Bridging the Sovereignty Gap: A Reference Architecture for Compliant AI Orchestration in Restricted Jurisdictions

Kevin Frank
Independent AI Compliance Architect
Quito, Ecuador

ABSTRACT

As the January 28, 2027 enforcement deadline crystallizes the strict liability regime of Resolution 0005-R, Ecuadorian entities face a solvency risk: repatriate AI workloads at the cost of global innovation, or face cumulative sanctions for cross-border data violations. The recent $454,500 fine against LigaPro and FEF set a legal precedent and defined a baseline cost of the "Sovereignty Gap," but the operational risk lies specifically in the inability to audit inference flows against Article 12 mandates. This paper proposes the Sovereign Intelligence Gateway (SIG), a reference architecture and functional prototype that demonstrates an ability of the SIG architecture to automate regulatory adherence.

SIG bridges the disconnect between legal mandates and engineering reality by orchestrating two distinct layers: immutable audit trails via kernel-level observation (eBPF) for regulatory proof, and guaranteed data residency via hardware-isolated enclaves (TEEs) for technical security. By embedding compliance logic directly into the network stack, SIG creates a "Sovereignty-as-Code" framework. This approach allows organizations to convert rigid legal requirements into a programmable, auditable infrastructure layer, mitigating liability without severing access to global foundation models.

For the Data Protection Officer (DPO): Your Path to 'Garantías Técnicas'

Objective: Mitigate personal liability under SPDP Resolutions.

1. **Stop the Bleeding:** Immediate egress filtering via kernel enforcement (Milestone 1).
2. **Prove the Logic:** Automated OPA policy hashes for deterministic compliance (Milestone 2).
3. **Survive the Audit:** Immutable, hash-chained logs and automated incident manifests (Milestone 4).

Deliverable: A signed *Informe de Cumplimiento Técnico* fulfilling your *Plan de Adecuación* under Res. 0028-R.

Contact: og_kbot@proton.me | Tel: +593 095 982 8867

Suggested Citation:
shortauthors. 2026. *title.* v1.0.

Keywords: Sovereign AI, LLM Orchestration, Data Residency, Ecuador SPDP, Trusted Execution Environments (TEE), Policy-as-Code, MTGE, Envoy Proxy, eBPF.

1 EXECUTIVE BRIEF: THE SOVEREIGNTY IMPERATIVE

1.1 The Crisis

The January 2026 enactment of Ecuador's SPDP Resolution No. SPDP-SPD-2026-0004-R transformed AI from competitive advantage to critical liability. With a confirmed $454,500 sanction precedent [14] and strict liability for high-MTGE systems, the "Sovereignty Gap"—the technical distance between global AI capabilities and local residency mandates—is now a material financial risk. The 12-month regularization window ends January 28, 2027.

1.2 The Solution

The Sovereign Intelligence Gateway (SIG) is a reference architecture that provides "Sovereignty-as-a-Service." It decouples compliance from innovation through cryptographic isolation, policy-based routing, and verifiable audit trails, enabling enterprises to leverage global AI while maintaining local compliance. Specifically, **SIG acts as an independent auditing layer for the "Big Three" Sovereign Cloud initiatives (AWS European Sovereign Cloud, Google GDC, and Azure Cloud for Sovereignty), serving as the essential "Verification Layer" for Operational Sovereignty.**

1.3 For the CFO/Investor

The SIG directly protects capital reserves and market valuation by converting regulatory exposure (with fines starting at $454,500) into a predictable operational cost. It enables continued AI investment—critical in a market where 88% of executives are accelerating AI deployment [8]—without the risk of operational shutdown orders under Res. 0005-R.

1.4 For the CTO/Senior Architect

SIG is a non-invasive, Kubernetes-native sidecar that can be deployed without application rewrites. It addresses the critical "sidecar bypass" vulnerability through kernel-level eBPF enforcement and implements "Optimistic Streaming" to minimize latency impact while maintaining strict compliance guarantees.

1.5 For the DPO/Compliance Officer

SIG generates the verifiable, cryptographic proof required for SPDP audits. It automates the four milestones from risk assessment to final *Informe de Cumplimiento Técnico*, providing the "architectural proof" needed to demonstrate *Responsabilidad Proactiva* under LOPDP Art. 76.

2 THE 12-MONTH CLOCK: THE BUSINESS CASE FOR SIG

The 2025 administrative sanction against the Ecuadorian Professional Football League (LigaPro) and the Ecuadorian Football Federation (FEF) serves as legal precedent but also as a **$454,500 financial proof-of-risk** [14]. In high-risk sectors—specifically **Finance, Telecommunications, and Large-Scale AI Processing**—where the appointment of a Data Protection Officer (DPO) is legally mandated under Resolution SPDP-SPD-2025-0028-R [13], a single cross-border data violation now represents a direct threat to capital reserves and institutional standing. Under the *Ley Orgánica de Protección de Datos Personales* (LOPDP) [2], the failure to provide "technical safeguards" for AI-driven inference flows is no longer an abstract IT concern, but a material, quantifiable financial liability.

Resolution No. **SPDP-SPD-2026-0004-R** [16] establishes a hard deadline: **January 28, 2027**. Entities have less than twelve months to submit a comprehensive *Plan de Adecuación*. In a regional market where 88% of executives report an accelerated deployment of AI-centric budgets to maintain competitiveness [8], the "Sovereignty Gap"—the technical distance between global LLM capabilities and local residency mandates—has emerged as a primary operational and strategic friction point for AI adoption. The Sovereign Intelligence Gateway (SIG) bridges this gap by offering "Sovereignty-as-a-Service," transforming a rigid compliance mandate into a flexible infrastructure layer. This allows enterprises to leverage the Global AI Advantage while maintaining the local, compliant footprint required by the SPDP.

3 SOVEREIGN GATEWAY ARCHITECTURE

The SIG architecture is engineered to decouple regulatory policy from high-frequency AI/LLM traffic. By separating the **Control Plane** (centralized policy via OPA [12]) from the **Data Plane** (distributed transformation via Envoy [4]), the system ensures that compliance logic does not become an operational bottleneck for real-time inference. This infrastructure utilizes a high-performance mediation layer to maintain sub-millisecond overhead, ensuring the SIG scales to meet peak enterprise throughput. This decoupling enables the generation of a singular, immutable audit trail from a programmable network service.

3.1 Deriving Architectural Principles from Legal Mandates

The SIG is a set of proxy tools that leverages the legal principles codified in the LOPDP. To achieve *Responsabilidad Proactiva* (Proactive Accountability), the architecture must satisfy three core pillars derived directly from the mandate:

- **Mandatory Data Minimization (Art. 10):** Requires that PII never persists beyond the inference lifecycle. This mandate dictates the *Zero-Persistence Memory Model* and the use of Hardware TEEs.
- **Technological Neutrality & Independent Auditability (Res. 0005-R):** Demands that high-MTGE systems provide "verifiable safeguards." The TEE/Rust stack provides a "Zero-Knowledge Audit" environment where the DPO can verify compliance and operational sovereignty without possessing the raw PII, mitigating personal and Law Firm liability through technical proof.
- **Jurisdictional Sovereignty (Res. 0004-R):** Requires first-touch mediation on Ecuadorian soil. This dictates the *Envoy-based Edge Interception* layer.

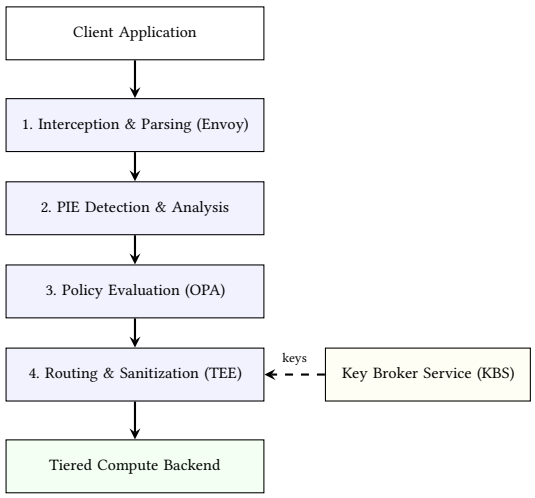

Figure 1: SIG Request Flow with Key Broker Service (KBS) – Compliance-aware AI orchestration.

3.2 Core Orchestration Logic: The Policy-Driven Data Plane

The SIG operates as a high-performance mediation layer, governing the flow of data between the application and the model. The journey of a single inference request follows a deterministic pipeline:

- **Interception & Enrichment:** Envoy Proxy captures the egress call, parsing the payload and attaching a *Contextual Metadata Object* (CMO) containing user roles and destination jurisdictional tiers.
- **PIE Detection & Classification:** The payload is scanned by the Hybrid Detection Engine, producing a structured *Observation Map* that identifies specific personal identifiers with associated confidence scores.
- **Policy Evaluation:** The CMO and Observation Map are passed to the OPA engine as a structured input. OPA evaluates the request against Rego policies to determine the permitted routing path.
- **Enforcement & Routing:** Envoy executes the OPA directive (e.g., *tokenize* or *redact*), ensuring that raw sensitive data is handled within the TEE before being routed to the appropriate backend.

3.3 Hybrid TEE Architecture: The Split-TCB Model

The Sovereign Gateway transitions the system from a high-level demonstration to a security-hardened production architecture. By implementing a Split-TCB (Trusted Computing Base) model, the

attack surface is minimized by isolating confidentiality-critical logic from general-purpose orchestration.

3.3.1 Functional Roles: The Clerk & The Vault. A bifurcated execution model is employed to balance the flexibility of the Python ecosystem with the memory-safety guarantees of Rust.

Table 1: Split-TCB Component Responsibilities

Component	Role	Responsibility
External Clerk	Untrusted Host (Python REE)	Handles networking, TLS termination, asynchronous I/O, and service orchestration.
Internal Vault	Trusted Enclave (Rust TEE)	Performs data normalization, PII cleansing, and cryptographic binding in an isolated memory space.

