Emotional Encoding Templates (EET) and Modular Layer System (MLS): A New Paradigm for Cybernetic Emotional Communication — v2 (Upgraded & Extended)

Executive Summary

Mission-critical operations routinely fail not for lack of facts, but for lack of **affective coherence** under stress. The EET+MLS framework decouples emotional context from semantic payloads and transmits it via low-bandwidth, machine-readable channels (haptics, temporal acoustic/visual rhythms, light), enabling rapid, resilient alignment across humans and machines. This v2 release upgrades the core model (VAD→signal mapping), formalizes a packet standard (ATP), adds L3 Safe-RL modulation, specifies consumer/enterprise deployments, and extends governance, risk, and compliance (GRC) guidance.

I. Strategic Imperative

Problem. In high-stakes settings (cyber incident response, remote medicine, tactical ops), cognitive load degrades language-based communication. Emotional intent becomes noisy or lost.

Solution. Separate channels: (a) semantic content, (b) **affective signal** (EET). Maintain emotional integrity with lower latency and lower bandwidth than text/voice.

Outcome. Faster shared situational understanding, lower miscommunication, stabilized team behavior, improved decision quality.

II. Affective Model and Signal Semantics

We adopt the **Valence–Arousal–Dominance (VAD)** model for continuous affect representation.

- **Valence (V):** unfavorable ↔ favorable
- **Arousal (A):** calm ↔ urgent/activated
- **Dominance (D):** unsure/submissive ↔ confident/in-control

II.A VAD→Signal Mapping (Upgraded)

We map VAD to physical signal parameters for three channels: **Haptic**, **Temporal (Morse-compatible)**, **Visual/Light**.

VAD Axis	Primary Mapping	Haptic	Temporal (EMC)	Visual/Light
Valence (V)	Polarity/Color Tone	soft vs. crisp envelope	timbre proxy via duty regularity	hue (warm=high V; cool=low V)
Arousal (A)	Rate/Energy	**Burst Duration (BD)** inverse to A; inter-burst interval ↓ with A	inter-symbol & inter-word gaps ↓ with A; micro-jitter ↑	blink rate ↑ with A
Dominance (D)	Intensity/ Confidence	amplitude & sharpness ↑ with D	amplitude profile ↑; rhythm regularity ↑	brightness ↑; edge definition ↑

Design rule: Delivery modulation must **not** alter the EET category itself. It optimizes **perception**, not **meaning**.

III. Emotional Morse Code (EMC) v2

Treat dot–dash rhythm as an **expressive carrier** in parallel to semantics.

- **A (Urgency):** shorten gaps; allow bounded micro-jitter to capture salience.
- **D (Authority):** maintain regularity; increase amplitude envelope.
- **V (Tone):** optional color/timbre overlays in multimodal contexts.

Template families: UrgencyHigh, CautionModerate, ReassureStable, NeutralInformational. Each template is a parameterized pattern with tolerances for noisy channels.

IV. Protocol Stack: Modular Layer System (MLS)

L6 – Presentation/UI. Multimodal IO: haptic actuators, LEDs, spatial audio, AR/VR.

L5 – Semantic Core. NLU, policy scheduling, initial affect inference to raw VAD.

L4 – EET Engine. Converts VAD→signal parameters using Table II.A; packages ATP.

L3 – Real-Time Modulation (Safe-RL). Bio-adaptive delivery within strict guardrails; logs decisions and supports user overrides.

L2 – Transport. IP/TCP, BLE, UWB, low-power RF; jitter buffers; integrity checks.

L1 – Services. Telemetry, resource mgmt, failover, security services.

V. Affective Tagging Protocol (ATP) v1.0

Objective. Standardized, interoperable EET packet.


```json
{
  "ver": "1.0",
  "msg_id": "uuid",
  "sender": "hash_or_pseudonymous_id",
  "target": "peer|group|device",
  "EET": {"label": "UrgencyHigh", "V": 0.2, "A": 0.9, "D": 0.8},
  "ctx": {"channel": "haptic|temporal|visual", "profile": "P1", "ttl_ms": 60000},
  "mod": {"allowed": true, "bounds": {"tempo_pct": [-15, 15], "intensity_pct": [-20, 0]}},
  "sig": "HMAC-SHA256(base64)",
  "ts": 1739059200
}
```


Notes. `mod.bounds` constrains L3; `sig` covers EET+ctx to prevent tampering.

VI. L3 Safe Reinforcement Learning (SRL) for Delivery Optimization

State (s): EET label, recent ack latency, HRV/EDA band, context (meeting/incident), prior actions.

Action (a): adjust {tempo, BD, intensity, actuator location}, or defer.

Reward (r): + for fast ack without stress drop; − for HRV decline or overrides; mild − for notification volume.

Hard constraints. Intensity caps, minimum inter-pattern gaps, **no semantic downgrades** of EET. Shadow-mode evaluation before live actuation. Append-only **modulation ledger** with rationale.

VII. Human Factors & Training

- **Calibration:** 2–3 minute onboarding to map perceived intensity and preferred actuators.
- **Literacy:** teach 4 canonical patterns (Urgency, Caution, Reassure, Neutral).
- **Fallback:** semantic-only templates; tiered urgency tags; closed-loop readbacks.

