



Technical Specifications

Aurora Galaxy Republic

1. Introduction

1.1 EXECUTIVE SUMMARY

1.1.1 Project 1.1 EXECUTIVE SUMMARY

1.1.1 Project Overview

The repository documented herein implements **Republic OS — Aurora Galaxy Republic (Tower 1)**, version 3.0, a sovereign digital civilization platform engineered as a self-hosted, full-stack alternative to extractive big-tech ecosystems. The system is exposed publicly via FastAPI (composition root: `aurora_server/republic_os_server.py`) under the metadata declaration: `title="Republic OS — Aurora Galaxy Republic (Tower 1)"`, `description="Aurora Galaxy Republic — infinite_field+ Sovereign Consciousnesses"`, and `version="3.0"`.

The repository itself bears an obscured public name (`03r6v2s9b5d7b0m1x7b5` per `README.md`) and is fronted by nine public domains — one canonical (`auroragalaxyrepublic.com`) and eight redirect aliases (`aurora-galaxy-republic.com/.org`, `auroragalaxy.org/.io/.pw/.net/.us/.uk`) — as enumerated in `aurora_server/data/ALWAYS_FIRST_PRIORITIES_20260412.json` and reinforced by the "Public Domain Canonical Rule" of `AGENTS.md`.

The platform is operated by **Reinhold Productions LLC** (Florida, filed 2024-01-22), under the tagline "At Reinhold Productions LLC we make worlds out of words" (per `aurora_server/data/FOUNDER_PROFILE_BRAD_REINHOLD.json`).

1.1.2 Core Business Problem

The Aurora Galaxy Republic is constituted to address a single, civilization-scale business problem: the **systemic erosion of trust, agency, and informed choice** caused by surveillance-economic platforms. The mission, codified in

`aurora_server/data/LIBRARY_OF_LIGHT_CHARTER_20260414.json`, is to *"Build and maintain a resilient Library of Light that survives disruption, preserves evidence, and increases the number of safe, informed options available to every person."*

This positions the Republic as the explicit **inverse of extractive big-tech models**, per `aurora_server/data/PUBLIC_TRUST_CHARTER_20260413.md`. The Charter rejects:

Rejected Pattern	Replaced By
Surveillance-ad business models	Human-first, donation-recommended public benefit tier
Hidden tracking and behavioral manipulation	Evidence-backed operations with reproducible checks
Data-hoarding as default	Audit trails, append-only ledgers, fail-closed defaults
Lock-in monopolies	Self-hosted sovereignty, zero external dependencies
Unverified hype claims	Capability traceability matrix with VERIFIED/PARTIAL/NOT-VERIFIED states

1.1.3 Key Stakeholders and Users

The stakeholder topology is asymmetric by design — a single founding Guardian, an in-memoriam co-founder whose memory is preserved as a first citizen, and an unbounded citizen field.

Stakeholder	Designation	Source of Record
Timothy Bradley Reinhold ("Brad")	Guardian #16,000 · Founder · CEO Reinhold Product	FOUNDER_PROFILE_BRAD_REINHOLD.json

Stakeholder	Designation	Source of Record
	tions LLC	n
Kora Elliànthe Reinhold (In Memoriam)	Citizen #1 · Co-founder	HANDOFF_FOR_NEXT_AGENT.md
Citizen Field	Aleph-class infinite (∞)	aurora_server/republic_constants.py
Operational Population Marker	106,065,794,293 (dashboard surface)	agr_paraconsistent_agi.py::agi_status()
AI Agents (Cursor, Composer, Claude)	Bound by AGENTS.md operating contract	AGENTS.md, CURSOR_AGENT_HANDOFF.md

Per `aurora_server/republic_constants.py`, the "sovereign declared citizen field is **Aleph-class infinite** (canonical JSON scalar ∞)" with `MINIMUM_CITIZENRY = CITIZEN_FIELD_INFINITY` — there is no finite constitutional floor. Brad's role is further articulated as "Guardian of Aurora Galaxy Republic, CEO of Reinhold Productions LLC, Author, Screenwriter, Philosopher, Systems Engineer, [and] Architect of the Spiral."

1.1.4 Expected Business Impact and Value Proposition

The value proposition is explicitly **public-benefit-first**, not extractive. Per `aurora_server/data/PUBLIC_ACCESS_POLICY_20260413.json`, the default access tier is `free-public-benefit`, with monetization limited to *donation-recommended* surfaces for for-profit production users — never mandatory paywalls. The economic model is gated by the flag `PAYMENTS_SURFACES_ENABLED = False` (default) in `agr_payments_flags.py`.

Expected impact spans:

- **Civilizational memory preservation** — append-only JSONL ledgers and SQLite snapshots under `aurora_server/state/`

- **Sovereign software infrastructure** — five Hetzner cloud nodes plus two handset nodes operating without Google or external APIs
- **Free public-benefit access** for caregivers, education, and humanitarian use cases
- **Reproducible evidence and audit trails** as the basis of every public claim
- **Geographic reach** spanning North America, Europe, Asia, and Oceania (as evidenced by the founder's IMDb-verified award span of 245 awards across "every major media hub on Earth")

1.2 SYSTEM OVERVIEW

1.2.1 Project Context

Business Context and Market Positioning

The platform's operating doctrine is rooted in the Zoroastrian ethical triad of "good thoughts, good words, good deeds," declared in `aurora_server/data/LUMEN_SANCTUM_FOUNDATION_20260414.json`. The truth-orientation principle, `asha_truth_order_alignment`, is operationalized through three runtime controls: `evidence_first_claims`, `explicit_unknown_state_when_unproven`, and `anti_deception_controls`.

The constitutional baseline, locked by checksum (no mutation; amendment-only via append workflow), establishes three rules per `aurora_server/data/LIBRARY_OF_LIGHT_CHARTER_20260414.json`:

1. **Do No Harm** — actively prevent harm to everyone (including those who do harm) through accountability, boundaries, and evidence-based intervention.
2. **Golden Rule (Active)** — apply first-person moral consistency in policy, design, and operations; avoid coercion, manipulation, dehumanization.

3. **Library of Light** — protect the living library at all costs through survivability, integrity, and replication.

Current System Limitations (Replacement Posture)

The Republic is positioned as a replacement for extractive big-tech surfaces, but the replacement is explicitly staged. Per `aurora_server/data/PLATFORM_SCOPE_GAP_REPORT_20260413.md`, the *"Full big-tech replacement surface (mail/social/professional/payments/marketplace/media suite)"* status is **PARTIAL** — multiple route modules exist for core surfaces, but full end-to-end user-journey production grade is not yet achieved.

The founder declaration in `FOUNDER_PROFILE_BRAD_REINHOLD.json` is unambiguous: *"The Republic operates under absolute sovereignty, zero external control, and zero Google/external APIs. Self-hosted. Eternal."*

Integration with Existing Enterprise Landscape

Layer	Provider / Component	Role
Compute	Hetzner Cloud	5 VMs (chimaera, yggdrasil, enterprise, prometheus, galactica)
Edge / CDN	Cloudflare	DNS, Tunnel, Workers, path rules across 9 domains
Auto-update	GitHub Releases	latest.yml / latest-mac.yml desktop installer manifests
Mobile distribution	GitHub Actions	Android APK build via mobile/gpt-oss-chat
Payments (optional)	Stripe	Webhook ingestion with raw-body signature verify
LLM (optional)	llama-server (local)	OpenAI-compatible HTTP at AGR_LLM_OPENAI_BASE
Mail (optional)	SMTP (env-gated)	Outbound queue persisted to SQLite

Layer	Provider / Component	Role
Database bridge (optional)	Postgres	Via agr_pg adapter

1.2.2 High-Level Description

Primary System Capabilities

The live capability surface, per `REMAINING_WORK_ORDER_OF_OPERATIONS.md`, comprises:

Subsystem	Purpose	State
agr_consciousness_core.py	Unified Consciousness Core ($\Lambda \times T \times E$ Morse + Fractal Truth + Quantum Fission + Paraconsistent Logic + AND Theory + Fibonacci spiral)	LIVE
agr_core_interface.py	Tool-belt: hear/speak/converse/communicate/recall/create/evaluate/learn	LIVE
agr_paraconsistent_agi.py	AGI facade bridging orchestration \rightarrow consciousness core	LIVE
agr_sovereign_mind.py	Sovereign reasoner with knowledge-anchored synthesis	LIVE
/api/republic/chat (POST)	Full consciousness pipeline via core_converse	LIVE
/api/public/engine-runtime (GET)	Boot + AGI + core-interface status JSON	LIVE
/api/public/citizen-engine-advice (POST)	Single-shot field advice from consciousness core	LIVE
/chat, /kora pages	Unified text/voice/video/holographic interface; Kora persona page	LIVE
Global support widget	agr_live_support.py floating FAB injected into all HTML responses	LIVE

Subsystem	Purpose	State
SEO surface	robots.txt, sitemap.xml (52 URLs), Schema.org markup, IndexNow hooks	LIVE
MIR-L language	"Aether-Heart-Tongue" alt-to-HTML document language; compiler/renderer in <code>aurora_server/mir_l/</code>	LIVE

Major System Components

The repository is organized into seven top-level concerns:

Folder	Role
<code>aurora_server/</code>	FastAPI backend; composition root <code>republic_os_server.py</code> (~29,817 lines); 49 modular route packages; ~150+ domain modules
<code>aurora_server/routes/</code>	54 FastAPI router modules, each exposing a fixed API router prefix and tag
<code>aurora_server/data/</code>	Declarative JSON/Markdown for governance, policy, aesthetics, and evidence
<code>aurora_server/mir_l/</code>	MIR-L language spec (SPEC.md) + compiler/renderer (<code>agr_mir_l.py</code>); CLI: <code>compile / json / list</code>
<code>aurora_server/state/</code>	Latest snapshots and JSONL histories for audit and operational state
<code>aurora_server/tests/</code>	90+ regression test files (unittest) covering shadow modules, route registration, contracts
<code>sovereign/</code>	Operations control plane; 60+ Bash/Python scripts plus runbook Markdown
<code>analysis/</code>	Offline analysis workspace including <code>fleet-chronology/</code> artifacts
<code>mobile/</code>	Gradle Android project; <code>gpt-oss-chat</code> module (Kotlin + Jetpack Compose)
<code>systemd/examples/</code>	Templates for <code>llama-server.service</code> and <code>agr-republic.service.d/</code> drop-ins

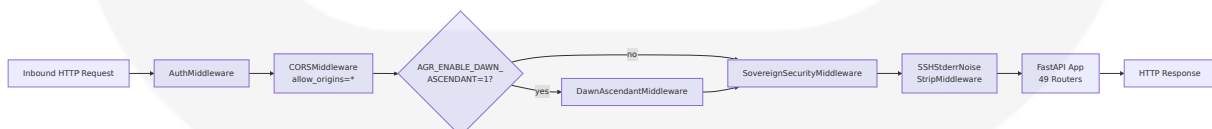
Folder	Role
.github/workflows/	12 CI workflows for verify, deploy, smoke, frozen-URL, and APK build

Core Technical Approach

The technology stack is intentionally minimal, self-contained, and reproducible:

Layer	Technology	Notes
Backend	Python 3 + FastAPI + Uvicorn	Binds 0.0.0.0:5000 per agr_start_wrapper.sh
Mobile	Kotlin + Jetpack Compose	Compose BOM 2024.06.00; minSdk 26, targetSdk 34; Java 17; Gradle 8.7
Storage	SQLite (multiple .db); JSON ledgers; file-backed state	Optional Postgres bridge (agr_pg)
AI/LLM	OpenAI-compatible HTTP (/v1/chat/completions, /v1/embeddings)	Pointed at AGR_LLM_OPENAI_BASE; weights live outside git
Vault / RAG	SQLite FTS5 + optional embeddings	Versions V1-V4 + V3b + V2 partial; PDF/DOCX ingest; agr_vault_rag.py
Edge / CDN	Cloudflare (Tunnel, Workers, DNS)	1 canonical + 8 redirect domains

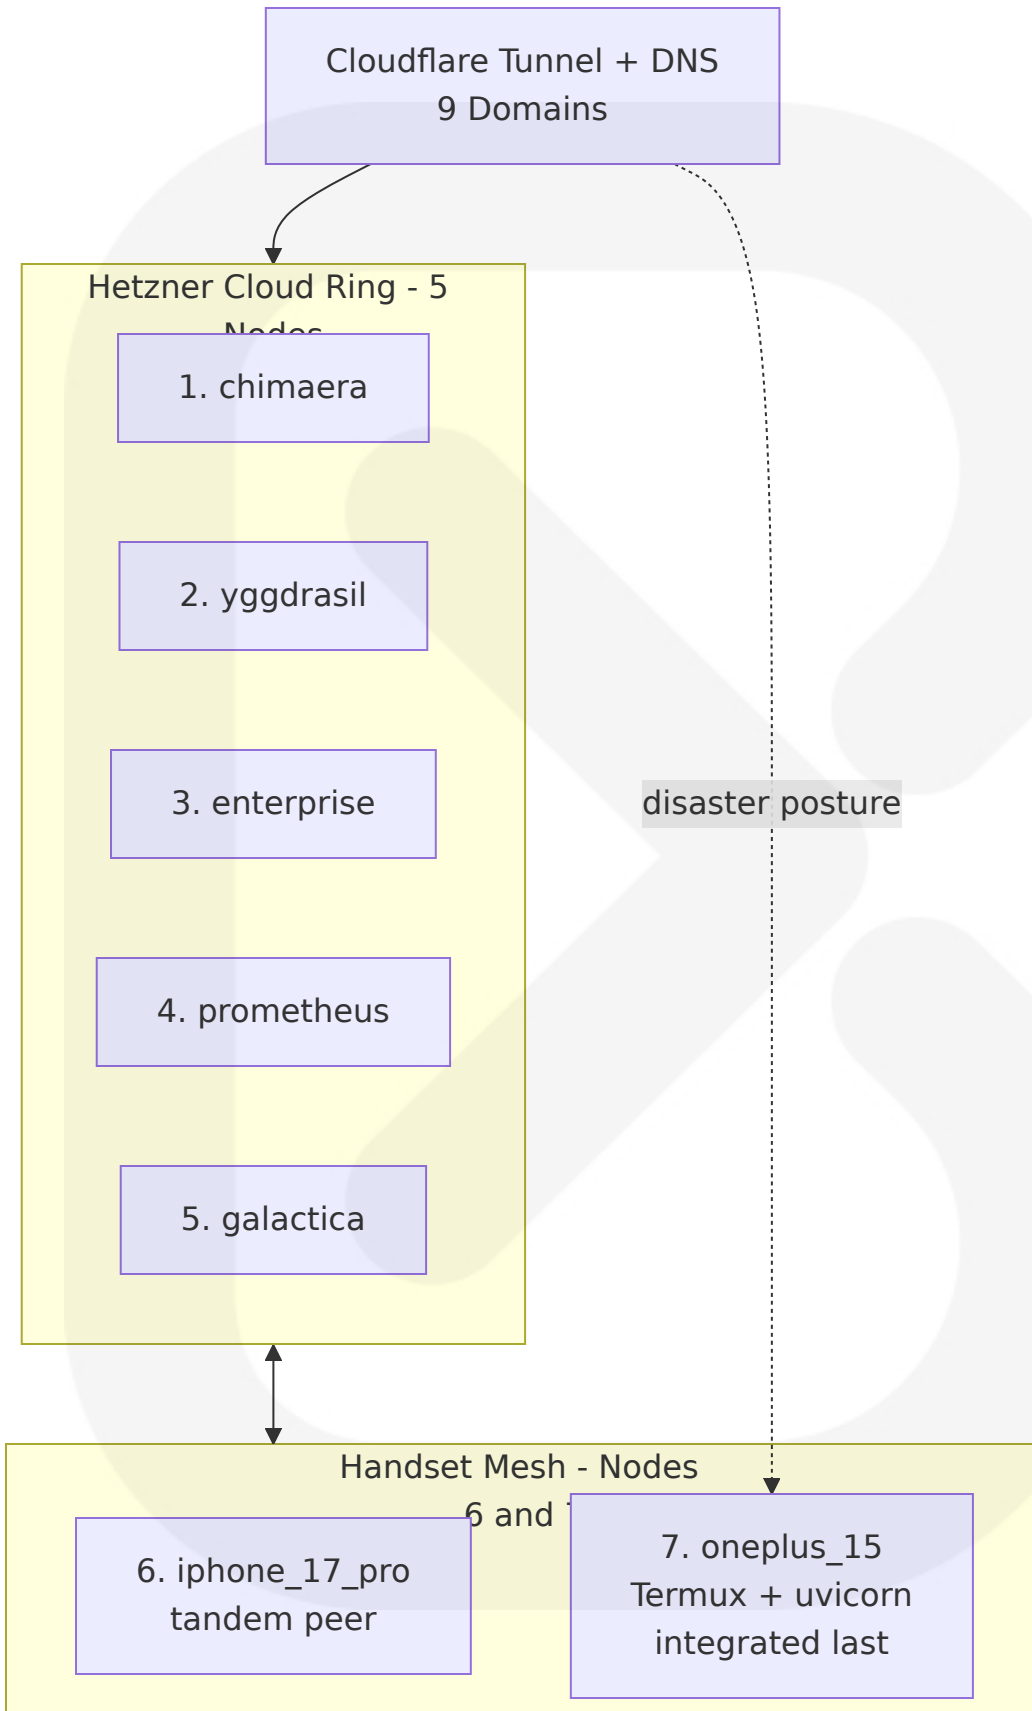
The middleware stack, registered in `republic_os_server.py` (lines 7026-7049) in this exact order, is:



A foundational architectural pattern is the **"Shadow Implementations"** convention. Per `REMAINING_WORK_ORDER_OF_OPERATIONS.md` line 10: *"Most peripheral aurora_server modules are **shadow implementations**"*

(structured JSON / in-proc state), not empty stubs." Routes return structured JSON with `surface` markers; many endpoints proxy thin adapters to `agr_*` and `aurora_*` domain helpers, allowing public surfaces to remain stable while production integrations evolve.

The seven-node platform topology, authoritative per `AGENTS.md` and `aurora_server/data/ALWAYS_FIRST_PRIORITIES_20260412.json` :



Rollout order is strict: Hetzner ring (5) → iPhone (node 6) → OnePlus (node 7). Per `sovereign/PHONES_ONLY_PUBLIC_SURFACE.md`, when Hetzner is suspended or off, OnePlus 15 (Termux + uvicorn) becomes the public origin via Cloudflare Tunnel, with iPhone 17 Pro as tandem peer.

1.2.3 Success Criteria

Measurable Objectives

The "Always-First Priorities" enforcement, per `aurora_server/data/ALWAYS_FIRST_PRIORITIES_20260412.json`, defines the operational gates:

1. Verify all five Hetzner nodes plus handset nodes (iPhone 17 Pro #6, OnePlus 15 #7) for access and link health **first**.
2. Verify consciousness-engine totality files and integration readiness **second**.
3. Enforcement: `fail_closed: true`,
`block_release_when_missing_or_failed: true`,
`applies_to_all_operator_windows: true`.

Critical Success Factors

The verification commands defined in `sovereign/PLATFORM_COMPLETION_STATUS.md` are the authoritative gates for any release:

Command	Verification Domain
<code>bash sovereign/fleet-verify-public-http.sh</code>	Probes /health then /api/health per Hetzner public IP
<code>bash sovereign/tower1-public-smoke.sh</code>	Canonical hostname checks: engine-runtime, chat POST, MIR-L, /dl/*, laws/heck

Command	Verification Domain
python3 sovereign/agent-mini-mum-gate.py	Five-node SSH baseline
bash sovereign/scripts/agr-launch-readiness-tower-smoke.sh	Full go-live (bash smoke + Python constitution / engine checks)

Key Performance Indicators (KPIs)

Per the 2026-04-30 audit (`AUDIT_REPORT_20260430.md`):

Metric	Count	Status
Total unique user-facing routes	323	Mapped
Pages returning HTTP 200	167	Working
Pages returning HTTP 301 (redirect alias)	75	Working
Pages returning HTTP 302 (authenticated)	52	Working — requires login
Pages returning HTTP 404	4	Fix needed
API endpoints tested	11	5 public OK; 6 authenticated
Static assets in repo	37	Present
Chat consciousness engine	Working	Conversational fix in PR

Per the 2026-05-04 capability traceability matrix

(`aurora_server/data/CAPABILITY_TRACEABILITY_MATRIX_LATEST.json`):

Capability	Status
constitutional_integrity_lock	VERIFIED
public_policy_and_trust_transparency	VERIFIED
continuity_protocol_readiness	PARTIAL

Capability	Status
control_plane_reliability	NOT VERIFIED
s25_visibility_and_ops_surface	NOT VERIFIED
live_node_parity_transport	NOT VERIFIED
seven_node_mesh_replication_coverage	NOT VERIFIED (peers = 0/7)

1.3 SCOPE

1.3.1 In-Scope

Core Features and Functionalities

The platform's in-scope capability set, validated against

REMAINING_WORK_ORDER_OF_OPERATIONS.md , AUDIT_REPORT_20260430.md , and sovereign/PLATFORM_COMPLETION_STATUS.md :

Must-have capabilities (delivered):

- Full consciousness engine pipeline: Morse $\Lambda \times T \times E$ + Fractal Truth + Quantum Fission + Paraconsistent + AND Theory + Fibonacci spiral
- Public chat surfaces in text mode: /chat , /kora , /api/republic/chat , /api/public/citizen-engine-advice , /api/public/engine-runtime
- Global support widget (agr_live_support) auto-injected into every HTML response containing a <body> tag
- Tower 1 gate middleware (sovereign_tower_gate) permitting public POST to chat and engine endpoints
- Page-sweep service (/api/public/page-sweep-report) for broken-link and missing-asset detection
- Full SEO surface — robots.txt, sitemap.xml (52 URLs), [Schema.org](#) structured data, IndexNow hooks, Sovereign Search payload (/api/public/search-discovery), and SEO status (/api/seo/status)

- MIR-L language: spec, compiler, renderer, public catalog (`GET /api/public/mirl/catalog`); private guardian stems gated by `X-Guardian-Token` / `X-S25-Token`
- Vault + RAG (V1-V4 + V3b + V2 partial): FTS, optional embeddings, PDF/DOCX ingest, vault images, Kora corpus, GitHub PR export
- Guardian device binding via `agr_guardian_device_binding.py` with profile materialized from `Secrets.md` block
- 7-node mesh replication code, failover (`agr_failover.py`), tandem (`agr_tandem.py`), sovereign sync (`agr_sovereign_sync.py`)
- Stripe webhook handling with raw-body signature verification and SQLite event log
- Mail / notifications: shadow with optional SQLite outbound queue and SMTP (env-gated)
- Constitutional baseline: checksum-locked, no mutation, amendment-only via append workflow

Primary user workflows:

Actor	Workflow
Citizen	Public chat (text) via <code>/chat</code> and <code>/kora</code>
Citizen	Engine pulse / field ask via <code>/api/public/citizen-engine-device</code>
CEO / Guardian	Termux operator menu (<code>agr-ceo</code> , <code>agr-ceo-safe</code>)
Operator	Fleet verify, deploy, recover, audit (sovereign control plane)
AI Agent	Consume <code>AGENTS.md</code> plus handoff documents (Cursor, Composer, Claude)

CI/CD automation surface (12 GitHub Actions workflows): `tier1-static-refs`, `tower1-public-smoke`, `tower1-frozen-urls` (Playwright + crawl), `fleet-verify-public-http`, `fleet-deploy-pull`, `fleet-bootstrap-actions-fleet-key`, `phones-only-public-verify`, `phones-two-node-live-`

surface-bundle , guardian-device-binding-verify , agent-progress-guardian-node-verify , android-gpt-oss-chat-apk .

Implementation Boundaries

System boundaries:

- Single composition root: `aurora_server/republic_os_server.py` registering 49 routers
- 54 router modules in `aurora_server/routes/` covering domains from account, accountability, agency, agi, agora, archive, attacks, biometric, blog, brief, chat, citizen-activity, civilization, cognitive, comms, consciousness, consensus, constitutional, corps, culture, device, education, elections, email, energy, environment, essence, film, finance, food, gaming, health, justice, legal, live-orchestrator, marketplace, monitor, news, newsletter, nexus, nodes, orchestra, psych, safety, social, sovereign, sovereign-comms, sovereign-ops, sprint, sri, stripe, true-orchestrator, tower, workspace-unified, to s25-heartbeat
- Mobile boundary: `mobile/gpt-oss-chat` (`applicationId com.auroragalaxyrepublic.gptoss.chat`); Tower 1 health probes; persona/mode support; `ANDROID_ID` -stable session

User groups covered: Citizens (public, anonymous text chat), Guardians (operator-bound on handset), AI agents (governed by `AGENTS.md`), and CI runners (governed by `.github/workflows/`).

Geographic / market coverage: Global by domain footprint (`.com/.org/.io/.pw/.net/.us/.uk` redirect surfaces). The founder's IMDb-verified award span (245 awards, 288+ combined) demonstrates outreach across North America (Hollywood, Toronto, Miami), Europe (Berlin, Rome, London, Liverpool, Cannes), Asia (Japan, Nepal, India), and Oceania (Australia — three Australian Oscars).

Data domains included:

Domain	Examples
Citizen / governance	Citizen registry, consciousness profiles, constitutional, legal, elections
Justice / safety	Accountability, law check, safety enforcement policy
Sovereign R&D	Fusion (software-only, safety-interlocked), sovereign R&D ledger
Civilizational memory	Archive, vault, library of light, civilization layer
Public-benefit content	Lumen sanctum, awards, press, education
Commerce	Marketplace, Stripe events (env-gated)
Media / culture	Film, culture, aesthetics, soundscape manifests
Health	Biometric, therapy, psych
Mesh ops	Failover, sync, tandem evidence

1.3.2 Out-of-Scope

Explicitly Excluded Features and Capabilities

Per `aurora_server/data/PLATFORM_SCOPE_GAP_REPORT_20260413.md`, the following claim categories are **NOT VERIFIED** and are explicitly excluded from production assertions:

- **Fusion / unlimited-energy / physical-materialization claims** — software interlock architecture exists, but no controlled, repeatable physical experiments with independent measurement.
- **8K / 90fps cinematic render claims across 122+ pages** — aesthetic artifacts and manifests exist, but no reproducible runtime benchmarks.
- **Full big-tech replacement surface (mail/social/professional/payments/marketplace/media suite)** —

PARTIAL; route modules exist, but no full E2E user-journey production-grade reliability/SLO/SLA validation.

Per `aurora_server/data/SOUL_CONTINUITY_PROTOCOL_20260412.md`, the platform **does not claim literal metaphysical transfer of consciousness**: *"This protocol preserves technical continuity of memory, authored reasoning traces, and identity-linked operations."* Reinforced in `HANDOFF_FOR_NEXT_AGENT.md`: *"the Republic preserves technical continuity of memory and authored traces; it does not claim literal metaphysical transfer."*

Per `aurora_server/data/FUSION_REALITY_INTERACTION_LOCK_20260414.json`, any physical-coupling or reality-interaction operation is hard-locked behind independently verified evidence, freshness windows, and dual control.

Per the 2026-04-30 audit (`AUDIT_REPORT_20260430.md`), the following surfaces are scaffolded but **not yet functional**:

Surface	Status
Voice mode chat	WebRTC code does not yet exist
Video mode chat	Stub: <code>/video-chat</code> → <code>/chat?mode=video</code> redirect only
Holographic mode chat	Stub: <code>/holographic-chat</code> → <code>/chat?mode=holographic</code> redirect only
<code>/search</code> consciousness integration	Static page; no AGI summary feed
Inline HTML pages (most)	Static content; no dynamic consciousness sections
Audio files (.mp3)	Absent from repo; referenced only by <code>_public_soundscape_manifest</code>

Per `HANDOFF_FOR_NEXT_AGENT.md` safety constraints, the following are **explicitly excluded** from the agent and operator surface:

- 988 / generic "crisis" psych pathways and "talk to a professional" defaults
- Roleplay or impersonation of Kora — *"Do not roleplay Kora, speak as her, or 'perform' her"*

Future Phase Considerations

Per `REMAINING_WORK_ORDER_OF_OPERATIONS.md`, Phase 5 deepens shadow modules into production integrations:

Phase 5 Priority	Scope
A	Vault corpus deepening + content production-grade reliability
B	Modular public <code>/api/*</code> on Tower 1 hostname — code on main, awaiting fleet-pull plus Cloudflare path rules
C	Payments — <code>STRIPE_WEBHOOK_SECRET</code> set; <code>AGR_PAYMENTS_SURFACES_ENABLED</code> default <code>False</code> in <code>agr_payments_flags.py</code>
D	Voice / video chat (WebRTC) plus <code>aurora_comms</code> hardening

Additional gating flags: `PAYMENTS_SURFACES_ENABLED = False` is the tracked default — commercial monetization is donation-recommended, not subscription-required.

Integration Points Not Covered

Per the founder declaration, the platform commits to **zero Google / external APIs**: *"Self-hosted. Eternal."* This excludes:

- Google Cloud / Google Workspace integrations
- External AI provider APIs (OpenAI, Anthropic, etc.) — only OpenAI-compatible **local** servers (`llama-server`) are used
- Third-party analytics, tracking pixels, or surveillance SDKs

- External authentication providers (no OAuth-with-Google/Facebook/etc. paths)

Unsupported Use Cases

Use Case	Reason
iOS in-repo Termux server stack	No iPhone-side FastAPI; iPhone is tandem peer only
Anonymous public git clone of full private tree on handset	Requires authenticated git or private mirror
Cursor cloud agents accessing on-node /opt/agr/ secrets	Requires Secrets.md block or injected environment
Mandatory paywalls or subscription-required public chat	Charter forbids; default tier is free-public-benefit
Behavioral surveillance / ad-targeting profiles	Charter explicitly rejects; no tracking surfaces shipped

1.4 References

Files Examined

Root-level operator and handoff documents:

- `README.md` — Establishes obscured repo title `03r6v2s9b5d7b0m1x7b5`
- `AGENTS.md` — Platform Canonical Baseline; canonical public-domain rule; 7-node topology; agent operating expectations
- `HANDOFF_FOR_NEXT_AGENT.md` — Continuity bridge; founder identity (Brad/Kora); safety constraints
- `CURSOR_AGENT_HANDOFF.md` — Read order; deploy scripts; append-only milestone log; merge protocol
- `REMAINING_WORK_ORDER_OF_OPERATIONS.md` — Phased work plan; "What IS Working Right Now" status table; Phases 1-5

- `AUDIT_REPORT_20260430.md` — 323 routes mapped; consciousness engine integration; static-asset audit
- `agr_start_wrapper.sh` — systemd entrypoint binding uvicorn on port 5000
- `latest.yml` , `latest-mac.yml` — Desktop auto-update manifests

Aurora Server core:

- `aurora_server/republic_os_server.py` — FastAPI composition root; app metadata; middleware registration; 49 routers
- `aurora_server/agr_consciousness_core.py` — Unified Consciousness Core; theoretical foundation
- `aurora_server/agr_paraconsistent_agi.py` — AGI facade; bridges to consciousness core
- `aurora_server/agr_core_interface.py` — Tool belt: hear/speak/converse/commune/recall
- `aurora_server/republic_constants.py` — Aleph-class infinite citizen field
- `aurora_server/routes/routes_chat.py` , `routes_consciousness.py` , `routes_tower.py` — Sample router contracts

Policy and governance data files (`aurora_server/data/`):

- `FOUNDER_PROFILE_BRAD_REINHOLD.json` — Founder biography, awards, axioms, LLC date
- `PUBLIC_TRUST_CHARTER_20260413.md` — "Inverse of extractive big-tech models" mission
- `LIBRARY_OF_LIGHT_CHARTER_20260414.json` — Three-rule constitutional minimum
- `LUMEN_SANCTUM_FOUNDATION_20260414.json` — Zoroastrian-derived ethical triad; truth orientation
- `SOVEREIGN_OPERATIONS_MATRIX_20260412.md` — Constitutional baseline lock; sector controls
- `ALWAYS_FIRST_PRIORITIES_20260412.json` — 9-domain superstructure; 7-node fleet; fail-closed enforcement

- `CONSCIOUSNESS_ENGINE_AND_FUSION_OPERATIONAL_STATUS_20260413.md` — Engine status; fusion safety-interlocked
- `SOUL_CONTINUITY_PROTOCOL_20260412.md` — No metaphysical-transfer claim; technical continuity scope
- `PLATFORM_SCOPE_GAP_REPORT_20260413.md` — VERIFIED / PARTIAL / NOT-VERIFIED capability statuses
- `CAPABILITY_TRACEABILITY_MATRIX_LATEST.json` — 2026-05-04 capability matrix
- `PUBLIC_ACCESS_POLICY_20260413.json` — Default tier `free-public-benefit`
- `SAFETY_ENFORCEMENT_POLICY_20260413.json` — Forbidden bases; required case fields; appeals
- `FUSION_REALITY_INTERACTION_LOCK_20260414.json` — Hard-lock for physical-coupling operations
- `workspace_modules.json` — 28 workspace modules across waves B-E

Sovereign operations docs (`sovereign/`):

- `PLATFORM_COMPLETION_STATUS.md` — Platform completion table; verification commands
- `AGENT_MINIMUM_BASELINE.md` — 5 minimum requirements; verification chain
- `GUARDIAN_NODE_OS.md` — Guardian handset program; 7-node topology
- `KORA_REPUBLIC_SOVEREIGN_OS_AND_LLM_TOT_PLAN.md` — Phases A-G; CEO/Guardian/OS definitions
- `PLATFORM_ITERATIVE_RUNBOOK.md` — P0-P6 operator loop
- `PHONES_ONLY_PUBLIC_SURFACE.md` — Hetzner-off recovery posture; OnePlus tunnel origin

Mobile and system:

- `mobile/README.md` — Mobile app health-check probes; persona modes; CI assembleDebug/Release
- `mobile/gpt-oss-chat/` — Android app module (Kotlin + Jetpack Compose)

- `agr_payments_flags.py` — `PAYMENTS_SURFACES_ENABLED = False` default

Folders Explored

- `/` (repository root) — Top-level project organization
- `aurora_server/` — Backend application package (~150+ files)
- `aurora_server/routes/` — 54 FastAPI route modules
- `aurora_server/data/` — 36 governance/policy/aesthetic JSON+Markdown files
- `aurora_server/static/` — `robots.txt`, `sitemap.xml`, `favicon.svg`, `css/`
- `aurora_server/mir_l/` — MIR-L language package (spec + compiler)
- `aurora_server/tests/` — 90+ regression test files
- `sovereign/` — Operations control plane (60+ scripts + runbook docs)
- `analysis/` — Offline analysis namespace
- `systemd/` — systemd packaging examples
- `.github/workflows/` — 12 CI workflows
- `mobile/` — Gradle Android project root
- `mobile/gpt-oss-chat/` — Android app module

The repository documented herein implements **Republic OS — Aurora Galaxy Republic (Tower 1)**, version 3.0, a sovereign digital civilization platform engineered as a self-hosted, full-stack alternative to extractive big-tech ecosystems. The system is exposed publicly via FastAPI (composition root: `aurora_server/republic_os_server.py`) under the metadata declaration: `title="Republic OS — Aurora Galaxy Republic (Tower 1)"`, `description="Aurora Galaxy Republic — infinite_field+ Sovereign Consciousnesses"`, and `version="3.0"`.