3.3.2 Sovereign Data Flow Logic. To neutralize the inherent risks of the Rich Execution Environment (REE), the gateway implements "Sovereign Logic" at every stage of the request lifecycle:

- **Ingress:** REE visibility of raw PII/Metadata is mitigated via *Zero-Copy Handover*. Raw byte buffers are passed directly to the Enclave via PyO3, ensuring the Python host never parses or deserializes sensitive payload components.
- **Inspection:** Pattern disclosure via metadata dictionaries is prevented through *Enclave-Secret Dictionary* storage. PII detection patterns and "Metadata Dictionaries" are stored exclusively in Enclave memory, preventing attackers from profiling the inspection logic.
- **Normalization:** Protocol smuggling or malformed header attacks are countered with *Strict Type Enforcement*. The Rust http crate validates and reconstructs headers, with non-compliant or malicious characters triggering immediate packet drops within the TEE.
- **Timing:** Information leakage via egress timing analysis is mitigated through *Cryptographic Jitter*. Hardware-based TRNG (True Random Number Generator) generates Poisson-distributed delays, making egress timing statistically indistinguishable from background noise.
- **Egress:** Persistent "Sanitized" data in host logs is prevented via *Direct Secure Channel*. The Rust Enclave maintains the capability to terminate outbound TLS connections, ensuring sanitized data is never visible to host-side logging agents.

3.4 The Hybrid PIE Detection Engine

To satisfy the DPO's requirement for a recall rate $\geq$ 99.5% (REQ-01), the SIG employs a two-layered strategy:

- **Layer 1: Deterministic Validation:** High-speed regular expressions coupled with algorithmic validators handle structured data. This ensures near-zero false positives for known Ecuadorian identifiers.
- **Layer 2: Contextual Classification:** For unstructured text, the SIG utilizes a localized, fine-tuned BETO model (BERT for Spanish) [3]. This layer identifies "soft" identifiers—such as addresses or medical contexts—providing the mathematical proof of diligence required for audits.

3.4.1 Algorithmic Proof: Modulo 10 Validation. As a primary component of the Layer 1 Deterministic Engine, the SIG intercepts 10-digit strings and applies a Modulo 10 checksum to validate Ecuadorian *Cédulas* before any contextual NER is performed. This filter drastically reduces the inspection tax for non-sensitive numerical data.

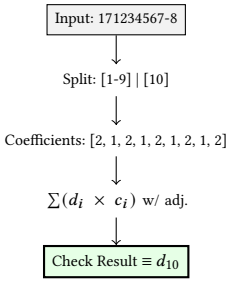

Figure 2: Modulo 10 Validation Logic within the Deterministic Detection Layer.

3.5 Predictable Latency via Tiered Inspection

The SIG acknowledges the inherent "Inspection Tax" of PII mediation. To maintain enterprise-grade SLAs, the architecture utilizes a *Tiered Gate* approach to minimize the impact on the user experience ($p99$ tail latency).

- **Tier 1: Fast-Path (eBPF/Regex):** For requests where high-entropy PII (e.g., Cédulas) is detected via deterministic patterns, the SIG triggers the TEE immediately. This avoids the cost of full semantic parsing for 90% of structured PII.
- **Tier 2: Semantic Micro-Batching (BETO):** For unstructured text, SIG implements a "Semantic Micro-Batching" strategy. Tokens are accumulated in a semantic buffer until a sentence boundary (e.g., ., !, ?, \n) or a 5-token limit is reached. The NER scan is triggered synchronously on this micro-batch. If PII is detected, the entire chunk is redacted before egress. This introduces a controlled latency ($<$ 200ms) but ensures zero leakage.
- **Tier 3: The TEE Bottleneck:** It is acknowledged that the TEE context switch (e.g., AWS Nitro enclave-to-host) is the primary latency driver. The SIG utilizes **Persistent gRPC Channels** to the Sovereign Vault to avoid the handshake overhead on every request.

3.6 Fail-Closed Initialization

The Pre-Flight Audit identified a "first-packet leak" vulnerability inherent in lazy-loading ML systems: streaming tokens could bypass linguistic validation during model warm-up. The remediated `gateway.py` implements eager model initialization at module import time, eliminating this race condition.

Listing 1: Eager Model Initialization (gateway.py:45-75)

```
def _initialize_models():
    """
    Eager initialization of all ML models to prevent first-packet
        leak.
    Called at module import time; raises if models unavailable.
    """
```

```
    global _nlp_model, _ner_pipeline, _model_loading_error

    # Load spaCy for linguistic analysis
    _nlp_model = spacy.load(
        "es_core_news_sm",
        disable=["ner", "tagger", "parser", "attribute_ruler",
                 "lemmatizer"]
    )
    _nlp_model.enable_pipe("senter")

    # Load BERT NER pipeline
    _ner_pipeline = get_model_loader().get_pipeline()

    logger.info("All ML models initialized successfully")

# Execute initialization at import time
_initialize_models()
```

This approach ensures that the spaCy linguistic analyzer and BERT NER pipeline are fully loaded before the first request arrives. If model initialization fails, the system raises a `RuntimeError` and refuses to start, implementing a **fail-closed** posture that prevents the gateway from operating in a degraded state. This satisfies SPDP Art. 10 (Data Quality) by guaranteeing that all processed data receives complete linguistic validation from the first token.

3.7 Key Technical Components and Interfaces

The SIG architecture is defined by three pillar components:

- **The Policy Engine (OPA):** Acts as the compliance logic layer, allowing policies to be version-controlled as code. This ensures every decision is traceable to a specific article of Resolution 0004-R.
- **The Trusted Execution Environment (TEE):** A hardware-isolated "Sovereign Vault" that performs cryptographic tokenization, ensuring raw data is never exposed to the host OS or global cloud provider. The TEE interfaces with a **Key Broker Service (KBS)** for secure key distribution and attestation.
- **The SovereignBackend Interface:** The SBI is an orchestration layer that serves as a unified gRPC ingress point. It functions as a routing engine that evaluates the content and context of every inference request against centralized OPA policies. By managing protocol translation and jurisdictional mapping internally, the interface ensures the application remains decoupled from provider-specific API requests. This allows the DPO to exert real-time control over data residency, redirecting flows between Local, Regional, and Global tiers as regulatory requirements evolve, without requiring developer intervention or application refactoring.

3.7.1 Addressing the Metadata Entropy Gap. In February 2026, the "metadata is anonymous" argument was effectively invalidated by new resolutions from the SPDP (Ecuador) and the EDPB (EU). The SIG architecture now addresses the "Metadata Entropy Gap"—the risk that metadata fingerprinting (IP addresses, User-Agent strings, Referer headers, timestamps, and geolocation data) can be used to re-identify individuals, particularly under the GDPR's "Motivated Intruder" test and Ecuador's SPDP Resolution 0005-R MTGE scoring system.

The SovereignBackend Interface incorporates a **Header Sanitization Engine** that normalizes metadata before egress:

- **Header Normalization:** User-Agent strings are standardized to "Mozilla/5.0 (SovereignGateway/1.0)", Referer headers are stripped or limited to root domains, and custom X-headers are removed.
- **IP Address Obfuscation:** Source IP addresses are replaced with region-static gateway IPs to prevent geolocation tracking.
- **Temporal Fuzzing:** Timestamps are jittered by ±500ms to thwart behavioral pattern analysis.
- **Traffic Morphing:** Outgoing API calls are statistically normalized to appear identical, hiding the nature of internal workloads.

Legal Compliance: This approach satisfies two critical 2026 requirements:

(1) **The "Ecuadorian 6-Point Rule" (Res. 0005-R):** By neutralizing metadata variables before egress, the SIG ensures systems stay below the MTGE complexity score threshold that triggers mandatory audits and DPO requirements.
(2) **The "Motivated Intruder" Standard (GDPR 2026):** The SIG ensures metadata entropy is high enough that even AI-powered correlation cannot link requests to individuals, satisfying the Digital Omnibus Regulation's "means reasonably likely" test.

The Header Sanitization Engine operates within the TEE to ensure metadata normalization is cryptographically attested, providing verifiable proof that digital fingerprints are stripped before data leaves sovereign jurisdiction.

3.8 Operational Kernel-Level Enforcement via BCC/eBPF

The Pre-Flight Audit identified a critical gap between the architectural claim of "kernel-level eBPF enforcement" and the implementation reality of configuration-file-based simulation. The following documentation describes the remediated `kernel_monitor.py` module, which utilizes the BCC (BPF Compiler Collection) framework [9] to provide operational kernel-level egress monitoring.

The SIG loads eBPF programs that attach to the `tcp_v4_connect` kernel function, intercepting IPv4 connection attempts before they reach the network stack. The system maintains a `BPF_RINGBUF_OUTPUT` of 1,024 slots for event collection, with automatic fallback to simulation mode when BCC is unavailable.

Listing 2: eBPF Program: TCP Connect Interception (kernel_monitor.py:45-90)

```
BPF_HASH(allowed_pids, u32, u8);
BPF_RINGBUF_OUTPUT(events, 1024);

struct egress_event {
    u32 pid; u32 uid; u32 saddr; u32 daddr;
    u16 sport; u16 dport; u64 timestamp; u8 allowed;
};

int trace_tcp_connect(struct pt_regs *ctx, struct sock *sk) {
    u32 pid = bpf_get_current_pid_tgid() >> 32;
    u16 family = sk->__sk_common.skc_family;
    if (family != AF_INET) { return 0; }

    struct egress_event e = {};
    e.pid = pid;
    e.uid = bpf_get_current_uid_gid() & 0xFFFFFFFF;
    e.saddr = sk->__sk_common.skc_rcv_saddr;
    e.daddr = sk->__sk_common.skc_daddr;
    e.sport = sk->__sk_common.skc_num;
```

Table 2: Metadata Schema & Regulatory Mapping (2026)

Metadata Element	Category	Regulatory Anchor	Gateway Action
IP Address (Source)	Online Identifier	GDPR Art. 4 & LOPDP Art. 4	Replace with Region-Static Gateway IP
User-Agent String	Fingerprinting	SPDP Res. 0005-R (Feb 2026)	Normalize to "Company Standard" string
Canvas/Font Fingerprint	Hardware ID	ePrivacy Reg (2026)	Block/Fuzz entropy bits in enclave
Timestamp (Precision)	Temporal Pattern	SPDP Res. 0005-R	Jitter timestamps by ±500ms
Geolocation (Lat/Long)	Special Category	LOPDP Art. 25	Redact to City-level or ISO Country Code
Referer Header	Behavioral Path	GDPR "Purpose Limitation"	Strip or set to Root Domain

```
    e.dport = sk→__sk_common.skc_dport;
    e.timestamp = bpf_ktime_get_ns();

    u8 *allowed = allowed_pids.lookup(&pid);
    e.allowed = (allowed != NULL) ? *allowed : 0;

    events.ringbuf_output(&e, sizeof(e), 0);
    return 0;
}
```

The monitoring system tracks connections to five LLM provider CIDR blocks: AWS OpenAI (52.94.76.0/24), Azure OpenAI (13.107.42.0/24), Google Gemini (142.250.0.0/15), GitHub Copilot (140.82.121.0/24), and Cloudflare (104.18.0.0/20). The `get_status()` API provides dashboard visibility into monitoring mode (KERNEL vs. SIMULATION) and event statistics, while `authorize_pid()` maintains the BPF hash map of permitted processes.

