reading the Subcategories: Intent and Evidence



- Four levers: scope collection, limit to purpose, contextualize purpose, control retention
- CT.PO-P1/P2: authorizations plus review, transfer, alteration, deletion, retention processes
- CT.DM-P8 blind spot: logs and audit records must be minimized yet support investigation
- CT.DP-P4/P5: selective collection via system configuration; attribute references over raw values
- Evidence reviewers expect: inventory entries, field-level purpose justification, retention settings, tests, deletion samples

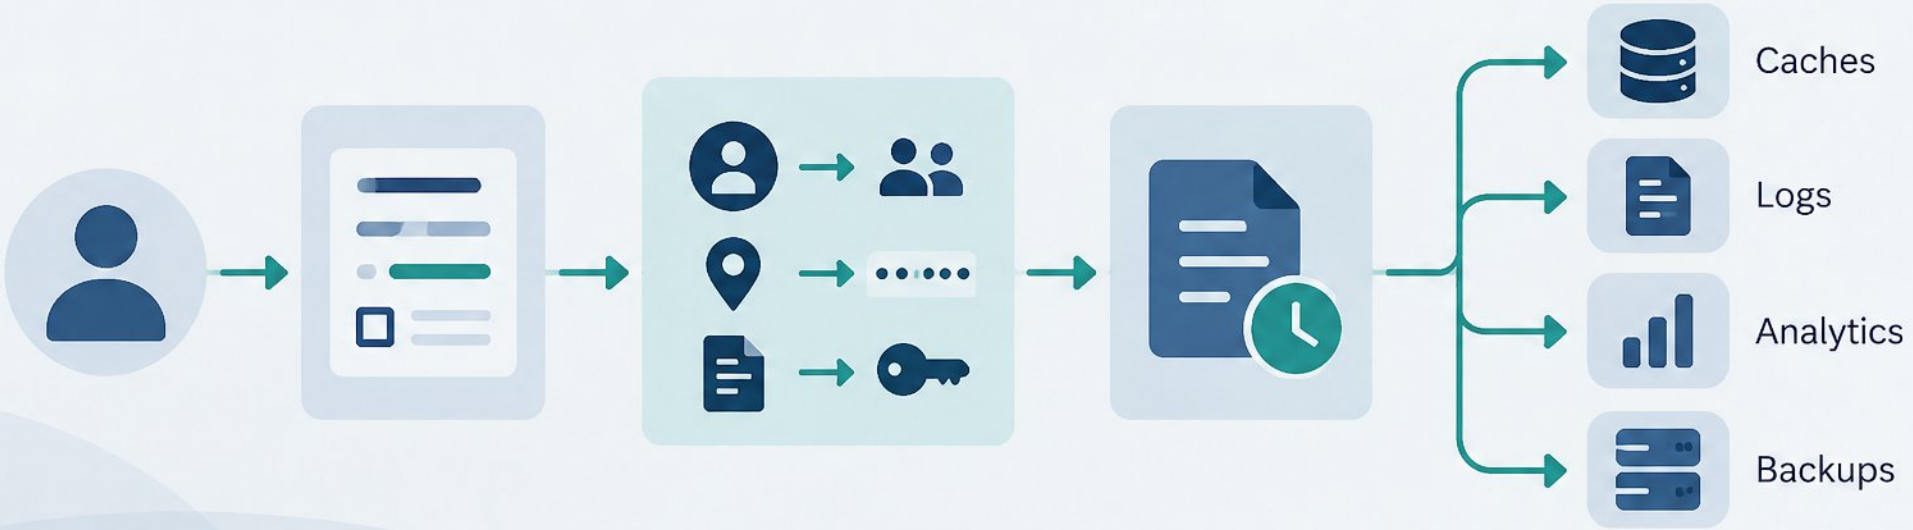
An Operating Model for Minimization

- Embed minimization at intake, design review, and lifecycle gates
- RACI: privacy owns purpose and risk; governance owns inventory and retention
- Product decides fields; security owns access, logging, destruction
- Reuse core controls: data inventory, purpose register, minimal-field standards
- Align with PIA, security review, and SDLC (CT.PO-P4)
- Define escalation criteria when product wants more data