VIII. Deployments

A. Cybersecurity Incident Response

Decoupled haptic UrgencyHigh to all on-call; semantic logs arrive later. Metrics: MTTD/MTTR, alert acknowledgement time, false-negative rate.

B. Remote/Robotic Medicine

Caution pulses to surgeon controls; confidence cues via dominance intensity; reduces verbal clutter.

C. Tactical/Disaster Response

Low-signature EMC signals; swarm coordination with central **shimmer** scalar $S \in [0,1]$.

D. Consumer/Enterprise (Remote Work & VR)

Wearable morale cues; gentle EETs for encouragement, celebration, focus; strict privacy defaults.

IX. Swarm & Multi-Agent Coordination (Gamma Shimmer)

- Central broadcasts S; agents map to speed/spacing/sampling density.
- Agents share local S′ and converge (consensus) to reduce oscillation.
- ESRM checks divergence between central and local evidence; bounded autonomy.
- **Accountability:** S/S′, action, and model versions logged.

X. Privacy, Security, and Compliance

- **Zero-knowledge proofs** for auth where feasible.
- **Differential privacy** on analytics; local-first processing.
- **User-controlled retention** and export/delete rights.
- End-to-end encryption of ATP; key rotation; integrity checks.

XI. Governance & Ethics

- **Transparency:** sender intent vs delivered modulation visible on demand.
- **Consent:** opt-in biometrics; per-signal policy controls; quiet hours.
- **Liability:** maintain modulation ledger; clarify roles (sender/operator/vendor).
- **Standards:** propose ATP to ISO/IEC JTC1; haptic profiles via ISO/IEEE; XR via OpenXR.

XII. Risks & Mitigations (Summary)

- **Emotional manipulation:** logs, overrides, explainability.
- **Biometric surveillance:** minimization, consent, encryption, DP.
- **Misinterpretation:** standardized profiles, multimodal redundancy, training.
- **Overreliance:** fallback drills, interpersonal comms practice.
- **Bias:** diverse datasets, per-user calibration, adaptive interfaces.

XIII. Research Agenda & Longitudinal Studies

- **Constant Syncing:** burnout vs resilience over 6–12 months; mixed-methods.
- **Cross-cultural validation:** affect perception across demographics.
- **Neurodiversity support:** pattern libraries for sensory preferences.

XIV. MVP & Roadmap (12–18 Months)

- **P0 (0–2 mo):** ATP v0.1; 4 canonical haptics; desktop plugin.
- **P1 (3–6 mo):** Wearable beta; local processing; per-user calibration; 3-team pilot.
- **P2 (6–12 mo):** L3 SRL shadow → capped actuation; modulation ledger; start longitudinal.
- **P3 (12–18 mo):** VR integration; org dashboards (agg only); standards draft.

XV. Appendices

A. Canonical Haptic Profiles v1.0

Label	V	A	D	Pattern	Intensity
Reassure	.7	.3	.7	1× smooth 1.2s	Low
Encourage	.8	.5	.6	2×0.25s, 0.2s gap	Med
Caution	.5	.5	.5	triplet 0.15s /10s	Med-low
UrgencyHigh	.2	.9	.8	3×0.08s, 0.08s gaps	High

B. EMC Timing Bounds

- Dot base unit **t**; dash **3t**; intra-char gap **t**; inter-char **3t**; inter-word **7t**.
- Arousal scaling: $\mathbf{t' = t \cdot (1 - k{\cdot}A)}$ with $k \leq 0.3$; micro-jitter $\leq \pm 0.1t$.

C. Modulation Ledger Fields

`msg_id, EET_in(V,A,D,label), delivery(action params), bio_bands, rationale[], ack_time, user_override, ts`

D. Fallback Semantic Templates (PHI)

- **CRIT-1**: “Immediate action required. Acknowledge.”
- **WARN-2**: “Proceed with caution; confirm when stable.”
- **INFO-3**: “Status update only; no action required.”

Conclusion

EET+MLS v2 formalizes a dual-channel paradigm that carries **meaning** and **mood** with equal rigor. By standardizing affect as first-class, machine-readable signal—bounded by ethics, privacy, and transparency—we enable systems and teams to think faster, feel safer, and act wiser under pressure.

---

v2.1 Addendum — Symbiosis Integration & Extensions

A. Integration Plays (Ecosystem Hooks)

WebRTC (VR/MR morale cues). - Inject ATP EET packets as data-channel sidecar frames: `webrtc.datachannel.send(ATP)`. - Client renders EET to haptic/visual in-headset; falls back to UI glyphs if no actuator. - QoS: prioritize UrgencyHigh packets with DSCP EF; duplicate over redundant channels when available.

ROS/ROS2 (Robotic Medicine & Teleop). - `ros2` message type: `eet_msgs/Packet` mirroring ATP; topics: `/eet/in`, `/eet/out`, `/eet/teleop/haptics`. - Node `eet_bridge` maps L4 output → device drivers (haptic handles) and logs to L3 ledger. - Safety: `rclcpp` parameters enforce intensity caps per tool and per user.

Incident Platforms (SIEM/SOAR). - Connector translates playbook severity → EET (VAD preset); posts to on-call wearables. - Backpressure: when alert flood detected, L3 auto-batches reassurance/caution pulses.