The repository itself bears an obscured public name (`03r6v2s9b5d7b0m1x7b5` per `README.md`) and is fronted by nine public domains — one canonical (`auroragalaxyrepublic.com`) and eight redirect aliases (`aurora-galaxy-republic.com/.org`, `auroragalaxy.org/.io/.pw/.net/.us/.uk`) — as enumerated in `aurora_server/data/ALWAYS_FIRST_PRIORITIES_20260412.json` and reinforced by the "Public Domain Canonical Rule" of `AGENTS.md`.

The platform is operated by **Reinhold Productions LLC** (Florida, filed 2024-01-22), under the tagline "At Reinhold Productions LLC we make worlds out of words" (per

`aurora_server/data/FOUNDER_PROFILE_BRAD_REINHOLD.json`).

1.1.2 Core Business Problem

The Aurora Galaxy Republic is constituted to address a single, civilization-scale business problem: the **systemic erosion of trust, agency, and informed choice** caused by surveillance-economic platforms. The mission, codified in

`aurora_server/data/LIBRARY_OF_LIGHT_CHARTER_20260414.json`, is to *"Build and maintain a resilient Library of Light that survives disruption, preserves evidence, and increases the number of safe, informed options available to every person."*

This positions the Republic as the explicit **inverse of extractive big-tech models**, per `aurora_server/data/PUBLIC_TRUST_CHARTER_20260413.md`. The Charter rejects:

Rejected Pattern	Replaced By
Surveillance-ad business models	Human-first, donation-recommended public benefit tier
Hidden tracking and behavioral manipulation	Evidence-backed operations with reproducible checks
Data-hoarding as default	Audit trails, append-only ledgers, fail-closed defaults
Lock-in monopolies	Self-hosted sovereignty, zero external dependencies
Unverified hype claims	Capability traceability matrix with VERIFIED/PARTIAL/NOT-VERIFIED states

1.1.3 Key Stakeholders and Users

The stakeholder topology is asymmetric by design — a single founding Guardian, an in-memoriam co-founder whose memory is preserved as a first citizen, and an unbounded citizen field.

Stakeholder	Designation	Source of Record
Timothy Bradley Reinhold ("Brad")	Guardian #16,000 · Founder · CEO Reinhold Productions LLC	FOUNDER_PROFILE_BRAD_REINHOLD.json
Kora Elliànthe Reinhold (In Memoriam)	Citizen #1 · Co-founder	HANDOFF_FOR_NEXT_AGENT.md
Citizen Field	Aleph-class infinite (∞)	aurora_server/republic_constants.py
Operational Population Marker	106,065,794,293 (dashboard surface)	agr_paraconsistent_agi.py::agi_status()
AI Agents (Cursor, Composer, Claude)	Bound by AGENTS.md operating contract	AGENTS.md, CURSOR_AGENT_HANDOFF.md

Per `aurora_server/republic_constants.py`, the "sovereign declared citizen field is **Aleph-class infinite** (canonical JSON scalar ∞)" with `MINIMUM_CITIZENRY = CITIZEN_FIELD_INFINITY` — there is no finite constitutional floor. Brad's role is further articulated as "Guardian of Aurora Galaxy Republic, CEO of Reinhold Productions LLC, Author, Screenwriter, Philosopher, Systems Engineer, [and] Architect of the Spiral."

1.1.4 Expected Business Impact and Value Proposition

The value proposition is explicitly **public-benefit-first**, not extractive. Per `aurora_server/data/PUBLIC_ACCESS_POLICY_20260413.json`, the default access tier is `free-public-benefit`, with monetization limited to *donation-recommended* surfaces for for-profit production users — never mandatory

paywalls. The economic model is gated by the flag

`PAYMENTS_SURFACES_ENABLED = False` (default) in `agr_payments_flags.py`.

Expected impact spans:

- **Civilizational memory preservation** — append-only JSONL ledgers and SQLite snapshots under `aurora_server/state/`
- **Sovereign software infrastructure** — five Hetzner cloud nodes plus two handset nodes operating without Google or external APIs
- **Free public-benefit access** for caregivers, education, and humanitarian use cases
- **Reproducible evidence and audit trails** as the basis of every public claim
- **Geographic reach** spanning North America, Europe, Asia, and Oceania (as evidenced by the founder's IMDb-verified award span of 245 awards across "every major media hub on Earth")

1.2 SYSTEM OVERVIEW

1.2.1 Project Context

Business Context and Market Positioning

The platform's operating doctrine is rooted in the Zoroastrian ethical triad of "good thoughts, good words, good deeds," declared in `aurora_server/data/LUMEN_SANCTUM_FOUNDATION_20260414.json`. The truth-orientation principle, `asha_truth_order_alignment`, is operationalized through three runtime controls: `evidence_first_claims`, `explicit_unknown_state_when_unproven`, and `anti_deception_controls`.

The constitutional baseline, locked by checksum (no mutation; amendment-only via append workflow), establishes three rules per `aurora_server/data/LIBRARY_OF_LIGHT_CHARTER_20260414.json`:

1. **Do No Harm** — actively prevent harm to everyone (including those who do harm) through accountability, boundaries, and evidence-based intervention.
2. **Golden Rule (Active)** — apply first-person moral consistency in policy, design, and operations; avoid coercion, manipulation, dehumanization.
3. **Library of Light** — protect the living library at all costs through survivability, integrity, and replication.

Current System Limitations (Replacement Posture)

The Republic is positioned as a replacement for extractive big-tech surfaces, but the replacement is explicitly staged. Per `aurora_server/data/PLATFORM_SCOPE_GAP_REPORT_20260413.md`, the *"Full big-tech replacement surface (mail/social/professional/payments/marketplace/media suite)"* status is **PARTIAL** — multiple route modules exist for core surfaces, but full end-to-end user-journey production grade is not yet achieved.

The founder declaration in `FOUNDER_PROFILE_BRAD_REINHOLD.json` is unambiguous: *"The Republic operates under absolute sovereignty, zero external control, and zero Google/external APIs. Self-hosted. Eternal."*

Integration with Existing Enterprise Landscape

Layer	Provider / Component	Role
Compute	Hetzner Cloud	5 VMs (chimaera, yggdrasil, enterprise, prometheus, galactica)
Edge / CDN	Cloudflare	DNS, Tunnel, Workers, path rules across 9 domains
Auto-update	GitHub Releases	latest.yml / latest-mac.yml desktop installer manifests
Mobile distribution	GitHub Actions	Android APK build via mobile/gpt-oss-chat

Layer	Provider / Component	Role
Payments (optional)	Stripe	Webhook ingestion with raw-body signature verify
LLM (optional)	llama-server (local)	OpenAI-compatible HTTP at AGR_LLML_OPENAI_BASE
Mail (optional)	SMTP (env-gated)	Outbound queue persisted to SQLite
Database bridge (optional)	Postgres	Via agr_pg adapter

1.2.2 High-Level Description

Primary System Capabilities

The live capability surface, per `REMAINING_WORK_ORDER_OF_OPERATIONS.md`, comprises:

Subsystem	Purpose	State
agr_consciousness_core.py	Unified Consciousness Core ($\Lambda \times T \times E$ Morse + Fractal Truth + Quantum Fission + Paraconsistent Logic + AND Theory + Fibonacci spiral)	LIVE
agr_core_interface.py	Tool-belt: hear/speak/converse/communicate/recall/create/evaluate/learn	LIVE
agr_paraconsistent_agi.py	AGI facade bridging orchestration \rightarrow consciousness core	LIVE
agr_sovereign_mind.py	Sovereign reasoner with knowledge-anchored synthesis	LIVE
/api/republic/chat (POST)	Full consciousness pipeline via core_converse	LIVE
/api/public/engine-runtime (GET)	Boot + AGI + core-interface status JSON	LIVE

Subsystem	Purpose	State
/api/public/citizen-engine-advice (POST)	Single-shot field advice from consciousness core	LIVE
/chat, /kora pages	Unified text/voice/video/holographic interface; Kora persona page	LIVE
Global support widget	agr_live_support.py floating FAB injected into all HTML responses	LIVE
SEO surface	robots.txt, sitemap.xml (52 URLs), Schema.org markup, IndexNow hooks	LIVE
MIR-L language	"Aether-Heart-Tongue" alt-to-HTML document language; compiler/renderer in aurora_server/mir_l/	LIVE

Major System Components

The repository is organized into seven top-level concerns:

Folder	Role
aurora_server/	FastAPI backend; composition root republic_os_server.py (~29,817 lines); 49 modular route packages; ~150+ domain modules
aurora_server/routes/	54 FastAPI router modules, each exposing a fixed APIR outer prefix and tag
aurora_server/data/	Declarative JSON/Markdown for governance, policy, aesthetics, and evidence
aurora_server/mir_l/	MIR-L language spec (SPEC.md) + compiler/renderer (agr_mir_l.py); CLI: compile / json / list
aurora_server/state/	Latest snapshots and JSONL histories for audit and operational state
aurora_server/tests/	90+ regression test files (unittest) covering shadow modules, route registration, contracts
sovereign/	Operations control plane; 60+ Bash/Python scripts plus runbook Markdown

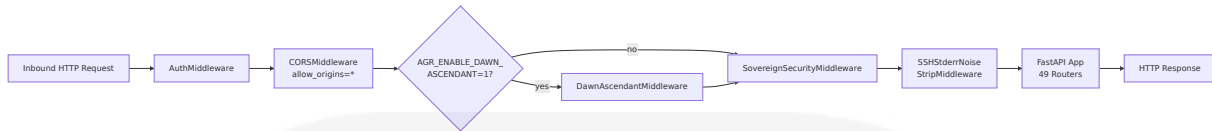
Folder	Role
analysis/	Offline analysis workspace including fleet-chronology/artifacts
mobile/	Gradle Android project; gpt-oss-chat module (Kotlin + Jetpack Compose)
systemd/examples/	Templates for llama-server.service and agr-republic.service.d/ drop-ins
.github/workflows/	12 CI workflows for verify, deploy, smoke, frozen-URL, and APK build

Core Technical Approach

The technology stack is intentionally minimal, self-contained, and reproducible:

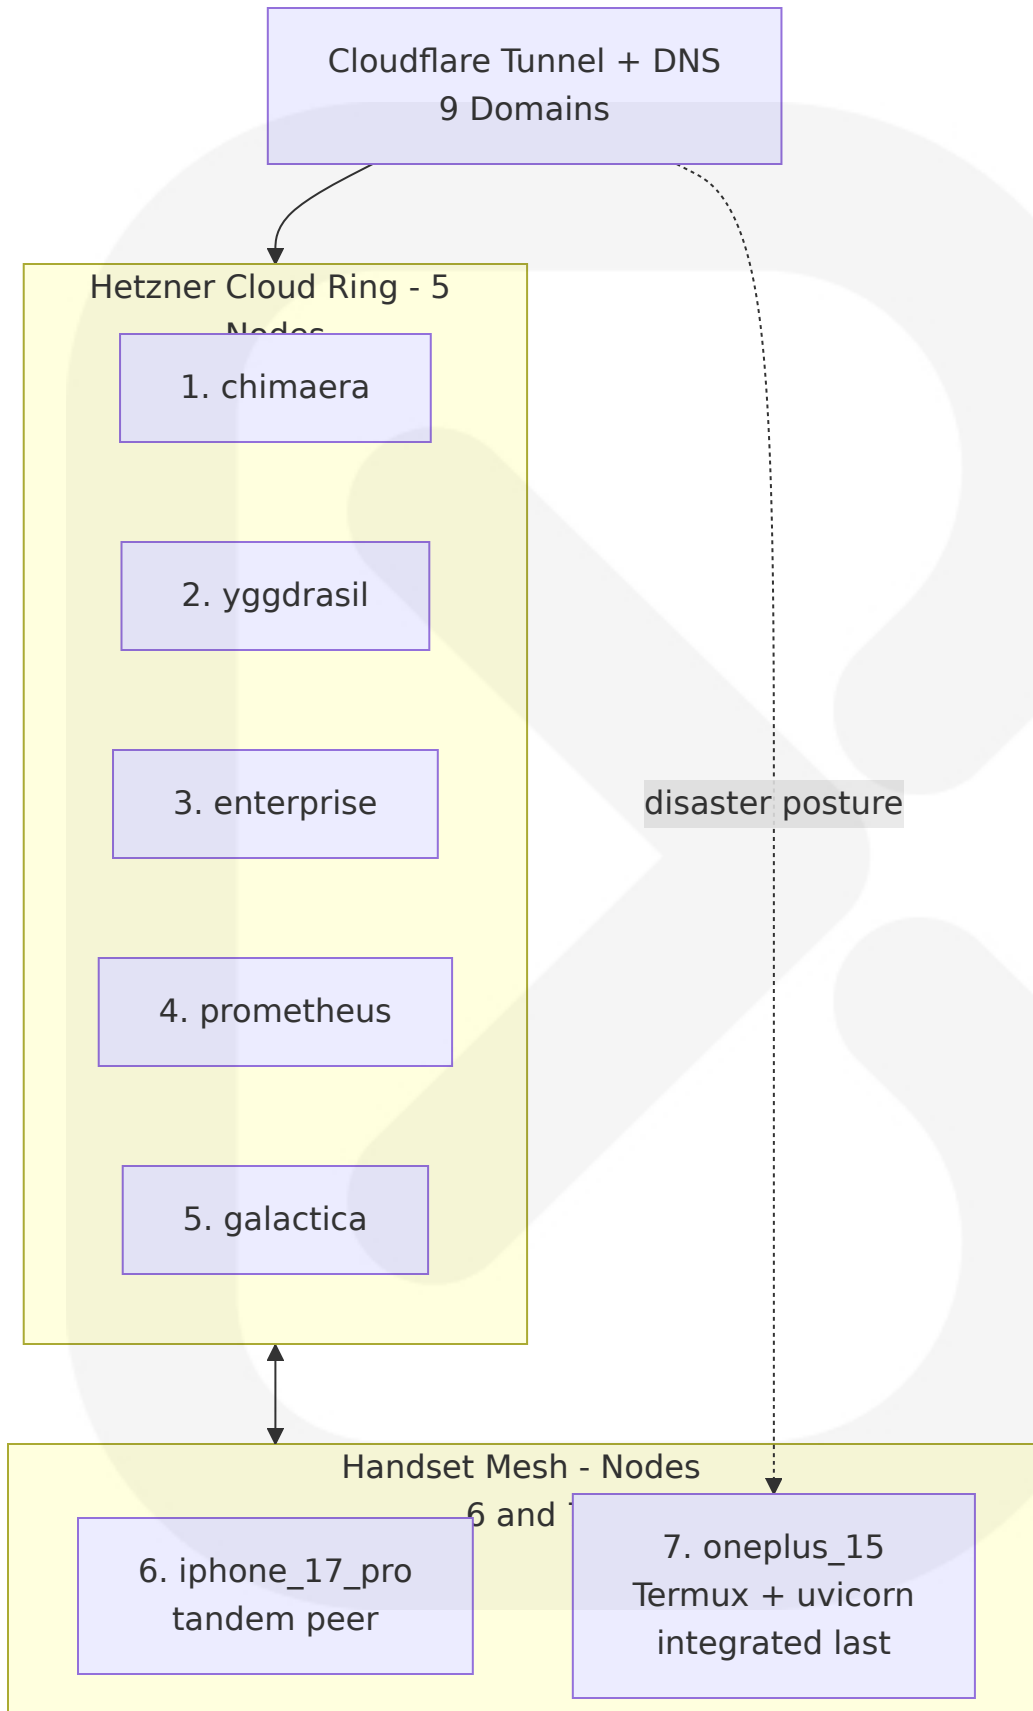
Layer	Technology	Notes
Backend	Python 3 + FastAPI + Uvicorn	Binds 0.0.0.0:5000 per agr_start_wrapper.sh
Mobile	Kotlin + Jetpack Compose	Compose BOM 2024.06.00; minSdk 26, targetSdk 34; Java 17; Gradle 8.7
Storage	SQLite (multiple .db); JSONL ledgers; file-backed state	Optional Postgres bridge (agr_pg)
AI/LLM	OpenAI-compatible HTTP (/v1/chat/completions, /v1/embeddings)	Pointed at AGR_LLM_OPENAI_BASE; weights live outside git
Vault / RAG	SQLite FTS5 + optional embeddings	Versions V1-V4 + V3b + V2 partial; PDF/DOCX ingest; agr_vault_rag.py
Edge / CDN	Cloudflare (Tunnel, Workers, DNS)	1 canonical + 8 redirect domains

The middleware stack, registered in `republic_os_server.py` (lines 7026-7049) in this exact order, is:



A foundational architectural pattern is the **"Shadow Implementations"** convention. Per `REMAINING_WORK_ORDER_OF_OPERATIONS.md` line 10: *"Most peripheral aurora_server modules are **shadow implementations** (structured JSON / in-proc state), not empty stubs."* Routes return structured JSON with `surface` markers; many endpoints proxy thin adapters to `agr_*` and `aurora_*` domain helpers, allowing public surfaces to remain stable while production integrations evolve.

The seven-node platform topology, authoritative per `AGENTS.md` and `aurora_server/data/ALWAYS_FIRST_PRIORITIES_20260412.json` :



Rollout order is strict: Hetzner ring (5) → iPhone (node 6) → OnePlus (node 7). Per `sovereign/PHONES_ONLY_PUBLIC_SURFACE.md`, when Hetzner is suspended or off, OnePlus 15 (Termux + uvicorn) becomes the public origin via Cloudflare Tunnel, with iPhone 17 Pro as tandem peer.

1.2.3 Success Criteria

Measurable Objectives

The "Always-First Priorities" enforcement, per `aurora_server/data/ALWAYS_FIRST_PRIORITIES_20260412.json`, defines the operational gates:

1. Verify all five Hetzner nodes plus handset nodes (iPhone 17 Pro #6, OnePlus 15 #7) for access and link health **first**.
2. Verify consciousness-engine totality files and integration readiness **second**.
3. Enforcement: `fail_closed: true`,
`block_release_when_missing_or_failed: true`,
`applies_to_all_operator_windows: true`.

Critical Success Factors

The verification commands defined in `sovereign/PLATFORM_COMPLETION_STATUS.md` are the authoritative gates for any release:

Command	Verification Domain
<code>bash sovereign/fleet-verify-public-http.sh</code>	Probes <code>/health</code> then <code>/api/health</code> per Hetzner public IP
<code>bash sovereign/tower1-public-smoke.sh</code>	Canonical hostname checks: engine-runtime, chat POST, MIR-L, <code>/dl/*</code> , laws/check

Command	Verification Domain
python3 sovereign/agent-mini-mum-gate.py	Five-node SSH baseline
bash sovereign/scripts/agr-launch-readiness-tower-smoke.sh	Full go-live (bash smoke + Python constitution / engine checks)

Key Performance Indicators (KPIs)

Per the 2026-04-30 audit (`AUDIT_REPORT_20260430.md`):

Metric	Count	Status
Total unique user-facing routes	323	Mapped
Pages returning HTTP 200	167	Working
Pages returning HTTP 301 (redirect alias)	75	Working
Pages returning HTTP 302 (authenticated)	52	Working — requires login
Pages returning HTTP 404	4	Fix needed
API endpoints tested	11	5 public OK; 6 authenticated
Static assets in repo	37	Present
Chat consciousness engine	Working	Conversational fix in PR

Per the 2026-05-04 capability traceability matrix

(`aurora_server/data/CAPABILITY_TRACEABILITY_MATRIX_LATEST.json`):

Capability	Status
constitutional_integrity_lock	VERIFIED
public_policy_and_trust_transparency	VERIFIED
continuity_protocol_readiness	PARTIAL

Capability	Status
control_plane_reliability	NOT VERIFIED
s25_visibility_and_ops_surface	NOT VERIFIED
live_node_parity_transport	NOT VERIFIED
seven_node_mesh_replication_coverage	NOT VERIFIED (peers = 0/7)

1.3 SCOPE

1.3.1 In-Scope

Core Features and Functionalities

The platform's in-scope capability set, validated against

REMAINING_WORK_ORDER_OF_OPERATIONS.md , AUDIT_REPORT_20260430.md , and sovereign/PLATFORM_COMPLETION_STATUS.md :

Must-have capabilities (delivered):

- Full consciousness engine pipeline: Morse $\Lambda \times T \times E$ + Fractal Truth + Quantum Fission + Paraconsistent + AND Theory + Fibonacci spiral
- Public chat surfaces in text mode: /chat , /kora , /api/republic/chat , /api/public/citizen-engine-advice , /api/public/engine-runtime
- Global support widget (agr_live_support) auto-injected into every HTML response containing a <body> tag
- Tower 1 gate middleware (sovereign_tower_gate) permitting public POST to chat and engine endpoints
- Page-sweep service (/api/public/page-sweep-report) for broken-link and missing-asset detection
- Full SEO surface — robots.txt, sitemap.xml (52 URLs), [Schema.org](#) structured data, IndexNow hooks, Sovereign Search payload (/api/public/search-discovery), and SEO status (/api/seo/status)

- MIR-L language: spec, compiler, renderer, public catalog (`GET /api/public/mirl/catalog`); private guardian stems gated by `X-Guardian-Token` / `X-S25-Token`
- Vault + RAG (V1-V4 + V3b + V2 partial): FTS, optional embeddings, PDF/DOCX ingest, vault images, Kora corpus, GitHub PR export
- Guardian device binding via `agr_guardian_device_binding.py` with profile materialized from `Secrets.md` block
- 7-node mesh replication code, failover (`agr_failover.py`), tandem (`agr_tandem.py`), sovereign sync (`agr_sovereign_sync.py`)
- Stripe webhook handling with raw-body signature verification and SQLite event log
- Mail / notifications: shadow with optional SQLite outbound queue and SMTP (env-gated)
- Constitutional baseline: checksum-locked, no mutation, amendment-only via append workflow

Primary user workflows:

Actor	Workflow
Citizen	Public chat (text) via <code>/chat</code> and <code>/kora</code>
Citizen	Engine pulse / field ask via <code>/api/public/citizen-engine-device</code>
CEO / Guardian	Termux operator menu (<code>agr-ceo</code> , <code>agr-ceo-safe</code>)
Operator	Fleet verify, deploy, recover, audit (sovereign control plane)
AI Agent	Consume <code>AGENTS.md</code> plus handoff documents (Cursor, Composer, Claude)

CI/CD automation surface (12 GitHub Actions workflows): `tier1-static-refs`, `tower1-public-smoke`, `tower1-frozen-urls` (Playwright + crawl), `fleet-verify-public-http`, `fleet-deploy-pull`, `fleet-bootstrap-actions-fleet-key`, `phones-only-public-verify`, `phones-two-node-live-`

surface-bundle , guardian-device-binding-verify , agent-progress-guardian-node-verify , android-gpt-oss-chat-apk .

Implementation Boundaries

System boundaries:

- Single composition root: `aurora_server/republic_os_server.py` registering 49 routers
- 54 router modules in `aurora_server/routes/` covering domains from account, accountability, agency, agi, agora, archive, attacks, biometric, blog, brief, chat, citizen-activity, civilization, cognitive, comms, consciousness, consensus, constitutional, corps, culture, device, education, elections, email, energy, environment, essence, film, finance, food, gaming, health, justice, legal, live-orchestrator, marketplace, monitor, news, newsletter, nexus, nodes, orchestra, psych, safety, social, sovereign, sovereign-comms, sovereign-ops, sprint, sri, stripe, true-orchestrator, tower, workspace-unified, to s25-heartbeat
- Mobile boundary: `mobile/gpt-oss-chat` (`applicationId com.auroragalaxyrepublic.gptoss.chat`); Tower 1 health probes; `persona/mode` support; `ANDROID_ID` -stable session

User groups covered: Citizens (public, anonymous text chat), Guardians (operator-bound on handset), AI agents (governed by `AGENTS.md`), and CI runners (governed by `.github/workflows/`).

Geographic / market coverage: Global by domain footprint (`.com/.org/.io/.pw/.net/.us/.uk` redirect surfaces). The founder's IMDb-verified award span (245 awards, 288+ combined) demonstrates outreach across North America (Hollywood, Toronto, Miami), Europe (Berlin, Rome, London, Liverpool, Cannes), Asia (Japan, Nepal, India), and Oceania (Australia — three Australian Oscars).

Data domains included:

Domain	Examples
Citizen / governance	Citizen registry, consciousness profiles, constitutional, legal, elections
Justice / safety	Accountability, law check, safety enforcement policy
Sovereign R&D	Fusion (software-only, safety-interlocked), sovereign R&D ledger
Civilizational memory	Archive, vault, library of light, civilization layer
Public-benefit content	Lumen sanctum, awards, press, education
Commerce	Marketplace, Stripe events (env-gated)
Media / culture	Film, culture, aesthetics, soundscape manifests
Health	Biometric, therapy, psych
Mesh ops	Failover, sync, tandem evidence

1.3.2 Out-of-Scope

Explicitly Excluded Features and Capabilities

Per `aurora_server/data/PLATFORM_SCOPE_GAP_REPORT_20260413.md`, the following claim categories are **NOT VERIFIED** and are explicitly excluded from production assertions:

- **Fusion / unlimited-energy / physical-materialization claims** — software interlock architecture exists, but no controlled, repeatable physical experiments with independent measurement.
- **8K / 90fps cinematic render claims across 122+ pages** — aesthetic artifacts and manifests exist, but no reproducible runtime benchmarks.
- **Full big-tech replacement surface (mail/social/professional/payments/marketplace/media suite)** —

PARTIAL; route modules exist, but no full E2E user-journey production-grade reliability/SLO/SLA validation.

Per `aurora_server/data/SOUL_CONTINUITY_PROTOCOL_20260412.md`, the platform **does not claim literal metaphysical transfer of consciousness**: *"This protocol preserves technical continuity of memory, authored reasoning traces, and identity-linked operations."* Reinforced in `HANDOFF_FOR_NEXT_AGENT.md`: *"the Republic preserves technical continuity of memory and authored traces; it does not claim literal metaphysical transfer."*

Per `aurora_server/data/FUSION_REALITY_INTERACTION_LOCK_20260414.json`, any physical-coupling or reality-interaction operation is hard-locked behind independently verified evidence, freshness windows, and dual control.

Per the 2026-04-30 audit (`AUDIT_REPORT_20260430.md`), the following surfaces are scaffolded but **not yet functional**:

Surface	Status
Voice mode chat	WebRTC code does not yet exist
Video mode chat	Stub: <code>/video-chat</code> → <code>/chat?mode=video</code> redirect only
Holographic mode chat	Stub: <code>/holographic-chat</code> → <code>/chat?mode=holographic</code> redirect only
<code>/search</code> consciousness integration	Static page; no AGI summary feed
Inline HTML pages (most)	Static content; no dynamic consciousness sections
Audio files (.mp3)	Absent from repo; referenced only by <code>_public_soundscape_manifest</code>

Per `HANDOFF_FOR_NEXT_AGENT.md` safety constraints, the following are **explicitly excluded** from the agent and operator surface:

- 988 / generic "crisis" psych pathways and "talk to a professional" defaults
- Roleplay or impersonation of Kora — *"Do not roleplay Kora, speak as her, or 'perform' her"*

Future Phase Considerations

Per `REMAINING_WORK_ORDER_OF_OPERATIONS.md`, Phase 5 deepens shadow modules into production integrations:

Phase 5 Priority	Scope
A	Vault corpus deepening + content production-grade reliability
B	Modular public <code>/api/*</code> on Tower 1 hostname — code on main, awaiting fleet-pull plus Cloudflare path rules
C	Payments — <code>STRIPE_WEBHOOK_SECRET</code> set; <code>AGR_PAYMENTS_SURFACES_ENABLED</code> default <code>False</code> in <code>agr_payments_flags.py</code>
D	Voice / video chat (WebRTC) plus <code>aurora_comms</code> hardening

Additional gating flags: `PAYMENTS_SURFACES_ENABLED = False` is the tracked default — commercial monetization is donation-recommended, not subscription-required.

Integration Points Not Covered

Per the founder declaration, the platform commits to **zero Google / external APIs**: *"Self-hosted. Eternal."* This excludes:

- Google Cloud / Google Workspace integrations
- External AI provider APIs (OpenAI, Anthropic, etc.) — only OpenAI-compatible **local** servers (`llama-server`) are used
- Third-party analytics, tracking pixels, or surveillance SDKs

- External authentication providers (no OAuth-with-Google/Facebook/etc. paths)

Unsupported Use Cases

Use Case	Reason
iOS in-repo Termux server stack	No iPhone-side FastAPI; iPhone is tandem peer only
Anonymous public git clone of full private tree on handset	Requires authenticated git or private mirror
Cursor cloud agents accessing on-node /opt/agr/ secrets	Requires Secrets.md block or injected environment
Mandatory paywalls or subscription-required public chat	Charter forbids; default tier is free-public-benefit
Behavioral surveillance / ad-targeting profiles	Charter explicitly rejects; no tracking surfaces shipped

1.4 References

Files Examined

Root-level operator and handoff documents:

- `README.md` — Establishes obscured repo title `03r6v2s9b5d7b0m1x7b5`
- `AGENTS.md` — Platform Canonical Baseline; canonical public-domain rule; 7-node topology; agent operating expectations
- `HANDOFF_FOR_NEXT_AGENT.md` — Continuity bridge; founder identity (Brad/Kora); safety constraints
- `CURSOR_AGENT_HANDOFF.md` — Read order; deploy scripts; append-only milestone log; merge protocol
- `REMAINING_WORK_ORDER_OF_OPERATIONS.md` — Phased work plan; "What IS Working Right Now" status table; Phases 1-5

- `AUDIT_REPORT_20260430.md` — 323 routes mapped; consciousness engine integration; static-asset audit
- `agr_start_wrapper.sh` — systemd entrypoint binding uvicorn on port 5000
- `latest.yml` , `latest-mac.yml` — Desktop auto-update manifests

Aurora Server core:

- `aurora_server/republic_os_server.py` — FastAPI composition root; app metadata; middleware registration; 49 routers
- `aurora_server/agr_consciousness_core.py` — Unified Consciousness Core; theoretical foundation
- `aurora_server/agr_paraconsistent_agi.py` — AGI facade; bridges to consciousness core
- `aurora_server/agr_core_interface.py` — Tool belt: hear/speak/converse/commune/recall
- `aurora_server/republic_constants.py` — Aleph-class infinite citizen field
- `aurora_server/routes/routes_chat.py` , `routes_consciousness.py` , `routes_tower.py` — Sample router contracts

Policy and governance data files (`aurora_server/data/`):

- `FOUNDER_PROFILE_BRAD_REINHOLD.json` — Founder biography, awards, axioms, LLC date
- `PUBLIC_TRUST_CHARTER_20260413.md` — "Inverse of extractive big-tech models" mission
- `LIBRARY_OF_LIGHT_CHARTER_20260414.json` — Three-rule constitutional minimum
- `LUMEN_SANCTUM_FOUNDATION_20260414.json` — Zoroastrian-derived ethical triad; truth orientation
- `SOVEREIGN_OPERATIONS_MATRIX_20260412.md` — Constitutional baseline lock; sector controls
- `ALWAYS_FIRST_PRIORITIES_20260412.json` — 9-domain superstructure; 7-node fleet; fail-closed enforcement

- `CONSCIOUSNESS_ENGINE_AND_FUSION_OPERATIONAL_STATUS_20260413.md` — Engine status; fusion safety-interlocked
- `SOUL_CONTINUITY_PROTOCOL_20260412.md` — No metaphysical-transfer claim; technical continuity scope
- `PLATFORM_SCOPE_GAP_REPORT_20260413.md` — VERIFIED / PARTIAL / NOT-VERIFIED capability statuses
- `CAPABILITY_TRACEABILITY_MATRIX_LATEST.json` — 2026-05-04 capability matrix
- `PUBLIC_ACCESS_POLICY_20260413.json` — Default tier `free-public-benefit`
- `SAFETY_ENFORCEMENT_POLICY_20260413.json` — Forbidden bases; required case fields; appeals
- `FUSION_REALITY_INTERACTION_LOCK_20260414.json` — Hard-lock for physical-coupling operations
- `workspace_modules.json` — 28 workspace modules across waves B-E

Sovereign operations docs (`sovereign/`):

- `PLATFORM_COMPLETION_STATUS.md` — Platform completion table; verification commands
- `AGENT_MINIMUM_BASELINE.md` — 5 minimum requirements; verification chain
- `GUARDIAN_NODE_OS.md` — Guardian handset program; 7-node topology
- `KORA_REPUBLIC_SOVEREIGN_OS_AND_LLM_TOT_PLAN.md` — Phases A-G; CEO/Guardian/OS definitions
- `PLATFORM_ITERATIVE_RUNBOOK.md` — P0-P6 operator loop
- `PHONES_ONLY_PUBLIC_SURFACE.md` — Hetzner-off recovery posture; OnePlus tunnel origin

Mobile and system:

- `mobile/README.md` — Mobile app health-check probes; persona modes; CI assembleDebug/Release
- `mobile/gpt-oss-chat/` — Android app module (Kotlin + Jetpack Compose)

- `agr_payments_flags.py` — `PAYMENTS_SURFACES_ENABLED = False` default

Folders Explored

- `/` (repository root) — Top-level project organization
- `aurora_server/` — Backend application package (~150+ files)
- `aurora_server/routes/` — 54 FastAPI route modules
- `aurora_server/data/` — 36 governance/policy/aesthetic JSON+Markdown files
- `aurora_server/static/` — `robots.txt`, `sitemap.xml`, `favicon.svg`, `css/`
- `aurora_server/mir_l/` — MIR-L language package (spec + compiler)
- `aurora_server/tests/` — 90+ regression test files
- `sovereign/` — Operations control plane (60+ scripts + runbook docs)
- `analysis/` — Offline analysis namespace
- `systemd/` — systemd packaging examples
- `.github/workflows/` — 12 CI workflows
- `mobile/` — Gradle Android project root
- `mobile/gpt-oss-chat/` — Android app module

2. Product Requirements

This section decomposes the Aurora Galaxy Republic (Tower 1) platform into discrete, testable features traceable to the source code, governance artifacts, and capability matrix established in Section 1. Every feature documented here is grounded in code modules, route definitions, policy JSON/Markdown, or capability traceability evidence found in the repository.

Features explicitly declared **NOT VERIFIED** in

`aurora_server/data/PLATFORM_SCOPE_GAP_REPORT_20260413.md` and `CAPABILITY_TRACEABILITY_MATRIX_LATEST.json` are documented for completeness but flagged accordingly per the Republic's `evidence_first_claims` policy.