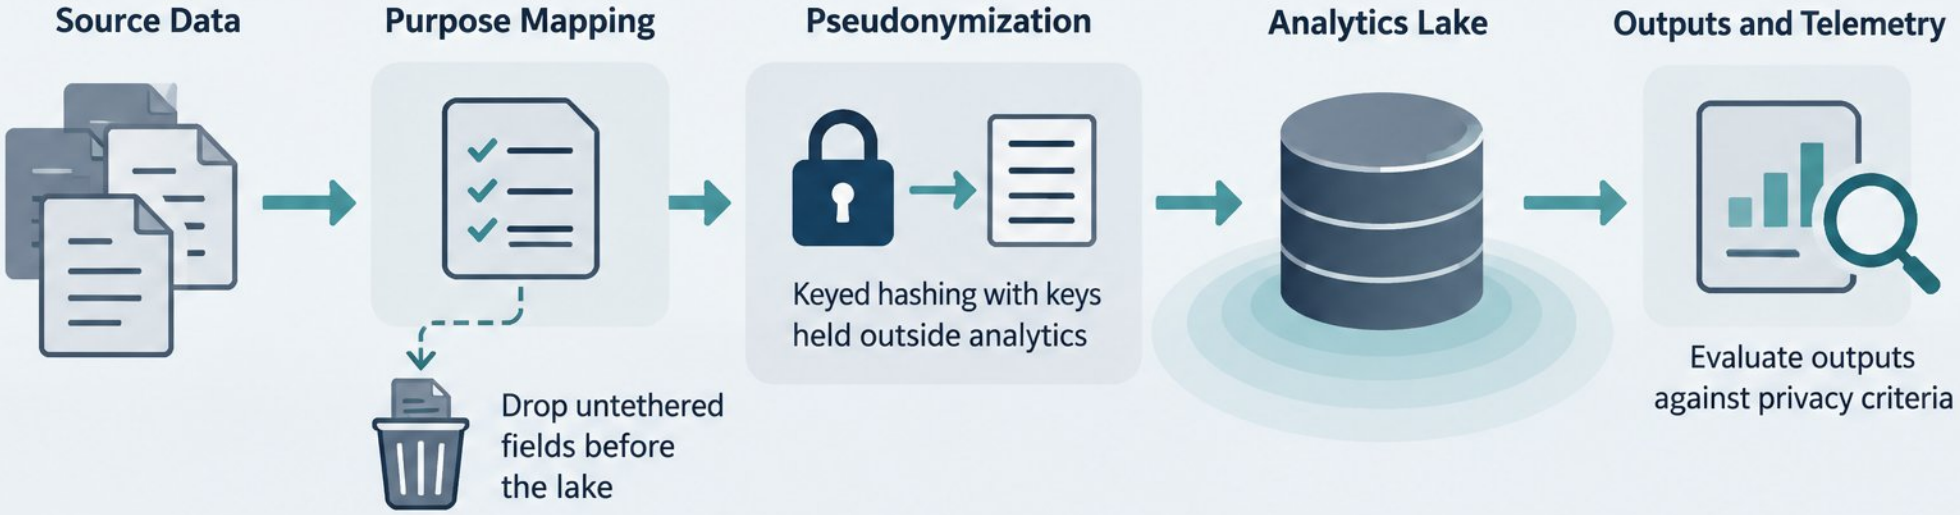


Product and Engineering Playbook

- **Collect just in time:** progressive profiling, optional fields, constrained inputs over free text
- **Derive and reference:** age bands not birth dates, truncated IP, tokens over copied identifiers
- **Know techniques:** anonymization is irreversible; pseudonymization keeps a key and stays personal data
- **Design retention in:** expiry rules per copy, idempotent deletion across caches, logs, analytics, backups
- **Crypto-shred immutable backups;** replay deletions before restored backups serve traffic



Minimization in AI and Analytics Pipelines



- PF 1.1 expands AI privacy risk; CT.DM-P10 embeds stakeholder preferences in algorithm design
- Map every training feature to a declared purpose; drop untethered fields before the lake
- Pseudonymized data is still personal data: keyed hashing with keys held outside analytics
- Aggregates leak in small groups; evaluate outputs against privacy criteria
- Minimize logs and telemetry: analytics events and debug logs are unmanaged collection channels

Metrics, Evidence, and Assurance

Leading indicators

- fields per feature,
- minimized-in-design rate,
- purpose justification coverage

Lagging indicators

- retention adherence,
- post-expiry copies,
- deletion request fulfillment time

Turn Subcategory achievement into auditable evidence for risk reporting and Tiers