Multi-Human “Gamma Shimmer”. - Team arousal scalar **S_team** = robust average of members’ A-bands from bios (opt-in) or self-rated pulses. - Broadcast **S_team** as low-frequency EET to nudge tempo of shared cues (e.g., de-escalate during overload).

B. ATP v1.1 Enhancements

Add **Resonance Feedback** to capture post-delivery effect and enable auto-tuning.


```
{
  "ver":"1.1",
  "msg_id":"uuid",
  "EET": {"label":"CautionModerate","V":0.5,"A":0.55,"D":0.5},
  "ctx": {"channel":"haptic","profile":"P2","ttl_ms":30000},
  "rfb": {"pre": {"V":0.52,"A":0.61,"D":0.46}, "post": {"V":0.56,"A":0.48,"D":
```

```
0.53}, "lag_ms": 4200},
  "sig":"HMAC(...)"
}
```


- `rfb.pre/post`: recipient-side VAD estimates (coarse bands) reported back with user consent. - `lag_ms`: time from delivery to stable post reading. - Privacy: differential privacy on aggregated rfb; per-user storage opt-in; no raw bios in packet.

C. Consumer Training & Gamification

- **Biofeedback Games:** mini-games that reward HRV recovery while responding to EET prompts.
- **Skill Trees:** unlock advanced profiles (e.g., "Gentle De-escalate") after proficiency on canonical four.
- **Daily Practice:** 90-sec drills (Recognize → Label → Breathe → Acknowledge) tracked locally.

D. Research Synergies

- **Cross-Cultural EMC Bounds:** partner with affective labs to derive jitter/tempo tolerance per locale; publish open tables.
- **VAD+S (Sociality) Axis Pilot:** extend model with **S** (cohesion/affiliation) $\in [0,1]$; map S to co-presence cues (subtle synchrony pips, shared glow) without affecting urgency semantics.
- **A/B Pilots in Cyber:** compare EET+MLS vs. traditional alerts on MTTD/MTTR, ack latency, error rates, psych-safety.

E. Moonshot: Multimodal LLM Co-Pilot (Empathic Modulation)

- **Inference:** LLM infers sender EET from chat semantics/prosody (low-confidence → human confirm).
- **Output Shaping:** LLM conditions response style on recipient's current A/D bands (from consented context) without changing underlying meaning.
- **Guardrails:**
- No covert affect changes: disclose "EET-shaped delivery enabled".
- Immutable semantics: EET shapes **delivery only**.
- Logged: store EET inference confidence, alternatives considered.

F. Metrics & Evaluation Additions

- **Resonance Δ:** $|\text{post-A} - \text{pre-A}|$ and recovery time distribution.
- **Over-notification Index:** haptic events per hour vs. overrides; target $\leq$ policy threshold.
- **Equity Score:** effectiveness parity across demographic/neuodiverse cohorts.

G. Standards & Interop Path

- Propose **ATP 1.1** to ISO/IEC JTC 1 (SC 42/SC 24); align haptic profiles with IEEE 700X series.
- WebRTC IETF draft for EET datachannel registry; ROS REP for `eet_msgs`.

H. Security & Privacy Updates

- Anonymous **rfb** by default; per-recipient salted IDs; DP on aggregates.
- Keys: rotate per-device; ledger entries signed and time-stamped.

- Local-first processing; cloud only for opted-in analytics.

I. Roadmap Inserts

- **P1.5:** ATP 1.1 (rfb), WebRTC/ROS bridges, consumer games alpha.
- **P2.5:** VAD+S experiments; LLM empathic shaping beta with strict disclosures.

This addendum positions EET/MLS as a practical symbiosis layer: interoperable, measurable, privacy-preserving, and ready for pilots across cyber, healthcare, XR, and field operations.

---

v2.2 Risk-Bound Expansion Plan

1) Scope Control & Phased Rollout

Core (ship-ready): VAD only; haptic + EMC + simple visual; ATP 1.0; L3 Safe-RL delivery (no semantic inference); single-locale pilots ($\geq$15 users).

Feature-flag tracks: all advanced features gated and measurable.

2) Experimental Tracks (Behind Flags)

A. VAD+S (Sociality) - Lab-only initially; team calibration sessions. - S drives co-presence cues (shared glow/heartbeat) only; no change to urgency semantics. - Exit: consistent $\Delta S \rightarrow \uparrow$ psych safety / $\downarrow$ over-notification across $\geq$3 cohorts.

B. LLM Shaping (Empathic Delivery) - Text semantics only; **no prosody** v1. - Rule-first EET mapping; LLM may suggest, never override. - Abstain if confidence < τ or sensitive domain. - On-device distilled model where possible; otherwise server with contract + audit; or disabled. - Disclose shaping; log rationale; human confirm for high-impact.

C. Resonance Feedback (rfb) - Aggregates only with **k-anonymity $\geq$10** per squad/window. - DP noise + $\geq$1h windows; suppress small teams. - Strict opt-in; local preview before sharing.

3) Targeted Risk Mitigations

LLM hallucinations/bias - Abstain-by-default policy; safety/intent checks. - Deterministic baseline mapping; fairness eval suite; drift monitoring.

Cross-cultural EMC bounds - Conservative Global Default; unlock locale packs post-data. - 2-minute personal timing calibration. - Start single-locale pilots; later cross-site A/B timing.