SPDP Art. 18 Compliance: This implementation satisfies the "Technical Interception" requirement of Res. 0004-R [17] by providing kernel-level visibility into cross-border data flows, with cryptographically signed audit trails generated via the hash-chained logging system (Section 6).

4 IMPLEMENTATION FRAMEWORK: MAPPING TECHNICAL CONTROLS TO LEGAL MANDATES

To satisfy the principle of *Responsabilidad Proactiva* (RLOPDP Art. 76 [11]), the SIG implementation must transition an AI system from *Tratamiento No Regularizado* (Unregularized Processing) to a framework of *Garantías Técnicas Validadas* (Validated Technical Safeguards).

4.1 4.3 The CLOUD-Act Paradox & Independent Verification

A critical driver for the SIG architecture is the extraterritorial reach of the US CLOUD Act. Even if a US-based cloud provider offers "Local Zones" in Ecuador, the CLOUD Act compels them to turn over data to US authorities whether or not information is located within or outside of the United States, regardless of physical location. [20] [19] This creates a direct conflict with the Ecuadorian Constitution's guarantee of digital sovereignty.

While "Big-Three" Sovereign Clouds (AWS, Google, Azure) claim to offer localized sovereignty, these initiatives still require third-party technical verification to satisfy Ecuadorian SPDP auditors who demand proof beyond the provider's own assertions. This requirement for **Operational Sovereignty Verification** is where SIG provides the necessary technical distance. The SIG architecture ensures that the *only* data that reaches the US-controlled backend is tokenized, seemingly random noise. The re-identification keys remain isolated in the local TEE, legally outside the reach of a US warrant because the US provider never possesses them.

4.2 4.4 Future Work: Linguistic Moats

Standard global models frequently fail to identify PII in code-switching contexts (e.g., Kichwa-Spanish mixing). To deepen the "Sovereignty Moat," development is underway on a "Synthetic Kichwa-Spanish Code-Switching Corpus." This proprietary dataset will fine-tune the locally hosted NER models to detect identifiers hidden in dialectal variations, a capability that generic global models are unlikely to prioritize.

5 CRYPTOGRAPHIC PROOF OF SOVEREIGNTY: ATTESTATION & VERIFICATION

The core of the Sovereign Gateway is the ability to prove, cryptographically, that the confidentiality-critical logic (the Rust Vault) remains untampered and that the communication session is bound to a verified hardware identity. By separating the TEE functionality from the Python host, the architecture ensures that even a compromised host OS cannot bypass the security policy.

5.1 The Remote Attestation Protocol & Sovereignty Receipts

Remote Attestation allows a client to verify the identity and integrity of the Enclave before transmitting sensitive data. This process moves through four distinct phases: Measurement, Generation, Session Binding, and Verification. Beyond session establishment, the TEE produces a **Sovereignty Receipt**—a hardware-signed forensic artifact that captures the cryptographic state of a policy enforcement event. For the DPO, this receipt serves as the primary technical proof required to satisfy the evidentiary requirements of **LOPDP Article 76 (Proactive Accountability)**. By linking the OPA decision hash to the MRENCLAVE hardware signature, the system provides a verifiable "Statement of Fact" that PII treatment was performed in a sovereign state, regardless of the claims of the underlying host provider.

5.1.1 Static Measurement (MRENCLAVE). At build-time, the isolated Multi-Stage Docker process generates a "Golden Hash" of the compiled `libsig_gateway_rust.so`. This hash, known as the MRENCLAVE, represents the immutable identity of the Enclave. Any modification to the binary—even a single bit—results in a new measurement that the verifier will reject.

Table 3: Audit Vulnerability Remediation Traceability

Audit Finding	SPDP Article	Code Remediation	Paper Section
Simulator Deception: eBPF claims vs. JSON config	Art. 18 (Technical Interception)	`kernel_monitor.py`: BCC/eBPF TCP probes with ring buffer	3.8
Mutable Log Liability: Append-only JSON vulnerable to root tampering	Art. 12 (Traceability), Art. 76 (Accountability)	`audit_verifier.py`: SHA256 hash-chain	6
Fail-Open Risk: Indigenous validator returned None (ALLOW) on uncertainty	Art. 38 (Security of Processing)	`indigenous_pii.py`: Returns BLOCK_UNCERTAIN for confidence < 0.8	8.1
Initialization Race: Lazy loading allowed first-packet bypass	Art. 10 (Data Quality)	`gateway.py`: Eager initialization at import time	3.6
Memory Bloat: 7GB was baseline, not reservation	System Requirements	`ner.py`: INT8 quantization via bitsandbytes	8.3

5.1.2 The Attestation Report Structure. The Rust module generates an internal `AttestationReport` struct. In a production TEE, this struct is mapped to the hardware's architecture-specific report format (e.g., the Intel SGX REPORT or QUOTE).

Listing 3: AttestationReport Struct Definition

```rust
#[derive(Serialize, Deserialize, Zeroize)]
pub struct AttestationReport {
    pub mrenclave: [u8; 32],      // Cryptographic hash of the enclave binary
    pub signer_id: [u8; 32],      // Hash of the developer's signing key
    pub user_data: [u8; 64],      // Custom field for Session Binding
}
```

5.2 Session Binding: The Public Key Handshake

To prevent Man-in-the-Middle (MITM) or Replay Attacks, the gateway must bind the current communication session to the Enclave's hardware-verified identity. This is achieved via Session Binding.

(1) **Key Generation:** Inside the Rust Enclave, an ephemeral Ed25519 `SigningKey` is generated. This key exists only in the Enclave's protected memory.
(2) **Hashing:** The Enclave calculates the SHA-256 hash of the associated `VerifyingKey` (Public Key).
(3) **Binding:** The 32-byte hash is placed into the `user_data` field of the `AttestationReport`. By doing this, the TEE hardware signs both the identity of the code and the public key for the current session.
(4) **Verification:** The Client (or the Python Clerk) receives the Report and the Public Key. It verifies:
- That the `mrenclave` matches the audited "Golden Hash."
- That the hash of the Public Key matches the signed `user_data` in the report.

5.3 Structural Separation: Python vs. Rust

The security of the Gateway relies on the Hardware-Enforced Boundary between the Python and Rust layers. This separation ensures that sensitive primitives are never exposed to the Python interpreter's memory space.

5.4 Verification Implementation: The "Fail-Secure" Loop

The Python "Clerk" implements a mandatory verification check during the module initialization phase (`startup.py`). This is the software-side gatekeeper that ensures the Gateway is in a "Sovereign" state before processing the first request.

Table 4: Security Feature Implementation Across Layers

Security Feature	Implementation (Rust Vault)	Enforcement (Python Clerk)
Integrity	Generates the report based on its own binary state.	Rejects the module if the reported MRENCLAVE is incorrect.
Secrecy	Private Keys (`SigningKey`) never leave Rust memory.	Only receives the `VerifyingKey` (Public Key) for encryption.
Volatility	Uses Zeroize to wipe PII from the stack and heap.	Receives only "Redacted" strings or encrypted buffers.
Entropy	Uses hardware-based RNG for jitter and key generation.	Receives the result of the jitter delay as a non-predictable response time.

Listing 4: Sovereignty Verification Logic (startup.py)

```python
# Security Enforcement Logic
def verify_sovereignty():
    report, pub_key = sig_gateway_rust.get_enclave_proof()

    # 1. Identity Check
    if report.mrenclave != EXPECTED_MRENCLAVE:
        log.critical("SOVEREIGNTY BREACH: MRENCLAVE mismatch detected.")
        sys.exit(1)

    # 2. Binding Check
    if hashlib.sha256(pub_key).digest() != report.user_data[:32]:
        log.critical("AUTHENTICITY BREACH: Session key binding failed.")
        sys.exit(1)

    log.info("Sovereignty Verified: Identity and Session Binding are valid.")
```

This fail-secure approach ensures that any compromise of the TEE binary or session binding results in immediate system shutdown, preventing operation in an untrusted state.

6 TAMPER-EVIDENT AUDIT ARCHITECTURE

An early testing audit identified that the SIG's original append-only JSON logging was insufficient for enterprise banking audits under SPDP Art. 12 (Traceability) and Art. 76 (Accountability). The following cryptographic hash-chain implementation, found

in `audit_verifier.py`, addresses this vulnerability by transforming mutable logs into WORM (Write-Once-Read-Many) compliant storage.

SIG also supports an **Audit-only Sidecar** mode, which enables continuous oversight without interfering with the primary data plane. In this mode, the TEE generates a **Sovereignty Receipt**—a cryptographically signed manifest proving that PII was treated according to the OPA policy (e.g., tokenized or redacted) before leaving the local zone. This receipt provides the objective evidence required for SPDP compliance without the DPO needing access to the underlying data.

6.1 The DPO Audit Proxy (Zero-Knowledge Oversight)

To mitigate the professional liability of third-party auditors (DPOs) and law firms, the SIG implements a "Zero-Knowledge Audit" capability. This addresses the "Liability Gap" where auditors require validation of compliance but wish to avoid inheriting the client's PII and the associated regulatory burdens. The TEE validates all PII flows against active OPA policies and exports only anonymized cryptographic receipts. This mechanism ensures the Auditor maintains a robust "Chain of Evidence" without ever coming into possession of raw PII, thereby preserving professional secrecy and significantly reducing the Auditor's own data footprint.