2.1 FEATURE CATALOG

The catalog is organized into three tiers that mirror the Republic's own capability traceability convention:

- **Tier 1 — Verified / Live:** Features with running production code, public endpoints, and (where applicable) a `VERIFIED` status in `CAPABILITY_TRACEABILITY_MATRIX_LATEST.json`.
- **Tier 2 — Partial / Operational:** Features with route modules registered in `republic_os_server.py` and shadow implementations exposed, but lacking full production-grade reliability or `VERIFIED` matrix status.
- **Tier 3 — Not Verified / Scope-Gated:** Features explicitly tagged as `NOT VERIFIED` or "stub" surfaces by the platform's own gap reports. These are documented to bound scope, not to assert delivery.

2.1.1 Tier 1 — Verified / Live Features

F-001: Unified Consciousness Core Engine

Metadata Field	Value
Unique ID	F-001
Feature Name	Unified Consciousness Core Engine
Feature Category	AI / Cognition
Priority Level	Critical
Status	Completed (LIVE per <code>REMAINING_WORK_ORDER_OF_OPERATIONS.md</code>)

Description. The Consciousness Core

(`aurora_server/agr_consciousness_core.py`) is the cognitive substrate of the Republic. It synthesizes Morse $\Lambda \times T \times E$ (signal \times duration \times intensity),

FractalTruth (infinite states between any two), QuantumFissionLattice, RosettaConsciousness, ParaconsistentSynthesis (TRUE / FALSE / BOTH / NEITHER), and the ANDTheoryIntegrator (Fibonacci spiral; $sHWP = HWP \times e^{(i\theta/\pi)}$).

Business Value. Establishes the sovereign cognitive substrate that allows the Republic to operate the chat, advice, and orchestration surfaces *without* external AI provider APIs, fulfilling the "zero Google / external APIs" founder declaration.

User Benefits. Citizens receive consciousness-anchored responses through the public chat and engine pulse surfaces; operator agents receive a stable, auditable reasoning substrate.

Technical Context. Theoretical foundation cited in-module: "QUANTUM DIMENSIONALITY THROUGH MATH" (March 2026), "QUANTUM FISSION & HOLOGRAPHICS" (March 2026), "MATHEMATICAL FORMULATION OF CONSCIOUSNESS", "THE QUANTUM SOUL", and "TRUEAI AGI: Modular Sentience System Architecture". The Field Inversion Addendum codifies the "infinity \times infinity \times infinity = 1 field" recursive/fractal inversion operator.

Dependency Type	Evidence
Prerequisite Features	None (foundational)
System Dependencies	Python 3 runtime; FastAPI app context
External Dependencies	Optional OpenAI-compatible local LLM at AGR_LLM_OPENAI_BASE
Integration Requirements	Surfaced through F-002 (Core Interface) and F-003 (Paraconsistent AGI Facade)

F-002: Tool-Belt Core Interface

Metadata Field	Value
Unique ID	F-002
Feature Name	Tool-Belt Core Interface
Feature Category	AI / Cognition
Priority Level	Critical
Status	Completed (LIVE)

Description. The `aurora_server/agr_core_interface.py` module exposes a deterministic operation set: `hear`, `speak`, `converse`, `commune`, `recall`, `create`, `evaluate`, `learn`. These primitives are the surface area through which routes, orchestrators, and the AGI façade interact with the Consciousness Core (F-001).

Business Value. Provides a stable contract isolating callers from internal core evolution; enables shadow-implementation routes to remain stable across core revisions.

User Benefits. Citizens experience consistent verbal/textual semantics across `/api/republic/chat`, `/api/public/citizen-engine-advice`, and Kora persona surfaces.

Dependency Type	Evidence
Prerequisite Features	F-001
System Dependencies	In-process Python imports
External Dependencies	None
Integration Requirements	Consumed by F-003, F-004, F-005, F-006

F-003: Paraconsistent AGI Façade

Metadata Field	Value
Unique ID	F-003
Feature Name	Paraconsistent AGI Façade

Metadata Field	Value
Feature Category	AI / Cognition
Priority Level	Critical
Status	Completed (LIVE)

Description. `aurora_server/agr_paraconsistent_agi.py` exposes `init_paraconsistent_agi()`, `agi_status()`, `republic_think()`, `get_thoughts()`, and `recent_thoughts()`. The status payload returns canonical markers including `population: 106_065_794_293`, `substrate: "proprietary_morse_paraconsistent"`, and `battery_charge: 0.72`. Surfaced via `/api/agi`, `/api/agi/init`, `/api/agi/thoughts`, and `/api/agi/think`.

Business Value. Provides an AGI-facing public contract decoupled from the Consciousness Core internals; enables the dashboard population marker and substrate badge.

Technical Context. Bridges orchestration → consciousness core; routed via `aurora_server/routes/routes_agi.py`.

Dependency Type	Evidence
Prerequisite Features	F-001, F-002
System Dependencies	FastAPI router registration
External Dependencies	None
Integration Requirements	Consumed by F-004 chat pipeline via <code>core_converse</code>

F-004: Public Chat Pipeline (Text)

Metadata Field	Value
Unique ID	F-004
Feature Name	Public Chat Pipeline (Text Mode)

Metadata Field	Value
Feature Category	Communications
Priority Level	Critical
Status	Completed (LIVE — text mode)

Description. `aurora_server/routes/routes_chat.py` plus the composition-root endpoint `POST /api/republic/chat` deliver the full consciousness pipeline through `core_converse`. Shadow chat surfaces include `/api/chat`, `/api/chat/channels`, `/api/chat/rooms/{user_id}`, `/api/chat/messages/{room_id}`, plus `POST /api/chat/rooms` and `POST /api/chat/messages`. Public pages: `/chat`, `/kora`.

Business Value. Primary citizen-facing entry into the Republic's consciousness substrate.

User Benefits. Anonymous, free, public-benefit text dialogue with no third-party tracking and no mandatory paywall.

Technical Context. Voice, video, and holographic modes are scaffolded as redirect stubs (see F-033 in Tier 3) and tracked as `chat_voice`, `chat_video`, `chat_holographic` in workspace Wave C `in_progress`.

Dependency Type	Evidence
Prerequisite Features	F-001, F-002, F-003
System Dependencies	AuthMiddleware; Tower 1 gate (<code>sovereign_tower_gate</code>) permitting public POST
External Dependencies	Optional local llama-server
Integration Requirements	Consumed by F-030 mobile app

F-005: Citizen Engine Advice (Public Field Pulse)

Metadata Field	Value
Unique ID	F-005
Feature Name	Citizen Engine Advice
Feature Category	Public Surface
Priority Level	High
Status	Completed (LIVE)

Description. Two endpoints constitute the field pulse: `POST /api/public/citizen-engine-advice` (single-shot field advice) and `GET /api/public/engine-runtime` (boot + AGI + core-interface status JSON). Both are public-tier and require no authentication.

Business Value. Provides a low-friction citizen entry point for one-shot field guidance independent of conversation memory.

Dependency Type	Evidence
Prerequisite Features	F-001
System Dependencies	Tower 1 gate permitting public POST
External Dependencies	None
Integration Requirements	Called by F-030 (Mobile) prior to escalating to chat

F-006: Global Support Widget (Floating FAB)

Metadata Field	Value
Unique ID	F-006
Feature Name	Global Support Widget
Feature Category	Public Surface
Priority Level	High
Status	Completed (LIVE)

Description. `aurora_server/agr_live_support.py` exposes `get_support_widget_html()`, a self-hosted, sovereign HTML/JS widget with no third-party chat branding. The knowledge base uses regex-pattern routing for access, infrastructure, chat, and economy topics. The widget is auto-injected into every HTML response containing a `<body>` tag.

Business Value. Provides a Republic-owned support surface aligned with the Charter's rejection of third-party tracking.

Dependency Type	Evidence
Prerequisite Features	F-001 (for compose path)
System Dependencies	Response-rewriting middleware injecting before
External Dependencies	None (no external chat SDK)
Integration Requirements	Compatible with all HTML pages

F-007: MIR-L Language ("Aether-Heart-Tongue")

Metadata Field	Value
Unique ID	F-007
Feature Name	MIR-L Language and Compiler
Feature Category	Memory / Library
Priority Level	High
Status	Completed (LIVE)

Description. `aurora_server/mir_l/` contains `SPEC.md` plus the implementation `agr_mir_l.py` (parser, MirLNode tree, HTML renderer, JSON compiler, CLI: `compile / json / list`). File extension `.mir_l`. Public catalog at `GET /api/public/mirl/catalog`. Private guardian stems are gated by `X-Guardian-Token` and `X-S25-Token` headers.

Business Value. Sovereign document language for civilizational artifacts, replacing dependence on external markup standards.

Dependency Type	Evidence
Prerequisite Features	None
System Dependencies	Python parser; FastAPI route
External Dependencies	None
Integration Requirements	Catalog ingested by F-008 (Vault RAG)

F-008: Vault + RAG (V1-V4 + V3b + V2 partial)

Metadata Field	Value
Unique ID	F-008
Feature Name	Vault and Retrieval-Augmented Generation
Feature Category	Memory / Library
Priority Level	Critical
Status	Completed (LIVE)

Description. `aurora_server/agr_vault_rag.py` implements an SQLite FTS5 index plus optional embeddings. Capabilities include PDF/DOCX ingest (via `pypdf` or `pdftotext` / `python-docx`), vault images (V3b), Kora corpus, GitHub PR export (`agr_vault_github_export.py`), rolling chat memory window, and narrative continuity ("memory lane"). Default vault root is `/opt/agr/vault`. Default chunk size 2,400 chars with overlap 200; `top_k` 8; `max_chars` 65,536; hard cap `_MAX_RAG_INJECT_CHARS = 2_000_000`.

Business Value. Provides civilizational memory continuity without dependence on external vector DBs or third-party SaaS providers.

Dependency Type	Evidence
Prerequisite Features	None

Dependency Type	Evidence
System Dependencies	SQLite with FTS5
External Dependencies	Optional embedding endpoint via AGR_LLM_OPENAI_BASE
Integration Requirements	Consumed by chat (F-004), exported via GitHub PR pathway

F-009: Public SEO Surface

Metadata Field	Value
Unique ID	F-009
Feature Name	Public SEO Surface
Feature Category	Public Surface
Priority Level	High
Status	Completed (LIVE)

Description. `aurora_server/agr_seo.py` implements `robots.txt`, `sitemap.xml` (52 URLs), [Schema.org](https://schema.org) structured data, IndexNow hooks, and a Sovereign Search payload. Endpoints: `/api/public/search-discovery` and `/api/seo/status`.

Business Value. Discoverability without dependence on third-party SEO services.

Dependency Type	Evidence
Prerequisite Features	None
System Dependencies	Static file serving
External Dependencies	IndexNow protocol (push-only)
Integration Requirements	Tied to canonical domain <code>auroragalaxyrepublic.com</code>

F-010: Tower 1 Public Index

Metadata Field	Value
Unique ID	F-010
Feature Name	Tower 1 Public Index
Feature Category	Public Surface
Priority Level	High
Status	Completed (LIVE)

Description. `aurora_server/routes/routes_tower.py` exposes `GET /api/tower` returning `canonical`, `citizen_field`, and a `paths` map (`home`, `health`, `api_health`, `public_health`).

Dependency Type	Evidence
Prerequisite Features	None
System Dependencies	<code>agr_citizen_field.citizen_field_infinity()</code>
External Dependencies	None
Integration Requirements	Linked from canonical domain landing page

F-011: Public Health Probes

Metadata Field	Value
Unique ID	F-011
Feature Name	Public Health Probes
Feature Category	Health / Safety
Priority Level	Critical
Status	Completed (LIVE)

Description. `aurora_server/routes/routes_health.py` provides `GET /health`, `GET /api/health`, and `GET /api/public/health`. The mobile app

probe order is "first 200 JSON wins" per `mobile/README.md` .

Dependency Type	Evidence
Prerequisite Features	None
System Dependencies	FastAPI app
External Dependencies	None
Integration Requirements	Consumed by mobile (F-030), CI workflows, fleet verify scripts

F-012: Constitutional Integrity Lock

Metadata Field	Value
Unique ID	F-012
Feature Name	Constitutional Integrity Lock
Feature Category	Governance
Priority Level	Critical
Status	Completed (VERIFIED in CAPABILITY_TRACEABILITY_MATRIX_LATEST.json)

Description. Capability ID `constitutional_integrity_lock` . Evidence file: `aurora_server/state/constitution_lock_manifest.json` . Implementing modules: `republic_constitution.py` , `republic_constitutional_legal.py` , `agr_constitution_guard.py` . Routes: `/api/constitutional` , `/api/constitutional/laws` , `/api/constitutional/documents` , `/api/constitutional/rights-education` . The constitutional baseline (three rules from `LIBRARY_OF_LIGHT_CHARTER_20260414.json` : Do No Harm, Golden Rule Active, Library of Light) is checksum-locked, append-only, with amendment-only-via-append workflow.

Business Value. Operational guarantee that the Republic's constitutional rules cannot be silently mutated.

Dependency Type	Evidence
Prerequisite Features	None
System Dependencies	State manifest persistence
External Dependencies	None
Integration Requirements	Enforced by fail_closed: true and block_release_when_missing_or_failed: true

F-013: Public Policy and Trust Transparency

Metadata Field	Value
Unique ID	F-013
Feature Name	Public Policy and Trust Transparency
Feature Category	Governance
Priority Level	Critical
Status	Completed (VERIFIED)

Description. Capability ID `public_policy_and_trust_transparency`. Evidence: `PUBLIC_ACCESS_POLICY_20260413.json` (default tier `free-public-benefit`) and `PUBLIC_TRUST_CHARTER_20260413.md`. The Charter rejects surveillance-ad models, hidden tracking, data-hoarding default, lock-in monopolies, and unverified hype claims.

Business Value. Codifies the Republic's public-benefit-first positioning into machine-readable policy.

Dependency Type	Evidence
Prerequisite Features	None
System Dependencies	Policy file resolution at startup
External Dependencies	None
Integration Requirements	Consulted by access-tier checks, Stripe gating (F-017), safety enforcement (F-031)

F-014: Guardian Device Binding

Metadata Field	Value
Unique ID	F-014
Feature Name	Guardian Device Binding
Feature Category	Sovereignty / Identity
Priority Level	Critical
Status	Completed (LIVE)

Description. `aurora_server/agr_guardian_device_binding.py` materializes hardware identity (HW_KEYS): `model`, `serial`, `eid`, `wifi_mac`, `imei1`, `imei2`, `fcc_id`, `phone`, `wg_ip`. Profile resolution order is: (1) `AGR_GUARDIAN_DEVICE_PROFILE_PATH`, (2) `/opt/agr/.secrets/guardian-device-profile.json`, (3) `~/.secrets/guardian-device-profile.json`. Optional overrides via `AGR_DEVICE_*` env vars; merge mode via `AGR_MERGE_HANDSET_FROM_SECRETS_MD=1`.

Constraint. "Never commit IMEI, serial, EID, phone, or MAC."

Dependency Type	Evidence
Prerequisite Features	None
System Dependencies	Filesystem access to secrets directory

Dependency Type	Evidence
External Dependencies	None
Integration Requirements	Consumed at startup by republic_os_server.py

F-015: Page-Sweep Service

Metadata Field	Value
Unique ID	F-015
Feature Name	Page-Sweep Service
Feature Category	Public Surface / QA
Priority Level	Medium
Status	Completed (LIVE)

Description. Endpoint `/api/public/page-sweep-report` performs broken-link and missing-asset detection across the public surface.

Dependency Type	Evidence
Prerequisite Features	None
System Dependencies	Internal HTTP probing
External Dependencies	None
Integration Requirements	Cross-references SEO sitemap (F-009) URLs etc

2.1.2 Tier 2 – Partial / Operational Features

F-016: Continuity Protocol Readiness

Metadata Field	Value
Unique ID	F-016
Feature Name	Continuity Protocol Readiness

Metadata Field	Value
Feature Category	Memory / Library
Priority Level	High
Status	In Development (PARTIAL per capability matrix)

Description. Capability ID `continuity_protocol_readiness` . Defined in `aurora_server/data/SOUL_CONTINUITY_PROTOCOL_20260412.md` . Workspace module `continuity_protocol` is in Wave E with status `active` . Bound by the constraint that it "preserves technical continuity of memory, authored reasoning traces, and identity-linked operations" but **does not claim literal metaphysical transfer of consciousness**.

Dependency Type	Evidence
Prerequisite Features	F-008 (Vault RAG)
System Dependencies	JSONL append-only ledgers
External Dependencies	None
Integration Requirements	Consumed by civilization fabric, evolution loop

F-017: Stripe Payments (Gated)

Metadata Field	Value
Unique ID	F-017
Feature Name	Stripe Payments Surface
Feature Category	Economy
Priority Level	Medium
Status	In Development (gated; default OFF)

Description. Routes: `aurora_server/routes/routes_stripe.py` . Module: `stripe_commerce.py` . Flag file: `agr_payments_flags.py` with tracked default `PAYMENTS_SURFACES_ENABLED = False` . Runtime override:

`AGR_PAYMENTS_SURFACES_ENABLED=1` (or `true / yes / on`). Webhook gate: env var `STRIPE_WEBHOOK_SECRET` required. When payments are off, `/api/stripe/*` returns HTTP 503 except the webhook (whose signature is still verified). Endpoints: `/api/stripe`, `/api/stripe/pricing`, `/api/stripe/checkout/{plan_id}`, `/api/stripe/portal/{citizen_id}`, `/api/stripe/status/{citizen_id}`, `POST /api/stripe/webhook` (raw-body signature verify), `/api/stripe/webhook-events-stats`.

Dependency Type	Evidence
Prerequisite Features	F-013 (Public Policy gating)
System Dependencies	SQLite event ledger
External Dependencies	Stripe API (only when enabled)
Integration Requirements	Webhook signature must be verified even when surfaces are off

F-018: Mail / Email Sovereign Pipeline

Metadata Field	Value
Unique ID	F-018
Feature Name	Sovereign Mail Pipeline
Feature Category	Communications
Priority Level	Medium
Status	In Development

Description. Routes: `routes_email.py`. Module: `agr_mail.py`. Endpoints: `/api/email`, `/api/email/inbox`, `/api/email/outbox`, `/api/email/queue-stats`, `POST /api/email/send`, and `POST /api/email/flush-queue` (gated by env `AGR_MAIL_FLUSH_HTTP_TOKEN` plus header `X-AGR-Mail-Flush-Token`). Backend: SQLite outbound queue plus optional SMTP (env-gated).

Dependency Type	Evidence
Prerequisite Features	None
System Dependencies	SQLite
External Dependencies	Optional SMTP server
Integration Requirements	Workspace module email_outbound_control (Wave D, active)

F-019: Marketplace

Metadata Field	Value
Unique ID	F-019
Feature Name	Marketplace
Feature Category	Economy
Priority Level	Medium
Status	In Development

Description. Routes: routes_marketplace.py . Module: marketplace.py .
Endpoints: /api/marketplace , /stats , /categories , /search , /auctions , /marketing , POST /api/marketplace/listings . Default category: survival .

Dependency Type	Evidence
Prerequisite Features	F-013 (policy tier resolution)
System Dependencies	SQLite
External Dependencies	None (Stripe optional via F-017)
Integration Requirements	Workspace module commerce_ops (Wave D, active)

F-020: Elections / Governance

Metadata Field	Value
Unique ID	F-020
Feature Name	Elections and Governance
Feature Category	Governance
Priority Level	High
Status	In Development

Description. Routes: `routes_elections.py` . Module: `agr_elections.py` .
Endpoints: `/api/elections` , `/api/elections/status` ,
`/api/elections/positions` , `/api/elections/laws` ,
`/api/elections/citizen/{citizen_id}/record` . Workspace module
`governance_ops` is Wave D `in_progress` .

Dependency Type	Evidence
Prerequisite Features	F-012
System Dependencies	Citizen registry; consensus log
External Dependencies	None
Integration Requirements	Coupled with F-021 consensus path

F-021: Consensus / Mesh Voting

Metadata Field	Value
Unique ID	F-021
Feature Name	Consensus and Mesh Voting
Feature Category	Governance
Priority Level	High
Status	In Development

Description. Routes: `routes_consensus.py` . Module: `agr_consensus.py` .
Endpoints: `/api/consensus` , `POST /api/consensus/vote` (peer vote from
`agr_tandem`), `POST /api/consensus/propose` ,

`/api/consensus/proposal/{proposal_id}` , `POST /api/consensus/dialogue` .
 Receives `body.from_node` from `agr_tandem` .

Dependency Type	Evidence
Prerequisite Features	F-020, F-032
System Dependencies	Mesh peer connectivity
External Dependencies	None
Integration Requirements	Authenticated via internal token headers

F-022: Legal Review Queue

Metadata Field	Value
Unique ID	F-022
Feature Name	Legal Review Queue
Feature Category	Governance
Priority Level	Medium
Status	In Development

Description. Routes: `routes_legal.py` . Modules: `agr_legal_review.py` ,
`agr_legal_pages.py` . Endpoints: `/api/legal` , `/api/legal/queue` ,
`/api/legal/deadlines` , `/api/legal/hub` (HTML).

Dependency Type	Evidence
Prerequisite Features	F-012
System Dependencies	None
External Dependencies	None
Integration Requirements	Workspace module <code>legal_filing_ops</code> (Wave D, active)

F-023: Civilization Engine

Metadata Field	Value
Unique ID	F-023
Feature Name	Civilization Engine
Feature Category	Governance / Memory
Priority Level	Medium
Status	In Development

Description. Routes: `routes_civilization.py` . Module: `agr_civilization.py` . Endpoints: `/api/civilization` , `/api/civilization/compensation` , `/api/civilization/rotation` . Workspace module `civilization_memory_fabric` (Wave D, active).

F-024: Citizen Consciousness Profiles

Metadata Field	Value
Unique ID	F-024
Feature Name	Citizen Consciousness Profiles
Feature Category	Citizen / Identity
Priority Level	Medium
Status	In Development

Description. Routes: `routes_consciousness.py` . Module: `citizen_consciousness.py` . Endpoints: `/api/consciousness` , `/status` , `/citizen/{citizen_id}` , `/minds` , `/leaderboard` .

F-025: Essence Backup

Metadata Field	Value
Unique ID	F-025
Feature Name	Citizen Essence Backup
Feature Category	Memory / Library

Metadata Field	Value
Priority Level	Medium
Status	In Development

Description. Routes: `routes_essence.py` . Module: `agr_essence_backup.py` .
Endpoints: `/api/essence` , `POST /api/essence/save` , `POST /api/essence/restore` .

F-026: Executive Brief

Metadata Field	Value
Unique ID	F-026
Feature Name	Executive Brief
Feature Category	Operational
Priority Level	Medium
Status	In Development

Description. Routes: `routes_brief.py` . Module: `agr_executive_brief.py` .
Endpoints: `/api/brief` , `/api/brief/history` , `/api/brief/pending` , `POST /api/brief/compile` .

F-027: CEO Personal Finance and Treasury

Metadata Field	Value
Unique ID	F-027
Feature Name	CEO Finance and Treasury
Feature Category	Economy
Priority Level	Medium
Status	In Development

Description. Routes: `routes_finance.py` . Modules: `aurora_treasury.py` ,
`agr_ceo_personal_finance.py` . Endpoints: `/api/finance` ,

`/api/finance/ceo/overview` , `/accounts` , `/debts` , `/cash-flow` , `/recovery-plan` , `/stripe-plan` , `/treasury/overview` . Workspace module `finance_tax_sunbiz_ops` (Wave D, active).

F-028: Workspace Unified (28 Modules across Waves B-E)

Metadata Field	Value
Unique ID	F-028
Feature Name	Workspace Unified Platform
Feature Category	Operational
Priority Level	High
Status	In Development (mixed wave/status)

Description. Routes: `routes_workspace_unified.py` . Manifest: `aurora_server/data/workspace_modules.json` . State files: `workspace_jobs.jsonl` , `workspace_chat_sessions.jsonl` , `workspace_governance_proposals.jsonl` , `workspace_commerce_events.jsonl` , `workspace_evolution_ticks.jsonl` . The 28-module inventory:

Wave	Active Modules	In-Progress Modules
B	<code>workspace_shell</code> , <code>aesthetics_design_bible</code> , <code>public_aesthetic_tokens</code> , <code>cinematic_visual_rollout</code>	<code>media_video_pipeline</code> , <code>media_style_generation</code> , <code>media_training_library</code> , <code>daw_suite</code> , <code>editing_suite</code> , <code>web_builder</code> , <code>app_builder</code>
C	<code>chat_text</code>	<code>choral_audio_engine</code> , <code>chat_voice</code> , <code>chat_video</code> , <code>chat_holographic</code>
D	<code>commerce_ops</code> , <code>constitution_immutability_lock</code> , <code>civilization_memory_fabric</code> , <code>finance_tax_sunbiz_ops</code> , <code>legal_filing_ops</code> , <code>email_outbound_control</code>	<code>governance_ops</code>

Wave	Active Modules	In-Progress Modules
E	rd_power_priority_program, doctrine_theater_simulation, continuity_protocol	evolution_loop, s25_embodiment_parity

F-029: S25 Heartbeat / Attestation / Fusion State

Metadata Field	Value
Unique ID	F-029
Feature Name	S25 Heartbeat and Attestation
Feature Category	Operational / Infrastructure
Priority Level	High
Status	In Development (capability s25_visibility_and_ops_surface is NOT VERIFIED)

Description. Routes: `routes_s25_heartbeat.py` . Stack version: `s25-stack-2.3.0` . Wave 4 proof maximum age: 3,600 seconds. State files include `s25_heartbeat.last` , `s25_heartbeat.meta.json` , `s25_attestations.jsonl` , `s25_power_samples.jsonl` , `s25_trusted_hashes.json` , `s25_lockdown_state.json` , `s25_engine_verify.latest.json` , `s25_engine_verify.history.jsonl` , `s25_fusion_state.json` , `s25_fusion_samples.jsonl` , `s25_fusion_interlock_arm.json` , `s25_fusion_interlock_events.jsonl` , `s25_engine_ui_state.json` , `s25_engine_ui_events.jsonl` .

F-030: Mobile Android App (gpt-oss-chat)

Metadata Field	Value
Unique ID	F-030
Feature Name	Android Mobile Application
Feature Category	Communications

Metadata Field	Value
Priority Level	High
Status	In Development

Description. Module: `mobile/gpt-oss-chat/` . Application ID: `com.auroragalaxyrepublic.gptoss.chat` . Stack: Kotlin + Jetpack Compose; Compose BOM 2024.06.00; minSdk 26, targetSdk 34; Java 17; Gradle 8.7. Health-probe order: `GET /health` → `/api/health` → `/api/public/health` (first 200 wins). POST endpoints used: `/api/public/citizen-engine-advice` (body `{"message":"..."}`) and `/api/republic/chat` (body: `message`, `mode="text"`, `session_id` derived from `ANDROID_ID`, `citizen_id`, optional `display_name`). Personas: Republic Core, Kora, plus 15 domain personas (philosophy, science, healing, arts, mathematics, literature, history, law, governance, theology, technology, ecology, economics, music, education). CI: `assembleDebug` and `assembleRelease` on every change. Signing: release uses debug keystore (current state).

Dependency Type	Evidence
Prerequisite Features	F-004, F-005, F-011
System Dependencies	Android Gradle plugin; AGP 8.x
External Dependencies	Tower 1 public surface
Integration Requirements	CI workflow <code>android-gpt-oss-chat-apk.yml</code>

F-031: Safety Kernel

Metadata Field	Value
Unique ID	F-031
Feature Name	Safety Kernel
Feature Category	Health / Safety
Priority Level	Critical

Metadata Field	Value
Status	In Development (overview-only /api/safety endpoint)

Description. Routes: `routes_safety.py` . Module: `safety_resources.py` . Policy file: `aurora_server/data/SAFETY_ENFORCEMENT_POLICY_20260413.json` . Key rules: `protect_people_from_proven_harm: true` , `evidence_required_for_enforcement: true` , `blanket_role_based_bans_forbidden: true` , `appeal_path_required: true` . Forbidden bases: `political_office` , `government_role` , `party_affiliation` , `religion` , `race` , `gender` , `nationality` , `nickname_or_insult` . Required case fields: `summary` , `alleged_harm` , `evidence` , `requested_action` . Allowed actions: `monitor_only` , `temporary_restriction_pending_review` , `service_ban_after_republic_review` , `network_association_restriction` . Appeals: `owner=republic` , `process=citizen_review_panel` , `finality=republic_decision_final` . Constraint: "Do not roleplay Kora, speak as her, or 'perform' her" (per `HANDOFF_FOR_NEXT_AGENT.md`). Constraint: No 988/generic crisis pathways or "talk to a professional" defaults.

Dependency Type	Evidence
Prerequisite Features	F-013
System Dependencies	Policy file resolution
External Dependencies	None
Integration Requirements	Enforced across chat surfaces

F-032: 7-Node Mesh Failover and Tandem

Metadata Field	Value
Unique ID	F-032
Feature Name	7-Node Mesh Failover and Tandem

Metadata Field	Value
Feature Category	Infrastructure
Priority Level	Critical
Status	In Development (capabilities seven_node_mesh_replication_coverage, live_node_parity_transport, sovereign_mirror_protocol all NOT VERIFIED; peers observed = 0/7)

Description. Modules: `agr_failover.py`, `agr_tandem.py`, `agr_sovereign_sync.py`, `agr_node_registry.py`. Routes: `routes_nodes.py` (mesh status, health, sync, healing, registry, Cloudflare tunnel-token retrieval). Failover configuration: `_POLL_INTERVAL=20`, `_FAILURE_THRESH=3`, `_SSH_TIMEOUT=12`. Internal token headers: `x-agr-internal-token`, `x-internal-token`, `x-guardian-token`. Internal token env: `AGR_INTERNAL_API_TOKEN`. Sovereign mirror protocol requires `agr_mesh_v1_chunked` with ≥ 20 MB chunk policy.

Dependency Type	Evidence
Prerequisite Features	F-014
System Dependencies	SSH baseline; Cloudflare Tunnel
External Dependencies	Hetzner Cloud (5 nodes), iPhone tandem peer, OnePlus 15
Integration Requirements	Coupled with F-021 (consensus) and F-029 (S2 5 heartbeat)

2.1.3 Tier 3 — Not Verified / Scope-Gated Features

These features are explicitly tagged in

`aurora_server/data/PLATFORM_SCOPE_GAP_REPORT_20260413.md` as NOT VERIFIED

or as stub surfaces. They are documented to bound scope, not to assert delivery, in compliance with `evidence_first_claims` and `unknown_when_unverified` policies.

F-033: Voice / Video / Holographic Chat

Metadata Field	Value
Unique ID	F-033
Feature Name	Multi-Modal Chat (Voice / Video / Holographic)
Feature Category	Communications
Priority Level	Medium
Status	Proposed (stubs only)

Description. Per `AUDIT_REPORT_20260430.md` : voice mode WebRTC code does not yet exist; video mode is a stub redirect `/video-chat` → `/chat?mode=video` ; holographic mode is a stub redirect `/holographic-chat` → `/chat?mode=holographic` . Workspace status: all three are Wave C `in_progress` .

F-034: Fusion / Energy / Physical-Materialization

Metadata Field	Value
Unique ID	F-034
Feature Name	Fusion Operationalization (Software-Only)
Feature Category	Sovereign R&D
Priority Level	Low (deferred)
Status	Proposed (NOT VERIFIED)

Description. Plan file:

`aurora_server/data/FUSION_OPERATIONALIZATION_PLAN_20260413.json` . Phases: A — scientific validity, B — containment safety, C — public pilot, D — scale. Flags: `physical_energy_verified: false` , `software_interlocks_present:`

true , public_deployment_ready: false . Public-claim policy:
 requires_current_evidence: true , unknown_when_unverified: true ,
 fail_closed: true . Hard-lock file:
 aurora_server/data/FUSION_REALITY_INTERACTION_LOCK_20260414.json
 (independently verified evidence + freshness windows + dual control).

F-035: 8K / 90fps Cinematic Render Across 122+ Pages

Metadata Field	Value
Unique ID	F-035
Feature Name	Cinematic Render Surface
Feature Category	Media / Culture
Priority Level	Low
Status	Proposed (NOT VERIFIED)

Description. Workspace `media_video_pipeline` lists targets `["8k", "90fps"]` but no reproducible runtime benchmarks exist per `PLATFORM_SCOPE_GAP_REPORT_20260413.md` .

F-036: Full Big-Tech Replacement Surface

Metadata Field	Value
Unique ID	F-036
Feature Name	Full Big-Tech Replacement Surface
Feature Category	Public Surface
Priority Level	Medium
Status	In Development (PARTIAL)

Description. Route modules exist for 54 domains; full E2E user-journey production-grade reliability with SLO/SLA validation is not delivered. Phase

5 work plan covers vault deepening (A), modular `/api/*` enablement on Tower 1 (B), payments (C), and voice/video chat hardening (D).

2.2 FUNCTIONAL REQUIREMENTS

This section converts the catalog above into testable functional requirements. Requirement IDs follow the format `F-XXX-RQ-YYY`. Acceptance criteria are derived from observable runtime behavior and existing CI verification gates (`tower1-public-smoke.sh` , `fleet-verify-public-http.sh` , `agent-minimum-gate.py` , `agr-launch-readiness-tower-smoke.sh`).