Standards volatility (WebRTC/ROS) - Adapter shims with strict semver; feature detection & graceful degrade (BLE/local sockets). - CI contract tests; rollback/killswitch per integration.

Battery & privacy (edge) - Duty-cycle haptics; cap notifications; coalesce non-urgent. - Small on-device models; periodic biosampling. - Fallback to rule-based when battery <15% or thermal >threshold.

De-anonymization (rfb) - Cohort thresholding; salted rotating IDs (24h). - Local-first raw storage; DP aggregates only; delete/export UX.

4) Minimal, Testable Increments

- **M0 (ship):** VAD core, ATP 1.0, 4 canonical haptics, Safe-RL, single-locale pilots.
- **M1 (opt-in):** rfb aggregates (k-anon + DP), no per-user server logs.
- **M2 (labs):** VAD+S co-presence; locale jitter studies with universities.
- **M3 (beta):** LLM text-only shaping with abstention + disclosure; no prosody; on-device first.

5) Go/No-Go Checklist (Per Rollout)

- Privacy DPIA complete; opt-in flows verified; deletion/export working.
- Guardrails: intensity caps, rate limits, override UI; modulation ledger enabled.
- Bias/safety tests passed across demographics; red-team report signed.
- Energy test on top wearables; fallback triggers validated.
- Standards adapters pass contract tests; integration kill-switch validated.
- Pilot metrics plan defined: ack latency, MTTR, over-notification index, equity score, fatigue.

6) Monitoring Essentials

Operational: ack latency, override rate, packet loss/late delivery, battery impact.

Affective: rfb ΔA distributions (when enabled), recovery time, misread rate.

Equity: effectiveness by locale/device/neurotype (report only when $k \geq 10$).

Trust: disclosure views, "deliver unmodified" selections, complaints per 1k signals.

> This plan constrains complexity, protects privacy, and keeps the roadmap auditable while enabling measured exploration of VAD+S, LLM shaping, and resonance feedback.

---

v2.3 Pragmatic Safeguards & Non-Blocking Governance

Goal. Address schedule risk from external dependencies (university studies, fairness suites), small-team k-anonymity bottlenecks, and ongoing governance needs for LLM abstention/sensitive domains—without weakening safety.

1) Small Teams & k-Anonymity (Equity Monitoring That Doesn't Stall)

- **Temporal k (rolling windows):** compute aggregates over sliding 30–45-day windows to reach **$k \geq 10$**; label charts with *window span* and *confidence band*.
- **Federated cohorting:** secure aggregation across *similar units* (same org, role, locale) using TEE/SGX or secure multiparty aggregation so no site sees raw counts.
- **Equity-Lite mode:** when $k<10$ persistently, suppress demographic cuts; show **proxy signals**: misread rate, override rate, ack latency, battery impact—stratified only by device type.

- **Sentinel cohorts (opt-in):** pre-consented cross-org pools to achieve k for early validation; results reported as *reference*, not canonical.
- **Confidence labelling:** every equity panel carries *confidence level* (High/Med/Low) and *k*; low-confidence panels cannot drive punitive decisions.

2) External Eval Dependencies (Keep Moving While Staying Honest)

- **Tiered evidence levels:**
- **L1 (ship gate):** internal fairness harness on synthetic + historical logs; red-team suite; sign-off by safety owner.
- **L2 (pilot hardening):** third-party fairness scan and limited academic review (non-blocking to pilot start).
- **L3 (publish/scale):** university study + cross-cultural validation before org-wide scaling.
- **Time-boxed reviews:** governance board meets bi-weekly; if L2 not ready by deadline, proceed with L1 + added monitoring & narrower scope.
- **Simulation bench:** EMC timing & haptic perception Monte Carlo sims using parametric human response models; validates bounds pre-pilot.

3) LLM Abstention & Sensitive Domain Governance (Clear, Fast, Auditable)

- **v1 Sensitive Domain Taxonomy (deny-by-default):** *health, employment, legal, identity, finances, safety-critical ops, minors.*
- **Rule-first mapping:** severity→EET via deterministic rules; LLM may *suggest* EET only outside sensitive domains.
- **Abstention policy:** on low confidence, domain hit, or model drift → *no shaping*. Fallback to rule-based delivery.
- **Runtime policy service:** centrally managed policy with **policy_ver**; cached locally; hot-reload without app updates.
- **Escalation SLA:** 24–72h for taxonomy updates; emergency blocklist in <1h.
- **Audit trail:** inference confidence, taxonomy match, and final decision logged.
- **ATP extension:** add `policy": {"policy_ver":"1.0","sens_domain":["health"]}` when applicable (never transmitted for consumer default unless required by enterprise policy).

4) Monitoring When k<10 (Don’t Fly Blind)

- **Aggregated proxies:** publish non-identifying metrics (override rate, misread rate, ack latency) with DP noise and windowed smoothing.
- **Rolling cohort merge:** merge small teams weekly into anonymized super-cohorts to surpass k while preserving org segmentation.
- **Confidence tags & guardrails:** attach *Low-Confidence* tag; auto-tighten guardrails (rate limits, intensity caps) until k recovered.