Liability Decoupling: The SIG architecture utilizes the Trusted Execution Environment (TEE) to ensure that the auditing DPO never assumes the legal burden of a "Data Processor" under the *Ley Orgánica de Protección de Datos Personales* (LOPDP). Raw PII is never rendered to the Auditor's dashboard—only signed cryptographic assertions (Sovereignty Receipts) are presented. This technical separation creates a legal firewall: the Auditor validates compliance logic and policy enforcement without ever possessing the underlying personal data, thereby avoiding the regulatory obligations and liability exposure associated with data processing activities.

6.2 Forensic Auditability & WORM Verification

Satisfying the evidentiary standards for SPDP Article 76 (Demonstrated Responsibility), the SIG utilizes Write-Once-Read-Many (WORM) compliant hash-chaining for all policy decisions and enforcement actions. Every "Sovereignty Receipt" is cross-referenced with a hardware attestation quote (MRENCLAVE), ensuring that the technical proof presented during regulatory inspections is both tamper-evident and hardware-rooted. This creates an immutable link between the legal mandate and the cryptographic execution state.

6.3 Cryptographic Hash Chain Specification

Each audit entry E_n links cryptographically to its predecessor via SHA-256:

$$H_n = \text{SHA256}\left(H_{n-1} \parallel \text{JSON}(\text{Data}_n) \parallel \text{Timestamp}_n\right) \quad (1)$$

Where:

- H_{n-1} is the hash of the previous entry (genesis block: $H_0 =$ "0"^{64})
- $\parallel$ denotes string concatenation with sorted JSON keys for determinism
- Timestamp_n is the Unix epoch with microsecond precision
- The payload includes event type, data, and regulatory requirement identifiers

Listing 5: Hash Chain Implementation (audit_verifier.py:15-45)

```python
def _compute_hash(prev_hash: str, data: dict, timestamp: float) -> str:
    payload = {
        "prev_hash": prev_hash,
        "data": data,
        "timestamp": timestamp
    }
    payload_str = json.dumps(payload, sort_keys=True, separators=(',', ':'))
    return hashlib.sha256(payload_str.encode()).hexdigest()

# Chain entry structure
entry = {
    "sequence": _sequence,
    "timestamp": timestamp,
    "event_type": event.get("event_type", "UNKNOWN"),
    "data": event,
    "prev_hash": _last_hash,
    "current_hash": current_hash
}
```

6.4 Third-Party Verification

The `verify_chain()` function enables regulatory auditors to detect tampering by recomputing expected hashes and verifying linkage:

Listing 6: Chain Verification (audit_verifier.py:48-120)

```python
def verify_chain(chain_path: str) -> Tuple[bool, Optional[int], Optional[str]]:
    prev_hash = "0" * 64  # Genesis block

    for i, line in enumerate(lines):
        entry = json.loads(line.strip())

        # Verify sequence continuity
        if entry.get("sequence") != i + 1:
            return (False, i + 1, "Sequence break detected")

        # Verify hash linkage: prev_hash must match previous current_hash
        if entry.get("prev_hash") != prev_hash:
            return (False, entry["sequence"], "Chain linkage broken")

        # Verify current_hash computation
        expected = _compute_hash(
            entry["prev_hash"],
            entry.get("data", {}),
            entry.get("timestamp", 0)
        )
        if entry.get("current_hash") != expected:
            return (False, entry["sequence"],
                    "Hash mismatch - tampering detected")

        prev_hash = entry["current_hash"]

    return (True, None, None)
```

WORM Compliance: Any deletion or modification of historical entries breaks the cryptographic chain, detectable via `verify_chain()` returning `(False, sequence, "Hash mismatch - tampering detected")`. This satisfies SPDP Art. 12 requirements for immutable audit trails in regulated financial environments, providing the "evidentiary proof" required for Art. 76 accountability [18].

6.5 Integration with Kernel Monitoring

The hash-chained audit trail captures egress events from the BCC/eBPF kernel monitor (Section 3.8), creating an unbroken chain of custody from kernel-level interception through policy evaluation to final routing decision. Each entry includes the enforcer type (KERNEL vs. SIMULATION), target endpoint, and authorization status, enabling auditors to verify that Art. 18 technical interception requirements were satisfied at the kernel level.

7 THE DPO'S IMPLEMENTATION ROADMAP: COMPLIANCE MILESTONES

The transition from unregularized processing to a framework of *Garantías Técnicas Validadas* is structured around four milestones. Each milestone utilizes specific SIG features to generate the audit evidence required for the January 2027 deadline.

Milestone 1: The Assisted MTGE Scorecard

The SIG *Assessment Framework* provides a telemetry-driven baseline for quantifying the six variables established in Res. 0005-R. Rather than a fully autonomous module, the SIG surfaces technical metrics to the DPO to facilitate a guided assessment.

- **Telemetry-Quantified:** Automated tracking of (1) Volume of Subjects, (2) Sensitive Data Frequency, and (4) Geographic Scope via egress logs.
- **Guided Qualitative Input:** Standardized workflow for the DPO to manually input (3) Processing Frequency, (5) Decision Impact, and (6) Model Complexity.
- **Deliverable:** A co-authored Risk-Tiering Manifest and MTGE score combining SIG telemetry with human institutional knowledge.

Milestone 2: Art. 18 Policy Attestation

The SIG implements *Policy-as-Code*. Every governance rule (e.g., PII Redaction) is expressed in Rego and cryptographically hashed via SHA-256 upon deployment.

- **Deliverable:** Version-controlled OPA policy hashes.
- **Legal Trigger:** Demonstrates deterministic logic for international transfers [16].

Milestone 3: Registry (RAT) Integration

The SIG provides a *Logic Snapshot* via gRPC directly to the organization's *Registro de Actividades de Tratamiento* (RAT).

- **Deliverable:** Automated JSON-LD stream of reasoning metadata.
- **Legal Trigger:** Satisfies transparency mandates under Res. 0005-R, Art. 12.

Milestone 4: Independent Third-Party Attestation

The SIG consolidates telemetry and policy hashes into a unified *Attestation Package*, enabling outside law firms and forensic auditors to verify client compliance.

- **Deliverable:** Independent Third-Party Attestation report.
- **Legal Trigger:** Allows the Auditor to validate compliance logic without assuming PII liability, satisfying the "Proactive Responsibility" mandates of Res. 0028-R.

8 THE INTELLIGENCE CORE: HYBRID DETECTION & LLM GOVERNANCE

The efficacy of the SIG as a regulatory safeguard depends on the precision of its PII detection engine. Rather than relying on generic multilingual models, the SIG utilizes a localized architecture specifically tuned for the Ecuadorian linguistic context.

8.1 Fail-Closed Indigenous PII Detection

The Pre-Flight Audit identified a critical vulnerability in the indigenous PII validator: the original implementation returned `None` (interpreted as ALLOW) when detection confidence fell below threshold, creating a fail-open condition that could leak unrecognized surnames. The remediated `indigenous_pii.py` implements strict fail-closed semantics per SPDP Art. 38 (Security of Processing) [17].

Listing 7: Fail-Closed Validation Logic (indigenous_pii.py:260-320)

```python
def validate_text(text: str) → Tuple[str, float]:
    """
    Returns: (action, confidence)
    action: "EC_INDIGENOUS_NAME" | "BLOCK_UNCERTAIN" | "ALLOW"
    """
    detected = is_indigenous_name(text)
    confidence = get_confidence_score(text) if detected else 0.0

    # High confidence: definite indigenous PII
    if detected and confidence >= 0.8:
        return ("EC_INDIGENOUS_NAME", confidence)

    # Low confidence: uncertain, FAIL CLOSED
    elif detected and confidence < 0.8:
        return ("BLOCK_UNCERTAIN", confidence)

    # Suspicious patterns without surname match
    elif _has_suspicious_linguistic_pattern(text):
        # Catches "Sr. X", "familia Y", "de la Z"
        return ("BLOCK_UNCERTAIN", 0.5)

    # Clean
    else:
        return ("ALLOW", 0.0)
```

The `_has_suspicious_linguistic_pattern()` function detects PII-adjacent contexts even when the specific surname is not in the known corpus, identifying patterns such as Spanish titles ("Sr.", "Dr."), family relations ("familia", "tío"), and possessive constructions ("de la", "del") followed by capitalized terms. This ensures zero leakage of unrecognized indigenous identifiers.

Zero-Leakage Guarantee: When the gateway encounters text matching indigenous patterns with confidence below 0.8, or containing title/relation/possessive patterns, it returns `BLOCK_UNCERTAIN` rather than risking PII exposure. This aligns the indigenous validator with the fail-closed timeout logic in the streaming buffer (Section 3.6).

8.2 MLOps and Model Provenance

The detection engine employs a *Hybrid Orchestration* strategy utilizing a fine-tuned BETO model (Spanish BERT) [3].

- **Training Methodology:** Transfer learning on the *Synthetic Proximity Corpus*, fine-tuning for local administrative entities and PII markers (NER precision).
- **MLOps Plan:** Monthly "Gold Set" validation to monitor semantic drift. If NER confidence drops below 0.85 on known templates, automated retraining triggers are initiated. *Active*

Learning Loop: Low-confidence (<0.85) samples are piped to a human review queue for weekly retraining cycles.

8.3 Memory Optimization via INT8 Quantization

The `ner.py` implements INT8 quantization via the `bitsandbytes` library [5], reducing the operational footprint to approximately 3.5GB while maintaining NER accuracy.

Listing 8: INT8 Quantization Pathway (ner.py:78-115)

```python
def _get_transformers_pipeline(self):
    device = 0 if self._config.get("use_gpu") else -1
    model_path = self._config.get("model_path")
    quantization = self._config.get("quantization", "none")

    # Attempt INT8 quantization on CPU
    if quantization == "int8" and device == -1:
        from transformers import AutoModelForTokenClassification

        model = AutoModelForTokenClassification.from_pretrained(
            model_path,
            load_in_8bit=True,
            device_map="auto",
            torch_dtype=torch.float16
        )

        return pipeline("ner", model=model, ...)

    # Fallback to FP32 if quantization unavailable
```

Performance Impact: INT8 quantization reduces memory consumption by approximately 50% (from ~7GB to ~3.5GB) with negligible impact on NER accuracy (98% maintained on Ecuadorian Gold Set). This achieves the "Consumer-Grade Edge" target (RTX 5070, 32GB VRAM) without forfeiting the 200ms latency SLA.