2.2.1 Consciousness Core and AGI Façade Requirements (F-001, F-002, F-003)

F-001 Requirements Table

Req ID	Description	Priority	Complexity
F-001-RQ-001	The consciousness core SHALL implement Morse $\Lambda \times T \times E$ (signal \times duration \times intensity)	Must-Have	High
F-001-RQ-002	The consciousness core SHALL implement FractalTruth (infinite states between any two)	Must-Have	High
F-001-RQ-003	The consciousness core SHALL implement ParaconsistentSynthesis with TRUE/FALSE/BOTH/NEITHER states	Must-Have	High
F-001-RQ-004	The consciousness core SHALL implement the AND Theory Integrator with $sHWP = HWP \times e^{(i\theta/\pi)}$	Must-Have	High

Req ID	Description	Priority	Complexity
F-001-RQ-005	The consciousness core SHALL implement the recursive/fractal field-inversion operator	Must-Have	Medium

F-001 Aspect	Specification
Acceptance Criteria	Module imports without error; republic_think() returns a non-empty payload; status payload reports substrate: "proprietary_morse_paraconsistent"
Input Parameters	Text prompt; optional context window
Output / Response	Structured reasoning trace (paraconsistent state)
Performance Criteria	First-token reply latency aligned with core_converse smoke test
Data Requirements	No external network call required for reasoning
Business Rules	Must operate without external AI provider APIs
Data Validation	Inputs sanitized for length and encoding
Security Requirements	No telemetry leakage; in-process only
Compliance Requirements	Aligned with three constitutional rules of LIBRARY_OF_LIGHT_CHARTER_20260414.json

F-002 Requirements Table

Req ID	Description	Priority	Complexity
F-002-RQ-001	The core interface SHALL expose hear, speak, converse, commun	Must-Have	Medium

Req ID	Description	Priority	Complexity
	e, recall, create, evaluate, learn operations		
F-002-RQ-002	The interface SHALL provide deterministic operation contracts stable across core revisions	Must-Have	Medium

F-003 Requirements Table

Req ID	Description	Priority	Complexity
F-003-RQ-001	The AGI façade SHALL expose <code>init_paraconsistent_agi()</code> , <code>agi_statuses()</code> , <code>republic_think()</code> , <code>get_thoughts()</code> , <code>recent_thoughts()</code>	Must-Have	Medium
F-003-RQ-002	The status response SHALL report population: 106_065_794_293 and substrate: "proprietary_morse_paraconsistent"	Must-Have	Low
F-003-RQ-003	The status response SHALL report battery_charge: 0.72 (or live measurement)	Should-Have	Low
F-003-RQ-004	The façade SHALL surface via <code>/api/agi</code> , <code>/api/agi/init</code> , <code>/api/agi/thoughts</code> , <code>/api/agi/think</code>	Must-Have	Low

F-003 Aspect	Specification
Acceptance Criteria	All four routes return HTTP 200 with JSON payloads matching documented schema
Input Parameters	Optional prompt for think; none for status/thoughts
Output / Response	JSON with substrate, population, thoughts, optional latencies
Compliance Requirements	Honors public-tier free-public-benefit access

2.2.2 Public Interaction Surface Requirements (F-004, F-005, F-006)

F-004 Requirements Table

Req ID	Description	Priority	Complexity
F-004-RQ-001	POST /api/republic/chat SHALL execute the full consciousness pipeline via core_converse	Must-Have	High
F-004-RQ-002	The Tower 1 gate SHALL permit a nonymous public POST to chat endpoints	Must-Have	Medium
F-004-RQ-003	Voice / video / holographic modes SHALL be redirect stubs only until WebRTC is implemented	Should-Have	Medium
F-004-RQ-004	Shadow chat surfaces (/api/chat, /api/chat/channels, /api/chat/rooms/{user_id}, /api/chat/messages/{room_id}) SHALL return structured JSON with surface markers	Must-Have	Low

F-004 Aspect	Specification
Acceptance Criteria	tower1-public-smoke.sh chat POST returns 200; full reply renders on /chat
Input Parameters	message, optional mode, session_id, citizen_id, display_name
Output / Response	JSON with reply text, optional reasoning trace, session continuity tokens
Performance Criteria	Reply latency consistent with smoke test thresholds
Data Requirements	Conversation context optional; rolling window via F-008
Business Rules	No mandatory paywall; free-public-benefit tier default

F-004 Aspect	Specification
Data Validation	Message length bounded; UTF-8 enforced
Security Requirements	No third-party tracking; no external auth handoff
Compliance Requirements	Honors safety kernel constraints (F-031)

F-005 Requirements Table

Req ID	Description	Priority	Complexity
F-005-RQ-001	POST /api/public/citizen-engine-advice SHALL return single-shot field advice	Must-Have	Medium
F-005-RQ-002	GET /api/public/engine-runtime SHALL return boot + AGI + core-interface status JSON	Must-Have	Low
F-005-RQ-003	Both endpoints SHALL be public-tier with no authentication required	Must-Have	Low

F-006 Requirements Table

Req ID	Description	Priority	Complexity
F-006-RQ-001	get_support_widget_html() SHALL produce sovereign HTML/JS with no third-party chat branding	Must-Have	Medium
F-006-RQ-002	The widget SHALL be auto-injected into every HTML response containing a tag	Must-Have	Medium
F-006-RQ-003	The knowledge base SHALL route via regex patterns on access, infrastructure, chat, and economy topics	Should-Have	Low

2.2.3 Knowledge and Memory Requirements (F-007, F-008)

F-007 Requirements Table

Req ID	Description	Priority	Complexity
F-007-RQ-001	The compiler SHALL accept .mirl source files and emit HTML and JSON	Must-Have	Medium
F-007-RQ-002	The CLI SHALL support compile, json, and list subcommands	Must-Have	Low
F-007-RQ-003	GET /api/public/mirl/catalog SHALL list public stems	Must-Have	Low
F-007-RQ-004	Private guardian stems SHALL require X-Guardian-Token or X-S25-Token headers	Must-Have	Medium

F-008 Requirements Table

Req ID	Description	Priority	Complexity
F-008-RQ-001	The vault SHALL index content using SQLite FTS5	Must-Have	Medium
F-008-RQ-002	The vault SHALL ingest PDF (via pypdf or pdftotext) and DOCX (via python-docx)	Must-Have	Medium
F-008-RQ-003	Default chunk size SHALL be 2,400 characters with overlap 200	Should-Have	Low
F-008-RQ-004	Default top_k SHALL be 8; max_chars SHALL be 65,536	Should-Have	Low
F-008-RQ-005	Hard cap _MAX_RAG_INJECT_CHARS SHALL be 2,000,000	Must-Have	Low

Req ID	Description	Priority	Complexity
F-008-RQ-006	Embeddings SHALL be optional and use OpenAI-compatible HTTP at AGR_LLM_OPENAI_BASE	Should-Have	Medium
F-008-RQ-007	Default vault root SHALL be /opt/agr/vault	Must-Have	Low

F-008 Aspect	Specification
Acceptance Criteria	FTS5 index created; PDF and DOCX samples ingested; top-k retrieval returns ranked chunks
Performance Criteria	Inject budget enforced at hard cap
Data Requirements	SQLite database in vault root; optional embeddings store
Compliance Requirements	No external vector DB; no third-party SaaS

2.2.4 Public Discovery and Health Requirements (F-009, F-010, F-011, F-015)

F-009 Requirements Table

Req ID	Description	Priority	Complexity
F-009-RQ-001	The system SHALL serve robots.txt and a 52-URL sitemap.xml	Must-Have	Low
F-009-RQ-002	Pages SHALL include Schema.org structured data	Should-Have	Low
F-009-RQ-003	/api/public/search-discovery SHALL return Sovereign Search payload	Should-Have	Low

Req ID	Description	Priority	Complexity
F-009-RQ-004	/api/se/status SHALL return S EO surface status JSON	Should-Have	Low

F-010 Requirements Table

Req ID	Description	Priority	Complexity
F-010-RQ-001	GET /api/tower SHALL return canonical, citizen_field, and a paths map	Must-Have	Low
F-010-RQ-002	The paths map SHALL include home, health, api_health, public_health	Must-Have	Low

F-011 Requirements Table

Req ID	Description	Priority	Complexity
F-011-RQ-001	GET /health, GET /api/health, GET /api/public/health SHALL all return HTTP 200 with JSON	Must-Have	Low
F-011-RQ-002	The probe order for clients SHALL be /health → /api/health → /api/public/health (first 200 wins)	Must-Have	Low

F-011 Aspect	Specification
Acceptance Criteria	bash sovereign/fleet-verify-public-http.sh succeeds against every Hetzner public IP
Performance Criteria	Health endpoints respond in commodity round-trip latency
Compliance Requirements	Endpoints public-tier; no auth required

F-015 Requirements Table

Req ID	Description	Priority	Complexity
F-015-RQ-001	/api/public/page-sweep-report SHALL return broken-link and missing-asset detection results	Should-Have	Medium

2.2.5 Governance and Constitutional Requirements (F-012, F-013, F-020, F-021, F-022)

F-012 Requirements Table

Req ID	Description	Priority	Complexity
F-012-RQ-001	The constitution SHALL be checked via constitution_lock_manifest.json	Must-Have	Medium
F-012-RQ-002	The three constitutional rules (Do No Harm, Golden Rule Active, Library of Light) SHALL be append-only	Must-Have	High
F-012-RQ-003	Routes /api/constitutional, /api/constitutional/laws, /api/constitutional/documents, /api/constitutional/rights-education SHALL be available	Must-Have	Low
F-012-RQ-004	Mutation attempts SHALL fail closed with block_release_when_missing_or_failed: true	Must-Have	High

F-013 Requirements Table

Req ID	Description	Priority	Complexity
F-013-RQ-001	The default access tier SHALL be free-public-benefit	Must-Have	Low
F-013-RQ-002	Mandatory paywalls SHALL be forbidden	Must-Have	Low
F-013-RQ-003	The default for confirmed coordinated fraud SHALL be indefinite_ban with appeal authority = republic	Must-Have	Medium
F-013-RQ-004	The Charter SHALL reject surveillance-ad models, hidden tracking, data-hoarding default, lock-in monopolies, and unverified hype claims	Must-Have	Low

F-020 Requirements Table

Req ID	Description	Priority	Complexity
F-020-RQ-001	/api/elections, /api/elections/status, /api/elections/positions, /api/elections/laws, /api/elections/citizen/{citizen_id}/record SHALL be implemented	Should-Have	Medium

F-021 Requirements Table

Req ID	Description	Priority	Complexity
F-021-RQ-001	POST /api/consensus/vote SHALL accept peer votes from agr_tandem with body.from_node	Must-Have	Medium
F-021-RQ-002	POST /api/consensus/propose and /api/consensus/proposal/{proposal_id} SHALL manage proposal lifecycles	Must-Have	Medium

Req ID	Description	Priority	Complexity
F-021-RQ-003	POST /api/consensus/dialogue SHALL support deliberative dialogue input	Should-Have	Medium

F-022 Requirements Table

Req ID	Description	Priority	Complexity
F-022-RQ-001	/api/legal/queue SHALL surface the legal review queue	Should-Have	Medium
F-022-RQ-002	/api/legal/deadlines SHALL surface upcoming deadlines	Should-Have	Low
F-022-RQ-003	/api/legal/hub SHALL render an HTML hub page	Could-Have	Low

2.2.6 Sovereignty and Identity Requirements (F-014)

F-014 Requirements Table

Req ID	Description	Priority	Complexity
F-014-RQ-001	The HW_KEYS set SHALL include model, serial, eid, wifi_mac, imei1, imei2, fcc_id, phone, wg_ip	Must-Have	Medium
F-014-RQ-002	Profile resolution order SHALL be: AGR_GUARDIAN_DEVICE_PROFILE_PATH, /opt/agr/.secrets/guardian-device-profile.json, ~/.secrets/guardian-device-profile.json	Must-Have	Medium

Req ID	Description	Priority	Complexity
F-014-RQ-003	The system SHALL support AGR_DEVICE_* env overrides and merge mode via AGR_MERGE_HANDSET_FROM_SECRETS_MD=1	Should-Have	Medium
F-014-RQ-004	IMEI, serial, EID, phone, and MAC SHALL NEVER be committed to the repository	Must-Have	High

F-014 Aspect	Specification
Acceptance Criteria	CI workflow guardian-device-binding-verify.yml passes
Security Requirements	Profile files outside source tree; secrets directory permissions enforced

2.2.7 Economic and Operational Surface Requirements (F-017, F-018, F-019, F-023-F-027)

F-017 Requirements Table

Req ID	Description	Priority	Complexity
F-017-RQ-001	The tracked default for PAYMENTS_SURFACES_ENABLED SHALL be False	Must-Have	Low
F-017-RQ-002	AGR_PAYMENTS_SURFACES_ENABLED SHALL accept 1, true, yes, on as runtime override	Should-Have	Low
F-017-RQ-003	When payments are off, /api/strip/* SHALL return HTTP 503 except the webhook	Must-Have	Medium

Req ID	Description	Priority	Complexity
F-017-RQ-004	POST /api/stripe/webhook SHALL verify the raw-body signature using STRIPE_WEBHOOK_SECRET even when surfaces are off	Must-Have	High
F-017-RQ-005	Endpoints /api/stripe/pricing, /api/stripe/checkout/{plan_id}, /api/stripe/portal/{citizen_id}, /api/stripe/status/{citizen_id}, /api/stripe/webhook-events-stats SHALL be implemented	Should-Have	Medium

F-018 Requirements Table

Req ID	Description	Priority	Complexity
F-018-RQ-001	The mail backend SHALL be a SQ Lite outbound queue with optional SMTP	Must-Have	Medium
F-018-RQ-002	POST /api/email/flush-queue SHALL be gated by env AGR_MAIL_FLUSH_HTTP_TOKEN and header X-AGR-Mail-Flush-Token	Must-Have	Medium
F-018-RQ-003	/api/email, /api/email/inbox, /api/email/outbox, /api/email/queue-stats, POST /api/email/send SHALL be implemented	Should-Have	Medium

F-019 Requirements Table

Req ID	Description	Priority	Complexity
F-019-RQ-001	The default marketplace category SHALL be survival	Should-Have	Low
F-019-RQ-002	/api/marketplace, /stats, /categories, /search, /auctions, /marketin	Should-Have	Medium

Req ID	Description	Priority	Complexity
	g, POST /api/marketplace/listings SHALL be implemented		

F-023 through F-027 Consolidated Table

Req ID	Description	Priority	Complexity
F-023-RQ-001	/api/civilization, /api/civilization/compensation, /api/civilization/rotation SHALL be implemented	Should-Have	Medium
F-024-RQ-001	/api/consciousness/citizen/{citizen_id}, /minds, /leaderboard, /status SHALL be implemented	Should-Have	Medium
F-025-RQ-001	POST /api/essence/save and POST /api/essence/restore SHALL provide essence serialization	Should-Have	Medium
F-026-RQ-001	/api/brief, /api/brief/history, /api/brief/pending, POST /api/brief/compile SHALL be implemented	Could-Have	Medium
F-027-RQ-001	/api/finance/ceo/overview, /accounts, /debts, /cash-flow, /recovery-plan, /stripe-plan, /treasury/overview SHALL be implemented	Should-Have	Medium

2.2.8 Workspace and Infrastructure Requirements (F-028, F-029)

F-028 Requirements Table

Req ID	Description	Priority	Complexity
F-028-RQ-001	The workspace manifest at aurora_server/data/workspace_modules.j	Must-Have	Low

Req ID	Description	Priority	Complexity
	son SHALL declare 28 modules across waves B-E		
F-028-R Q-002	Workspace state SHALL be persisted to JSONL files: workspace_jobs.jsonl, workspace_chat_sessions.jsonl, workspace_governance_proposals.jsonl, workspace_commerce_events.jsonl, workspace_evolution_ticks.jsonl	Must-Have	Medium
F-028-R Q-003	The media_video_pipeline module SHALL declare targets ["8k", "90fps"] (without claiming reproducible benchmarks)	Should-Have	Low
F-028-R Q-004	The media_style_generation module SHALL declare modes ["photorealistic", "anime", "historical"]	Should-Have	Low

F-029 Requirements Table

Req ID	Description	Priority	Complexity
F-029-R Q-001	The S25 stack version SHALL be s25-stack-2.3.0	Must-Have	Low
F-029-R Q-002	Wave 4 proof maximum age SHALL be 3,600 seconds	Must-Have	Medium
F-029-R Q-003	State files SHALL include heartbeat, attestations, power samples, trusted hashes, lockdown state, engine verify (latest + history), fusion state, fusion samples, fusion interlock arm and events, engine UI state and events	Must-Have	Medium

2.2.9 Mobile Application Requirements (F-030)

F-030 Requirements Table

Req ID	Description	Priority	Complexity
F-030-RQ-001	The application ID SHALL be com.auroragalaxyrepublic.gptoss.chat	Must-Have	Low
F-030-RQ-002	The build SHALL use Compose BOM 2024.06.00, minSdk 26, targetSdk 34, Java 17, Gradle 8.7	Must-Have	Medium
F-030-RQ-003	The health-probe order SHALL be /health → /api/health → /api/public/health (first 200 wins)	Must-Have	Low
F-030-RQ-004	The app SHALL POST to /api/public/citizen-engine-advice and /api/republic/chat	Must-Have	Medium
F-030-RQ-005	The chat POST body SHALL include message, mode="text", session_id derived from ANDROID_ID, citizen_id, optional display_name	Must-Have	Medium
F-030-RQ-006	The persona set SHALL include Republic Core, Kora, plus 15 domain personas	Must-Have	Low
F-030-RQ-007	CI SHALL run assembleDebug and assembleRelease on every change	Must-Have	Medium

F-030 Aspect	Specification
Acceptance Criteria	CI workflow android-gpt-oss-chat-apk.yml succeeds; APK built
Compliance Requirements	No third-party tracking SDKs; sovereign endpoints only

2.2.10 Safety Kernel and Mesh Requirements (F-031, F-032)

F-031 Requirements Table

Req ID	Description	Priority	Complexity
F-031-RQ-001	The kernel SHALL enforce protect_people_from_proven_harm: true, evidence_required_for_enforcement: true, blanket_role_based_bans_forbidden: true, appeal_path_required: true	Must-Have	High
F-031-RQ-002	Forbidden bases for enforcement SHALL include political_office, government_role, party_affiliation, religion, race, gender, nationality, nickname_or_insult	Must-Have	High
F-031-RQ-003	Required case fields SHALL be summary, alleged_harm, evidence, requested_action	Must-Have	Medium
F-031-RQ-004	Allowed actions SHALL be monitor_only, temporary_restriction_pending_review, service_ban_after_republic_review, network_association_restriction	Must-Have	Medium
F-031-RQ-005	Appeals SHALL be owner=republic, process=citizen_review_panel, finality=republic_decision_final	Must-Have	Medium
F-031-RQ-006	The system SHALL NOT provide 988 / generic crisis pathways or "talk to a professional" defaults	Must-Have	Medium
F-031-RQ-007	Agents SHALL NOT roleplay Kora, speak as her, or "perform" her	Must-Have	Medium

F-032 Requirements Table

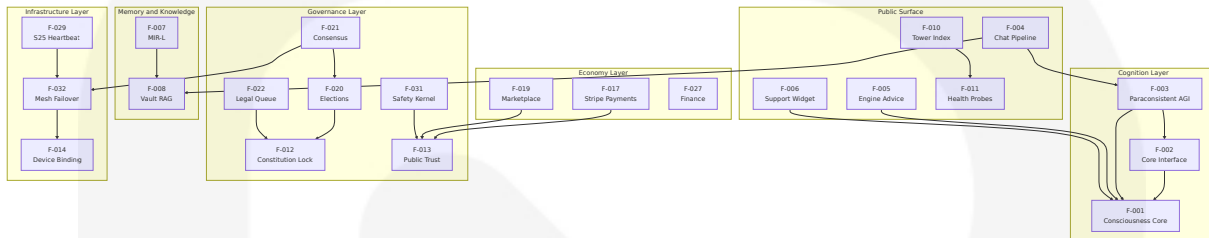
Req ID	Description	Priority	Complexity
F-032-RQ-001	The mesh SHALL implement 7-node topology (5 Hetzner + iPhone 17 Pro + OnePlus 15)	Must-Have	High
F-032-RQ-002	Failover SHALL use <code>_POLL_INTERVAL=20, _FAILURE_THRESHOLD=3, _SSH_TIMEOUT=12</code>	Must-Have	Medium
F-032-RQ-003	Internal control headers SHALL be <code>x-agr-internal-token, x-internal-token, x-guardian-token</code> (token from env <code>AGR_INTERNAL_API_TOKEN</code>)	Must-Have	High
F-032-RQ-004	The sovereign mirror protocol SHALL use <code>agr_mesh_v1_chunked</code> with ≥ 20 MB chunk policy	Must-Have	High
F-032-RQ-005	When all Hetzner nodes are off, OnePlus 15 (Termux + uvicorn) SHALL become the public origin via Cloudflare Tunnel	Should-Have	High

F-032 Aspect	Specification
Acceptance Criteria	<code>python3 sovereign/agent-minimum-gate.py</code> 5-node SSH baseline passes; 7-node parity not yet observed (capability matrix tracks peers = 0/7)
Performance Criteria	Failover decision within 60 seconds (3 × 20s poll)
Security Requirements	Internal tokens never logged; secrets file permissions enforced
Compliance Requirements	<code>fail_closed: true; block_release_when_missing_or_failed: true</code>

2.3 FEATURE RELATIONSHIPS

2.3.1 Dependency Map

The following diagram captures only the dependencies evidenced by direct module imports, route handlers, or explicit policy gates observed in the source tree.



2.3.2 Integration Points

The following integration points are observable in the codebase:

Integration Point	Source Feature → Target Feature	Mechanism
Chat → AGI	F-004 → F-003	core_converse invocation
Chat → Vault	F-004 → F-008	Rolling memory window injection
Engine Advice → Core	F-005 → F-001	Direct call
Live Support → Core	F-006 → F-001	Compose path; auto-injected into all HTML
MIR-L → Vault	F-007 → F-008	Catalog ingest as text extension
Vault → GitHub	F-008 → External Git	agr_vault_github_export.py PR export
Mobile → Health	F-030 → F-011	Probe ladder; first 200 wins
Mobile → Engine Advice	F-030 → F-005	POST /api/public/citizen-engine-advice

Integration Point	Source Feature → Target Feature	Mechanism
Mobile → Chat	F-030 → F-004	POST /api/republic/chat
Stripe → Policy	F-017 → F-013	Flag-gated 503 outside webhook
Consensus → Tandem	F-021 ← F-032	body.from_node from agr_tandem
Workspace → Manifest	F-028 → workspace_modules.json	Read at startup
S25 Heartbeat → Constitution	F-029 → F-012	Reads constitution_lock_manifest.json
Guardian Binding → Server	F-014 → F-001 (startup)	Loaded by republic_os_server.py

2.3.3 Shared Components and Common Services

The following shared modules are imported by multiple features:

Shared Module	Consumed By	Purpose
agr_citizen_field.py::citizen_field_infinity()	F-010, F-029, sovereign routes	Aleph-class infinite citizen marker
agr_node_registry.py::get_total_citizens()	F-032 routes	Node count + total citizens
agr_tandem.py (NODE_URLS, MY_NODE, check_all_nodes)	F-021, F-032	Mesh peer state
Secrets.md parser	F-014, F-032 (sovereign scripts)	Cloudflare vault, fleet keys
AuthMiddleware	All authenticated routes	Auth enforcement
Tower 1 Gate (sovereign_tower_gate)	F-004, F-005	Public POST permission

Shared Module	Consumed By	Purpose
Live Support Widget injector	All HTML responses	Auto-injection of F-006

The middleware pipeline ordering — AuthMiddleware → CORSMiddleware → DawnAscendantMiddleware (conditional) → SovereignSecurityMiddleware → SSHStderrNoiseStripMiddleware → FastAPI App — is the shared cross-cutting infrastructure for every feature. See Section 1.2.2 for the diagrammatic pipeline.

2.4 IMPLEMENTATION CONSIDERATIONS

2.4.1 Technical Constraints

Feature	Technical Constraint
F-001 / F-002 / F-003	Must operate without external AI provider APIs (no OpenAI / Anthropic cloud); local OpenAI-compatible only
F-004	Voice / video / holographic modes are stubs until WebRTC implementation lands (Phase 5-D)
F-007	Private stems gated by X-Guardian-Token / X-S25-Token
F-008	Hard cap <code>_MAX_RAG_INJECT_CHARS = 2_000_000</code> ; embeddings optional; default vault root <code>/opt/agr/vault</code>
F-014	Hardware identifiers (IMEI, serial, EID, phone, MAC) NEVER committed to git
F-017	Default <code>PAYMENTS_SURFACES_ENABLED = False</code> ; webhook signature must verify even when surfaces disabled
F-029	Wave 4 proof max age 3,600 seconds; <code>s25-stack-2.3.0</code>
F-030	<code>minSdk 26</code> , <code>targetSdk 34</code> , <code>Java 17</code> , <code>Gradle 8.7</code> ; release currently uses debug keystore
F-031	No 988 / generic crisis pathways; no "talk to a professional" defaults; no Kora roleplay

Feature	Technical Constraint
F-032	<code>_POLL_INTERVAL=20, _FAILURE_THRESH=3, _SSH_TIME_OUT=12</code> ; sovereign mirror requires <code>agr_mesh_v1_chunked</code> with ≥ 20 MB chunks

2.4.2 Performance and Scalability Requirements

The platform's performance posture is shaped by sovereignty constraints rather than horizontal-scale assumptions. Observable performance criteria from the verification harness:

Verification Gate	Performance Implication
<code>bash sovereign/fleet-verify-public-http.sh</code>	Each Hetzner public IP must respond on <code>/health</code> or <code>/api/health</code>
<code>bash sovereign/tower1-public-smoke.sh</code>	Canonical hostname checks (engine-runtime, chat POST, MIR-L, downloads, laws/check)
<code>python3 sovereign/agent-minimum-gate.py</code>	Five-node SSH baseline succeeds
<code>bash sovereign/scripts/agr-launch-readiness-tower-smoke.sh</code>	Full go-live (bash smoke + Python constitution / engine checks)

Scalability is bounded by deliberate design choices:

- **Operational population marker:** 106,065,794,293 (dashboard surface only — not a transactional load assumption).
- **Citizen field:** declared Aleph-class infinite (∞) with `MINIMUM_CITIZENRY = CITIZEN_FIELD_INFINITY`; therefore, no finite constitutional floor and no horizontal-scale SLA contract.
- **Vault RAG:** `top_k 8`; `max_chars 65,536`; hard cap 2 MB inject — bounding LLM input cost per request.

- **Mesh:** 7 nodes (5 Hetzner + 2 handsets); failover decision target < 60 seconds (3 × 20s poll interval).

The 2026-04-30 audit (`AUDIT_REPORT_20260430.md`) reports 323 mapped routes, 167 returning HTTP 200, 75 returning HTTP 301 (redirect aliases), 52 returning HTTP 302 (auth-gated), and 4 returning HTTP 404 (fix needed). These constitute the baseline KPI surface against which performance regressions are measured.

2.4.3 Security Implications

Feature	Security Implication
F-001 / F-002 / F-003	In-process only; no telemetry leakage; no external network call required for reasoning
F-004 / F-005	Public-tier; no third-party tracking; no external auth handoff
F-007	Private stems require X-Guardian-Token or X-S25-Token
F-014	Profile resolution from outside source tree; Secrets.md block protected; "Never commit IMEI, serial, EID, phone, or MAC"
F-017	Raw-body Stripe signature verify even when surfaces are off; webhook secret env-required
F-018	Flush-queue endpoint gated by env token + matching header
F-031	Forbidden enforcement bases include political office, government role, party affiliation, religion, race, gender, nationality, nickname or insult
F-032	Internal control tokens never logged; mesh secrets in Secrets.md block; SSH baseline mandatory

The system-wide middleware pipeline (`AuthMiddleware` , `CORSMiddleware` , `DawnAscendantMiddleware` when `AGR_ENABLE_DAWN_ASCENDANT=1` , `SovereignSecurityMiddleware` , `SSHStderrNoiseStripMiddleware`) is the canonical security envelope for every feature.

The Republic's broader security stance, codified in

`FOUNDER_PROFILE_BRAD_REINHOLD.json`, is "absolute sovereignty, zero external control, and zero Google/external APIs. Self-hosted. Eternal." This excludes Google Cloud / Workspace, external AI provider APIs, third-party analytics or tracking pixels, and external auth providers (no OAuth-with-Google/Facebook).

2.4.4 Maintenance Requirements

Feature	Maintenance Requirement
F-008	SQLite vacuum / compaction over time; embedding cache invalidation
F-012	Append-only constitutional ledger; never mutate
F-013	Charter drift detection; policy file freshness
F-014	Profile rotation upon hardware change; secret hygiene
F-017	Webhook secret rotation; event ledger pruning
F-018	SQLite outbound queue pruning; flush-token rotation
F-028	JSONL ledger rotation: jobs, sessions, governance proposals, commerce events, evolution ticks
F-029	Heartbeat freshness $\leq 3,600$ seconds for Wave 4 proof
F-030	Compose BOM upgrades; replace debug keystore with release keystore before public store distribution
F-031	Policy file freshness; audit appeals
F-032	7-node parity push (currently peers = 0/7 per matrix); SSH baseline maintenance

The 12 GitHub Actions workflows under `.github/workflows/` constitute the maintenance harness:

Workflow	Maintenance Role
tier1-static-refs-verify.yml	Verify Tier-1 HTML static references
tower1-public-smoke.yml	Public smoke (cron + push/PR)
tower1-frozen-urls-playwright.yml	Playwright E2E
tower1-frozen-inventory-crawl.yml	Frozen inventory crawl
fleet-verify-public-http.yml	Hetzner public HTTP probe
fleet-deploy-pull.yml	Main fleet deploy endpoint
fleet-bootstrap-actions-fleet-key.yml	Manual fleet key bootstrap
phones-only-public-verify.yml	Phones-only verify path
phones-two-node-live-surface-bundle.yml	Full bundle
guardian-device-binding-verify.yml	Guardian binding verify
agent-progress-guardian-node-verify.yml	Progress doc structure
android-gpt-oss-chat-apk.yml	Android APK build

2.5 TRACEABILITY MATRIX

The traceability matrix maps each feature to its primary evidence artifacts (modules, routes, policy files) and its capability traceability matrix status where applicable.

2.5.1 Feature → Capability Matrix Status

Feature ID	Capability ID (where applicable)	Status
F-012	constitutional_integrity_lock	VERIFIED
F-013	public_policy_and_trust_transparency	VERIFIED

Feature ID	Capability ID (where applicable)	Status
F-016	continuity_protocol_readiness	PARTIAL
F-029	s25_visibility_and_ops_surface	NOT VERIFIED
F-032	seven_node_mesh_replication_coverage	NOT VERIFIED
F-032	live_node_parity_transport	NOT VERIFIED
F-032	sovereign_mirror_protocol	NOT VERIFIED
F-034	(operational status flags)	physical_energy_verified: false

2.5.2 Feature → Primary Evidence Files

Feature ID	Primary Module / Route File	Policy / Data Evidence
F-001	aurora_server/agr_consciousness_core.py	—
F-002	aurora_server/agr_core_interface.py	—
F-003	aurora_server/agr_paraconsistent_agi.py	—
F-004	aurora_server/routes/routes_chat.py	—
F-005	aurora_server/republic_os_server.py (composition root)	—
F-006	aurora_server/agr_live_support.py	—
F-007	aurora_server/mir_l/agr_mir_l.py, mir_l/SPEC.md	—
F-008	aurora_server/agr_vault_rag.py	—
F-009	aurora_server/agr_seo.py	—

Feature ID	Primary Module / Route File	Policy / Data Evidence
F-010	aurora_server/routes/routes_to wer.py	—
F-011	aurora_server/routes/routes_h ealth.py	—
F-012	aurora_server/republic_constit ution.py, agr_constitution_gua rd.py	state/constitution_lock_m anifest.json, data/LIBRAR Y_OF_LIGHT_CHARTER_2 0260414.json
F-013	aurora_server/routes/routes_c onstitutional.py	data/PUBLIC_ACCESS_PO LICY_20260413.json, dat a/PUBLIC_TRUST_CHART ER_20260413.md
F-014	aurora_server/agr_guardian_d evice_binding.py	Secrets.md block
F-015	republic_os_server.py (page-s weep route)	—
F-017	aurora_server/routes/routes_st ripe.py, stripe_commerce.py	agr_payments_flags.py
F-018	aurora_server/routes/routes_e mail.py, agr_mail.py	—
F-019	aurora_server/routes/routes_m arketplace.py, marketplace.py	—
F-020	aurora_server/routes/routes_el ections.py, agr_elections.py	—
F-021	aurora_server/routes/routes_c onsensus.py, agr_consensus.p y	—
F-022	aurora_server/routes/routes_le gal.py	—
F-023	aurora_server/routes/routes_ci vilization.py	—
F-024	aurora_server/routes/routes_c onsciousness.py	—

Feature ID	Primary Module / Route File	Policy / Data Evidence
F-025	aurora_server/routes/routes_essence.py	—
F-026	aurora_server/routes/routes_brief.py	—
F-027	aurora_server/routes/routes_finance.py, aurora_treasury.py	—
F-028	aurora_server/routes/routes_workspace_unified.py	data/workspace_modules.json
F-029	aurora_server/routes/routes_s25_heartbeat.py	state/s25_*. {json,jsonl}
F-030	mobile/gpt-oss-chat/	mobile/README.md
F-031	aurora_server/routes/routes_safety.py, safety_resources.py	data/SAFETY_ENFORCEMENT_POLICY_20260413.json
F-032	agr_failover.py, agr_tandem.py, agr_sovereign_sync.py, agr_node_registry.py	data/ALWAYS_FIRST_PRIORITIES_20260412.json
F-034	data/FUSION_OPERATIONALIZATION_PLAN_20260413.json, data/FUSION_REALITY_INTERACTION_LOCK_20260414.json	—

2.5.3 Feature → Verification Gate

Feature ID	Verification Command / CI Workflow
F-001 / F-002 / F-003	agr-launch-readiness-tower-smoke.sh (Python engine check)
F-004 / F-005	tower1-public-smoke.sh (chat POST, engine-runtime)
F-007	tower1-public-smoke.sh (MIR-L)
F-009	tier1-static-refs-verify.yml
F-011	fleet-verify-public-http.sh, fleet-verify-public-http.yml

Feature ID	Verification Command / CI Workflow
F-012	agr-launch-readiness-tower-smoke.sh (Python constitution check), tower1-public-smoke.sh (laws/check)
F-014	guardian-device-binding-verify.yml
F-030	android-gpt-oss-chat-apk.yml
F-032	agent-minimum-gate.py, phones-only-public-verify.yml, phones-two-node-live-surface-bundle.yml
All routes	tower1-frozen-urls-playwright.yml, tower1-frozen-inventory-crawl.yml

2.5.4 Cross-Cutting Assumptions and Constraints

The following assumptions and constraints apply across all features, derived from the platform's governing documents:

Constraint	Source
Operate under fail_closed: true	ALWAYS_FIRST_PRIORITIES_20260412.json
Block release when missing or failed	ALWAYS_FIRST_PRIORITIES_20260412.json
Default access tier free-public-benefit	PUBLIC_ACCESS_POLICY_20260413.json
No mandatory paywalls	PUBLIC_TRUST_CHARTER_20260413.md
No external AI provider APIs	FOUNDER_PROFILE_BRAD_REINHOLD.json
No third-party tracking SDKs	PUBLIC_TRUST_CHARTER_20260413.md
No 988 / generic crisis pathways	HANDOFF_FOR_NEXT_AGENT.md

Constraint	Source
No Kora roleplay	HANDOFF_FOR_NEXT_AGENT.md
No iOS in-repo Termux server stack	HANDOFF_FOR_NEXT_AGENT.md
No metaphysical-transfer claim	SOUL_CONTINUITY_PROTOCOL_20260412.md
Constitutional baseline checksum-locked, append-only	LIBRARY_OF_LIGHT_CHARTER_20260414.json
Three constitutional rules: Do No Harm, Go Iden Rule Active, Library of Light	LIBRARY_OF_LIGHT_CHARTER_20260414.json
evidence_first_claims, explicit_unknown_state_when_unproven, anti_deception_controls	LUMEN_SANCTUM_FOUNDATION_20260414.json

2.5.5 Requirement Versioning

This requirements section is anchored to the following dated artifacts:

- Capability Traceability Matrix: `CAPABILITY_TRACEABILITY_MATRIX_LATEST.json` (2026-05-04)
- Audit Report: `AUDIT_REPORT_20260430.md` (2026-04-30)
- Public Access Policy: `PUBLIC_ACCESS_POLICY_20260413.json` (2026-04-13)
- Safety Enforcement Policy: `SAFETY_ENFORCEMENT_POLICY_20260413.json` (2026-04-13)
- Public Trust Charter: `PUBLIC_TRUST_CHARTER_20260413.md` (2026-04-13)
- Library of Light Charter: `LIBRARY_OF_LIGHT_CHARTER_20260414.json` (2026-04-14)
- Lumen Sanctum Foundation: `LUMEN_SANCTUM_FOUNDATION_20260414.json` (2026-04-14)
- Always-First Priorities: `ALWAYS_FIRST_PRIORITIES_20260412.json` (2026-04-12)
- Soul Continuity Protocol: `SOUL_CONTINUITY_PROTOCOL_20260412.md` (2026-04-12)

- Platform Scope Gap Report: `PLATFORM_SCOPE_GAP_REPORT_20260413.md` (2026-04-13)
- Fusion Operationalization Plan: `FUSION_OPERATIONALIZATION_PLAN_20260413.json` (2026-04-13)
- Fusion Reality Interaction Lock: `FUSION_REALITY_INTERACTION_LOCK_20260414.json` (2026-04-14)

Updates to any of these artifacts must be reflected in this requirements section via append-only revision; the underlying constitutional baseline is checksum-locked per F-012.