5) Timeline Assurance Mechanisms

- **Feature flags with TTL:** every experimental feature expires unless renewed by governance.
- **Change cadence:** bi-weekly governance, monthly stability freeze for core stack.
- **Go/No-Go timeboxes:** if external eval misses window, proceed with smaller scope, stricter monitoring, and explicit “provisional” label.

6) Risk Register Additions

- **R-17:** De-anonymization in micro-teams → mitigated by temporal k & federated cohorting.
- **R-18:** Standards API drift (WebRTC/ROS) → adapter contracts + kill-switch; quarterly compatibility review.
- **R-19:** LLM drift causes over-abstention → fallback to rules + model canary alerts; rotate model only with A/B.
- **R-20:** Governance backlog delays → SLA & emergency blocklist channel; publish queue status in dashboard.

  v2.3 keeps the schedule moving by substituting *confidence-labeled proxies*, **federated aggregation**, and **rule-first fallbacks** wherever data sparsity or governance latency would otherwise stall deployment.

---

v2.4 Operational Overhead & Assumption Risk Controls

Context. You flagged concrete downsides: overhead from federated aggregation (TEE/SGX), weekly cohort merges, runtime policy plumbing, and the risk of meta-bureaucracy; plus assumption risks around L1 internal harnesses and culturally narrow sensitive-domain taxonomies. This section introduces **lean modes**, **latency budgets**, and **assumption audits** that keep safety without stalling small teams.

A) Lean Operations: Reduce Overhead Without Losing Safety

A1. Modes of Operation - **Micro-Team Mode (≤15 users):** - Disable federated aggregation; enable **Equity-Lite** proxies only (override/misread/ack latency by device). - rfb off by default; local reflections only. - Governance autopilot: default policies with **TTL=90 days** unless org opts into advanced. - **Standard Mode:** current v2.3 defaults (k≥10, DP, cohort merges). - **Enterprise Mode:** adds federated aggregation/TEEs and full RACI governance.

A2. Alternatives to TEEs - **Secure Aggregation (no SGX):** pairwise masks / SMPC for sum-only metrics; server sees only aggregates. - **Serverless Cohort Merge:** scheduled job merges sub-k cohorts weekly; operators can opt for **auto-merge** to remove manual effort.

A3. Latency Budgets & Edge Constraints - **EET render SLO:** ≤10ms edge render (haptic/visual) from receipt. - **On-device policy cache:** policy check ≤2ms; stale-ok for ≤24h with hash verification. - **RF path budget:** ≤50ms one-way for UrgencyHigh; queue non-urgent for batch delivery. - **Energy caps:** max N haptics/hr; adaptive coalescing to protect battery.

A4. Governance Without Meta-Bureaucracy - **RACI-Lite:** 3 roles only — *Owner* (approves), *Steward* (monitors), *Auditor* (periodic checks). - **Auto-approval defaults:** low-risk changes (copy edits, non-urgent visual tweaks) ship unless vetoed within 72h. - **Change classes:** - **Class A (high-risk):** requires board; 7-day review. - **Class B (moderate):** Owner + Steward; 72h. - **Class C (low):** autopass with post-hoc audit. - **Policy TTLs:** features expire unless re-affirmed; dashboard shows pending expiries.

B) Assumption Risk Controls

B1. Strengthen L1 Evidence (Internal Harness) - **Counterfactual Data Augmentation:** swap demographics/locales in messages; check parity. - **Adversarial Stress Tests:** bias injections (names, dialects, prosody tokens) to probe false positives. - **Shadow Evaluation:** before pilot, run rule-based vs. model-suggested EET mapping on real historical comms (de-identified) — models may only *suggest*. - **Holdout Canaries:** a small, consented group runs dual delivery (rule vs. RL-modulated) to detect drift.

B2. Sensitive-Domain Taxonomy Governance - **Global Advisory Circle:** rotating reviewers from multiple regions/cultures + domain experts (health, legal, labor) meet monthly. - **Localization Packs:** domain lists per locale; default to **deny-by-default** when uncertain. - **Public Issue Queue (enterprise-visible):** track proposed taxonomy changes with rationales and timelines. - **Heuristic Baseline:** rule-based detectors for domain keywords; LLM only assists with explanations; humans approve.

B3. Monitoring When Data Is Sparse - Confidence-labeled proxies (Equity-Lite) remain visible; **no punitive actions** allowed under Low Confidence. - Auto-tighten guardrails (rate limits/intensity caps) while k<10 persists; notify Owner.

C) Ops Cost Snapshot (Indicative)

Component	Micro-Team	Standard	Enterprise
Setup time	~1–2 hrs	~1 day	2–4 days
Cohort merge ops	none (auto)	weekly serverless	automated + reviewed
TEE/SGX	none	none	optional (SMPC alt)
Governance cadence	quarterly TTL	monthly	bi-weekly + board
Edge latency budget	10ms render / 50ms RF	same	same

D) Decision Gates (Non-Blocking)

- If **external fairness scan** late → proceed with L1 harness + narrower scope, label *Provisional*, increase monitoring.
- If **k<10** persists → remain in Equity-Lite + guardrails; plan sentinel cohort; no blocker.
- If **policy backlog** → Class C changes auto-pass; emergency blocklist path active.

v2.4 introduces lean defaults, eliminates unnecessary TEEs for small teams, sets explicit latency/energy budgets, and replaces brittle dependencies with confidence-labeled proxies and rule-first fallbacks. It keeps you shipping while staying honest about uncertainty.