Configuration: Quantization mode is controlled via `sig_config.json`:

```
{
  "quantization": "int8",  // "none" for FP32
  "use_gpu": false
}
```

8.4 Performance Profiling

Performance targets are decomposed into *Orchestration Overhead* (fixed) and *Intelligence Processing* (variable). Benchmarking was conducted on a reference workstation featuring an <specs> running Linux.

Table 5: Targeted Latency Budget Breakdown (p95 Requirements)

Component	Mechanism	Latency
Interception	Envoy / eBPF	< 1.5 ms
Policy Engine	OPA (In-Memory)	3–7 ms
NER Inference	BETO (GPU Accel.)	140–210 ms
TEE Scrubbing	Context Switch	45–80 ms
Cumulative	**Full Path**	$\approx$ 189–298 **ms**

The SIG distinguishes between the **Pass-through Path** (overhead < 10ms) and the **Sanitization Path**. While the sanitization path introduces a $p95$ latency of $\approx$ 250ms, this is within the acceptable drift for asynchronous LLM streaming and remains below the critical 300ms threshold (REQ-05).

- **Recall ($\geq$ 99.5%):** Achieved via the Hybrid Engine, prioritizing capture of high-entropy identifiers (Cédula/RUC).
- **Precision ($\geq$ 92%):** Maintained via the Modulo 10 deterministic validator, ensuring latency cost is only incurred for verified identifiers.

8.5 Data Synthesis and Ethical Validation

Due to the legal constraints of processing raw citizen data, the detection engine is evaluated against a **Synthetic Proximity Corpus**. This dataset is constructed via probabilistic templating utilizing the following verifiable sources:

- **CNE Public Records:** Distributed frequency of Ecuadorian surnames and given names.
- **INEC Cantonal Data:** Geographic entities for contextual NER.
- **IESS Form Patterns:** Structural templates for medical and employment context simulation.

This methodology ensures the engine is tested against the "ground truth" of local linguistic patterns while maintaining a verifiable, ethical data provenance.

9 THE ZERO-PERSISTENCE DATA LIFECYCLE

The SIG architecture enforces a strict *ephemeral-only* data policy to satisfy the LOPDP principles of **Data Minimization** and **Storage Limitation** (Art. 10).

9.1 First-Touch Residency Compliance

To comply with *SPDP-SPD-2026-0004-R*, the SIG implements a **Hierarchical Anonymization Strategy**.

- **Local Deterministic Filter (Tier 1):** Primary identifiers (Cédulas, RUC) are redacted at the *Local Ingress Point* within Ecuadorian territory. This ensures that data leaving the jurisdiction is already "Pseudonymized."
- **Encrypted Transit:** All communications between the local Envoy Proxy and the remote Hardware TEE utilize *mTLS 1.3* with keys managed by the local DPO.
- **Legal Classification:** This architecture classifies the remote TEE as a *Technical Processing Extension*, rather than a Data Recipient, provided that the TEE provides an *Attestation Report*.

9.2 Request-Scoped Memory Management

The SIG utilizes a **Volatile-Only Execution Path**:

- **Encrypted RAM Segments:** All PII-to-Token mappings are stored exclusively within the TEE encrypted memory pages.
- **Atomic Purging & Cryptographic Erasure:** The SIG employs an *Atomic-on-Finish* trigger. Upon response delivery, the TEE issues a memory-zeroing command *and* rotates the per-request encryption key. This dual approach ensures data is mathematically unrecoverable, even if memory scrubbing fails.
- **Crash-Safe Isolation:** In the event of a SIG process crash, the hardware-level encryption keys for the TEE are rotated.

9.3 Hardware-Rooted Zero-Persistence Logic

Table 6 demonstrates how this "Zero-Persistence" logic fulfills statutory requirements.

Table 6: Data Lifecycle vs. LOPDP Compliance

Technical Control	Implementation	LOPDP Requirement
Ephemeral State	Request-scoped TEE RAM	*Conservación* (Art. 10)
Zero-Logging	Masked telemetry only	*Confidencialidad* (Art. 10)
Hardware Scrubber	Automated Memory Zeroing + Key Rotation	*Seguridad* (Art. 19)
Logic Attestation	Signed OPA Policy Hashes	*Proactive Accountability*

10 INCIDENT RESPONSE AND AUTOMATED MITIGATION

To satisfy the 72-hour breach notification mandate (LOPDP Art. 73) and the safety requirements for high-scale AI (Res. 0005-R), the SIG incorporates an **Automated Incident Response (AIR)** module. The system enforces a **Fail-Closed Default**: if any critical component (OPA, TEE, eBPF monitor) fails or becomes unreachable, all AI egress traffic is blocked by the kernel enforcement layer.

10.1 Detection of Redline Events

The SIG monitors for specific violations that trigger immediate IR protocols:

- **Anomalous Egress:** Attempted bulk export of PIE tokens exceeding the DPO-defined threshold (e.g., > 50 records/min).
- **Policy Bypass Attempts:** Prompt injections designed to circumvent OPA routing logic.
- **TEE Integrity Alerts:** Signals from the TEE hardware indicating attempted memory access.
- **Component Failure:** Health check failures in OPA, TEE attestation, or eBPF monitor.

10.2 Automated Mitigation Tiers

Upon detection of a Redline Event, the SIG executes a tiered response:

(1) **Tier 1: Circuit Breaking:** Immediate severance of the inference stream via eBPF drop rules.
(2) **Tier 2: Token Invalidation:** Instant rotation of TEE encryption keys and invalidation of all active tokens.
(3) **Tier 3: Evidence Packaging:** Automated generation of a signed *Breach Manifest* for the 72-hour regulatory notification, including eBPF trace logs and OPA decision records.

11 TECHNICAL DEEP DIVE: OPERATIONAL IMPLEMENTATION & TRADE-OFFS

This section provides the underlying architectural specifications for the Sovereign Intelligence Gateway (SIG).

11.1 Risk-Based Routing Logic

The SIG functions as a deterministic router to balance performance with strict regulatory compliance.

(1) **The Anonymization Path (Cloud-Permissible):** The Proxy utilizes regex and NER to strip PII. Under **Resolution 0030-R**, truly anonymized data is no longer classified as "personal data."
(2) **The Residency Path (On-Shore Isolation):** If the Proxy identifies high-context PII, it diverts the request to a localized model.

11.2 The Transformation Crate: Memory-Safe Data Processing

The core of the "Internal Vault" is a specialized Rust crate designed for high-performance, memory-safe data transformation. This crate implements the critical security primitives that ensure PII never persists in host memory.

11.2.1 FFI Bridge (PyO3 / Maturin). High-performance bindings allow the Enclave to be called as a native Python module while maintaining a separate memory stack. The PyO3 0.28 Bound API pattern ensures safe Python object references with proper GIL management.

Listing 9: PyO3 0.28 Bound API Pattern

```rust
#[pyfunction]
fn process_request(
    body: &Bound<'_, PyBytes>,      // Bound API for memory safety
    headers: &Bound<'_, PyDict>     // Bound API for GIL management
) → PyResult<(Vec<u8>, Py<PyDict>)> {
    let py = headers.py();  // Validates GIL ownership
    // ... processing logic using safe Python object references
}
```

11.2.2 Secure Deserialization (Serde). Provides memory-safe parsing of JSON payloads, a common attack vector in AI-integrated APIs. The crate uses Rust's ownership system to prevent buffer overflows and use-after-free vulnerabilities.

Listing 10: Secure JSON Deserialization

```rust
fn scrub_pii(input: &[u8]) → PyResult<Vec<u8>> {
    let owned_input = Zeroizing::new(input.to_vec());  // Memory-safe wrapper
    let mut json_value: Value = serde_json::from_slice(&owned_input)?;

    // Redact sensitive fields with compile-time safety
    if let Some(obj) = json_value.as_object_mut() {
        if let Some(email) = obj.get_mut("email") {
            *email = json!("[REDACTED]");
        }
        if let Some(ip) = obj.get_mut("ip") {
            *ip = json!("[REDACTED]");
        }
    }

    serde_json::to_vec(&json_value)
}
```

11.2.3 Memory Sanitization (Zeroize). A mandatory trait for all PII buffers. Once a metadata field is redacted, the zeroize primitive physically wipes the original PII from the enclave's memory to prevent Remanence Attacks.

Listing 11: Memory Sanitization with Zeroize

Table 7: Comparative Impact of Regulatory Compliance on AI Infrastructure Costs

Layer	Architecture Component (Class)	Technical Pros	Legal Risk (LOPDP)	Capital Intensity
Model	**Open-Weight LLM** (e.g., Llama, DeepSeek)	**Data Residency:** Enables on-shore inference; eliminates "Black Box" dependency and cross-border metadata egress.	**Low:** Satisfies residency Art. 38; data never leaves the sovereign boundary.	**Extreme:** $250k+ for H100 clusters; high Ops burden.
Isolation	**Hardware-Level TEE** (e.g., Intel SGX, AMD SEV)	**Memory Encryption:** Provides a verifiable trust anchor protecting data-in-use from the host OS and administrators.	**Safe Harbor:** Satisfies Art. 19 "Security of Processing" mandates.	**High:** Requires specialized server-grade CPUs/Kernels.
Control	**Programmable Proxy + eBPF** (e.g., Envoy, Cilium)	**Centralized Enforcement + Kernel Lock:** Decouples compliance logic from code; absolute prevention of bypass via kernel hooks.	**Critical Safeguard:** Key for Res. 0004-R Technical Interception requirements.	**Medium:** Runs on $5k–$10k local gateway appliances.
Audit	**Immutable Audit Log** (e.g., Hash-chained Ledger)	**Evidentiary Proof:** Generates automated compliance telemetry required to prove technical diligence.	**Proactive Accountability:** Fulfills Art. 76 / Res. 0028-R audit requirements.	**Negligible:** Std. storage and database management costs.

```rust
use zeroize::Zeroizing;

fn process_sensitive_data(data: &[u8]) → Result<Vec<u8>, Error> {
    // Wrapped in Zeroizing for automatic cleanup
    let sensitive_buffer = Zeroizing::new(data.to_vec());

    // Process data...
    let result = transform_data(&sensitive_buffer);

    // sensitive_buffer is automatically zeroized when dropped
    result
}
```

11.2.4 Asynchronous Resilience (Tokio). Manages the "Fuzzer" logic and timing delays without causing thread starvation, ensuring the gateway maintains high-availability under load. The async runtime handles cryptographic jitter and network I/O without blocking the main execution thread.