Maturity ladder

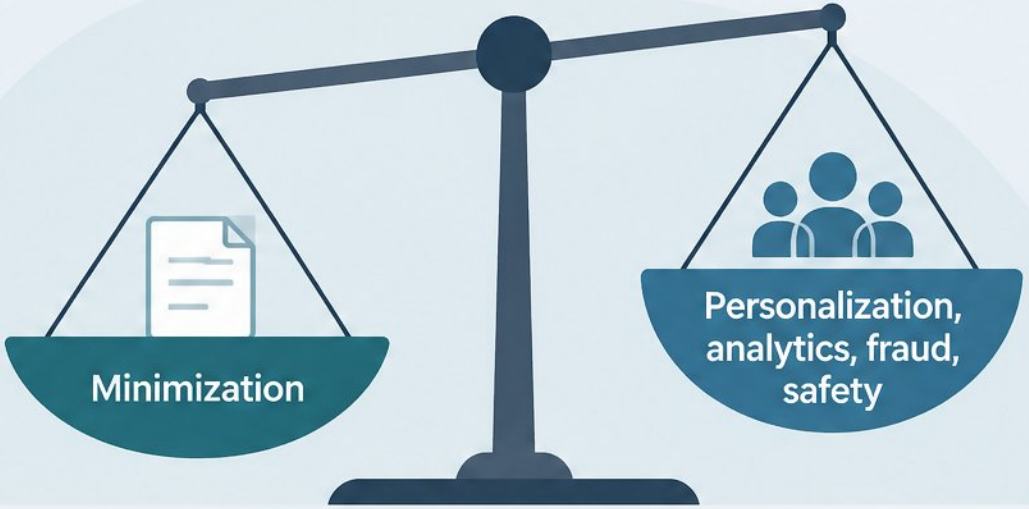
ad hoc → governed → measured → continuously proven

Tiers support budget talks

Avoid overclaiming: no "compliance with the Privacy Framework"; don't game field counts

Common Pitfalls and Cross-Functional Trade-offs

- Minimization vs. personalization, analytics, fraud, safety: negotiate on evidence and risk
- Failure modes: scope creep, orphaned stores, purpose drift, over-retention by default
- Security tension: scoped logging and access, not blanket retention (CT.DM-P8)
- Governance failure: policy without enforcement, expired exceptions never closed
- Turn documented failure patterns into teaching cases (GV.MT-P6)
- CT.PO-P2 and CT.PO-P4 anchor retention and life cycle controls



Common Failure Modes


scope creep


orphaned stores


purpose drift


over-retention by default



Workshop: Applying the Framework to a Product Scenario

Scenario:
the customer analytics feature with broad collection and third-party SDKs

→ **Map to Framework**

CT.DM-P | CT.DP-P | CT.PO-P | GV.PO-P | GV.MT-P

Step 2:
find gaps against a Current Profile; draft a Target Profile excerpt

→ **Current vs. Target Profile (excerpt)**

Current Profile	Target Profile excerpt

Step 3:
assign controls, owners, and dates across privacy, data governance, product, and security

→ **Controls, Owners, and Dates**

Control	Owner	Date



Debrief:
defend minimization to product and security; decide when to escalate

Adoption Roadmap: First 90 Days and First Year

A phased path from focused adoption to broader implementation, built on governance and guided by what's ahead.

Days 0–90

- Days 0–90: one product line, one data domain, a minimal evidence set, and a Current versus Target Profile excerpt

Months 3–12

- Months 3–12: expand to retention automation, intake gates, and metrics reporting into the risk register

Governance Hooks

- Governance hooks: policy updates, training, exception handling with expiry, and continuous monitoring (GV.MT-P)

Approach

- Sequence by risk and feasibility; small wins build credibility for broader enforcement

Watch for 2026

- Watch for 2026: final PF 1.1, implementation guidance, small-business quick start, Privacy Workforce Taxonomy mapping, PF 1.1 use case, and the Data Governance and Management Profile

Next Steps, Resources, and Q&A

- **Act now:** pick the scenario product, convene the four-team group, book the intake gate review
- **Watch 2026 releases:** final Privacy Framework 1.1, implementation guidance, small-business quick start guide
- **Track PF 1.1 shifts:** CSF 2.0 realignment, expanded AI privacy risk, Community Profiles
- **Stay current** via the NIST Privacy Framework site, Resource Repository crosswalks, and email list
- **Take home:** Profile template, field-purpose worksheet, retention rule template
- **Reminder:** framework is voluntary and risk-based — use it to prioritize, not certify



Take-home kit





Q&A

PersonWise.

PersonWise • AI digital-human course platform

Scan the code or click anywhere to watch this course live, with voice Q&A

personwise.ai/t/3e21b39d/public