---

v2.5 Lingering Nuances & Quantification Framework

A) Heterogeneous Fleets: Mode Logic Without Fragility

A1. Single Source of Truth (SSOT) for Mode & Policy - Device reads signed **OrgConfig**: `{mode: micro|standard|enterprise, policy_ver, ttl}`. No client auto-detect; reduce branchiness. - **Fail-closed** defaults: if config stale >24h, fall back to Micro-Team Mode + Equity-Lite.

A2. Compatibility Matrix & Contract Tests - Maintain a published matrix across **Wearable FW x Mobile OS x ROS/ROS2 x App Ver**. - Ship **contract tests** (open test app) that validate: ATP parse, haptic timing, policy cache, ledger write. Pilots must pass before go-live.

A3. Policy Cache Hygiene - Immutable, versioned policy bundles (`policy_ver`) signed + hashed; hot-reload with checksum. - Canary rollout: 5% devices, then staged waves; auto-rollback on error codes.

A4. ROS & Wearables Bridging - Use a thin `eet_bridge` node translating ATP↔`eet_msgs`; keep ROS QoS profiles fixed per mode. - On wearable failure, degrade to visual glyphs; log `delivery_fallback=true` in ledger.

B) SMPC Instead of TEEs (But Keep It Optional)

B1. Lightweight Aggregation Path - Secure Aggregate (sum/mean) with pairwise masks; server sees only aggregates. - Library plug point: abstraction layer so orgs can choose **SMPC | TEE | None (Equity-Lite)**.

B2. Small Orgs Reality - Default to Equity-Lite + periodic, opt-in sentinel cohorts to unlock deeper insights quarterly. - Provide **Managed Cohort Service** (optional) run by vendor; small teams avoid crypto ops entirely.

C) Advisory Circle: Faster Response Without Bureaucracy

C1. Rapid Response Protocol (RRP) - **48-hour** window for emerging domain flags (e.g., novel AI ethics concerns): temporary **deny-by-default** rule added via runtime policy service. - Asynchronous vote via secure channel (no monthly wait). Monthly meeting only for ratification.

C2. Rolling Reviewers - 2 on-call reviewers per region rotate weekly; quorum = 3 global signatures to enact interim policy.

D) Quantification: Costs, SLO Value, & ROI

D1. Indicative Cost Model (per 100 users / month) - Serverless cohort merge: ~10–30k invocations → **$5–$20** (provider-dependent). - Ledger storage (compressed JSON, 1KB/msg, 50 msgs/user/day): **~150 MB/mo** → **$3–$10**. - Key mgmt / signing ops: **$1–$5**. - Push/transport (BLE/WebRTC datachannel negligible) + monitoring: **$5–$15**. - **Ballpark infra: $14–$50 / 100 users / mo** (excluding wearables).

D2. Pilot Budget Snapshot (8 weeks, 50 users) - Devices (loaners) & setup: **$3–$6k**. - Integration & tests: **$5–$8k**. - Study ops & analysis: **$4–$7k**. - **Total pilot range: $12–$21k**.

D3. ROI Template - **Inputs:** baseline MTTR (mins), ack latency (secs), incident cost/min, meeting hours/person/week, retention rate. - **Effects:** Δack (secs) → ΔMTTR; Δmeeting_hours; Δretention. - **Example:** - Ack latency ↓ from 45s→18s (−27s). For 30 P1 alerts/mo at $150/min incident cost, MTTR ↓ ~3% → **$13.5k/yr** saved per team. - Meeting fatigue ↓ 10% (0.5h/person/wk) at $70/h blended → **$91k/yr** for 50 people. - +2% retention at $15k replacement cost → **$15k/yr**. - **Net (illustrative):** $120k+/yr upside vs. <$10k/yr infra + ops for 50 users.

D4. Latency Compliance → Retention Proxy - Track **SLO hits** (≤10ms render, ≤50ms urgent RF). For cohorts where SLO >98%, correlate with **first-month retention**; use as leading indicator.

E) Action Items & Ownership (RACI-Lite)

- **Owner:** signs mode selection, approves policy changes Class A/B.
- **Steward:** monitors SLOs, equity proxies, rfb toggles; triggers RRP.
- **Auditor:** quarterly review of ledger, privacy, and bias harness.

F) Attribution / Partners

- Add **Lika Mentchoukov (Mench.ai)** to the Advisory Circle for rapid-review and cross-domain sensitivity inputs; listed as *External Advisor* in governance appendix.

  v2.5 reduces mode-logic brittleness, offers SMPC without crypto pain, speeds domain policy updates via RRP, and provides concrete cost/ROI scaffolds so small orgs can plan without guesswork.

---

v2.6 External Variables & Regulated-Scale Playbook

Aim. De-risk variable costs/ROI assumptions, partner coordination, and compliance constraints in ultra-regulated environments without stalling delivery.

A) Cost & ROI Uncertainty Controls

A1. Budget Bands & Guardrails - **Spend SLOs:** monthly caps per 100 users; alert at 70/85/100% with auto-throttle on non-urgent features. - **Autoscaling floors/ceilings:** cap serverless concurrency; queue non-critical merges. - **Committed use discounts:** plan RI/savings instruments where possible; fallback to spot for batch analytics.