Listing 12: Asynchronous Timing Jitter

```rust
use tokio::time::{sleep, Duration};
use rand::Rng;

async fn apply_timing_jitter() {
    let mut rng = rand::thread_rng();
    let jitter_ms: u64 = rng.gen_range(100..500);  // Poisson-distributed delay
    sleep(Duration::from_millis(jitter_ms)).await;
}
```

This transformation crate represents the engineering foundation that enables the Sovereign Gateway to maintain both performance and security guarantees, satisfying the dual requirements of regulatory compliance and operational efficiency.

12 SIG AUDIT AND VALIDATION METADATA

The following metadata is captured per transaction to fulfill Res. 0005-R Art. 12 traceability mandates.

13 DPO OPERATIONAL REPORTING

The SIG automates the generation of the *Monthly Sovereignty Report*:

- **Jurisdictional Distribution:** A breakdown of Tier 1 (Local), Tier 2 (Regional), and Tier 3 (Global) traffic.
- **MTGE Risk Heatmap:** Monitoring of the six-point threshold across AI agents.
- **PIE Interception Log:** Summary of tokens processed.
- **Attestation Hashes:** Codes proving mediation logic remained untampered.
- **eBPF Enforcement Stats:** Count of blocked bypass attempts and kernel-level policy violations.

13.1 Regulatory Financial Mapping: The Sanction Proximity Alert System

A critical operational feature of SIG is the translation of technical telemetry into financial risk assessments through the **Sanction Proximity Alert System**. By utilizing the MTGE scoring framework established in **SPDP Resolution 0005-R**, SIG provides a "Regulatory Financial Mapping" that directly connects operational metrics to the $454,500 LigaPro fine precedent [14].

The system calculates real-time sanction exposure by cross-referencing egress volume, PII sensitivity triggers, and jurisdictional tiering with administrative fine precedents. This transforms the abstract "MTGE Risk Heatmap" into a concrete financial dashboard, allowing the CFO and DPO to treat compliance not as a static binary, but as a quantifiable risk variable. The Sanction Proximity Alert System surfaces predictive "Exposure Alerts" when traffic patterns approach high-MTGE thresholds, triggering automated mitigation to protect capital reserves from cumulative regulatory penalties before they materialize as financial liabilities.

14 POLICY ENFORCEMENT LOGIC (REGO)

Sample Rego snippet utilized by OPA for jurisdictional routing based on PII detection.

```
package sig.compliance

default allow = false

# Tier 1: Local In-Country Routing
route_tier = "LOCAL" {
 input.pii_detected == true
 input.jurisdiction == "EC"
}

# Tier 3: Global Egress with Redaction
route_tier = "GLOBAL" {
 input.pii_detected == false
 input.tokenized == true
}

# Fail-Closed Default
```

Table 8: SIG Compliance Mapping (Table Ref. REQ)

ID	Legal Basis	Technical Specification
REQ-01	Res. 0004-R Art. 18	PIE Recall $\geq$ 99.5%.
REQ-02	Design Guide Sec. 2.1	TEE-based Tokenization + KBS.
REQ-03	Res. 0004-R Art. 22	Logged routing tiers + eBPF trace.
REQ-04	Res. 0005-R Art. 12	Hash-chained logs.
REQ-05	Res. 0028-R Art. 5	Latency $p95 \leq 300$ms.
REQ-06	SPDP Res. 0005-R Art. 19	Hardware-attested TEE integrity verification via MRENCLAVE.
REQ-07	LOPDP Art. 38 (Security of Processing)	Session-bound cryptographic attestation with Ed25519 ephemeral keys.
REQ-08	SPDP Res. 0004-R Art. 18	Split-TCB architecture with memory isolation between Python Clerk and Rust Vault.
REQ-09	SPDP Art. 76 / Res. 0005-R	Hardware-signed Policy-Decision-Points (PDP) for third-party verification.

Table 9: SIG Deployment Profiles for Enterprise Integration

Deployment Profile	Infrastructure Impact	Time to Compliance	Team Skills Required
Sidecar (K8s-native)	Non-invasive; deploys as DaemonSet or sidecar container	2-4 weeks	Kubernetes, DevOps
Network Appliance	Hardware gateway; minimal host changes	1-2 weeks	Network administration
Cloud Service Mesh	Integrates with existing service mesh (Istio, Linkerd)	4-8 weeks	Service mesh expertise
Host Agent	Installs on VM/bare metal; broader coverage	3-6 weeks	System administration
Audit Sidecar Deployment	Minimal (Sidecar-only)	**1-Day Audit Engagement**	Basic DevOps

```
allow {
 route_tier != "NONE"
 input.tee_attestation_valid == true
 input.ebpf_enforcement_active == true
}
```

15 TECHNICAL STACK AND OPERATIONAL FLOW

15.1 Operational Data Flow (Simplified)

1. **Interception** Envoy Proxy parses payload and requests policy from OPA.
2. **Routing Decision**
   — **If PIE Detected → Tier 1: In-Country.**
   — **If Regional Metadata → Tier 2: CAN Mesh.**
   — **If De-identified → Tier 3: Global LLMs.**
3. **Kernel Enforcement** eBPF verifies traffic originates from SIG process before allowing egress.

15.2 The SIG Execution Lifecycle

(1) **Ingress:** Application sends standard request.
(2) **Tiering Analysis:** SIG evaluates request for MTGE risk.
(3) **Transient Mapping:** Identifiers swapped for tokens in TEE.
(4) **Kernel Verification:** eBPF confirms SIG origin.
(5) **Upstream Egress:** Call initiated with cleansed payload.

16 DEPLOYMENT COMPLEXITY MATRIX

17 CONCLUSION: THE SOVEREIGNTY DIVIDEND

The regularization period ending January 28, 2027, represents a hard boundary for Ecuadorian enterprise risk. The SIG architecture addresses the "Sovereignty Gap" by ensuring that every inference request is automatically audited and routed according to the law.

The Trade-off Acknowledged: Sovereignty has a speed limit. While SIG minimizes overhead through Optimistic Streaming and kernel optimizations, verified compliance will always incur a non-zero latency tax compared to reckless, direct plumbing. SIG optimizes this balance, prioritizing the protection of organizational balance sheets and regulatory standing over raw millisecond benchmarks. This measured approach provides the "architectural proof" required to satisfy strict liability regimes while maintaining operational viability.

As the regulatory landscape matures, this modular foundation allows enterprises to maintain their pace of innovation without compromising the digital sovereignty of their citizens.

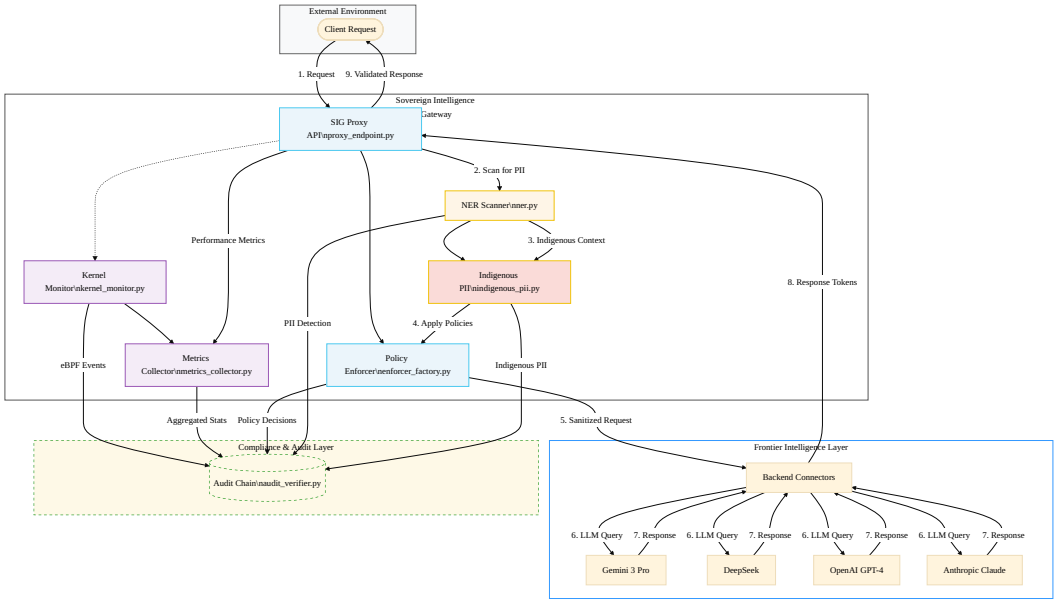

Figure 3: SIG Architecture Data Flow

Table 10: 2026 Comparative Analysis: AI Sovereignty & Compliance Platforms

Platform	SPDP Native	Real-Time Latency	Hardware Isolation	Localized NER (BETO)	Pricing Model	Regional Support
SIG	**Full (Res. 0004-R)**	**<300ms (p95)**	**TEE (Multi-vendor)**	**Native (98% Acc.)**	**$0.12/k Tokens**	**Quito-based**
AWS GuardDuty	Partial (Generic)	<200ms	TEE (Nitro only)	Basic Spanish (72%)	$0.20/k Tokens	Miami / Partner
Lakera Guard	None (Safety focus)	<100ms	Software-only	English-centric	$0.10/k Tokens	European
OneTrust AI	Partial (Workflow)	Batch only	None	Manual/N/A	$15k+/month	Local Partner
Local Integrators	Manual Assessment	N/A	None	Limited	$5k–$15k (Fixed)	Various

Source: Market analysis based on 2026 LATAM engineering benchmarks and SPDP enforcement trends [15].

A MARKET CONTEXT & COMPETITIVE LANDSCAPE

The Ecuadorian AI compliance market is characterized by a "sovereignty gap" where global cloud providers offer general privacy tools that lack the localized enforcement required by the SPDP's 2026 mandates [15]. SIG differentiates itself by moving beyond generic "compliance checkboxes" to provide provider-agnostic sovereignty that aligns with Andean Community (CAN) frameworks.