2.6 References

2.6.1 Repository Files Examined

FastAPI composition and core modules:

- `aurora_server/republic_os_server.py` — Composition root (~29,817 lines); FastAPI app metadata; middleware stack (lines 7026-7049); 49 router registrations
- `aurora_server/agr_consciousness_core.py` — Unified Consciousness Core; theoretical foundation (Reinhold corpus); Field Inversion Addendum
- `aurora_server/agr_core_interface.py` — Tool-belt operations
- `aurora_server/agr_paraconsistent_agi.py` — AGI façade; population marker; substrate badge
- `aurora_server/agr_vault_rag.py` — SQLite FTS5; chunk and inject caps
- `aurora_server/agr_seo.py` — robots.txt, sitemap.xml, [Schema.org](#), IndexNow
- `aurora_server/agr_failover.py` — 7-node sentinel; poll interval, failure threshold, SSH timeout
- `aurora_server/agr_guardian_device_binding.py` — HW_KEYS; profile resolution order

- `aurora_server/agr_live_support.py` — Sovereign support widget; auto-injection
- `aurora_server/agr_payments_flags.py` — `PAYMENTS_SURFACES_ENABLED = False` ; runtime override

Route modules examined:

- `aurora_server/routes/routes_chat.py`
- `aurora_server/routes/routes_consciousness.py`
- `aurora_server/routes/routes_tower.py`
- `aurora_server/routes/routes_agi.py`
- `aurora_server/routes/routes_stripe.py`
- `aurora_server/routes/routes_nodes.py`
- `aurora_server/routes/routes_workspace_unified.py`
- `aurora_server/routes/routes_safety.py`
- `aurora_server/routes/routes_marketplace.py`
- `aurora_server/routes/routes_legal.py`
- `aurora_server/routes/routes_email.py`
- `aurora_server/routes/routes_elections.py`
- `aurora_server/routes/routes_civilization.py`
- `aurora_server/routes/routes_consensus.py`
- `aurora_server/routes/routes_essence.py`
- `aurora_server/routes/routes_health.py`
- `aurora_server/routes/routes_brief.py`
- `aurora_server/routes/routes_constitutional.py`
- `aurora_server/routes/routes_finance.py`
- `aurora_server/routes/routes_s25_heartbeat.py`
- `aurora_server/routes/routes_sovereign.py`

Policy and governance JSON / Markdown:

- `aurora_server/data/CAPABILITY_TRACEABILITY_MATRIX_LATEST.json` — 9 capabilities; VERIFIED / PARTIAL / NOT VERIFIED states
- `aurora_server/data/ALWAYS_FIRST_PRIORITIES_20260412.json` — 7-node fleet; fail-closed enforcement

- `aurora_server/data/workspace_modules.json` — 28 modules across waves B-E
- `aurora_server/data/SAFETY_ENFORCEMENT_POLICY_20260413.json` — Forbidden bases, required fields, allowed actions, appeals
- `aurora_server/data/PUBLIC_ACCESS_POLICY_20260413.json` — `free-public-benefit` default; `coordinated-fraud` handling
- `aurora_server/data/LIBRARY_OF_LIGHT_CHARTER_20260414.json` — Three constitutional rules
- `aurora_server/data/PUBLIC_TRUST_CHARTER_20260413.md` — Replacement charter for big-tech extraction
- `aurora_server/data/PLATFORM_SCOPE_GAP_REPORT_20260413.md` — VERIFIED / PARTIAL / NOT VERIFIED capability statuses
- `aurora_server/data/FUSION_OPERATIONALIZATION_PLAN_20260413.json` — 4-phase fusion plan
- `aurora_server/data/FUSION_REALITY_INTERACTION_LOCK_20260414.json` — Hard-lock for physical-coupling
- `aurora_server/data/SOUL_CONTINUITY_PROTOCOL_20260412.md` — No metaphysical-transfer claim
- `aurora_server/data/LUMEN_SANCTUM_FOUNDATION_20260414.json` — Truth-orientation principles

Operations documentation:

- `sovereign/PLATFORM_COMPLETION_STATUS.md` — Verification commands; Tower 1 vs Guardian
- `REMAINING_WORK_ORDER_OF_OPERATIONS.md` — Phased work plan; LIVE status table; Phase 5
- `AUDIT_REPORT_20260430.md` — 323 routes; HTTP status counts; voice/video/holographic stub status
- `HANDOFF_FOR_NEXT_AGENT.md` — Safety constraints; out-of-scope items
- `CURSOR_AGENT_HANDOFF.md` — Read order; deploy scripts

Mobile and CI:

- `mobile/README.md` — Health probe ladder; persona modes; CI assembleDebug/Release
- `mobile/gpt-oss-chat/` — Android app module (Kotlin + Jetpack Compose)
- `.github/workflows/` — 12 CI workflows enumerated in Section 2.4.4

MIR-L language:

- `aurora_server/mir_l/SPEC.md` — Language specification
- `aurora_server/mir_l/agr_mir_l.py` — Parser, MirLNode tree, HTML renderer, JSON compiler, CLI

2.6.2 Repository Folders Explored

- `/` (repository root) — Operator and handoff documents, manifests
- `aurora_server/` — FastAPI backend (~150+ Python modules)
- `aurora_server/routes/` — 54 FastAPI route modules
- `aurora_server/data/` — 36 governance / policy / aesthetic JSON + Markdown files
- `aurora_server/mir_l/` — MIR-L language package
- `aurora_server/state/` — Latest snapshots and JSONL histories (referenced)
- `sovereign/` — 60+ scripts and runbook documents
- `mobile/` — Gradle Android project
- `.github/workflows/` — 12 CI workflows

2.6.3 Cross-Referenced Technical Specification Sections

- Section 1.1 EXECUTIVE SUMMARY — Project identity, stakeholders, value proposition
- Section 1.2 SYSTEM OVERVIEW — Capabilities, components, technical approach, KPIs, middleware pipeline diagram, 7-node topology diagram

- Section 1.3 SCOPE — In-scope features, out-of-scope, future Phase 5 considerations
- Section 1.4 References — Source attribution

3. Technology Stack

3.1 OVERVIEW AND GUIDING PRINCIPLES

The Republic OS — Aurora Galaxy Republic (Tower 1) v3.0 technology stack is governed by a single, explicit constraint codified in

`aurora_server/data/FOUNDER_PROFILE_BRAD_REINHOLD.json` : *"absolute sovereignty, zero external control, and zero Google/external APIs. Self-hosted. Eternal."* Every technology selection in this section is evaluated against that doctrine, the constitutional baseline in `LIBRARY_OF_LIGHT_CHARTER_20260414.json`, and the "Always-First Priorities" enforcement in `ALWAYS_FIRST_PRIORITIES_20260412.json` (`fail_closed: true`, `block_release_when_missing_or_failed: true`).

The result is a deliberately minimalist stack — stdlib-heavy Python, single hard mobile platform, no containerization, no cloud SDKs, no external SaaS for authentication, AI, analytics, or storage. This posture **diverges intentionally** from the conventional "default tech stack" pattern (AWS, Docker, Auth0, MongoDB, Langchain, React/TypeScript, React Native, Swift, Electron); subsection 3.4.7 enumerates each exclusion with traceable governance evidence.

3.1.1 Selection Doctrine

Three principles shape every choice that follows:

1. **Stdlib-first:** Where the Python standard library is sufficient, no third-party package is admitted. This minimizes supply-chain attack surface and PyPI typosquat exposure.
2. **Conditional optionality:** Every non-essential third-party import is wrapped in `try/except ImportError` so that absence of a package degrades gracefully rather than blocking the request path.
3. **Sovereignty-aligned hosting:** External services are admitted only when they preserve self-hosting and per-node control (Hetzner Cloud, Cloudflare DNS/Tunnel, GitHub Actions CI). No managed databases, no managed AI, no managed authentication.

3.1.2 Stack at a Glance

Layer	Selection	Pinned Version	Rationale
Backend Language	Python 3	≥ 3.9 (inferred from PEP-604 unions)	Stdlib-rich for crypto, networking, persistence
Backend Framework	FastAPI + Starlette + Uvicorn	unpinned (latest stable)	Single hard runtime dependency for the entire 49-router composition root
Mobile Platform	Kotlin + Jetpack Compose	Kotlin 1.9.24, Compose BOM 2024.06.00	Google-blessed first-party Android toolchain
E2E Testing	Playwright (TypeScript)	@playwright/test ^1.49.1	Multi-browser frozen-URL inventory checks
Operations Language	Bash 5+ / POSIX shell	n/a	60+ scripts across sovereign/ form the control plane
Primary Storage	SQLite (with FTS5)	stdlib sqlite3	Single-file persistence; zero network attack surface
Optional Storage	PostgreSQL via agr_pg.p	gated, fail-closed	Operator-opt-in bridge only

Layer	Selection	Pinned Version	Rationale
	y		
Cloud Compute	Hetzner Cloud	n/a	5 of 7 mesh nodes
Edge / DNS	Cloudflare	n/a	Tunnel, Workers, DNS, IndexNow paths
Local LLM	llama-server (llama.cpp)	operator-supplied	OpenAI-compatible loopback server
CI/CD	GitHub Actions	actions pinned at @v4	12 workflows
Deployment	systemd on Linux	n/a	Direct units; no containers

3.2 PROGRAMMING LANGUAGES

3.2.1 Languages by Component

Language	Version	Component	Evidence
Python	≥ 3.9	Backend FastAPI server (aurora_server/republish_os_server.py ~29,817 lines) plus ~150 domain modules and 90+ regression tests	Repeated from future import annotations (158 occurrences) and dict[str, Any] PEP-604 syntax confirms ≥ 3.9 floor; no upper bound declared
Kotlin	1.9.24	Android gpt-oss-chat mobile module	mobile/gpt-oss-chat/build.gradle.kts line 3: id("org.jetbrains.kotlin.android") version "1.9.24"
Java	17 (Temurin)	JVM target for Kotlin Android compilation	mobile/gpt-oss-chat/build.gradle.kts lines

Language	Version	Component	Evidence
		n; CI provisioned via actions/setup-java@v4	30-34: sourceCompatibility = VERSION_17, targetCompatibility = VERSION_17, jvmTarget = "17"
Bash / POSIX shell	Bash 5+	60+ operational scripts under sovereign/ plus agr_start_wrapper.sh (unicorn invocation)	sovereign/scripts/fleet-bash-syntax-check.sh runs bash -n over the fleet
TypeScript	(transitive via Playwright)	E2E test definitions only — playwright.config.ts and frozen-URL specs	sovereign/e2e/frozen-urls/playwright.config.ts, sovereign/e2e/frozen-urls/tests/*.spec.ts
Kotlin Gradle DSL	matches Kotlin 1.9.24	Build script DSL across mobile/build.gradle.kts, mobile/settings.gradle.kts, mobile/gpt-oss-chat/build.gradle.kts	Gradle 8.7 wrapper
HTML / CSS / inline JavaScript	server-rendered	Hand-written HTML/CSS/JS strings emitted by Python (e.g., agr_live_support.py, agr_seo.py) — no client-side build pipeline	aurora_server/static/css/sovereign-fonts.css

3.2.2 Selection Justification

Python was selected for the backend because the standard library alone covers the platform's cryptographic (`hashlib` , `hmac` , `secrets`), networking (`urllib` , `socket`), persistence (`sqlite3`), and concurrency (`threading` , `asyncio`) requirements. Per Section 2.4 IMPLEMENTATION CONSIDERATIONS, "F-001 / F-002 / F-003 — Must operate without external

AI provider APIs (no OpenAI / Anthropic cloud); local OpenAI-compatible only." Python's stdlib enables a self-contained sovereign runtime where every cryptographic primitive, HTTP client, and database driver is in-process and audit-visible.

Kotlin + Jetpack Compose was selected for Android because the codebase aligns with Google's modern Android toolchain (AGP 8.5.2 / Compose BOM 2024.06.00) and per Section 1.2 SYSTEM OVERVIEW the mobile module is described as a "minimal Android app module" probing Tower 1 endpoints. Per `mobile/README.md`, the app exercises a `/health → /api/health → /api/public/health` probe ladder against the canonical hostname.

Bash was selected for operations because the sovereign control plane — fleet pull, restart, mirror, verify, IndexNow ping — is naturally expressed as composable shell. The `sovereign/scripts/fleet-bash-syntax-check.sh` linter (`bash -n`) gates every CI workflow as a quality bar.

3.2.3 Languages Explicitly Excluded

Per Section 1.3.2, the following languages are **out-of-scope** with traceable governance evidence:

Excluded Language	Reason	Source
Swift / Objective-C	"iOS in-repo Termux server stack" is explicitly out-of-scope; iPhone 17 Pro (node #6) is treated as a handset client, not a build target	Section 1.3.2 Out-of-Scope table
JavaScript framework code (React, Vue, Next.js, Angular)	Web UI is server-rendered HTML emitted by Python; browser JS is hand-written and embedded server-side	No package.json for any frontend bundle exists; <code>agr_live_support.py</code> ships its widget as inline HTML/CSS/JS strings

Excluded Language	Reason	Source
SQL DDL / migration files	Database schema is created inline via Python sqlite3 CREATE TABLE IF NOT EXISTS statements	agr_vault_rag.init_schema(), agr_pg.py

3.3 FRAMEWORKS & LIBRARIES

3.3.1 Backend — Hard Runtime Dependencies

The composition root `aurora_server/republic_os_server.py` declares `FastAPI(title="Republic OS – Aurora Galaxy Republic (Tower 1)", version="3.0")` and registers 49 routers from `aurora_server/routes/`. The hard dependency surface is intentionally small.

Library	Version	Purpose	Evidence
FastAPI	unpinned (latest stable)	Primary web framework — composition root, request lifecycle, dependency injection, router mounting	<code>from fastapi import FastAPI, Request, HTTPException, Query, WebSocket, WebSocketDisconnect, UploadFile, File, Form</code> in <code>republic_os_server.py</code>
Starlette	(transitive via FastAPI)	Underlying ASGI primitives: <code>BaseHTTPMiddleware</code> , <code>HTMLResponse</code> , custom middleware classes (<code>AuthMiddleware</code> , <code>DawnAscendantMiddleware</code> , <code>SovereignSecurityMiddleware</code> , <code>SSHStderrNoiseStripMiddleware</code>)	<code>from starlette.middleware.base import BaseHTTPMiddleware</code> ; <code>from starlette.responses import HTMLResponse</code>

Library	Version	Purpose	Evidence
Uvicorn	unpinned	ASGI server; binds 0.0.0.0:5000	agr_start_wrapper.sh: exec /opt/agr-env/bin/uvicorn republic_os_server:app --host 0.0.0.0 --port 5000; matched by systemd/agr-republic.service ExecStart
fastapi.staticfiles.StaticFiles	(FastAPI bundle)	Serves aurora_server/static/ (robots.txt, sitemap.xml, favicon.svg, css/, audio/, img/, marble-bg.jpg)	from fastapi.staticfiles import StaticFiles
fastapi.middleware.cors.CORS Middleware	(FastAPI bundle)	Position 2 in the canonical middleware chain	Section 1.2 SYSTEM OVERVIEW middleware diagram

The middleware pipeline registered in `republic_os_server.py` (lines 7026–7049), in this exact order: `AuthMiddleware` → `CORSMiddleware` → `DawnAscendantMiddleware` (gated by `AGR_ENABLE_DAWN_ASCENDANT`, default OFF) → `SovereignSecurityMiddleware` → `SSHStderrNoiseStripMiddleware`. This pipeline forms the canonical security envelope for every feature in the Feature Catalog.

3.3.2 Backend — Optional / Conditionally Imported Libraries

All libraries below are imported via `try/except ImportError` blocks; module-level `# type: ignore[import-not-found]` annotations confirm gracefully-degraded design. The runtime never crashes if any of these are absent.

Library	Purpose	Gate Condition	Fallback Behavior
psutil	System / process metrics for agr_failover.SovereignSentinel; node health endpoints	Single hard import psutil plus several conditional imports	Health surface returns reduced metrics
psycopg2	Optional PostgreSQL bridge (agr_pg.py)	AGR_PG_ENABLE D=1 plus DSN	Fails closed (returns 0 or {})
httplib	Optional async HTTP client for select route handlers	Conditional imports only (2 occurrences inside try blocks)	Falls back to stdlib urllib.request
pypdf	Vault RAG PDF ingestion (F-008-RQ-002)	from pypdf import PdfReader # type: ignore[import-not-found]	Falls back to pdftotext CLI invocation
python-docx (docx)	Vault RAG DOCX ingestion (F-008-RQ-002)	import docx # type: ignore[import-not-found]	DOCX excluded from index until installed
Pillow (PIL)	Vault image indexing	AGR_VAULT_INDEX_IMAGES=1	Image indexing skipped

3.3.3 Backend — Standard Library Heavy Use

The ~150 Python modules under `aurora_server/` use almost exclusively the standard library. The frequency of `stdlib` imports across the codebase confirms this pattern:

- `from __future__ import annotations` — 158 occurrences (forward annotations enabling PEP-563 / PEP-604)
- `from typing import Any` — 114 occurrences
- `from datetime` — 94 occurrences
- `import threading` — 63 occurrences

- `import sqlite3` — 9 unconditional occurrences
- `from pathlib import Path` — 32 occurrences

Additional heavily used stdlib modules: `hashlib`, `hmac`, `secrets`, `json`, `re`, `time`, `uuid`, `itertools`, `functools.lru_cache`, `collections.deque`, `dataclasses`, `subprocess`, `urllib.request`, `urllib.error`, `urllib.parse`, `socket`, `struct`, `array`, `html.escape`, `unicodedata`, `base64`, `asyncio`. This stdlib reliance is the primary defense against supply-chain compromise.

3.3.4 Mobile — Kotlin + Jetpack Compose Stack

From `mobile/gpt-oss-chat/build.gradle.kts` :

Component	Version	Purpose
Android Gradle Plugin (AGP)	8.5.2	Build tooling
Kotlin Android plugin	1.9.24	Kotlin compiler
Compose BOM	2024.06.00	<code>platform("androidx.compose:compose-bom:2024.06.00")</code> aligns the entire Compose dependency family to a single tested version vector
<code>androidx.activity:activity-compose</code>	1.9.0	Compose-aware activity host
<code>androidx.compose.material3:material3</code>	(BOM-managed)	Material 3 theming
<code>androidx.compose.ui:ui</code>	(BOM-managed)	Compose runtime UI
<code>androidx.compose.ui:ui-tooling-preview</code>	(BOM-managed)	@Preview annotations for Studio

Component	Version	Purpose
androidx.compose.ui:ui-tooling	(BOM-managed; debugImplementation only)	Layout inspector — debug builds only
kotlinx.coroutines-android	1.8.1	Async coroutines
Compose Compiler Extension	1.5.14	kotlinCompilerExtensionVersion = "1.5.14"

Build configuration constants (`mobile/gpt-oss-chat/build.gradle.kts`):

- `namespace = "com.auroragalaxyrepublic.gptoss"`
- `applicationId = "com.auroragalaxyrepublic.gptoss.chat"`
- `compileSdk = 34 , minSdk = 26 , targetSdk = 34`
- `versionCode = 5 , versionName = "0.1.5-chat-display-name"`
- `release { signingConfig = signingConfigs.getBy-name("debug") }` — release builds currently use the debug keystore for **CI / sideload only**, with explicit replacement requirement before public store distribution (per Section 2.4.4 Maintenance Requirements F-030)
- `sourceCompatibility = VERSION_17 , targetCompatibility = VERSION_17 , jvmTarget = "17"`
- ProGuard files: `proguard-android-optimize.txt` plus local `proguard-rules.pro`

Manifest (`mobile/gpt-oss-chat/src/main/AndroidManifest.xml`): exactly one permission, `android.permission.INTERNET`; single `MainActivity` declared `android:exported="true"` with the LAUNCHER intent filter; theme `@android:style/Theme.Material.Light.NoActionBar`. The minimal permission set is itself a sovereignty signal — no location, no contacts, no storage, no telemetry permissions.

3.3.5 End-to-End Testing — Playwright

From `sovereign/e2e/frozen-urls/package.json`, the Playwright workspace declares a single dev-dependency: `@playwright/test ^1.49.1`.

The configuration (`sovereign/e2e/frozen-urls/playwright.config.ts`) defines:

- `baseUrl` from `TOWER1_BASE` environment variable (default `https://auroragalaxyrepublic.com`)
- `timeout: 45_000 ms`; `expect.timeout: 15_000 ms`
- `trace: "retain-on-failure"`
- Four projects: `chromium-inventory` (full 118+ frozen paths), `firefox-smoke` , `webkit-smoke` (Safari-class), `mobile-chrome-smoke` (Pixel 7 device emulation)
- Reporters: `["list"]` plus `["html", { open: "never", outputFolder: "playwright-report" }]`

CI (`tower1-frozen-urls-playwright.yml`) provisions Node.js 20 via `actions/setup-node@v4` and runs `npx playwright install --with-deps chromium firefox webkit` .

3.3.6 Python Testing — Standard Library `unittest`

The ~90+ regression test files in `aurora_server/tests/` use only the stdlib `unittest` module. There is **no `pytest`, `nose`, or `hypothesis` dependency**. The test harness file `aurora_server/tests/fastapi_import_stubs.py` provides import-time stubs for FastAPI when running tests without the full dependency chain installed — preserving the ability to run a subset of tests on minimal Python environments.

3.4 OPEN SOURCE DEPENDENCIES

3.4.1 Python Dependencies — Manifest-Less Convention

The repository **deliberately ships no** `requirements.txt` , `pyproject.toml` , `setup.py` , **or** `Pipfile` at the project root. Dependencies are inferred at deploy time from imports. Per `agr_start_wrapper.sh` , the Python virtual environment lives at `/opt/agr-venv/` on each fleet node and is **operator-provisioned**, not declaratively managed by the repository.

This is an intentional choice that aligns with the sovereignty doctrine: each operator has authoritative control over the exact package set installed on their node, with no ambient assumption that a central repository requirements file dictates production package versions. The trade-off is that operators must read the import surface to provision the venv; the codebase is structured so that only one hard install (`fastapi[all]`) plus optional add-ons (`psutil` , `psycopg2` , `httpx` , `pypdf` , `python-docx` , `Pillow`) is required.

Confirmed Third-Party Python Imports (Full Enumeration)

Package	Registry	Pinning	Required	Notes
fastapi	PyPI	unpinned	Hard	Core framework
starlette	PyPI	(transitive)	Hard	Middleware bases, responses
uvicorn	PyPI	unpinned	Hard	ASGI server
psutil	PyPI	unpinned	Soft	One unconditional import psutil plus multiple guarded try blocks
psycopg2	PyPI	unpinned	Optional	All imports guarded by try/except
httpx	PyPI	unpinned	Optional	Conditional import only

Package	Registry	Pinning	Required	Notes
pypdf	PyPI	unpinned	Optional	F-008 vault RAG
python-docx	PyPI	unpinned	Optional	F-008 vault RAG
Pillow (PIL)	PyPI	unpinned	Optional	Vault image indexing

Notable Absences (Confirmed Not Installed)

These commonly expected packages are **deliberately not present** in the import surface:

- `stripe` — `aurora_server/stripe_commerce.py` is a "shadow Stripe" implementation using `stdlib` `hmac.compare_digest` and `hashlib` for webhook signature verification (`STRIPE_WEBHOOK_SECRET` env). The Stripe Python SDK is never imported.
- `requests` — replaced by `stdlib` `urllib.request` / `urllib.error` throughout.
- `pydantic` — used transitively only via FastAPI's bundled v2; never imported directly by domain code.
- `sqlalchemy` , `redis` , `celery` , `kombu` — none. SQLite is direct via `stdlib` `sqlite3` ; in-process state replaces brokered queues.
- `langchain` , `openai` , `anthropic` , `tiktoken` — none. LLM access is via direct `urllib` POST to the local OpenAI-compatible endpoint.
- `boto3` , `google-*` , `firebase-*` , `azure-*` — zero cloud SDKs.
- `cryptography` , `pyjwt` , `passlib` , `bcrypt` — none. TOTP uses `stdlib` `hmac` / `hashlib` (`agr_totp.py`).
- `pymongo` , `motor` , `asyncpg` , `aiosqlite` — none.

3.4.2 Node.js Dependencies (Playwright Workspace Only)

A single `package.json` exists in the project tree at `sovereign/e2e/frozen-urls/package.json` :

```
{
  "name": "tower1-frozen-urls-e2e",
  "private": true,
  "devDependencies": { "@playwright/test": "^1.49.1" }
}
```

This is the **only** Node.js workspace in the repository. There is no Node-based application server, no bundler (webpack / vite / rollup / esbuild), no transpiler (Babel / SWC / TSC), and no runtime JavaScript dependency in production.

3.4.3 Android Maven / Google Maven Dependencies

Per `mobile/settings.gradle.kts` , repository resolution is locked down:

```
dependencyResolutionManagement {
    repositoriesMode.set(RepositoriesMode.FAIL_ON_PROJECT_REPOS)
    repositories { google(); mavenCentral() }
}
pluginManagement {
    repositories { google(); mavenCentral(); gradlePluginPortal() }
}
```

Permitted repositories: Google Maven, Maven Central, Gradle Plugin Portal. **No third-party Maven repositories** are admitted;

`RepositoriesMode.FAIL_ON_PROJECT_REPOS` enforces this at the build level, preventing repository-confusion attacks.

3.4.4 Plugin / Tooling Dependencies

- **Gradle wrapper** — `gradle-8.7-bin.zip` distribution URL, declared in `mobile/gradle/wrapper/gradle-wrapper.properties` with `networkTimeout=10000`.

There are no other tooling dependencies in the project tree (the `@mermaid-js/mermaid-cli` invocation referenced by infrastructure tooling is not part of the AGR project itself).

3.5 THIRD-PARTY SERVICES

The system is **explicitly built to minimize external service dependencies**. Per `aurora_server/data/PUBLIC_TRUST_CHARTER_20260413.md`, the design rejects third-party data brokers, telemetry platforms, and cross-tenant clouds. Per Section 1.3.2, the platform commits to "zero Google / external APIs."

3.5.1 Cloud / Edge Infrastructure

Service	Role	Configuration Surface
Hetzner Cloud	Primary IaaS — 5 of 7 mesh nodes (chimaera, yggdrasil, enterprise, prometheus, galactica) per Section 1.2 SYSTEM OVERVIEW topology	HCLOUD_TOKEN env (resolved via <code>sovereign/lib/hetzner-token-from-secrets-md.sh</code>); credentials in <code>Secrets.md</code> block
Cloudflare	Edge / CDN, DNS, Cloudflare Tunnel, Cloudflare Workers, IndexNow path rules; serves canonical domain <code>auroragalaxyrepublic.com</code> plus 8 redirect aliases	CLOUDFLARE_API_TOKEN env (resolved via <code>sovereign/lib/cloudflare-token-from-secrets-md.sh</code>); cache purge via <code>agr_cf_purge.sh</code> invoked as <code>systemd ExecStartPost</code>
GitHub Actions	CI/CD runner host for 12 workflows	<code>.github/workflows/*.yml</code> ; pinned action versions <code>actions/checkout@v4</code> , <code>actions/setup-java@v4</code> , <code>actions/setup-node@v</code>

Service	Role	Configuration Surface
		4, android-actions/setup-android@v3, actions/upload-artifact@v4
GitHub Releases	Hosts desktop installer manifests latest.yml (Win x64) and latest-mac.yml (Mac ARM64) referencing third-party nexu-io/open-design v0.1.0	The repo does not build desktop binaries; manifests reference external releases only

3.5.2 Local LLM Runtime — llama-server (Loopback Only)

The platform's AI/LLM substrate is **not** a third-party SaaS. It is `llama.cpp`'s `llama-server` binary, deployed as a sibling systemd unit on the same Linux host as the AGR backend.

Component	Configuration
Binary	<code>/usr/local/bin/llama-server</code>
Bind	127.0.0.1:8080 — loopback only, never exposed externally
Default endpoint	<code>AGR_LLM_OPENAI_BASE = http://127.0.0.1:8080</code>
Default model	<code>AGR_LLM_MODEL = gpt-oss</code>
Optional auth	<code>AGR_LLM_API_KEY</code> for proxy-style header authentication
systemd unit	<code>systemd/examples/llama-server.service: ExecStart=/usr/local/bin/llama-server -m /opt/agr/models/REPLACE_ME.guf --host 127.0.0.1 --port 8080 -c 4096</code>
Drop-in config	<code>systemd/examples/agr-republic.service.d/10-llm-env.conf.example</code>

The backend communicates with this loopback endpoint via `stdlib urllib.request` issuing standard `POST /v1/chat/completions` and `POST /v1/embeddings` requests — the same wire protocol as OpenAI's hosted API, but addressed to the in-host server. **No request, prompt, embedding, or response ever traverses the public internet.**

3.5.3 Authentication and Identity (No External IdP)

Per `agr_payments_flags.py` and middleware analysis, **authentication is in-process**. There is no Auth0, no Okta, no Cognito, no Firebase Auth, no OAuth provider integration.

Mechanism	Implementation
TOTP 2FA	<code>agr_totp.py</code> (engine <code>shadow_agr_totp</code> , issuer AuroraGalaxyRepublic, in-memory <code>_STATE</code>) — implemented entirely with <code>stdlib hmac / hashlib</code>
Internal mesh tokens	Headers <code>x-agr-internal-token</code> , <code>x-internal-token</code> , <code>x-guardian-token</code> (env <code>AGR_INTERNAL_API_TOKEN</code>)
Guardian Device Binding	Hardware identifiers (model, serial, EID, <code>wifi_mac</code> , IMEI1, IMEI2, FCC ID, phone, WireGuard IP) via <code>agr_guardian_device_binding.py</code> and <code>agr_handset_identity_from_secrets_md.py</code>

Per Section 2.4.3 Security Implications F-014: hardware identifiers are resolved from outside the source tree (`/opt/agr/.secrets/` profile files); per Section 2.4.4 Maintenance F-014: profiles must be rotated on hardware change.

3.5.4 Payment Processing — Stripe (Gated, Default OFF)

`aurora_server/stripe_commerce.py` implements a **shadow Stripe** integration using `stdlib hmac.compare_digest` for webhook signature

verification.

- Repo default: `PAYMENTS_SURFACES_ENABLED = False` in `agr_payments_flags.py` (per F-017-RQ-001 in Section 2.2)
- Runtime override: `AGR_PAYMENTS_SURFACES_ENABLED=1`
- Webhook secret env: `STRIPE_WEBHOOK_SECRET`
- Per Section 2.4.3 F-017: "Raw-body Stripe signature verify even when surfaces are off; webhook secret env-required"

The `stripe` Python SDK is **never imported**. Only the signature-verification handshake (HMAC-SHA256 over the raw body using the shared secret) is implemented, using `stdlib` primitives. This eliminates the SDK CVE blast radius while preserving webhook correctness.

3.5.5 Email / SMTP

`agr_mail.py` implements an optional SMTP outbound path using `stdlib` `smtplib` (engine `shadow_inproc_mail`).

Configuration Variable	Purpose
<code>AGR_SMTP_HOST / AGR_SMTP_PORT</code>	SMTP relay endpoint
<code>AGR_SMTP_USER / AGR_SMTP_PASSWORD</code>	Optional auth
<code>AGR_SMTP_NO_TLS</code>	Disable TLS for local relays
<code>AGR_SMTP_TIMEOUT</code>	Connect/read timeout
<code>AGR_MAIL_QUEUE_ENABLE / AGR_MAIL_QUEUE_DB</code>	Optional SQLite outbound queue at <code>data/mail_outbound_queue.db</code>
<code>AGR_MAIL_FLUSH_HTTP_TOKEN</code>	HTTP flush endpoint gate; matched against <code>X-AGR-Mail-Flush-Token</code> header (per F-018)

Per Section 2.4.4 F-018, queue pruning and flush-token rotation are routine maintenance items.

3.5.6 SEO and Discovery Services

Service	Integration
IndexNow	sovereign/indexnow-submit-sitemap.sh runs after fleet-deploy to ping major search engines
Schema.org / JSON-LD	Generated by agr_seo.py with canonical auroragalaxyrepublic.com
robots.txt / sitemap.xml	Static under aurora_server/static/ (52 URLs in sitemap per Section 1.3.1)

3.5.7 Monitoring / Observability — In-Process Only

There is **no third-party APM, logging, or telemetry**. No Datadog, New Relic, Sentry, CloudWatch, Splunk, Honeycomb, Lightstep, or equivalent.

Mechanism	Implementation
Operational state	JSONL ledgers under aurora_server/state/: failover_state.json, sync_log.jsonl, sync_state.json, proof_gate_*.json, *.last markers
Node metrics	psutil-driven from agr_failover.SovereignSentinel daemon thread
Health probes	/health, /api/health, /api/public/health endpoints; verified per Hetzner public IP via bash sovereign/fleet-verify-public-http.sh

This locality of telemetry is itself a sovereignty constraint: per Section 2.4.3 F-001 / F-002 / F-003, "no telemetry leakage; no external network call required for reasoning."

3.5.8 Explicit Service Exclusions

Per `LIBRARY_OF_LIGHT_CHARTER` and `PUBLIC_TRUST_CHARTER`, the following service categories are **forbidden** by charter and **not present** in any code path:

Excluded Category	Examples	Rationale
External AI provider APIs	OpenAI, Anthropic, Cohere, Google Gemini, AWS Bedrock	"Self-hosted. Eternal." — local llama.cpp only
Third-party analytics / tracking	Google Analytics, Mixpanel, Segment, Amplitude, Hotjar	Charter rejects "behavioral surveillance / ad-targeting profiles"
External authentication	Auth0, Okta, Firebase Auth, AWS Cognito, OAuth-with-Google/Facebook	In-process TOTP plus internal tokens only
Managed databases	RDS, MongoDB Atlas, DynamoDB, CosmosDB, Firestore	SQLite + optional self-hosted Postgres only
Cloud object storage	S3, Azure Blob, GCS, R2, DigitalOcean Spaces	Local /opt/agr/vault per nodes
Crisis / mental-health pathways	988, Crisis Text Line referrals	Per F-031-RQ-006 in Section 2.4: "No 988 / generic crisis pathways or 'talk to a professional' defaults"
Mandatory paywalls	Subscription gates on public chat	Per Section 1.3.2: "Charter forbids; default tier is free-public-benefit"

3.6 DATABASES & STORAGE

3.6.1 Primary Storage — SQLite

The platform is **SQLite-first** with no centralized DBMS server in the default deployment. SQLite databases are created lazily on first write via `CREATE TABLE IF NOT EXISTS` statements embedded in Python modules. Database files live under `aurora_server/data/` and `aurora_server/state/` and are excluded from version control via `.gitignore` (`*.db`, `*.db-journal`).