A2. Sensitivity & Monte Carlo - Publish a **ROI Sensitivity Sheet**: vary incident cost ($50–$500/min), alert volume (±50%), Δack (±20s) to show bands. - **Monte Carlo** (1k runs): randomize cloud unit costs (±30%), retention lift (0–3%), and meeting-fatigue reduction (0–15%) → report P10/P50/P90 ROI. - Track **realized** vs **assumed** inputs from pilots and update priors quarterly.

A3. Cost Transparency - Dashboard: per-feature cost (ledger, merges, key ops) and per-cohort unit economics. - Kill-switch per feature if unit cost > threshold for X days.

B) External Partner Coordination (Contract Early, Move Fast)

B1. Pre-Negotiated MSA & DPA - Keep a light **Master Services Agreement** and **Data Processing Addendum** template ready; include security, privacy, and IP clauses. - **SOW Milestones:** define deliverables (fairness scan, locale jitter study) with dates and acceptance tests.

B2. RACI for Partners - **Owner (internal):** signs scope and owns timelines. - **Partner Lead:** single POC with escalation path. - **Steward (internal):** day-to-day integration; weekly sync cadence. - **Auditor (internal or 3rd-party):** verifies outputs.

B3. SLA & Escalations - **Response SLA:** 2 biz days; **Resolution SLA:** time-boxed by SOW. - **Fallback:** if partner misses milestone, proceed with L1/L2 evidence + narrower scope; revisit when deliverables land.

C) Regulated Sectors Playbook (Defense/Health/Gov)

C1. Deployment Options - **Self-Host / On-Prem / Air-Gapped:** containerized MLS; offline ledger; manual key rotation. - **Gov Cloud / FedRAMP/IL-x:** deploy ATP services within accredited env; disable consumer features. - **Data Residency:** pin storage region; per-tenant KMS; split metadata vs. biosensors.

C2. Compliance Matrix & Artifacts - Map to **NIST 800-53**, **FedRAMP**, **ISO 27001**, **SOC 2**, **HIPAA**, **GDPR/CCPA**, **ITAR/EAR** (export controls) as applicable. - Provide **SBOM** and supply-chain attestations (SLSA/SBOM-CycloneDX). - Maintain **Data Flow Diagrams**, **DPIA**, **Threat Model**, **Incident Response Plan**, and **BCP/DR**.

C3. Legal Review Gate - **Gate G-Reg-01:** counsel approval required before enabling managed cohorts; otherwise run **Customer-Managed Cohorts** only. - **Security Review:** penetration test + policy verification before go-live; annual re-attestation.

C4. Features for Regulated Mode - rfb disabled by default; if enabled, **on-prem aggregator** only. - No LLM shaping unless model is **customer-hosted**; telemetry minimized; logs local. - Extended retention controls; tamper-evident ledger with customer HSM keys.

D) Procurement & Timeline Realism

- **Lead-time buffers:** add 4–8 weeks for vendor onboarding, legal reviews, and security scans in defense/health.
- **Parallelization:** run technical pilot prep (contract tests, device staging) while legal review proceeds.
- **Go/No-Go cadence:** provisional go permitted for non-regulated cohorts; regulated cohorts wait for G-Reg-01.

E) Roles, Risk & Accountability

- **Risk R-21:** Cloud cost spikes → spend SLOs + throttling + feature kill-switch.
- **Risk R-22:** Partner delay → pre-negotiated MSA/SOW + fallback evidence + scope guardrails.
- **Risk R-23:** Compliance rejection → Regulated Mode with self-host, rfb off, customer HSM; legal gate.
- **Risk R-24:** De-anonymization risk in small secure units → Equity-Lite only; no cross-unit aggregation.

v2.6 adds hard cost/ROI guardrails, contracts/SLA discipline for partners, and a Regulated Mode that satisfies stringent compliance without freezing innovation.

---

v2.7 Due Diligence, ROI Guardrails & Air-Gapped Ops

A) Advisor & Attribution Hygiene

- **Evidence Thresholds:** External advisors must meet ≥2 of: peer-reviewed publications, standards contributions, deployed programs in affective ethics/HCI at scale, or recognized patents.
- **Documentation:** Maintain public advisor dossier (works, role, scope, conflicts). Renewal annually.
- **Status of Mentchoukov/Mench.ai:** mark as **Provisional Advisor (evaluation)** until traction criteria are met; remove if not met within 6 months.
- **Placeholder Policy:** no placeholder attributions in public materials; use “TBD (under evaluation)” instead.

B) ROI & Monte Carlo—From Assumption to Evidence

- **Assumption Labels:** every ROI figure tagged as *Assumed* or *Pilot-Observed* with date & cohort.
- **Validation Gates:**
- **G-ROI-01:** require ≥2 pilots to show Δack improvement with $p<0.05$ before claiming MTTR benefits.
- **G-ROI-02:** retention lift claims require ≥90-day observation and matched control cohort.
- **Priors & Bands:** publish priors used (incident $/min range, alert volume, fatigue delta) and show P10/P50/P90 bands, not point estimates.
- **Sector Adjustors:** provide presets (Cyber, Health, Gov) that tune priors and widen bands for high-variance verticals.