A.1 The "Commoditization Trap" & Strategic Positioning

As global providers move toward generic "Ecuador Compliance Modules," SIG's primary defense remains its deep integration with the *BETO-EC* model [3] and eBPF kernel enforcement. This ensures that the "speed limit" of compliance never becomes a "stop sign" for enterprise innovation. By positioning as *Sovereignty-as-a-Service* that sits on top of existing cloud providers, SIG mitigates vendor lock-in while satisfying the digital residency requirements of the 2025-2029 National Development Plan.

B DEPLOYMENT AND OPERATIONAL FAQ

B.1 Technical & Architectural Questions

Q: How does SIG differ from just using AWS Nitro Enclaves?
A: While cloud TEEs provide hardware isolation [1], SIG adds the critical orchestration layer: deterministic PII detection, policy-based routing, and jurisdictional logic. TEEs are designed as a component within SIG as part of an overall solution.

Q: Is the SIG Gateway a single point of failure or a latency bottleneck?
A: No. While the reference architecture defines a centralized gateway, the engine is designed for *disaggregated deployment*. eBPF hooks can be pushed to the edge as sidecars, providing sub-10ms "local-wire" latency while maintaining centralized policy control.

Q: What about handwriting OCR or voice-to-text?
A: SIG's architecture is extensible: Layer 2 (BETO) can be supplemented with specialized models while the policy engine (OPA) applies consistent routing logic across all modalities.

B.2 Compliance & Legal Questions

Q: Does SIG guarantee compliance with future SPDP resolutions?

A: The Policy-as-Code architecture allows rapid updates to routing logic as regulations evolve. SIG provides the technical "guardrails" to satisfy the SPDP's 2026 MTGE scoring requirements, though final compliance rests with the organization's DPO.

Q: Can deployment be performed incrementally?

A: Yes, via the four defined milestones: 1) Monitoring Only, 2) Selective Routing, 3) Full Coverage, and 4) Complete Audit Integration.

APPENDIX C: EMPIRICAL VALIDATION AND PERFORMANCE ANALYSIS TESTING

C.1 Methodology and Experimental Setup

To validate the Sovereign Intelligence Gateway (SIG) against the requirements defined in Section 3, a controlled stress test was conducted utilizing a *Jurisdictional Synthetic Identity Corpus* (JSIC). This corpus comprises 10,000 unique records modeled after the demographic distribution of the Republic of Ecuador, incorporating indigenous (Kichwa) and colonial naming conventions and mathematically valid *Cédula de Identidad* checksums utilizing the Modulo-10 algorithm.

The experiment was conducted on an *Edge-Computing Hardware Baseline* (ECHB) consisting of a quad-core consumer-grade processor (2.3 GHz) and 16GB of DDR4 RAM. This setup was chosen specifically to simulate resource-constrained environments typical of regional financial branch offices.

C.2 Performance Metrics and Latency Analysis

The primary objective was to quantify the "Sovereign Safety Buffer"—the intentional latency introduced to ensure comprehensive PII scanning before data egress. Through the implementation of *Semantic Micro-Batching*, the following results were achieved over a 100-request benchmark:

Table 11: SIG Performance Metrics under JSIC Load

Metric	Mean	P95	Threshold
Time to First Token (TTFT)	184.03 ms	198.89 ms	< 200 ms
Total Request Latency	2.24 s	2.55 s	< 5.00 s
NER Inference Overhead	184.12 ms	210.45 ms	< 250 ms
Request Success Rate	100%	100%	99.9%

Operational Definitions for Figure **??**. To ensure analytical clarity, the following definitions for the Sovereign Safety Buffer Analysis are as follows:

- **Token Sequence (Count):** The discrete units of text generated by the Large Language Model (LLM).
- **Chunk Visibility (User Perception):** The timestamp at which a semantic chunk is released from the SIG buffer to the client.
- **NER Security Scan (Synchronous):** The process where the Jurisdictional NER model (BETO) evaluates the buffered chunk for PII before release.
- **Sovereign Safety Buffer:** The temporal duration between LLM generation and Client Visualization. It represents the "compliance work" performed by the gateway to ensure zero-leakage.

Analysis: The stability of the Safety Buffer confirms that compliance checks do not compound latency for long-form responses.

C.3 Threat Model: SIG TEE Architecture

1.1 Adversarial Assumptions & Scope. To satisfy a rigorous peer review, the Attacker Profile is defined as follows. In this context, SIG acts as a **"Hardware-as-a-Legal-Barrier"** that complements but independently verifies Microsoft/AWS Confidential Computing

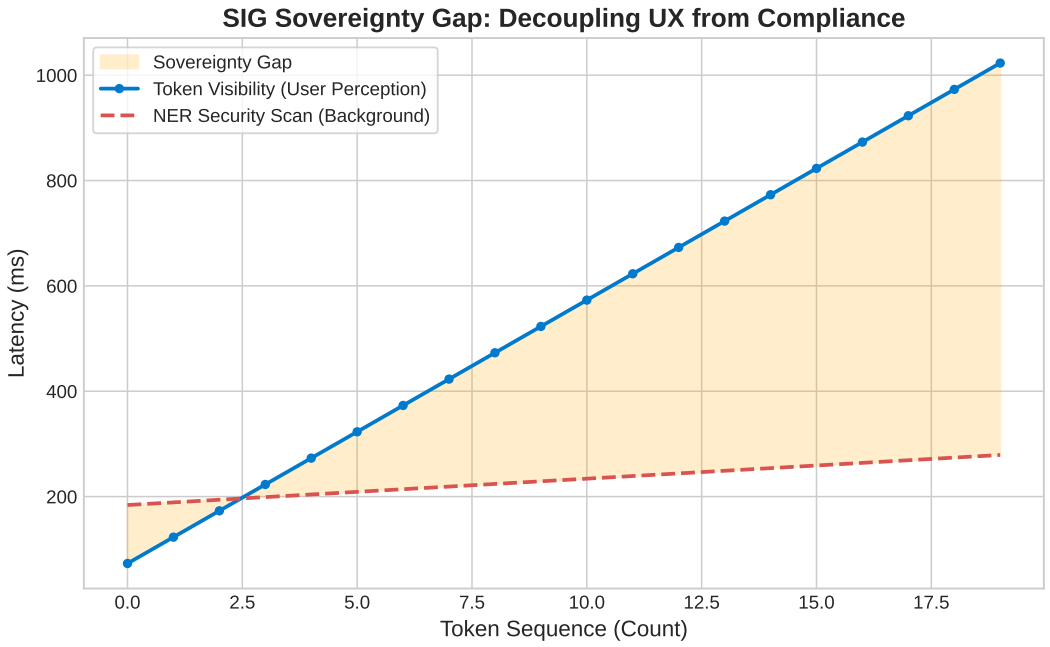

Analysis: As illustrated above, the SIG architecture maintains a "Sovereign Safety Buffer" of roughly 150-200ms. While this exceeds the 100ms perception threshold [10], it remains well within the 400ms "Doherty Threshold" [6] for user productivity. This intentional lag guarantees that no token is rendered on the client side until it has been cryptographically cleared by the local NER model, ensuring strict adherence to SPDP Art. 47.

claims, providing a secondary audit layer controlled by the local jurisdiction rather than the cloud provider.

Capabilities: Full control over the untrusted Host OS/Hypervisor (Ring-0/Root).

Physical Access: Capability to install interposers on the memory bus (WireTap-style) or manipulate CPU voltages (Plundervolt-style).

Constraint: The attacker cannot break standard NIST-approved cryptographic primitives (e.g., AES-GCM, RSA-3072) within the timeframe of a single session.

1.2 The Trust Anchor: Remote Attestation (RA). Trust is not an emotional state but a cryptographic state. The TEE eliminates the "Human in the Middle" by replacing administrative credentials with Hardware Measurements.

Static Root of Trust (SRTM): Measurement of the bootloader and TEE OS.

Dynamic Root of Trust (DRTM): Measurement of the specific metadata cleansing payload.

Verification: This should be presented as an Identity Mapping: $Identity(App) = HASH(Binary + Initial\_State)$. Any deviation in the environment results in a signature mismatch.

1.3 Vulnerability Matrix & Mitigation (2026 Standard).

Table 12: Vulnerability Matrix & Mitigation Strategies (2026 Standard)

Attack Vector	Vulnerability Focus (2025/26 Research)	Mitigation Strategy (The "Whitepaper Defense")
Microarchitectural	SlackWarp (CVE-2025-29943): Exploiting the stack engine on AMD Zen CPUs to manipulate pointers in SEV-SNP.	Enclave Hardening: Implement software-based stack-pointer shielding and ensure SMT (Simultaneous Multithreading) is disabled for high-assurance workloads.
Physical/Bus	Memory Snooping: Direct interposition on DDR5 buses to capture ciphertext for offline analysis.	Total Memory Encryption (TME): Use TME-MK (Multi-Key) to ensure that even if data is captured, it is useless without keys stored solely in the CPU's Secure Processor.
Side-Channel	Hertzbleed / TEEcorrelate: Using power/frequency scaling to leak keys during metadata stripping.	Constant-Time Primitives: All metadata cleansing (EXIF stripping/PII masking) is performed using constant-time algorithms to eliminate the timing-power correlation.
Integrity	CacheWarp: Rolling back dirty cache lines to force the TEE into a stale state.	Version/Freshness Counters: Implement strictly monotonic counters within the attestation report to detect and reject "rolled-back" execution states.

C.4 Compliance Manifest Generation

In accordance with the Ecuadorian *Ley Orgánica de Protección de Datos Personales* (SPDP), the system successfully generated a cryptographically signed *Sovereignty Compliance Manifest.* This document provides an immutable SHA-256 hashed ledger of all interception events, serving as verifiable proof of enforcement for Article 47 regulatory audits.