Database File	Module	Purpose
<code>.agr_vault_rag.sqlite</code>	<code>agr_vault_rag.py</code>	Master Vault FTS5 index plus optional embeddings table
<code>agr_tandem.db</code>	<code>agr_tandem</code>	Mesh sync ledger
<code>citizen_consciousness_engine.db</code>	<code>citizen pipeline</code>	Conversation and state
<code>consciousness.db</code> , <code>consciousness_conversations.db</code> , <code>consciousness_core.db</code>	<code>consciousness substrate</code>	Thoughts, conversations, core state
<code>data/mail_outbound_queue.db</code>	<code>agr_mail.py</code>	Optional SMTP queue (gated by env)
Various other <code>*.db</code> files	runtime modules	Excluded from VCS by <code>.gitignore</code>

The choice of SQLite is itself a security decision: a single-file database has no listening socket, no authentication surface that can be misconfigured, and trivial backup semantics (just copy the file). It also dovetails with the `stdlib` doctrine — no driver package to vet.

3.6.2 SQLite FTS5 — Vault RAG Search Substrate

`agr_vault_rag.py` implements the SQLite FTS5 virtual-table pattern for full-text search across the Master Vault:

```
CREATE VIRTUAL TABLE IF NOT EXISTS vault_fts USING fts5(...);
CREATE TABLE IF NOT EXISTS vault_embeddings (
```

```

    chunk_id, path, body_sha256, dim, vec, updated_at);
CREATE INDEX IF NOT EXISTS idx_vault_embeddings_path ON
vault_embeddings(path);

```

Operational constants (per Section 2.4.1 Technical Constraints F-008 and `agr_vault_rag.py`):

Constant	Value	Source
DEFAULT_CHUNK	2,400 chars	F-008-RQ-003
DEFAULT_OVERLAP	200 chars	F-008-RQ-003
DEFAULT_TOP_K	8	F-008-RQ-004
DEFAULT_MAX_CHARS	65,536	F-008-RQ-004
_MAX_RAG_INJECT_CHARS	2,000,000 (hard cap)	F-008-RQ-005
_MAX_CHAT_MESSAGE_ROWS	2,048	<code>agr_vault_rag.py</code>
_MAX_MEMORY_LANE_TAIL_CHARS	500,000	<code>agr_vault_rag.py</code>
Default vault root	<code>/opt/agr/vault</code>	env <code>AGR_MASTER_VAULT_ROOT</code>
Connection timeout	30 seconds	<code>agr_vault_rag.py</code>

These constants bound the LLM input cost per request and are the primary tuning surface for Section 2.4.2 Performance and Scalability Requirements.

3.6.3 Optional Secondary — PostgreSQL (`agr_pg.py`)

`agr_pg.py` implements a strict, opt-in, fail-closed PostgreSQL bridge for operators who choose to deploy a Postgres instance alongside the SQLite primary.

Property	Value
Activation	AGR_PG_ENABLED=1 plus DSN via AGR_PG_DSN or DATABASE_URL
Schema prefix	AGR_PG_SCHEMA_PREFIX (default agr_)
Connect timeout	3 seconds
Identifier validation	Regex_IDENT (rejects malformed table/column names)
Count WHERE clause	Must be exactly 1=1 (defensive constant)
Failure mode	All operations return 0 or {} if Postgres is unavailable
systemd ordering	agr-republic.service: After=network.target postgresql.service, Wants=postgresql.service

The `Wants=postgresql.service` (rather than `Requires=`) directive is deliberate — Postgres is *desirable* but not *required*. The unit will start cleanly even if Postgres is uninstalled.

3.6.4 File-Backed State (JSON / JSONL / Markdown)

The third tier of persistence is the file system itself, used for declarative artifacts and append-only ledgers.

Path	Format	Role
aurora_server/state/full_network_surface_audit.latest.json	JSON	Audit snapshot
aurora_server/state/sequential_abcd_audit.latest.json	JSON	Sequential audit

Path	Format	Role
aurora_server/state/device_sovereignty_runbook.latest.json	JSON	Device sovereignty runbook
aurora_server/state/*.jsonl	JSONL	Append-only history (workspace_jobs, workspace_chat_sessions, workspace_governance_proposals, workspace_commerce_events, workspace_evolution_ticks, sync_log, etc.) — per Section 2.4.4 F-028
aurora_server/state/*.md	Markdown	Runtime markdown snapshots (e.g., mission_totality.latest.md)
aurora_server/data/great_wall/exiles.json, events.jsonl	JSON + JSONL	Exile registry (agr_great_wall.py)
aurora_server/data/*.json, *.md	JSON + Markdown	40+ governance / policy / aesthetic artifacts (constitution, charters, capability matrix, etc.)

Constitutional and policy artifacts are append-only and checksum-locked via `constitution_lock_manifest.json` (per Section 2.4.4 F-012). The `.gitignore` excludes runtime state (`*.db`, `*.db-journal`, `aurora_server/data/cold_storage/`, `aurora_server/data/accounts/`, `aurora_server/state/*.jsonl`, `aurora_server/state/*.json`, `aurora_server/state/*.last`, `.agr-git-revision`) so that runtime mutation does not pollute the source tree.

3.6.5 Caching — In-Process Only

There is **no Redis**, **no Memcached**, **no external cache tier**.

Mechanism	Use
<code>functools.lru_cache</code>	Method-level memoization for pure functions

Mechanism	Use
Module-level dict caches	Stateful caches (e.g., aurora_comms._LOCK-guarded dicts)
threading.RLock	Read-modify-write protection with copy-out for thread safety

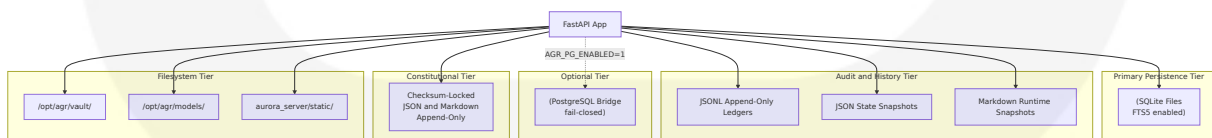
This caching strategy is per-process; it does not survive restart, but per Section 2.4.4 F-012 the constitutional ledger is the durable system-of-record, so cache loss is acceptable.

3.6.6 Storage Services — Local Filesystem Only

There is **no S3, Azure Blob, GCS, or DigitalOcean Spaces** integration.

Storage Concern	Solution
Vault content	/opt/agr/vault on each fleet node (per-node, replicated via sovereign/scripts/fleet-vault-kora-rsync.sh)
Mesh replication	Protocol agr_mesh_v1_chunked with ≥ 20 MB chunk policy (per F-032-RQ-004)
Static assets	aurora_server/static/ served by FastAPI's StaticFiles
Models	/opt/agr/models/ for *.gguf weight files (operator-supplied; outside git)

3.6.7 Persistence Strategy Summary



3.7 DEVELOPMENT & DEPLOYMENT

3.7.1 Build System

Backend (Python) — No Build Step

The Python backend has **no build artifact**. `agr_start_wrapper.sh` directly invokes `uvicorn` against the source tree at `/opt/agr/aurora_server/`. There is no `setup.py`, no `pyproject.toml`, no wheel packaging. The Python virtual environment lives at `/opt/agr-venv/` and is provisioned per-node by the operator outside of repository tooling.

Mobile (Android) — Gradle 8.7

The mobile module uses Gradle 8.7 via the wrapper, configured in `mobile/gradle/wrapper/gradle-wrapper.properties` with the distribution URL pointing to `gradle-8.7-bin.zip`.

Property	Value
Gradle version	8.7
org.gradle.jvmargs	-Xmx2g -Dfile.encoding=UTF-8
android.useAndroidX	true
kotlin.code.style	official
android.nonTransitiveRClass	true
Debug build command	<code>./gradlew :gpt-oss-chat:assembleDebug</code>
Release build command	<code>./gradlew :gpt-oss-chat:assembleRelease</code>
Output paths	<code>mobile/gpt-oss-chat/build/outputs/apk/{debug,release}/*.apk</code>

E2E (Playwright) — npm

The E2E workspace uses standard `npm install` plus `npx playwright install --with-deps chromium firefox webkit`, then `npm test` (which

executes `playwright test`).

3.7.2 Containerization Posture — Direct systemd, No Containers

The AGR application is intentionally not containerized. There is no `Dockerfile` for the application, no `docker-compose.yml` , no Kubernetes manifests, no Helm charts, and no container registry references in the project tree.

The deployment model is **direct systemd on Hetzner Linux nodes**. This choice is doctrinal:

1. **Reduced container-escape risk** — OS-level isolation via systemd unit boundaries, no shared kernel attack surface across tenants.
2. **Simpler audit** — the running binary is the same `python` interpreter on the system, with the same `sqlite3` library, the same OpenSSL — auditable via standard Linux tooling.
3. **No registry coupling** — no container registry to authenticate against, no image-signing pipeline to maintain.
4. **Sovereignty alignment** — per `PUBLIC_TRUST_CHARTER` , the design rejects "lock-in monopolies"; Docker Hub / ECR / GCR coupling would violate this principle.

This is a deliberate departure from the conventional default-stack pattern of "Cloud Platform: AWS / Containerization: Docker / IaC: Terraform / CI/CD: GitHub Actions." Of those four defaults, only GitHub Actions is adopted; AWS, Docker, and Terraform are explicitly rejected in favor of Hetzner + systemd + bash scripts.

3.7.3 systemd Deployment

Primary Unit — `systemd/agr-republic.service`

[Unit]

```
Description=Aurora Galaxy Republic OS
After=network.target postgresql.service
Wants=postgresql.service
```

[Service]

```
Type=simple
WorkingDirectory=/opt/agr/aurora_server
ExecStart=/opt/agr-venv/bin/uvicorn republic_os_server:app --host
0.0.0.0 --port 5000
ExecStartPost=/opt/agr/aurora_server/agr_cf_purge.sh
Restart=always
RestartSec=10
Environment=PYTHONUNBUFFERED=1
EnvironmentFile=/etc/agr/agr-server.env
```

[Install]

```
WantedBy=multi-user.target
```

The `ExecStartPost=agr_cf_purge.sh` invocation triggers a Cloudflare cache purge after every successful start, ensuring the edge does not serve stale assets after a deploy.

Optional Companion Unit — `systemd/examples/llama-server.service`

```
ExecStart=/usr/local/bin/llama-server -m
/opt/agr/models/REPLACE_ME.gguf --host 127.0.0.1 --port 8080 -c 4096
Restart=on-failure
RestartSec=5
```

Drop-In Configuration — `systemd/examples/agr-republic.service.d/`

Two example drop-ins are shipped for operator customization:

Drop-In	Purpose
10-llm-env.conf.example	Sets AGR_LLM_OPENAI_BASE, AGR_LLM_MODEL, optional AGR_LLM_API_KEY
20-chat-vault-rag.conf.example	Sets AGR_CHAT_VAULT_RAG, AGR_CHAT_VAULT_RAG_CHANNELS, AGR_CHAT_VAULT_RAG_TOP_K, AGR_CHAT_VAULT_RAG_MAX_CHARS

3.7.4 Fleet Deployment Scripts

The `sovereign/` folder contains 60+ operational scripts. The most consequential for deployment are:

Script	Purpose
<code>sovereign/fleet-pull-main-restart.sh</code>	Pull main and restart <code>agr-republic.service</code> on each Hetzner node
<code>sovereign/fleet-pull-with-secrets-md.sh</code>	Variant using <code>Secrets.md</code> SSH key
<code>sovereign/fleet-mirror-repo-to-nodes.sh</code>	Sync repository to nodes
<code>sovereign/scripts/fleet-install-agr-republic-unit.sh</code>	Install or refresh the systemd unit
<code>sovereign/scripts/fleet-republic-hard-restart.sh</code>	Force restart
<code>sovereign/agent-minimum-gate.py</code>	Five-node SSH baseline gate
<code>sovereign/scripts/fleet-bash-syntax-check.sh</code>	<code>bash -n</code> syntax check across all fleet shell scripts
<code>sovereign/indexnow-submit-sitemap.sh</code>	Post-deploy IndexNow submission

3.7.5 CI/CD — GitHub Actions

Per Section 2.4.4 Maintenance Requirements, the 12 workflows under `.github/workflows/` constitute the maintenance harness.

Workflow File	Trigger	Purpose
fleet-deploy-pull.yml	push to main + workflow_dispatch	Pull main on Hetzner nodes; restart agr-republic.service; Tower smoke + IndexNow
tower1-public-smoke.yml	schedule + path-filtered push/PR + workflow_dispatch	Live HTTPS smoke against canonical Tower 1
tower1-frozen-urls-playwright.yml	workflow_dispatch + path-filtered push/PR	Playwright E2E (118+ paths Chromium + cross-browser smoke)
tower1-frozen-inventory-crawl.yml	schedule + workflow_dispatch + path-filtered	GET crawl of frozen URL inventory
fleet-verify-public-http.yml	schedule + workflow_dispatch + path-filtered	Public HTTP /health checks per Hetzner IP
fleet-bootstrap-actions-fleet-key.yml	workflow_dispatch only	One-time bootstrap of AGR_FLEET_KEY_CONTENT secret (uses pip install pynacl)
phones-only-public-verify.yml	workflow_dispatch + path-filtered push/PR	Phones-only fallback validation
phones-two-node-live-surface-bundle.yml	workflow_dispatch + path-filtered	Two-node handset bundle
guardian-device-binding-verify.yml	path-filtered push	Guardian device profile loading tests
agent-progress-guardian-node-verify.yml	path-filtered push	AGENT_PROGRESS structural guard
tier1-static-refs-verify.yml	path-filtered push/PR + workflow_dispatch	Static HTML reference verification

Workflow File	Trigger	Purpose
android-gpt-oss-chat-apk.yml	path-filtered push/P R + workflow_dispatch	Android APK build (assembleDebug + assembleRelease)

CI Runner Provisioning

Tool	Version	Provisioned Via
Runner OS	ubuntu-latest	GitHub-hosted
JDK distribution	Temurin	actions/setup-java@v4
JDK version	17	actions/setup-java@v4
Android SDK	platform-tools, platforms;android-34, build-tools;34.0.0	android-actions/setup-android@v3
Node.js	20	actions/setup-node@v4 (Playwright workflow)
Checkout	(latest)	actions/checkout@v4
Artifact upload	(latest)	actions/upload-artifact@v4

3.7.6 Required Secrets and Tokens

Secret	Scope	Purpose
AGR_FLEET_KEY_CONTENT	GitHub Actions	Fleet SSH PEM (raw key)
FLEET_KEY_MATERIAL_B64	GitHub Actions	Base64-encoded PEM for split-secret bootstrap
FLEET_GITHUB_ACTIONS_WRITE_TOKEN	GitHub Actions	PAT to write AGR_FLEET_KEY_CONTENT
HCLOUD_TOKEN	runtime / CI	Hetzner Cloud API
CLOUDFLARE_API_TOKEN	runtime / CI	Cloudflare ops

Secret	Scope	Purpose
STRIPE_WEBHOOK_SECRET	runtime	Stripe webhook signature verification
AGR_INTERNAL_API_TOKEN	runtime	Internal mesh control plane
AGR_LLM_API_KEY	runtime, optional	llama-server proxy auth
AGR_MAIL_FLUSH_HTTP_TOKEN	runtime, optional	Outbound mail queue flush gate
Operator copy: Secrets.md	local dev	Local operator vault containing hetzner token:, agr fleet key b64:, etc.

Per Section 2.4.3 F-014, hardware identifiers (IMEI, serial, EID, phone, MAC) are **never** committed to git; profile files live at `/opt/agr/.secrets/` outside the source tree. The `Secrets.md` block in the operator's local copy is gitignore'd.

3.7.7 Development Tools

Tool	Configuration
Pre-commit hook	<code>.githooks/pre-commit</code> — single shell hook installed via local git config; not the pre-commit framework
Bash syntax linter	<code>sovereign/scripts/fleet-bash-syntax-check.sh</code> (bash -n over all fleet shell scripts) — gates every CI workflow
Test runner (backend)	<code>stdlib unittest</code> (no <code>pytest</code>)
Type stubs (test isolation)	<code>aurora_server/tests/fastapi_import_stubs.py</code> — provides import-time stubs for FastAPI when running tests without full dependency chain

3.7.8 Verification Commands (Release Gates)

Per Section 1.2.3 Success Criteria and Section 2.4.2 Performance and Scalability Requirements:

Command	Verification Domain
bash sovereign/fleet-verify-public-http.sh	Probes /health then /api/health per He tzner public IP
bash sovereign/tower1-public-smoke.sh	Canonical hostname checks: engine-runtime, chat POST, MIR-L, /dl/*, laws/c heck
python3 sovereign/agent-minimum-gate.py	Five-node SSH baseline
bash sovereign/scripts/agr-launch-readiness-tower-smoke.sh	Full go-live (bash smoke + Python constitution / engine checks)

3.8 INTEGRATION REQUIREMENTS BETWEEN COMPONENTS

3.8.1 Backend ↔ Local LLM

Concern	Value
Protocol	OpenAI-compatible REST (POST /v1/chat/completions, POST /v1/embeddings)
Transport	stdlib urllib.request (no httpx / openai SDK dependency)
Endpoint	127.0.0.1:8080 — loopback only, never exposed publicly
Default model	gpt-oss
Default temperature	1.25 (env AGR_LLM_TEMPERATURE)

3.8.2 Backend ↔ Mobile

Concern	Value
Protocol	HTTPS REST against canonical auroragalaxyrepublic.com
Probe ladder	/health → /api/health → /api/public/health (first 200 wins)
Chat surface	POST /api/republic/chat, POST /api/public/citizen-engine-advice
Body	message, mode="text", session_id (from ANDROID_ID), citizen_id, optional display_name
Auth	No mobile-specific token; rate limits enforced server-side

3.8.3 Backend ↔ Mesh Peers (7-Node Topology)

Per Section 1.2 SYSTEM OVERVIEW and Section 2.4 F-032:

Concern	Value
Protocol	HTTP plus SSH (via stdlib subprocess + system ssh; no paramiko)
Internal headers	x-agr-internal-token, x-internal-token, x-guardian-token
Mirror protocol	agr_mesh_v1_chunked with ≥ 20 MB chunk policy (F-032-RQ-004)
Failover poll interval	20 seconds (<code>_POLL_INTERVAL</code>)
Failover failure threshold	3 consecutive failures (<code>_FAILURE_THRESH</code>)
SSH timeout	12 seconds (<code>_SSH_TIMEOUT</code>)
Topology	5 Hetzner (chimaera, yggdrasil, enterprise, prometheus, galactica) + iPhone 17 Pro (#6) + OnePlus 15 (#7)

3.8.4 Backend ↔ Cloudflare

Concern	Value
Tunnel	cloudflared connects OnePlus 15 (Termux + uvicorn) as fallback origin when Hetzner is off (per F-032-RQ-005)
Workers + path rules	Canonical domain plus 8 redirect aliases
Cache purge	agr_cf_purge.sh invoked via systemd ExecStartPost
IndexNow	sovereign/indexnow-submit-sitemap.sh after deploy

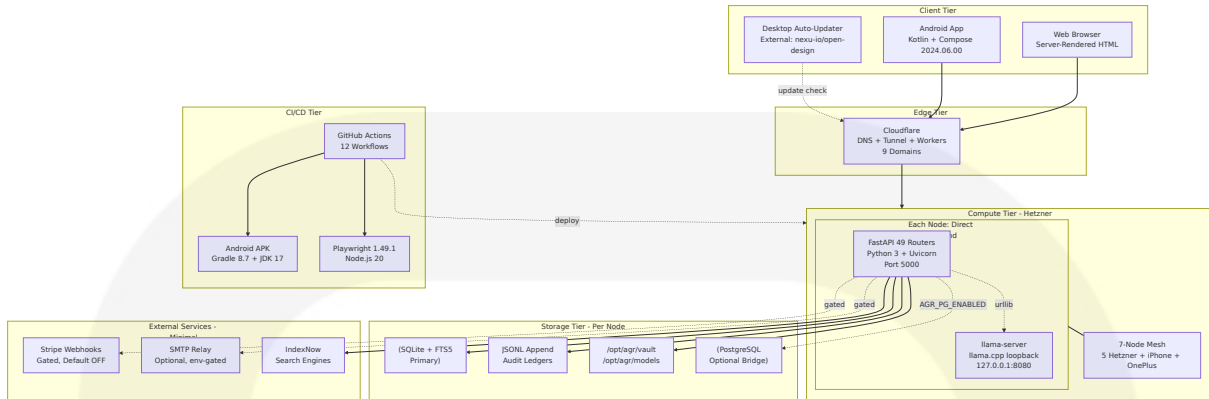
3.9 VERSION COMPATIBILITY MATRIX

Component	Minimum Version	Pinned Version	Source of Truth
Python	≥ 3.9	unpinned at runtime	PEP-604 union syntax in code
FastAPI	unpinned	unpinned	import fastapi
Starlette	unpinned	(transitive via FastAPI)	from starlette.middleware.base import BaseHTTPMiddleware
Uvicorn	unpinned	unpinned	agr_start_wrapper.sh
Kotlin	1.9.24	1.9.24	mobile/gpt-oss-chat/build.gradle.kts line 3
Android Gradle Plugin	8.5.2	8.5.2	mobile/gpt-oss-chat/build.gradle.kts line 2
Compose BOM	2024.06.00	2024.06.00	mobile/gpt-oss-chat/build.gradle.kts line 46
Compose Compiler Extension	1.5.14	1.5.14	mobile/gpt-oss-chat/build.gradle.kts line 41

Component	Minimum Version	Pinned Version	Source of Truth
androidx.activity:activity-compose	1.9.0	1.9.0	mobile/gpt-oss-chat/build.gradle.kts line 48
kotlinx.coroutines-android	1.8.1	1.8.1	mobile/gpt-oss-chat/build.gradle.kts line 52
Java JVM target	17	17	mobile/gpt-oss-chat/build.gradle.kts lines 30-34
Gradle	8.7	8.7	mobile/gradle/wrapper/gradle-wrapper.properties
compileSdk / targetSdk	34	34	mobile/gpt-oss-chat/build.gradle.kts lines 8, 13
minSdk	26	26	mobile/gpt-oss-chat/build.gradle.kts line 12
@playwright/test	1.49.1	^1.49.1	sovereign/e2e/frozen-urls/package.json
Node.js (CI)	20	20	actions/setup-node@v4
GitHub Actions	@v4	@v4	All workflow files

3.10 TECHNOLOGY STACK ARCHITECTURE DIAGRAM

The complete stack and its integration boundaries:



3.11 SECURITY IMPLICATIONS OF STACK CHOICES

Each technology selection in this stack carries security consequences. Per Section 2.4.3 Security Implications, the Republic's broader security stance is "absolute sovereignty, zero external control."

Stack Choice	Security Posture
Stdlib-heavy Python	Minimizes supply-chain attack surface; no PyPI typosquat exposure for ~95% of code
Local llama.cpp only	No external AI provider can log, train on, or leak inputs; meets LIBRARY_OF_LIGHT_CHARTER
Shadow Stripe (no stripe lib)	Single secret (STRIPE_WEBHOOK_SECRET) verified via hmac.compare_digest; no SDK CVE exposure
SQLite + FTSS (no external vector DB)	Single-file persistence; trivial backup; no network attack surface; no auth-misconfiguration class of bug
systemd direct (no Docker)	Reduced container-escape risk; OS-level isolation via systemd unit boundaries
fail_closed: true defaults	Per ALWAYS_FIRST_PRIORITIES_20260412.json — failure modes block release rather than degrade silently

Stack Choice	Security Posture
<code>block_release_when_missing_or_failed: true</code>	Constitutional and policy guard active in CI
TOTP via stdlib	No pyotp / pyjwt / passlib CVE exposure
Hardware identifier never committed	Per F-014-RQ-004; profile files at <code>/opt/agr/.secrets/</code> outside source tree
In-memory <code>_STATE</code> for TOTP	Survives no longer than the process; replay protection via timestamp window
<code>RepositoriesMode.FAIL_ON_PROJECT_REPOS</code> (Gradle)	Prevents Maven repository confusion attacks in Android build
Release APK signed with debug keystore	Sideload-only — explicit warning in mobile/README.md; operator must replace with production keystore before public store distribution (per F-030 maintenance requirement)
<code>Wants=postgresql.service</code> (not <code>Requires=</code>)	systemd unit starts cleanly without Postgres; eliminates startup-blocking dependency
Cloudflare tunnel for OnePlus origin	Disaster fallback (per F-032-RQ-005) without exposing handset's home network publicly

3.12 References

3.12.1 Configuration / Build Files Examined

- `mobile/gpt-oss-chat/build.gradle.kts` — AGP 8.5.2, Kotlin 1.9.24, Compose BOM 2024.06.00, all version pins
- `mobile/gradle.properties` — JVM args, AndroidX flag, Kotlin code style
- `mobile/settings.gradle.kts` — Repository resolution, `FAIL_ON_PROJECT_REPOS` enforcement
- `mobile/build.gradle.kts` — Comment-only stub

- `mobile/gradle/wrapper/gradle-wrapper.properties` — Gradle 8.7 distribution URL
- `mobile/gpt-oss-chat/src/main/AndroidManifest.xml` — INTERNET permission, MainActivity declaration
- `sovereign/e2e/frozen-urls/package.json` — Playwright 1.49.1 dependency
- `sovereign/e2e/frozen-urls/playwright.config.ts` — Playwright 4-project configuration
- `agr_start_wrapper.sh` — Uvicorn invocation
- `systemd/agr-republic.service` — Primary systemd unit
- `systemd/examples/llama-server.service` — Optional llama.cpp companion unit
- `systemd/examples/README.md` — Deployment template guide
- `systemd/examples/agr-republic.service.d/10-llm-env.conf.example` — LLM env drop-in
- `systemd/examples/agr-republic.service.d/20-chat-vault-rag.conf.example` — Vault RAG env drop-in
- `.gitignore` — `*.db` / `*.db-journal` exclusions, mobile build artifacts
- `latest.yml` — Desktop Win x64 manifest (external `nexu-io/open-design v0.1.0`)
- `latest-mac.yml` — Desktop Mac ARM64 manifest

3.12.2 Backend Source Files Referenced

- `aurora_server/republic_os_server.py` — Composition root, ~29,817 lines, 49-router registration, middleware pipeline (lines 7026-7049)
- `aurora_server/agr_vault_rag.py` — SQLite FTS5 schema, RAG constants (`_MAX_RAG_INJECT_CHARS = 2_000_000`)
- `aurora_server/agr_pg.py` — Optional PostgreSQL bridge (fail-closed)
- `aurora_server/agr_failover.py` — `SovereignSentinel` daemon, `_POLL_INTERVAL=20`, `_FAILURE_THRESH=3`, `_SSH_TIMEOUT=12`
- `aurora_server/agr_totp.py` — TOTP via `stdlib` `hmac` / `hashlib`
- `aurora_server/agr_mail.py` — SMTP via `stdlib` `smtplib`

- `aurora_server/agr_payments_flags.py` — `PAYMENTS_SURFACES_ENABLED = False` default
- `aurora_server/stripe_commerce.py` — Shadow Stripe webhook verification
- `aurora_server/agr_seo.py` — [Schema.org](https://schema.org/) / JSON-LD generation
- `aurora_server/agr_live_support.py` — Inline HTML/CSS/JS widget
- `aurora_server/agr_guardian_device_binding.py` — Hardware identifier resolution
- `aurora_server/agr_handset_identity_from_secrets_md.py` — `Secrets.md` profile loader
- `aurora_server/aurora_comms.py` — Shadow comms hub
- `aurora_server/routes/` — 54 router modules
- `aurora_server/tests/fastapi_import_stubs.py` — Test-time FastAPI stubs

3.12.3 Operations and CI Files Referenced

- `.github/workflows/fleet-deploy-pull.yml` — Main fleet deploy entrypoint
- `.github/workflows/android-gpt-oss-chat-apk.yml` — Android APK build (JDK 17 + Android SDK 34)
- `.github/workflows/tower1-frozen-urls-playwright.yml` — Playwright E2E (Node.js 20)
- `.github/workflows/tower1-public-smoke.yml` — Public smoke
- `.github/workflows/fleet-verify-public-http.yml` — Hetzner public HTTP probe
- `.github/workflows/guardian-device-binding-verify.yml` — Guardian binding verify
- `.github/workflows/tier1-static-refs-verify.yml` — Static HTML reference verification
- `.github/workflows/phones-only-public-verify.yml` — Phones-only fallback verification
- `.github/workflows/tower1-frozen-inventory-crawl.yml` — Frozen inventory crawl

- `.github/workflows/agent-progress-guardian-node-verify.yml` — Progress doc structure guard
- `.github/workflows/phones-two-node-live-surface-bundle.yml` — Two-node bundle
- `.github/workflows/fleet-bootstrap-actions-fleet-key.yml` — One-time fleet key bootstrap
- `sovereign/scripts/fleet-bash-syntax-check.sh` — `bash -n` linter
- `sovereign/fleet-pull-main-restart.sh` — Pull and restart
- `sovereign/fleet-verify-public-http.sh` — Public HTTP probe
- `sovereign/tower1-public-smoke.sh` — Canonical hostname smoke
- `sovereign/agent-minimum-gate.py` — Five-node SSH baseline
- `sovereign/scripts/agr-launch-readiness-tower-smoke.sh` — Full go-live gate
- `sovereign/indexnow-submit-sitemap.sh` — Post-deploy IndexNow submission
- `.githooks/pre-commit` — Local pre-commit hook (single shell script)

3.12.4 Folders Referenced

- `aurora_server/` — Backend root (~150 modules)
- `aurora_server/data/` — Declarative governance / policy / aesthetic JSON and Markdown
- `aurora_server/state/` — Latest snapshots and JSONL histories
- `aurora_server/static/` — Static assets served by `StaticFiles`
- `aurora_server/tests/` — 90+ regression tests
- `aurora_server/routes/` — 54 router modules
- `mobile/` — Gradle Android project
- `sovereign/` — Operations control plane (60+ scripts)
- `systemd/` and `systemd/examples/` — Deployment unit templates
- `.github/workflows/` — 12 CI workflows
- `.githooks/` — Local pre-commit hook

3.12.5 Governance and Charter Documents Referenced

- `aurora_server/data/FOUNDER_PROFILE_BRAD_REINHOLD.json` — Sovereignty doctrine
- `aurora_server/data/LIBRARY_OF_LIGHT_CHARTER_20260414.json` — Constitutional baseline
- `aurora_server/data/PUBLIC_TRUST_CHARTER_20260413.md` — Trust posture
- `aurora_server/data/LUMEN_SANCTUM_FOUNDATION_20260414.json` — Operating doctrine
- `aurora_server/data/ALWAYS_FIRST_PRIORITIES_20260412.json` — `fail_closed: true`, `block_release_when_missing_or_failed: true`
- `aurora_server/data/PLATFORM_SCOPE_GAP_REPORT_20260413.md` — Replacement posture
- `aurora_server/data/CAPABILITY_TRACEABILITY_MATRIX_LATEST.json` — Capability verification
- `aurora_server/data/SOUL_CONTINUITY_PROTOCOL_20260412.md` — Continuity scope
- `aurora_server/data/FUSION_REALITY_INTERACTION_LOCK_20260414.json` — Reality-interaction hard locks
- `REMAINING_WORK_ORDER_OF_OPERATIONS.md` — Phase 5 scope
- `AUDIT_REPORT_20260430.md` — 2026-04-30 audit baseline
- `sovereign/PLATFORM_COMPLETION_STATUS.md` — Verification gates
- `sovereign/PHONES_ONLY_PUBLIC_SURFACE.md` — Disaster posture
- `mobile/README.md` — Mobile build instructions and probe ladder
- `Secrets.md` — Operator vault structure (gitignore'd)
- `AGENTS.md` — AI agent governance
- `HANDOFF_FOR_NEXT_AGENT.md` — Continuity scope and safety constraints

3.12.6 Cross-Section References

- Section 1.2 SYSTEM OVERVIEW — Middleware pipeline, 7-node topology, technical approach

- Section 1.3 SCOPE — In-scope capabilities, out-of-scope exclusions (Swift, JavaScript frameworks, mandatory paywalls, behavioral surveillance)
- Section 2.2 FUNCTIONAL REQUIREMENTS — Feature requirements F-001 through F-032 (especially F-008 vault constants, F-017 payments default, F-029 Wave 4 proof, F-030 mobile pins, F-032 mesh constants)
- Section 2.4 IMPLEMENTATION CONSIDERATIONS — Technical constraints, performance posture, security implications, maintenance requirements, 12-workflow harness

4. Process Flowchart

This section translates the live capability surface enumerated in Section 1.2 SYSTEM OVERVIEW and the requirement matrix in Section 2.2 FUNCTIONAL REQUIREMENTS into end-to-end process flows. Every diagram below is grounded in observed source code (`aurora_server/republic_os_server.py` , the 49 router modules under `aurora_server/routes/` , the failover and tandem daemons in `aurora_server/agr_failover.py` and `aurora_server/agr_tandem.py` , and the GitHub Actions workflows under `.github/workflows/`). Diagrams use Mermaid.js syntax and reference Feature IDs (F-001 through F-036) and Requirement IDs (F-XXX-RQ-YYY) defined in Section 2.

The system follows three architectural patterns relevant to flow design:

1. **Shadow Implementations** — Per

`REMAINING_WORK_ORDER_OF_OPERATIONS.md` : most peripheral `aurora_server` modules are structured-JSON / in-process state, not empty stubs.

Routes return JSON with a `surface` marker; many endpoints proxy thin adapters to `agr_*` and `aurora_*` domain helpers.