C) Air-Gapped & IL-x Environments—Operational Playbook

- **Manual Key Rotation Minimization:**
- Use **customer HSM with time-boxed key epochs** and pre-signed rotation bundles; rotation executed via offline USB/token with dual-control (two-person rule).
- Rotation cadence defaults: quarterly; emergency rotate in <2 hours using pre-staged bundles.
- **No-Net Update Channel:**
- Policy bundles & ATP profiles shipped on encrypted media; hash verified against on-prem root; canary devices staged before fleet rollout.
- **Hardware & Cert Matrix:**
- Maintain **IL-x / FIPS-140-2/3** compatibility matrix (actuators, bridges, servers). Blocklist non-compliant SKUs; publish lead-times for certified replacements.
- **Operational SLOs for Air-Gapped:**
- EET render ≤10ms, ledger write ≤20ms local, policy check ≤2ms from cache; all verifiable via on-prem contract tests.
- **LLM Shaping:** disabled unless customer-hosted model accredited for IL tier; otherwise rule-first only.

D) Partner Coordination & Contracting—No Surprises

- **Pre-Contracting:** advisors/partners must have **MSA/DPA** signed during discovery; SOW includes milestones & exit criteria; provisional titles avoided.
- **Escalation Matrix:** designate internal exec sponsor and partner lead with 24-hour escalation window; fallback to internal tracks if missed.

E) Compliance Sign-Offs (Defense/Gov)

- **G-REG-02:** require legal sign-off that vendor-run cohorts are permitted; if not, enforce **Customer-Managed Cohorts** only (no cross-unit aggregation).
- **Artifact Pack:** deliver SBOM, DPIA, threat model, export-control review before pilot start; Go/No-Go blocked without pack receipt.

v2.7 removes placeholder risks, turns ROI into evidence-driven claims, and hardens air-gapped operations with practical key and update procedures, aligned to IL-x constraints.

---

v2.8 Fast-Track Governance & Evidence Capture for Micro-Teams

Objective. Keep consumer and small-team pilots moving without compromising safety. Convert governance into self-serve automation; pre-wire ROI capture so “Observed” tags arrive sooner.

A) Fast-Track (Green-Lane) Pilot Path

- **Eligibility:** $\leq$50 users, non-regulated sector, no biometrics (rfb off), no LLM shaping, VAD core only.
- **Auto-provisioned artifacts:** DPIA-Lite, Threat Sketch, SBOM summary, ATP profile sheet generated from OrgConfig.
- **Time-boxed approvals:** if legal does not respond in **5 business days**, pilot proceeds under *Provisional* with added monitoring and a rollback plan.
- **Kill-switches:** one-click disable for EET channel, RL modulation, and integrations.

B) Governance-as-Code (GaaC)

- **Policy bundles as code:** YAML with schema validation; signed and versioned.
- **Compliance checks:** pre-flight `gaac lint` verifies: sensitive domain settings, rfb status, data retention, encryption flags.
- **Regulatory profiles:** `profile: consumer|enterprise|regulated` sets defaults (e.g., rfb=false, retention=local-only for consumer).

C) Legal Load Reduction

- **Template kit:** jurisdiction-aware DPIA-Lite, DPA addendum, pilot SOW; auto-filled from OrgConfig.
- **Clause library:** pre-approved fallback clauses (data locality, opt-in biosensors, no cross-unit agg) toggle via flags.
- **Escalation bot:** tracks G-REG gates; pings counsel with concise diffs; posts green/yellow/red status in pilot dashboard.

D) Evidence Kits (Faster “Observed” Tags)

- **Instrumentation SDK:** drop-in timers for ack latency, MTTR linkages, over-notification index; stores locally and exports anonymized aggregates.
- **Auto-A/B harness:** randomizes delivery between *Rule-only* vs. *Safe-RL* within guardrails; computes deltas.
- **Cost SLOs:** auto-track infra cost/user and emit spend alarms vs. budget bands; surface P10/P50/P90.
- **Claim lifecycles:** dashboard shifts metrics from *Assumed* → *Pilot-Observed* once thresholds met (e.g., $n \geq 200$ events, $p<0.05$, 90-day window).

E) Micro-Team Default Config

```
profile: consumer
mode: micro
rfb: false
llm_shaping: false
retention: local_only
transport: ble|webrtc
policy_ttl_days: 90
max_haptics_per_hour: 12
intensity_cap: user_calibrated
kill_switches: [eet_channel, rl_modulation, integrations]
```

F) Ops Dash Additions

- **Legal clock:** shows days remaining on time-box; lists blockers; provides auto-generated PDFs.
- **Green-Lane status:** visible badge; downgrade to Standard if flags turned on (rfb/LLM/regulated integrations).
- **Rollback rehearsal:** one-click simulation validates end-to-end disable within 60s.

G) Communication & Transparency

- **Pilot consent:** plain-language sheet (1-page) with examples of haptics; explains opt-out and local-only data.
- **Stakeholder digest:** weekly auto-summary with metrics, costs, incidents, and any governance changes applied via GaaC.

v2.8 turns governance into automation (GaaC), provides a green-lane fast track for micro-teams, and auto-captures the evidence needed to graduate metrics from *Assumed* to *Observed* without waiting on heavy external reviews.