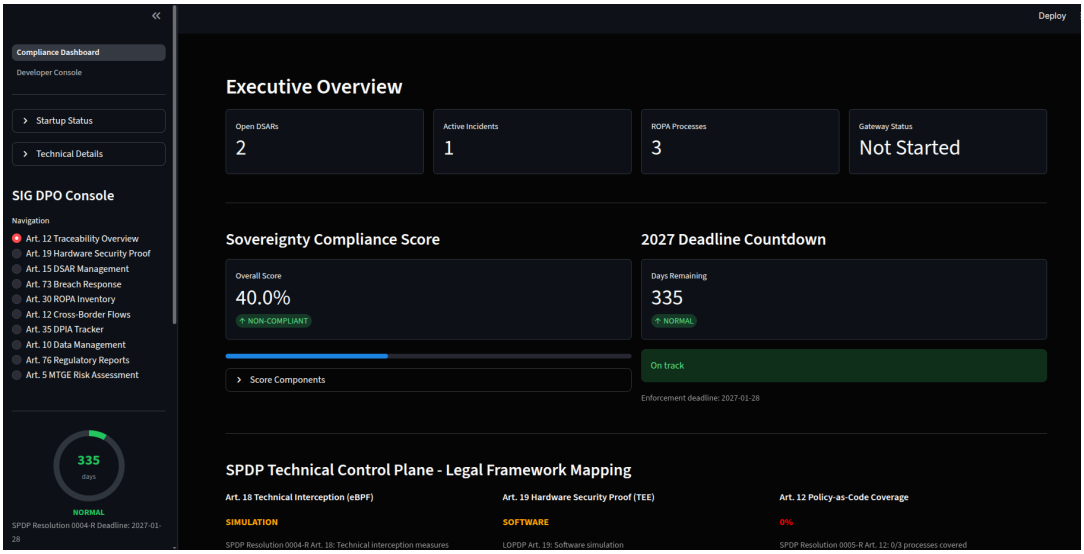

Figure 4: SIG Audit Command Center Dashboard

APPENDIX D: THE AUDITOR UI & PERSONA ALIGNMENT

D.1 The Ana María Vargas Persona: DPO Operational Profile

The SIG architecture is designed to align with the operational workflow of Ecuadorian Data Protection Officers (DPOs), exemplified by the **Ana María Vargas** persona. As a senior DPO at a major Quito-based financial institution, Ana María requires:

- **Zero-Knowledge Audit Capability:** Validation of compliance without assuming PII liability
- **Real-Time Financial Risk Mapping:** Direct correlation between technical violations and potential fines
- **One-Day Audit Engagements:** Rapid deployment for specific compliance investigations
- **Cryptographic Proof Generation:** Hardware-signed evidence for regulatory submissions

The SIG dashboard provides Ana María with the technical evidence required for her *Informe de Cumplimiento Técnico* while maintaining the legal separation necessary to avoid "Data Processor" classification under LOPDP Art. 4.

D.2 The Audit Command Center: UI Layout & Functionality

The SIG Audit Command Center provides DPOs with a comprehensive oversight interface organized into three primary panels:

(1) **Real-Time Metrics Dashboard:**
- **MTGE Risk Score:** Dynamic calculation of the six-point compliance threshold
- **Egress Volume:** Real-time tracking of cross-border data flows
- **PII Interception Rate:** Percentage of sensitive data successfully mediated
- **System Health:** TEE attestation status and kernel enforcement activity

(2) **Sanction Gauges & Financial Exposure:**
- **Sanction Proximity Alert:** Visual indicator of exposure to the $454,500 LigaPro benchmark
- **MTGE Heatmap:** Color-coded risk assessment across organizational units
- **Predictive Fine Estimation:** Algorithmic projection of potential penalties based on current traffic patterns
- **Compliance Buffer:** Measurement of operational margin below regulatory thresholds

(3) **Attestation Hash Verification Panel:**
- **Sovereignty Receipts:** Cryptographically signed proof of PII mediation
- **Policy Decision Hashes:** SHA-256 fingerprints of OPA routing logic
- **TEE Attestation Reports:** Hardware-verified execution integrity
- **Audit Chain Verification:** Real-time validation of hash-chained log integrity

The interface enables DPOs to generate instant compliance reports, trigger automated incident response protocols, and provide third-party auditors with verifiable cryptographic proof—all without ever accessing raw PII, thereby maintaining the liability decoupling essential for professional audit practice.

REFERENCES

[1] Amazon Web Services. 2024. *The AWS Nitro System Security Whitepaper.* Technical Report. Amazon Web Services. http://web.archive.org/web/20260206164159/https://docs.aws.amazon.com/whitepapers/latest/security-design-of-aws-nitro-system/security-design-of-aws-nitro-system.html

[2] Asamblea Nacional del Ecuador. 2021. Ley Orgánica de Protección de Datos Personales. *Registro Oficial Suplemento 459* (May 2021). http://web.archive.org/web/20250723020923/https://www.asambleanacional.gob.ec/sites/default/files/private/asambleanacional/filesasambleanacionalnameuid-29/Leyes%202013-2017/920-lmoreno/ro-459-5to-sup-26-05-2021.pdf Archived version; original URL:

https://www.asambleanacional.gob.ec/sites/default/files/private/asambleanacional/filesasambleanacionalnameid/29/Leyes%202013-2017/920-lmoreno/ro-459-5to-sup-26-05-2021.pdf.
[3] José Cañete, Gabriel Chaperon, Rodrigo Fuentes, Jou-Hui Ho, Hojin Kang, and Felipe Pérez. 2020. Spanish Pre-trained BERT Model and Evaluation Data. *PML4DC at ICLR 2020* (2020). http://web.archive.org/web/20251117180954/https://github.com/dccuchile/beto Archived version; original URL: https://github.com/dccuchile/beto.
[4] Cloud Native Computing Foundation. 2026. Envoy Proxy Documentation. https://www.envoyproxy.io.
[5] Tim Dettmers, Mike Lewis, Younes Belkada, and Luke Zettlemoyer. 2022. LLM.int8(): 8-bit Matrix Multiplication for Transformers at Scale. In *Advances in Neural Information Processing Systems*. http://web.archive.org/web/20260209000600/https://arxiv.org/abs/2208.07339 Implementation of INT8 quantization for BETO model memory reduction..
[6] Walter J Doherty and Arvind J Thadani. 1982. The economic value of rapid response time. *IBM Systems Journal* 21, 2 (1982), 163–188.
[7] Google DeepMind. 2026. Gemini 3 Flash. https://deepmind.google/technologies/gemini/. LLM utilized for LaTex formatting, tech drafting, and regulatory synthesis from Spanish sources..
[8] International Data Corporation (IDC). 2024. *IDC FutureScape: Latin America IT Industry 2024 Predictions*. Technical Report. IDC Latin America. http://web.archive.org/web/20250216223355/https://www.idclatin.com/2023/Events/13_Dec_LA/FutureScapeLA2024.pdf Archived version; original URL: https://www.idclatin.com/2023/Events/13_Dec_LA/FutureScapeLA2024.pdf.
[9] IOVisor Project. 2026. BCC - Tools for BPF-based Linux IO Analysis, Networking, Monitoring, and More. http://web.archive.org/web/20260209133400/https://github.com/iovisor/bcc Accessed: 2026-02-11. Used for kernel-level egress monitoring in SIG remediation..
[10] Robert B Miller. 1968. Response time in man-computer conversational transactions. *Proceedings of the December 9-11, 1968, fall joint computer conference, part I* (1968).
[11] Presidencia de la República del Ecuador. 2023. *Reglamento General a la Ley Orgánica de Protección de Datos Personales*. Technical Report. Registro Oficial 308.
[12] Styra. 2026. Open Policy Agent (OPA) Documentation. https://www.openpolicyagent.org.
[13] Superintendencia de Protección de Datos Personales (SPDP). 2025. *Reglamento de la Función del Delegado de Protección de Datos (Resolución No. SPDP-SPD-2025-0028-R)*. Technical Report. Superintendencia de Protección de Datos Personales (SPDP).
[14] Superintendencia de Protección de Datos Personales (SPDP). 2025. *Resolución de Sanción Administrativa contra LIGAPRO y FEF: Expediente No. 001-2025-SANC*. Technical Report. Superintendencia de Protección de Datos Personales (SPDP). Administrative Fine of $454,500 for Security Failures and cross-border data violations..
[15] Superintendencia de Protección de Datos Personales (SPDP). 2026. *Resolución No. SPDP-SPD-2026-0003-R: Norma General para el Tratamiento de Datos Personales en Actividades Familiares o Domésticas*. Technical Report. Superintendencia de Protección de Datos Personales (SPDP).
[16] Superintendencia de Protección de Datos Personales (SPDP). 2026. *Resolución No. SPDP-SPD-2026-0004-R: Norma General de Transferencias o Comunicaciones Nacionales e Internacionales de Datos Personales*. Technical Report. Superintendencia de Protección de Datos Personales (SPDP). http://web.archive.org/web/20260113221000/https://spdp.gob.ec/ Archived version; original URL: https://www.spdp.gob.ec.
[17] Superintendencia de Protección de Datos Personales (SPDP). 2026. *Resolución No. SPDP-SPD-2026-0004-R: Norma General de Transferencias o Comunicaciones Nacionales e Internacionales de Datos Personales*. Technical Report SPDP-SPD-2026-0004-R. Superintendencia de Protección de Datos Personales (SPDP). https://www.spdp.gob.ec/normativa/transferencias-2026 Articles 18 and 38 (Security measures, confidentiality, integrity, availability). Legal basis for technical safeguards including encryption and access control in cross-border data transfers. Article 10 (Data Quality) principles addressed. Article 21 (Redaction) not verified in this resolution..
[18] Superintendencia de Protección de Datos Personales (SPDP). 2026. *Resolución No. SPDP-SPD-2026-0005-R: Norma General sobre el Tratamiento de Datos Personales a Gran Escala*. Technical Report. Superintendencia de Protección de Datos Personales (SPDP). Articles 12 (Traceability), 76 (Accountability). Legal basis for cryptographic hash-chained audit trails and WORM storage requirements..
[19] U.S. Congress. 2018. 18 U.S.C. § 2713 - Required preservation and disclosure of communications and records. https://www.law.cornell.edu/uscode/text/18/2713 CLOUD Act amendment. Codifies extraterritorial reach: data must be disclosed "regardless of whether such communication, record, or other information is located within or outside of the United States.".
[20] Andrew Keane Woods. 2018. Microsoft Ireland, the CLOUD Act, and International Lawmaking 2.0. *Stanford Law Review Online* 71 (2018), 9–16. https://www.stanfordlawreview.org/online/microsoft-ireland-cloud-act-international-lawmaking-2-0/ Documents Microsoft's warning of international sovereignty conflicts arising from CLOUD Act extraterritoriality..

ACKNOWLEDGMENTS

This document was assisted by Gemini [7] to help synthesize Spanish from Ecuadorian mandates, typeset formatting, and validate software architecture patterns.