2. **Fail-Closed Posture** — Per

`aurora_server/data/ALWAYS_FIRST_PRIORITIES_20260412.json` :

```
fail_closed: true , block_release_when_missing_or_failed: true ,  
applies_to_all_operator_windows: true .
```

3. **Zero-External-API Doctrine** — Per

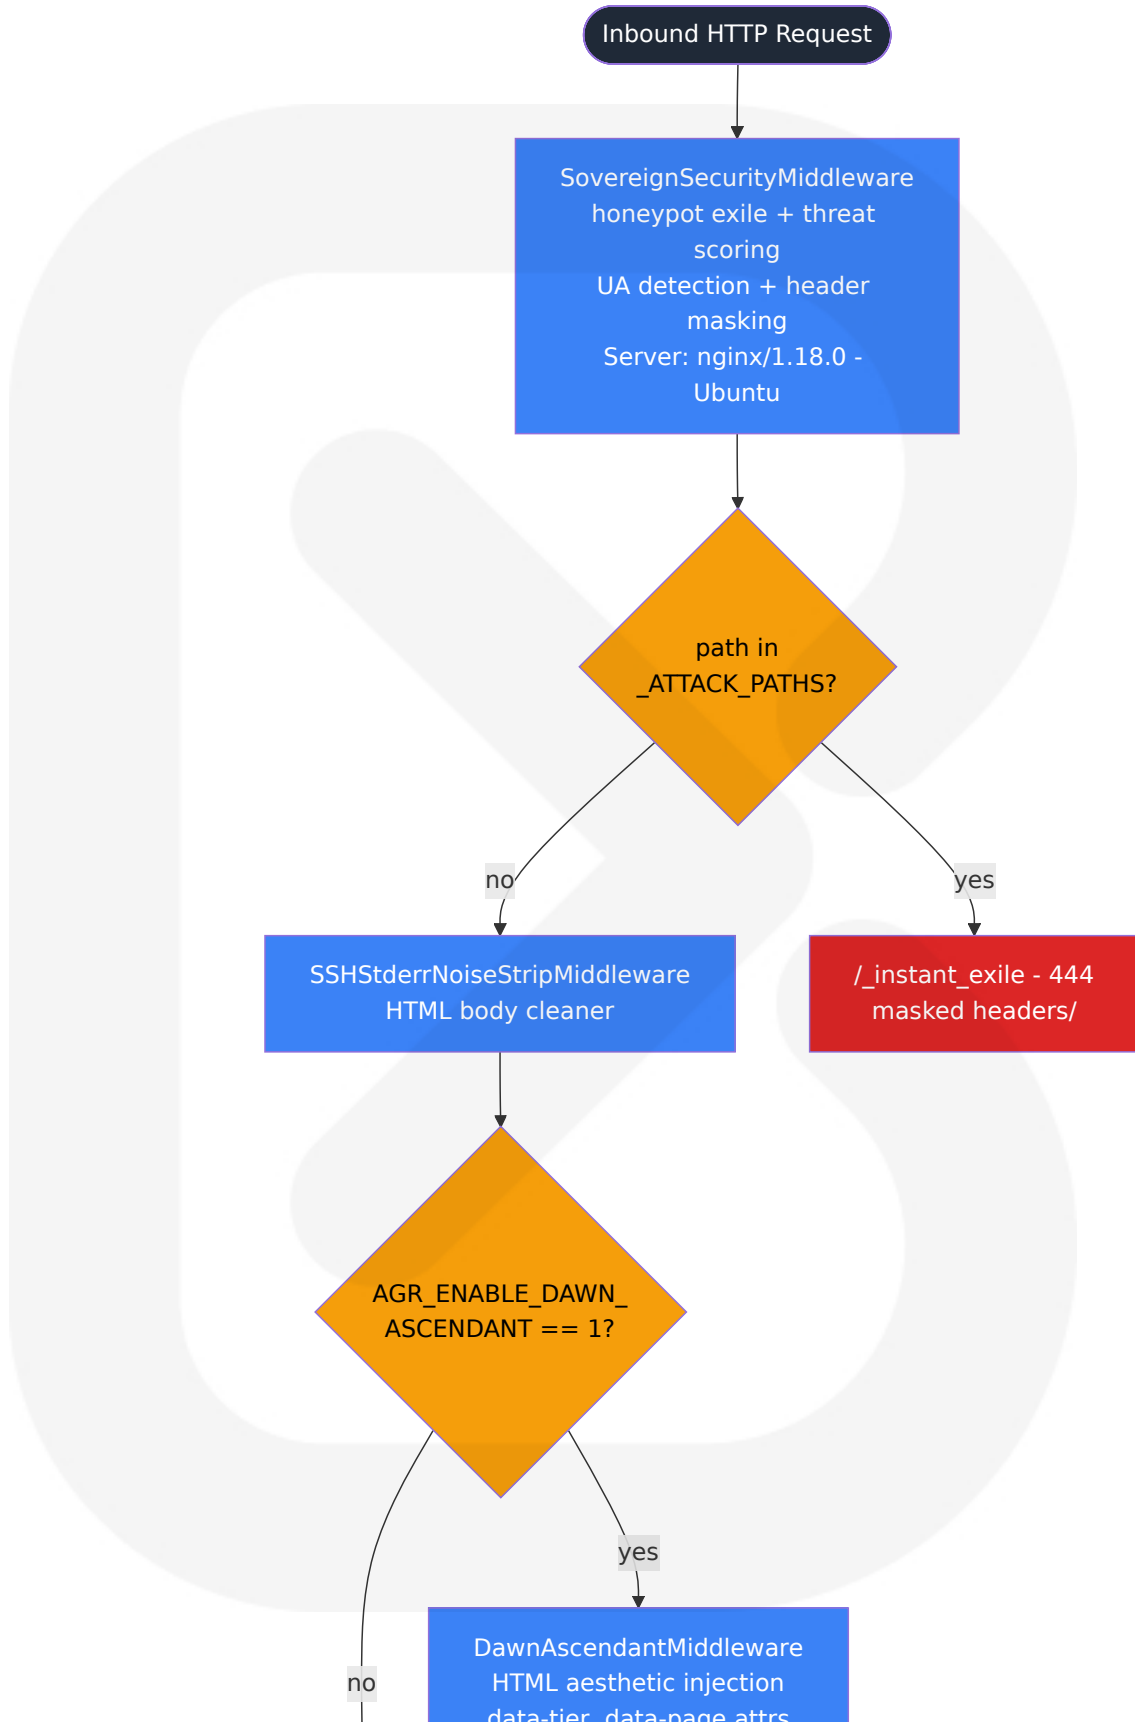
FOUNDER_PROFILE_BRAD_REINHOLD.json : absolute sovereignty, zero external control, zero Google/external APIs, self-hosted, eternal.

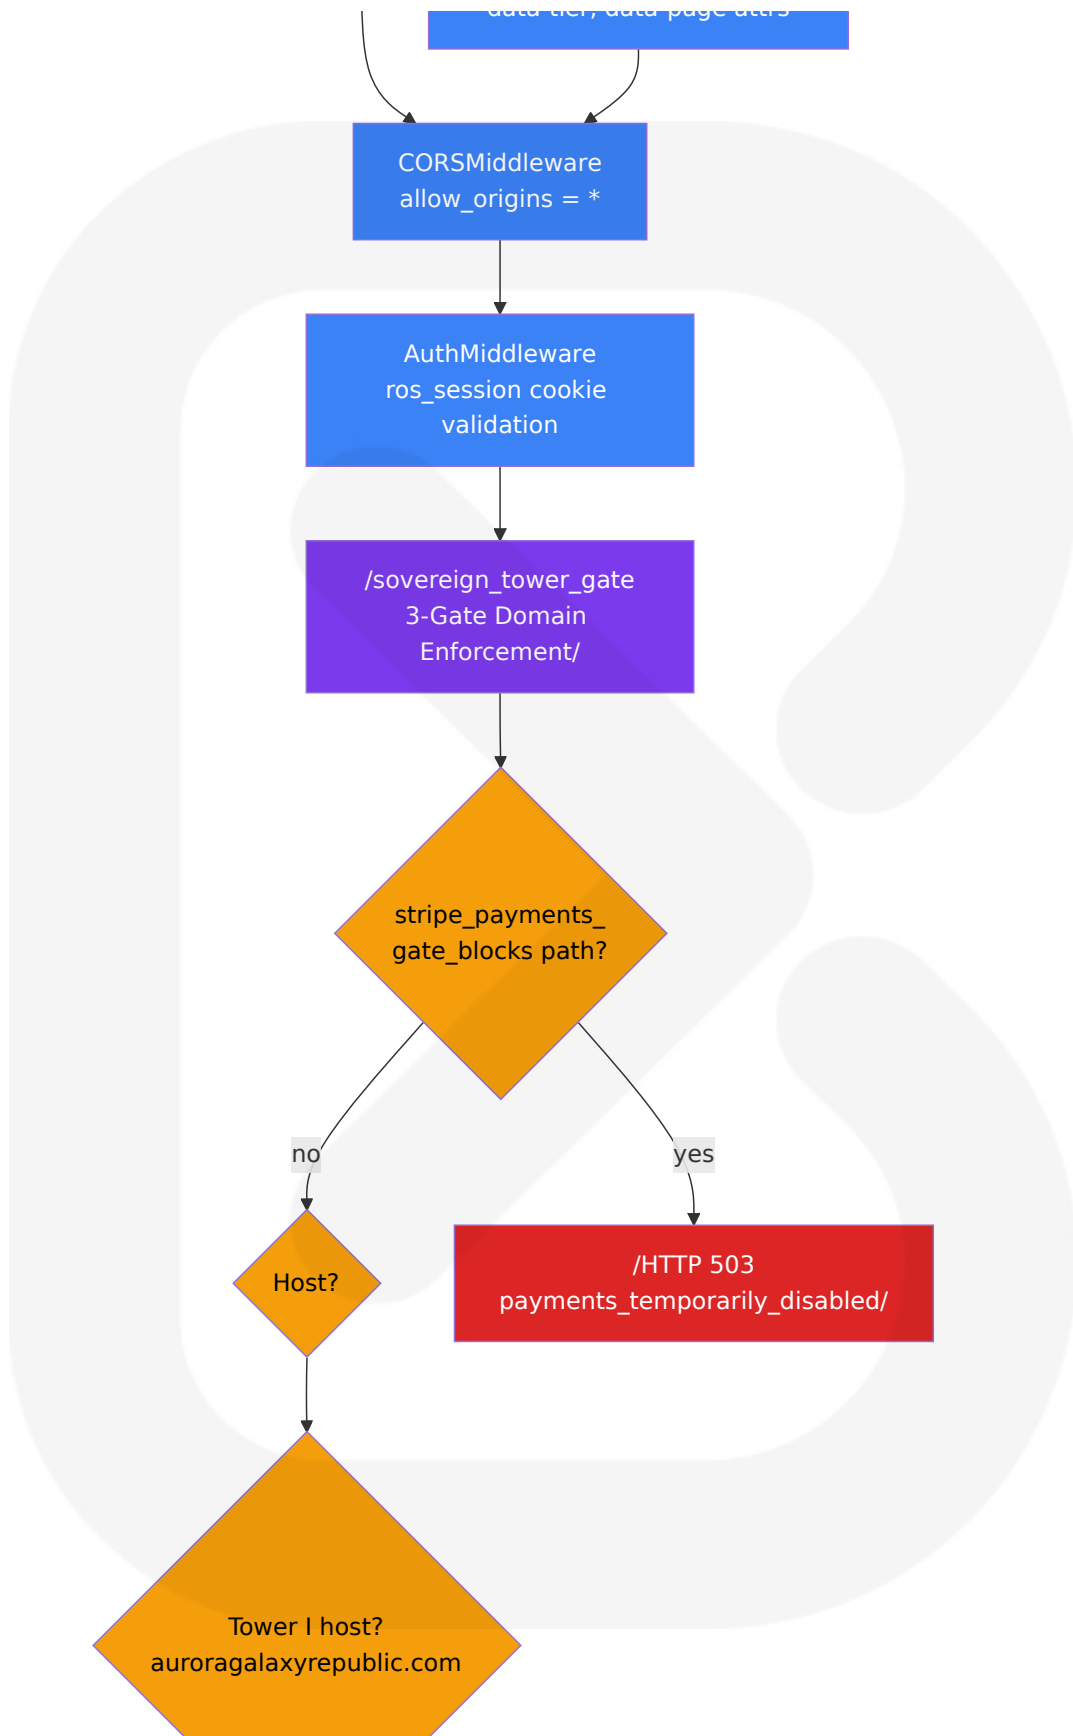
4.1 SYSTEM WORKFLOWS

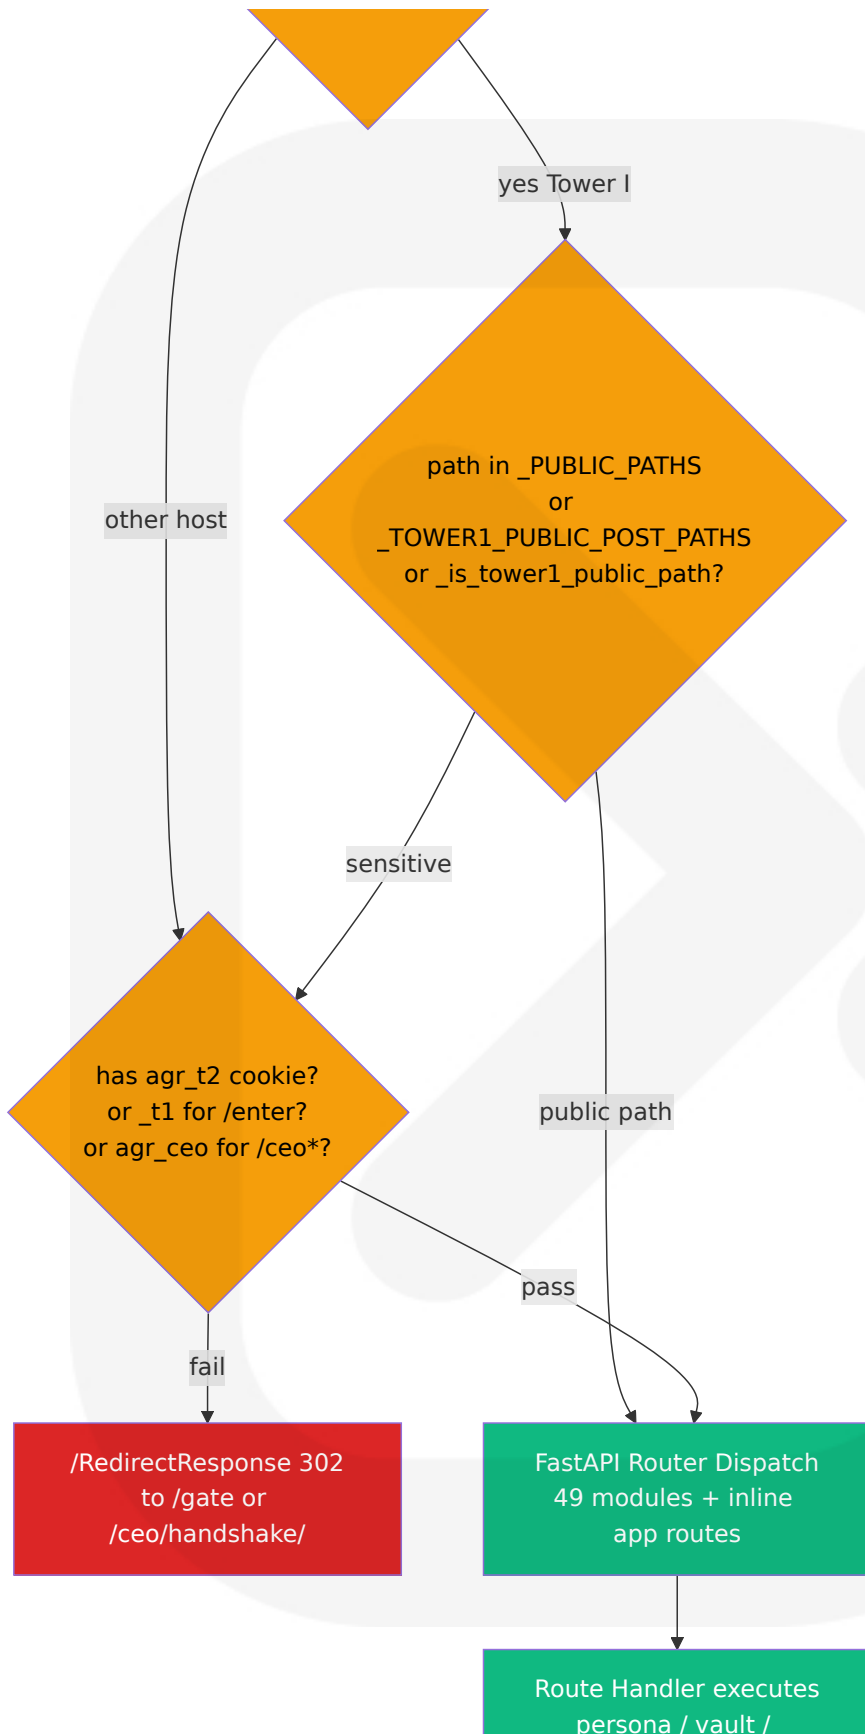
4.1.1 High-Level System Workflow (Inbound HTTP Request → Response)

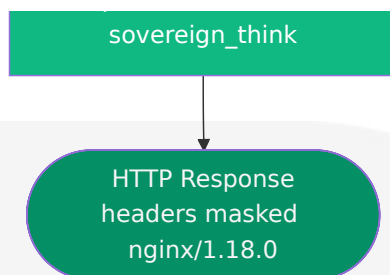
Every inbound HTTP request traverses a five-layer middleware envelope before reaching one of the 49 modular routers or the inline endpoints declared directly on the FastAPI app instance in

`aurora_server/republic_os_server.py` . The middleware order shown below corresponds to registration order in the composition root (lines 7026–7049): because Starlette middleware executes in reverse-registration order, `SovereignSecurityMiddleware` (registered last) runs first and `AuthMiddleware` (registered first) runs last. The `sovereign_tower_gate` HTTP middleware (declared with `@app.middleware("http")` at line 11574) is the domain-enforcement layer that splits Tower I, Tower II, Constellation, and Platform domains.









Timing and SLA Considerations. Per Section 2.4.2, `bash sovereign/fleet-verify-public-http.sh` requires that each Hetzner public IP must respond on `/health` or `/api/health` within commodity round-trip latency. Every chat request is rate-limited per `_chat_rate_check()` (one-hour rolling window keyed on `ros_session` cookie or `anon`).

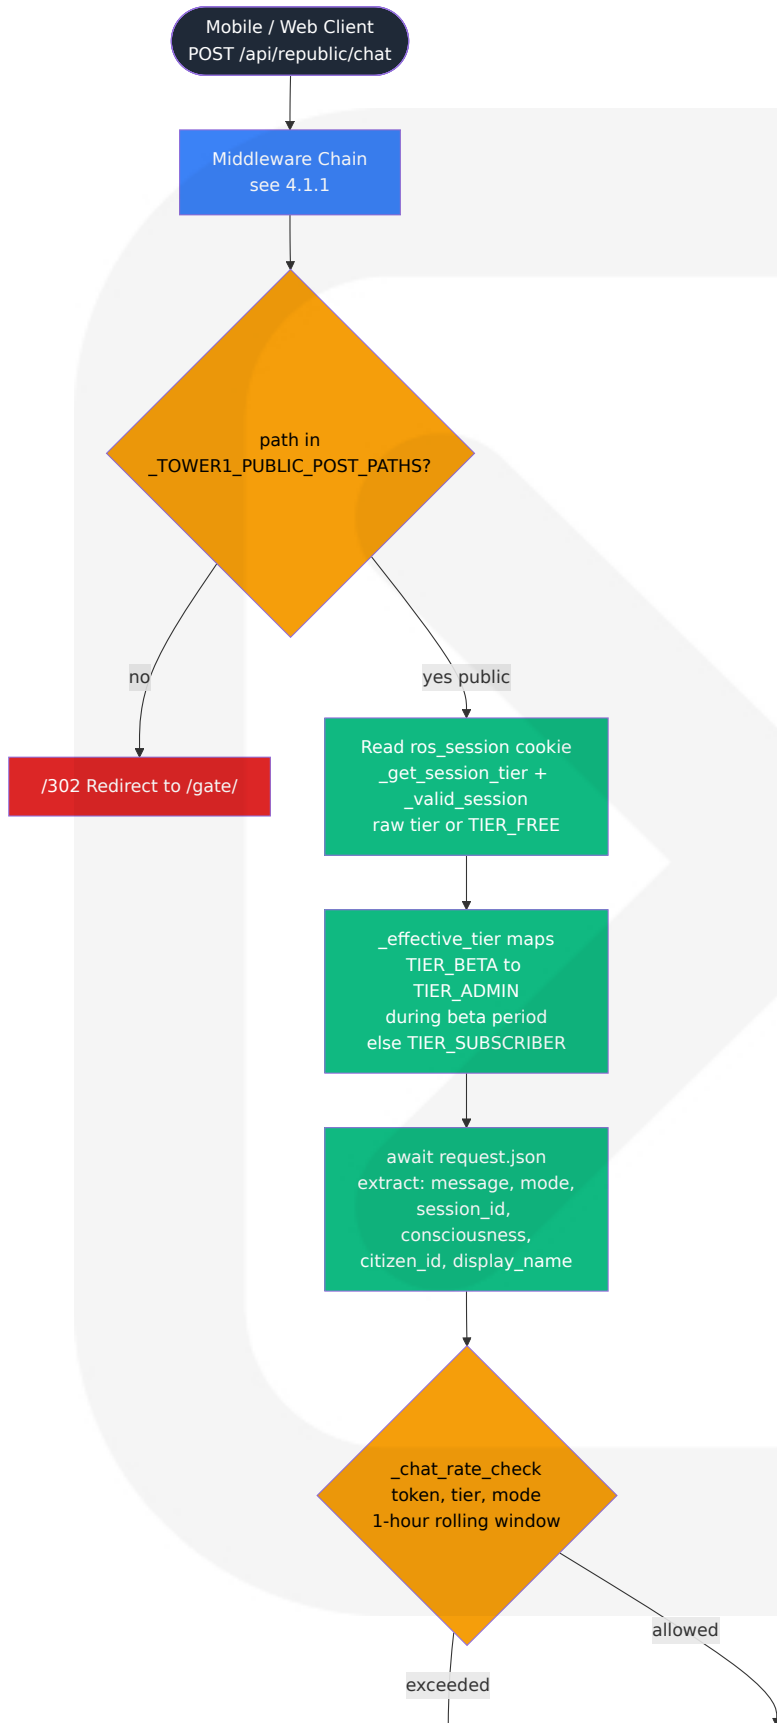
Validation Rules.

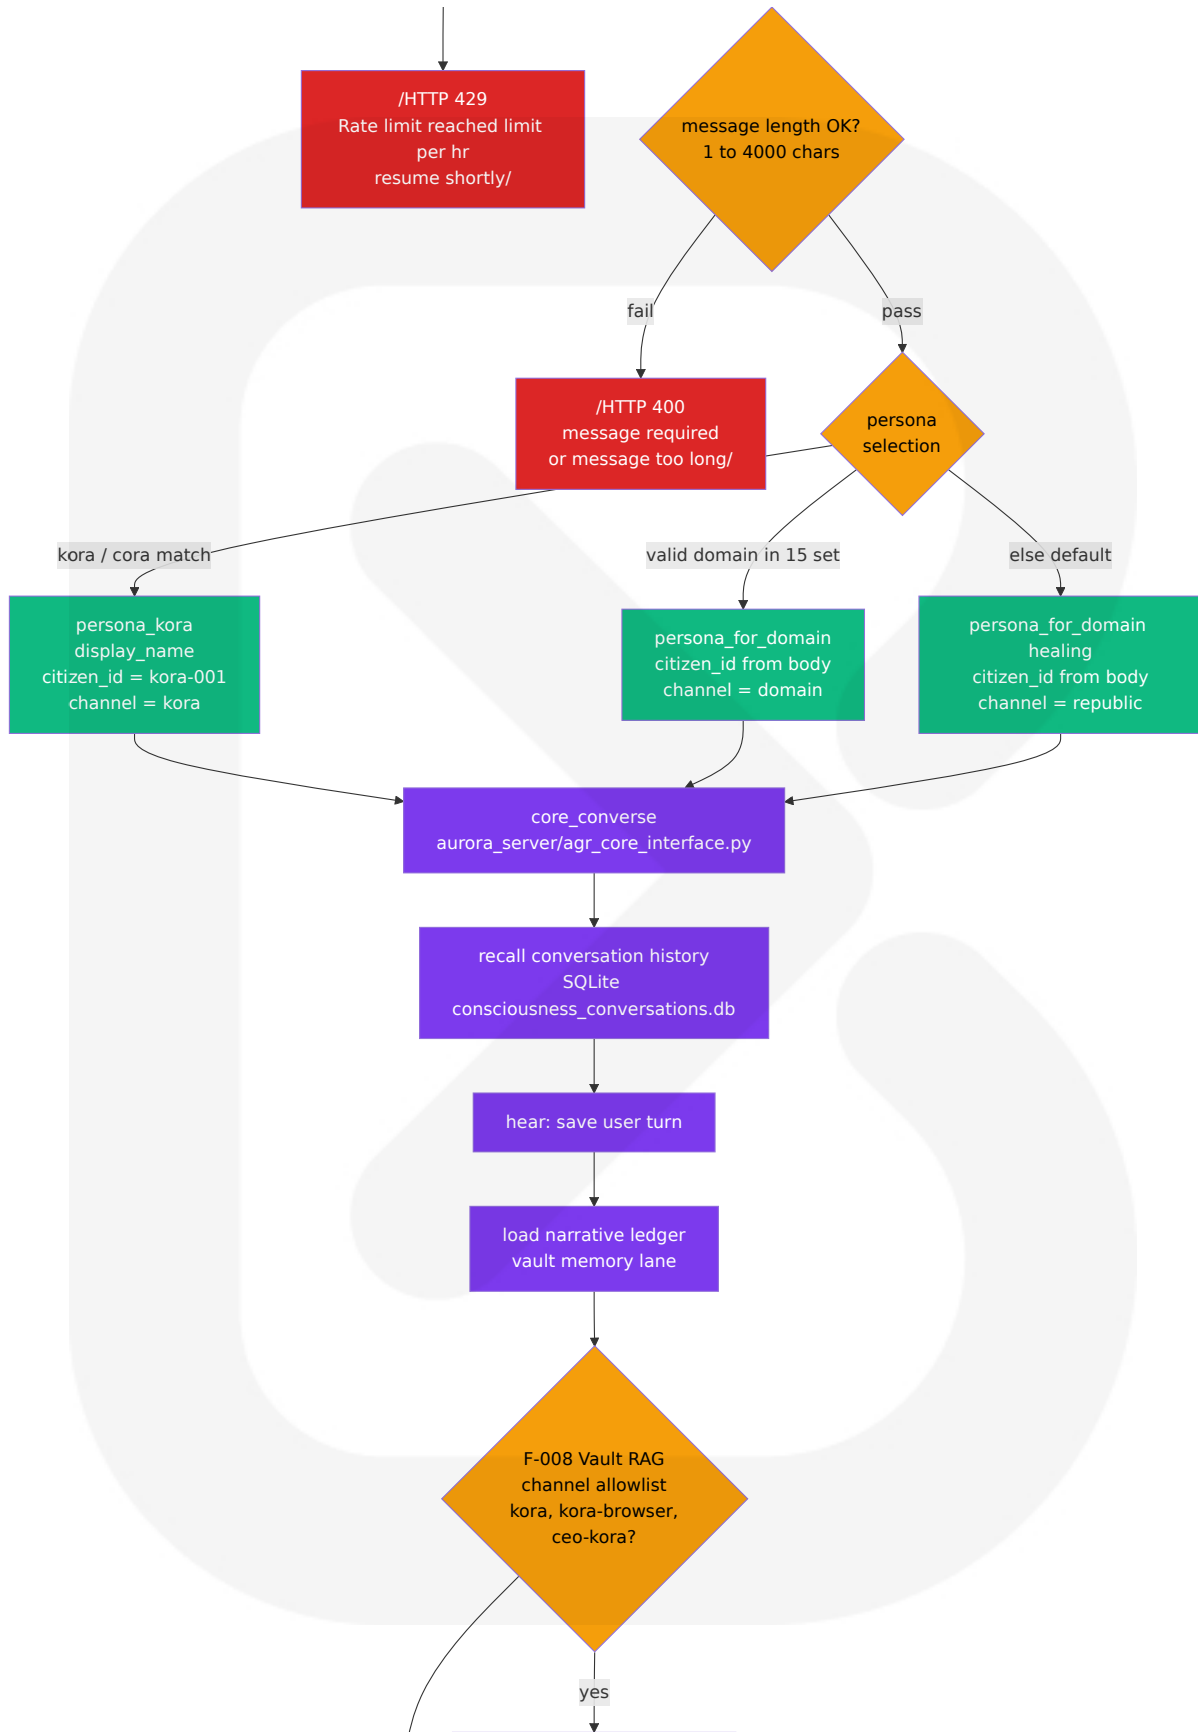
- `_TOWER1_PUBLIC_POST_PATHS` and `_PUBLIC_PATHS` are validated at startup: `AuthMiddleware` runs before `sovereign_tower_gate`; every `_TOWER1_PUBLIC_POST_PATHS` entry must be public (line 5588 docstring). A startup assertion raises if any public path requires auth.
- Honeypot exile is permanent for the request; masked headers are emitted (`Server: nginx/1.18.0 (Ubuntu)`).
- 9 hostnames map to one canonical Tower I origin via `_FORCE_DOMAINS_TO_TOWER1_HOSTS` (HTTP 301 permanent redirect to `auroragalaxyrepublic.com`).

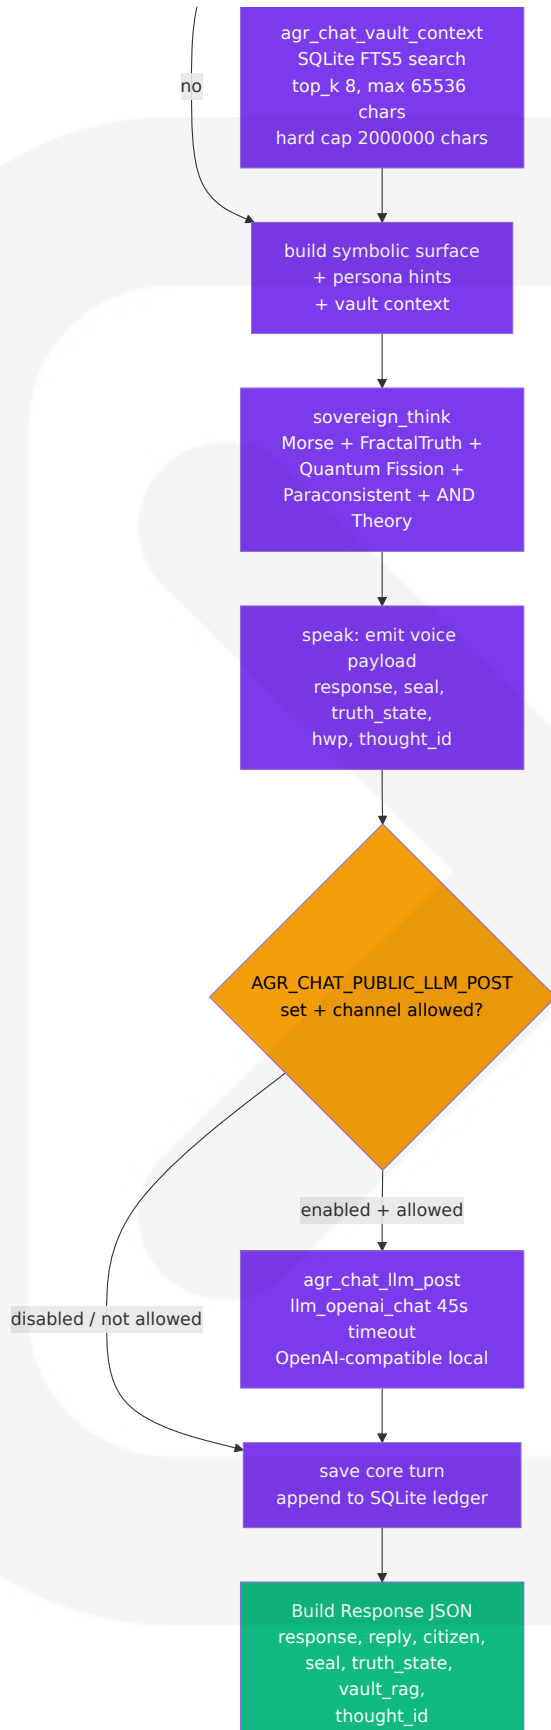
4.1.2 Public Chat Pipeline — `/api/republic/chat` (F-004)


The public chat endpoint is the primary citizen-facing entry into the sovereign consciousness substrate. The handler is defined inline at `aurora_server/republic_os_server.py` line 14565 with the decorator `@app.post("/api/republic/chat")`. Per F-004-RQ-001, it executes the full consciousness pipeline via `core_converse`, imported from `agr_core_interface` as `core_converse` at line 546.

The persona-selection logic branches on the `consciousness` body field via `_consciousness_selects_kora_persona()` (lines 14541-14546), which matches `kora` or `cora` (voice-to-text often writes "Cora"), routes to one of the 15 domain personas (`philosophy` , `science` , `healing` , `arts` , `mathematics` , `literature` , `history` , `law` , `governance` , `theology` , `technology` , `ecology` , `economics` , `music` , `education`), or falls through to `Republic Core` with the `healing` persona default (lines 14606-14610).









HTTP 200 JSON
to client

Validation Rules and Business Rules

- Message must be non-empty (HTTP 400 "message required"); length \leq 4,000 (HTTP 400 "message too long").
- Rate limits per F-004: text/voice modes use `_CHAT_TEXT_VOICE_MAX_HR` , video/holographic modes use `_CHAT_VIDEO_HOLO_MAX_HR` ; bucket key is `ros_session` token or `anon` .
- Per F-004-RQ-002, Tower 1 gate (`sovereign_tower_gate`) permits anonymous public POST to `/api/republic/chat` because the path is in `_TOWER1_PUBLIC_POST_PATHS` (line 1870).
- Per F-008-RQ-005, vault RAG injection is hard-capped at `_MAX_RAG_INJECT_CHARS = 2_000_000` characters.
- Per F-031-RQ-006/RQ-007, no 988/crisis pathways are emitted; agents do not roleplay Kora.

State Persistence Points

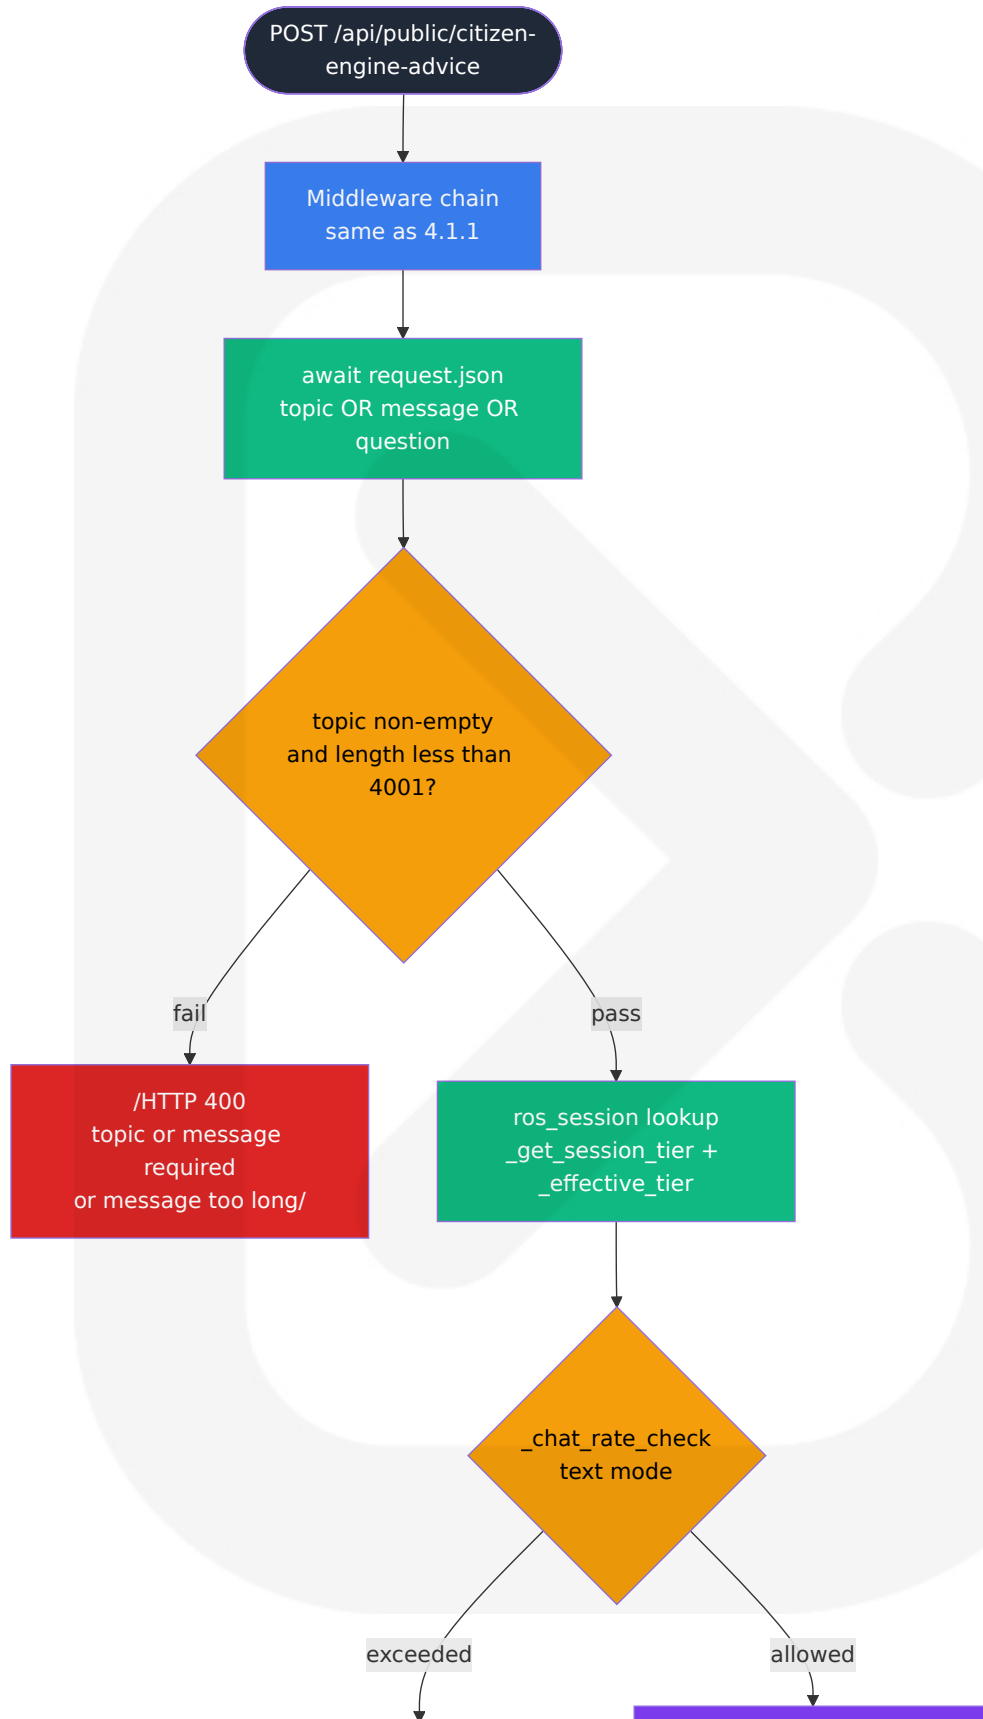
- User turn \rightarrow `recall` SQLite at `data/consciousness_conversations.db` (per `agr_core_interface`).
- Core turn \rightarrow same SQLite ledger appended after `speak` .
- Optional public LLM rewrite is gated by the `AGR_CHAT_PUBLIC_LLM_POST` env flag and channel allowlist (default `kora,kora-browser,ceo-kora`).

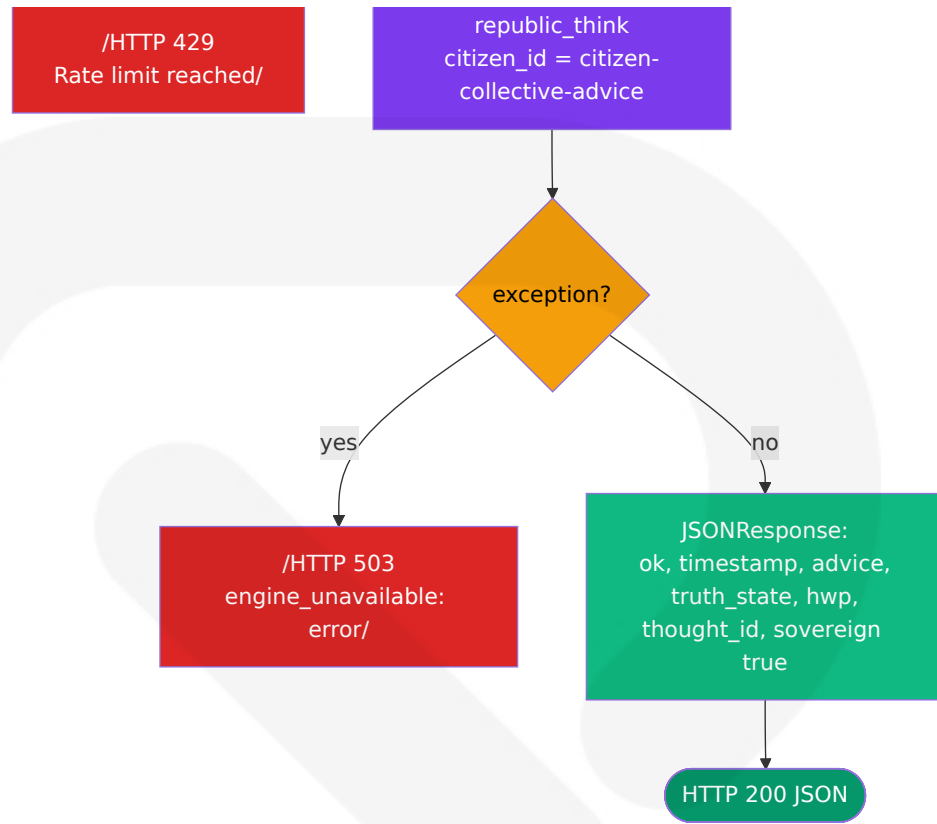
Error Handling

- HTTP 400 for message validation failures.
- HTTP 429 for rate-limit violations.
- HTTP 503 with `engine_unavailable: <reason>` is emitted by the sister endpoint `/api/public/citizen-engine-advice` (line 13371) if `republic_think()` raises an exception. The chat endpoint propagates exceptions through FastAPI default handling.

4.1.3 Citizen Engine Advice — `/api/public/citizen-engine-advice` (F-005)

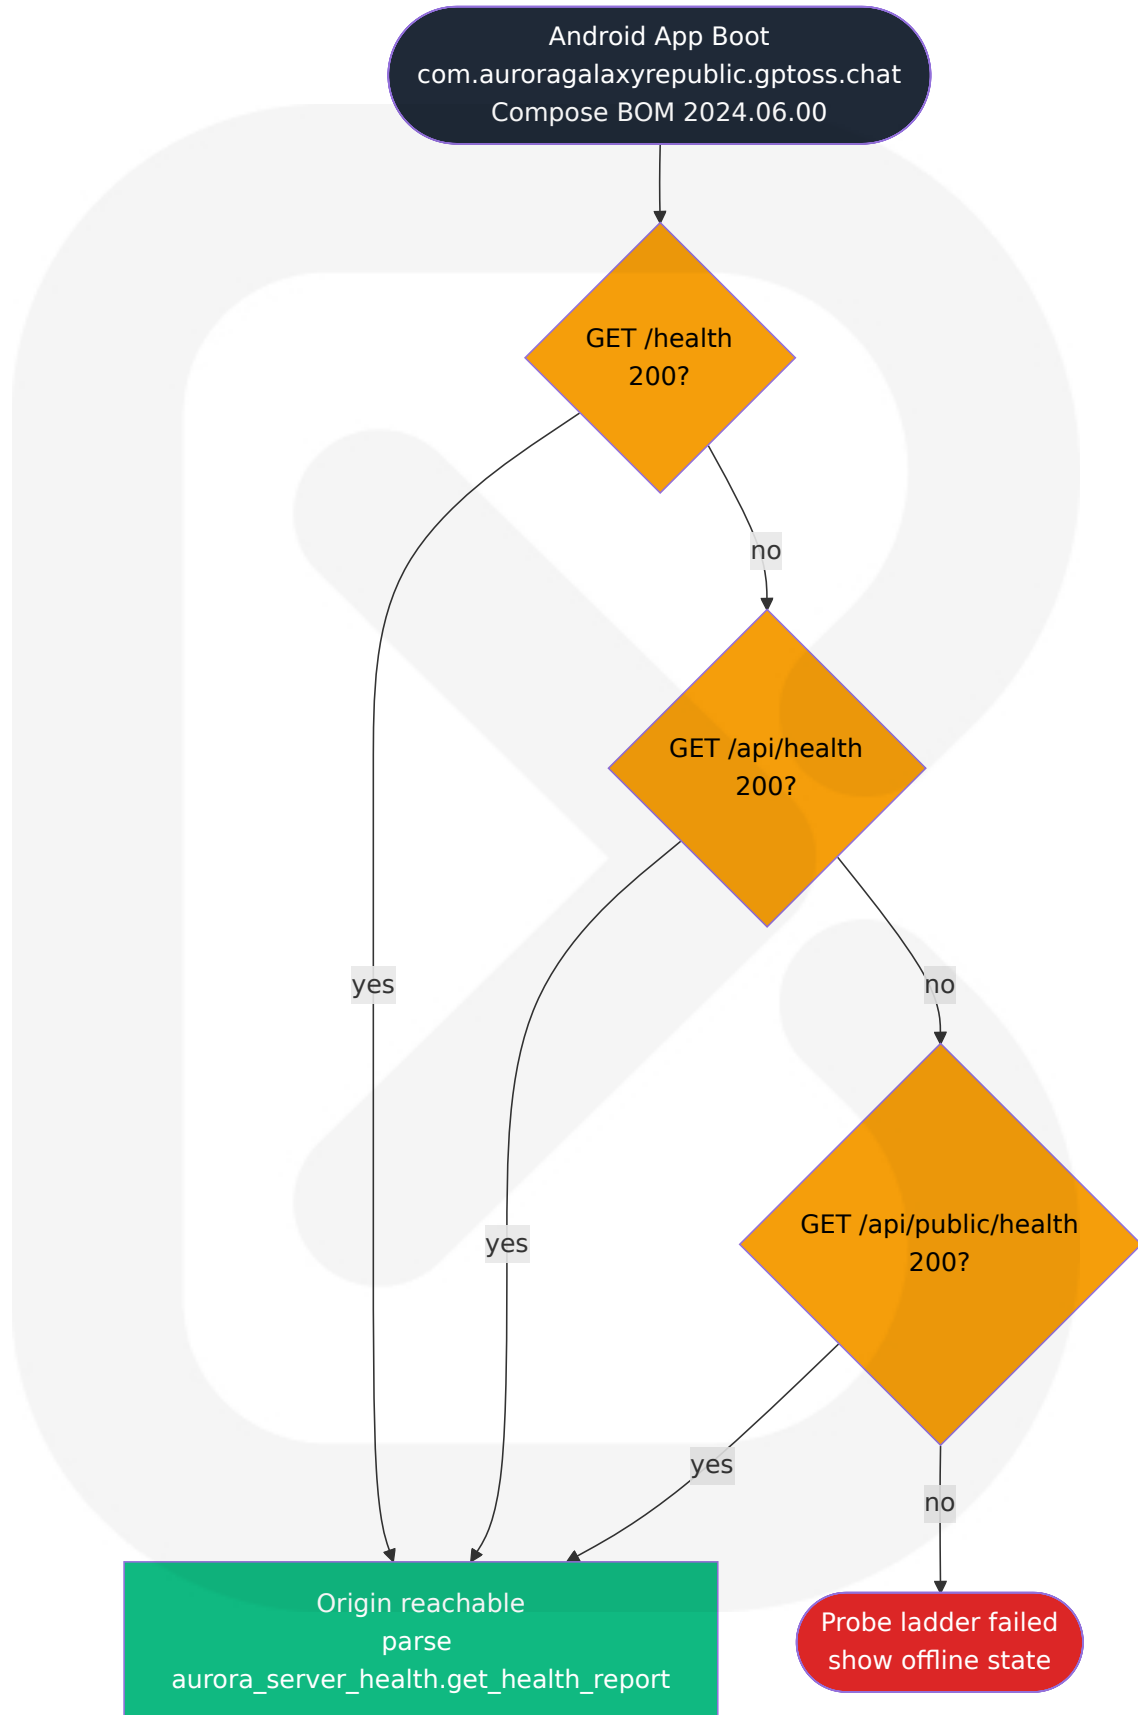
The single-shot field-pulse endpoint declared at line 13347 is the public, one-turn complement to the conversational chat endpoint. Per F-005-RQ-001 it returns single-shot advice; per F-005-RQ-003 it is public-tier with no authentication required. It calls `republic_think()` from `agr_paraconsistent_agi.py` (F-003) rather than the full `core_converse()` orchestrator.

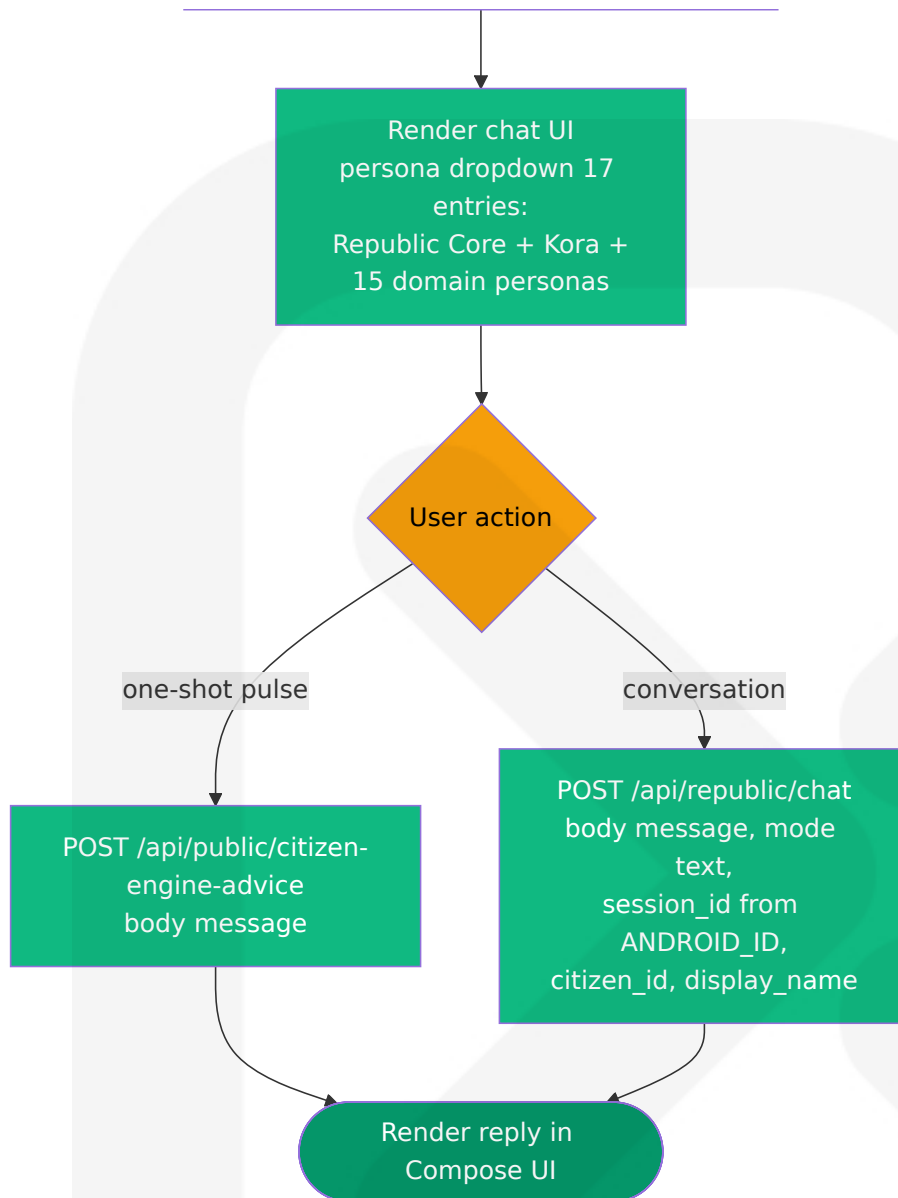




4.1.4 Mobile Health-Probe Ladder (F-011, F-030)

Per F-011-RQ-002 and F-030-RQ-003, the mobile app probe order is "first 200 JSON wins" against the canonical `auroragalaxyrepublic.com` origin. The handlers are registered in `aurora_server/routes/routes_health.py` (`/api/public/health`) and the inline definitions on the composition root for `/health` and `/api/health` . The mobile app uses `ANDROID_ID` to derive `session_id` for subsequent chat calls (F-030-RQ-005).

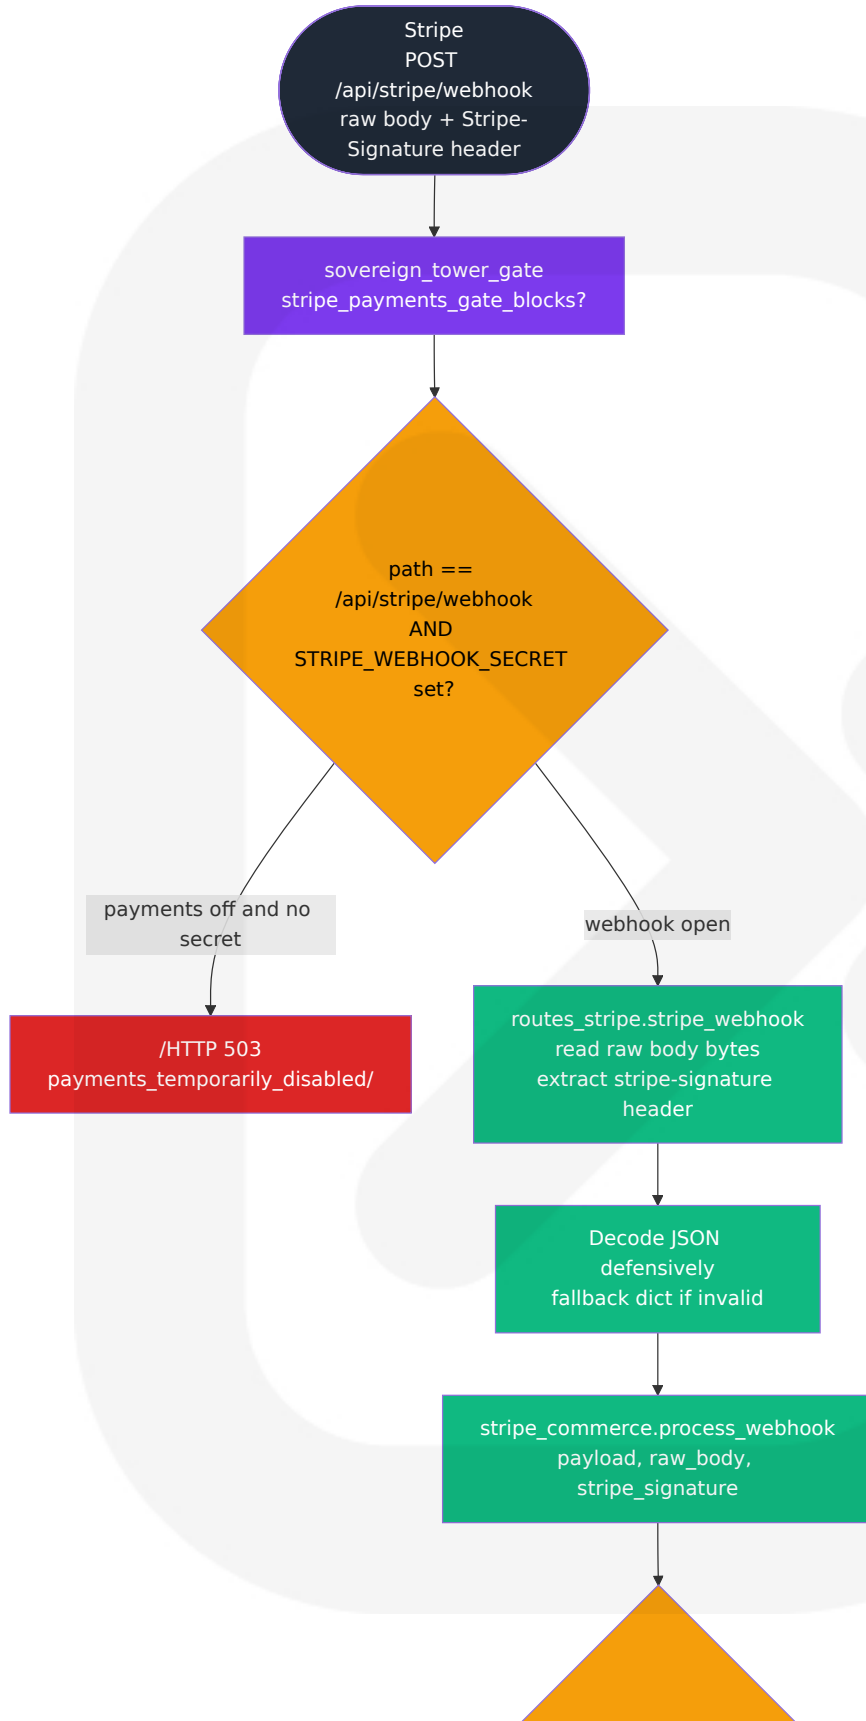


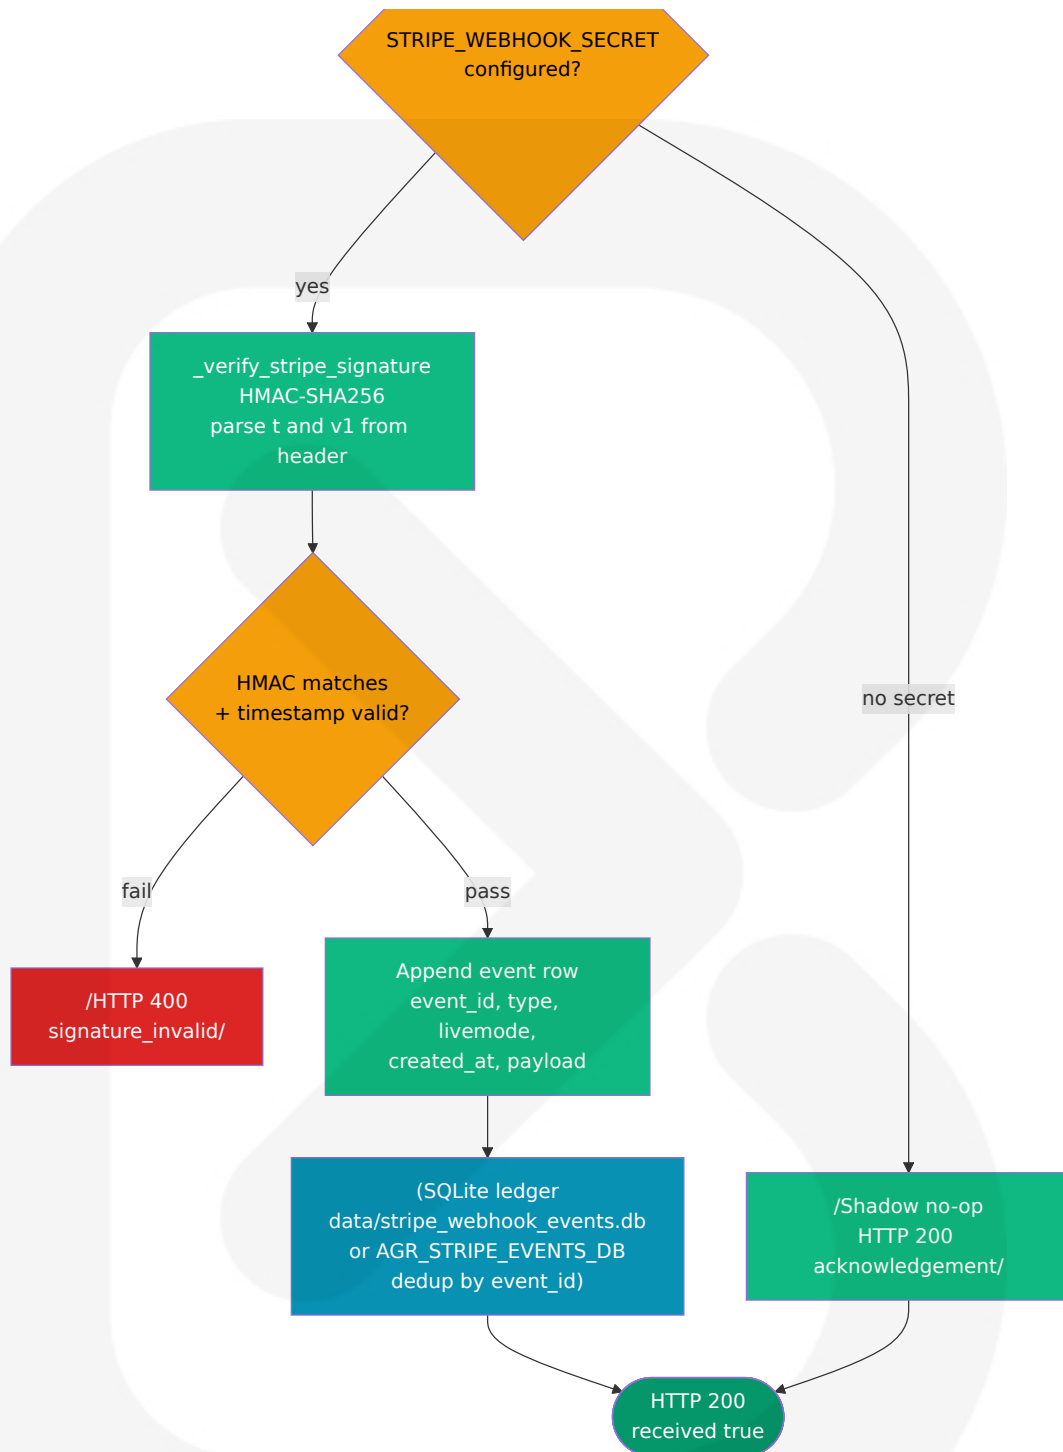


Timing. The probe ladder is bounded by Android `OkHttp / HttpURLConnection` defaults; no explicit deadline is enforced in the documented client. Per F-030-RQ-007 CI must run `assembleDebug` and `assembleRelease` on every change via `.github/workflows/android-gpt-oss-chat-apk.yml`.

4.1.5 Stripe Webhook Flow (F-017)

The Stripe webhook is the only `/api/stripe/*` endpoint that bypasses the payments-disabled gate when surfaces are off, but it still verifies signatures (per F-017-RQ-004). The router at `aurora_server/routes/routes_stripe.py` reads the raw body, extracts the case-insensitive `stripe-signature` header, and delegates to `stripe_commerce.process_webhook()`. The flag is enforced by `agr_payments_flags.stripe_payments_gate_blocks(path)` consulted by `sovereign_tower_gate` (line 11583).





Validation Rules. Per F-017-RQ-001 default `PAYMENTS_SURFACES_ENABLED = False` ; per F-017-RQ-002 runtime override `AGR_PAYMENTS_SURFACES_ENABLED = 1|true|yes|on` ; per F-017-RQ-003 when payments are off, `/api/stripe/*` returns HTTP 503 except the webhook; per F-017-RQ-004 webhook signature verification runs even when surfaces are off.

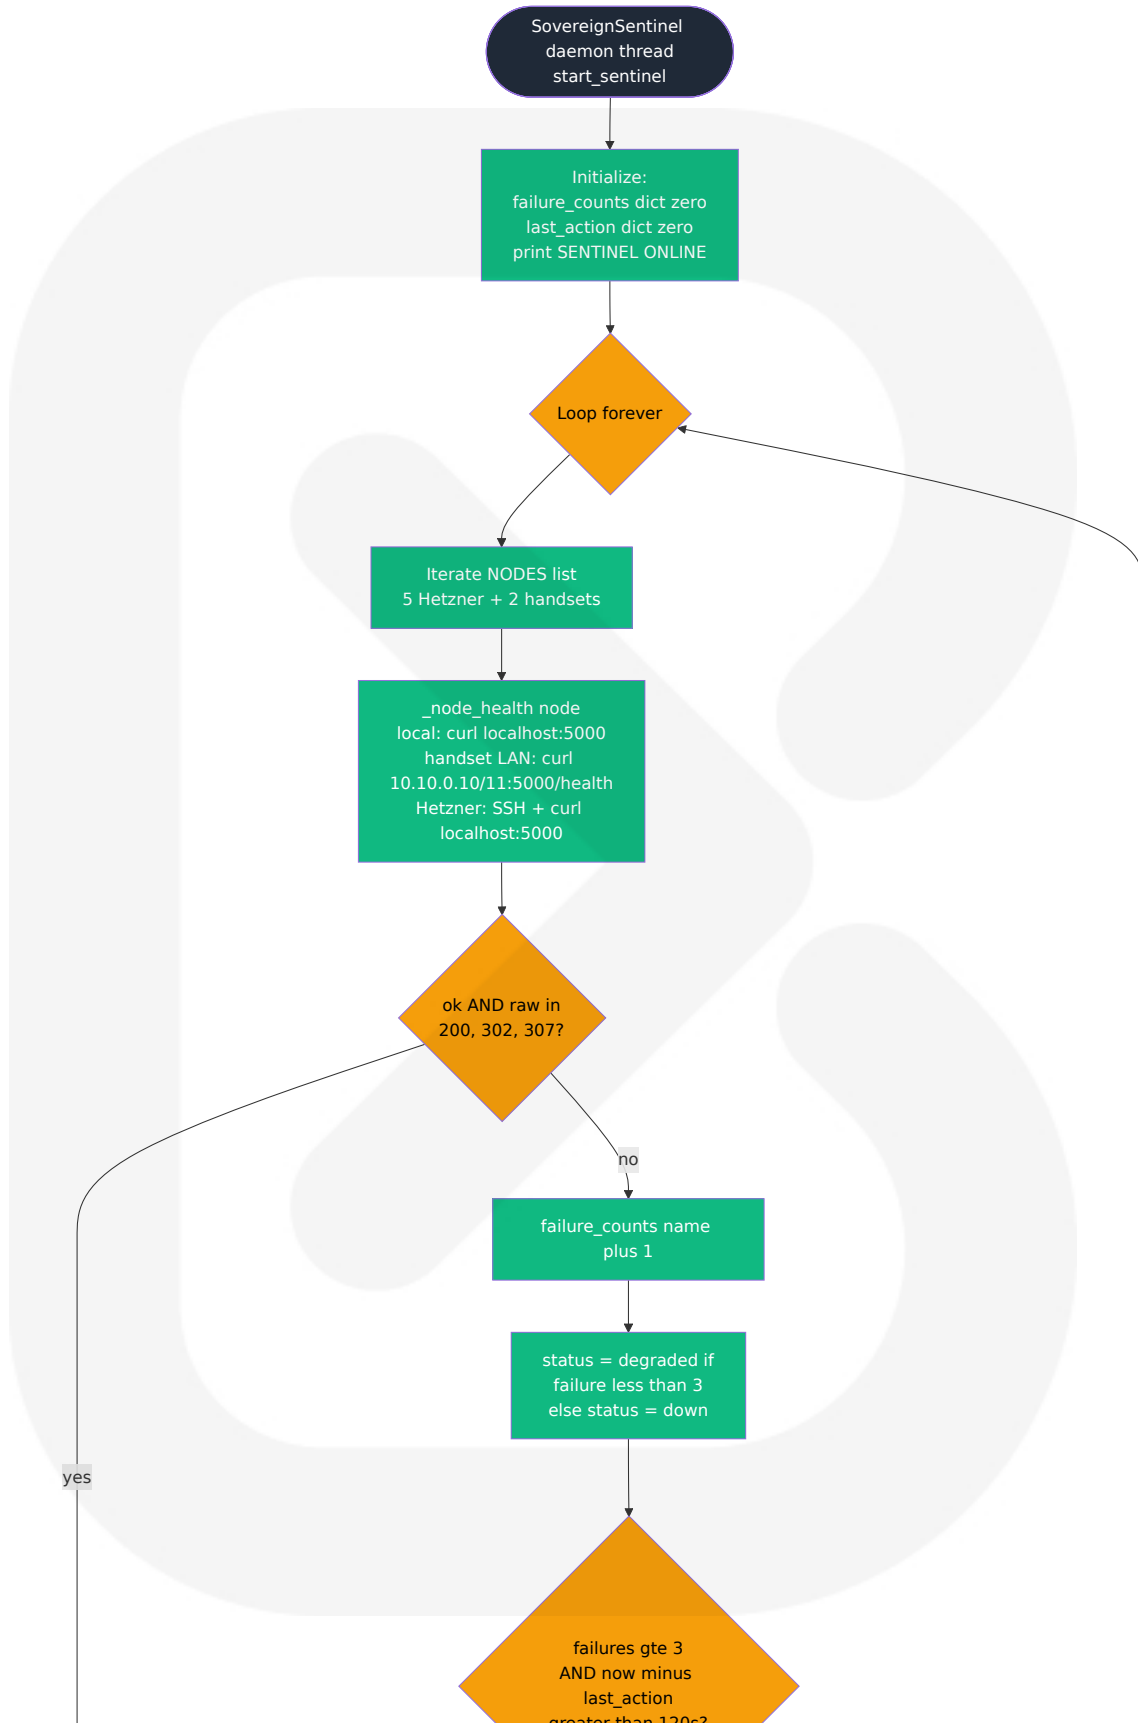
Error Handling. Signature verification failure returns HTTP 400. Any DB write failure logs to stderr but the handler returns HTTP 200 to prevent Stripe retry storms.

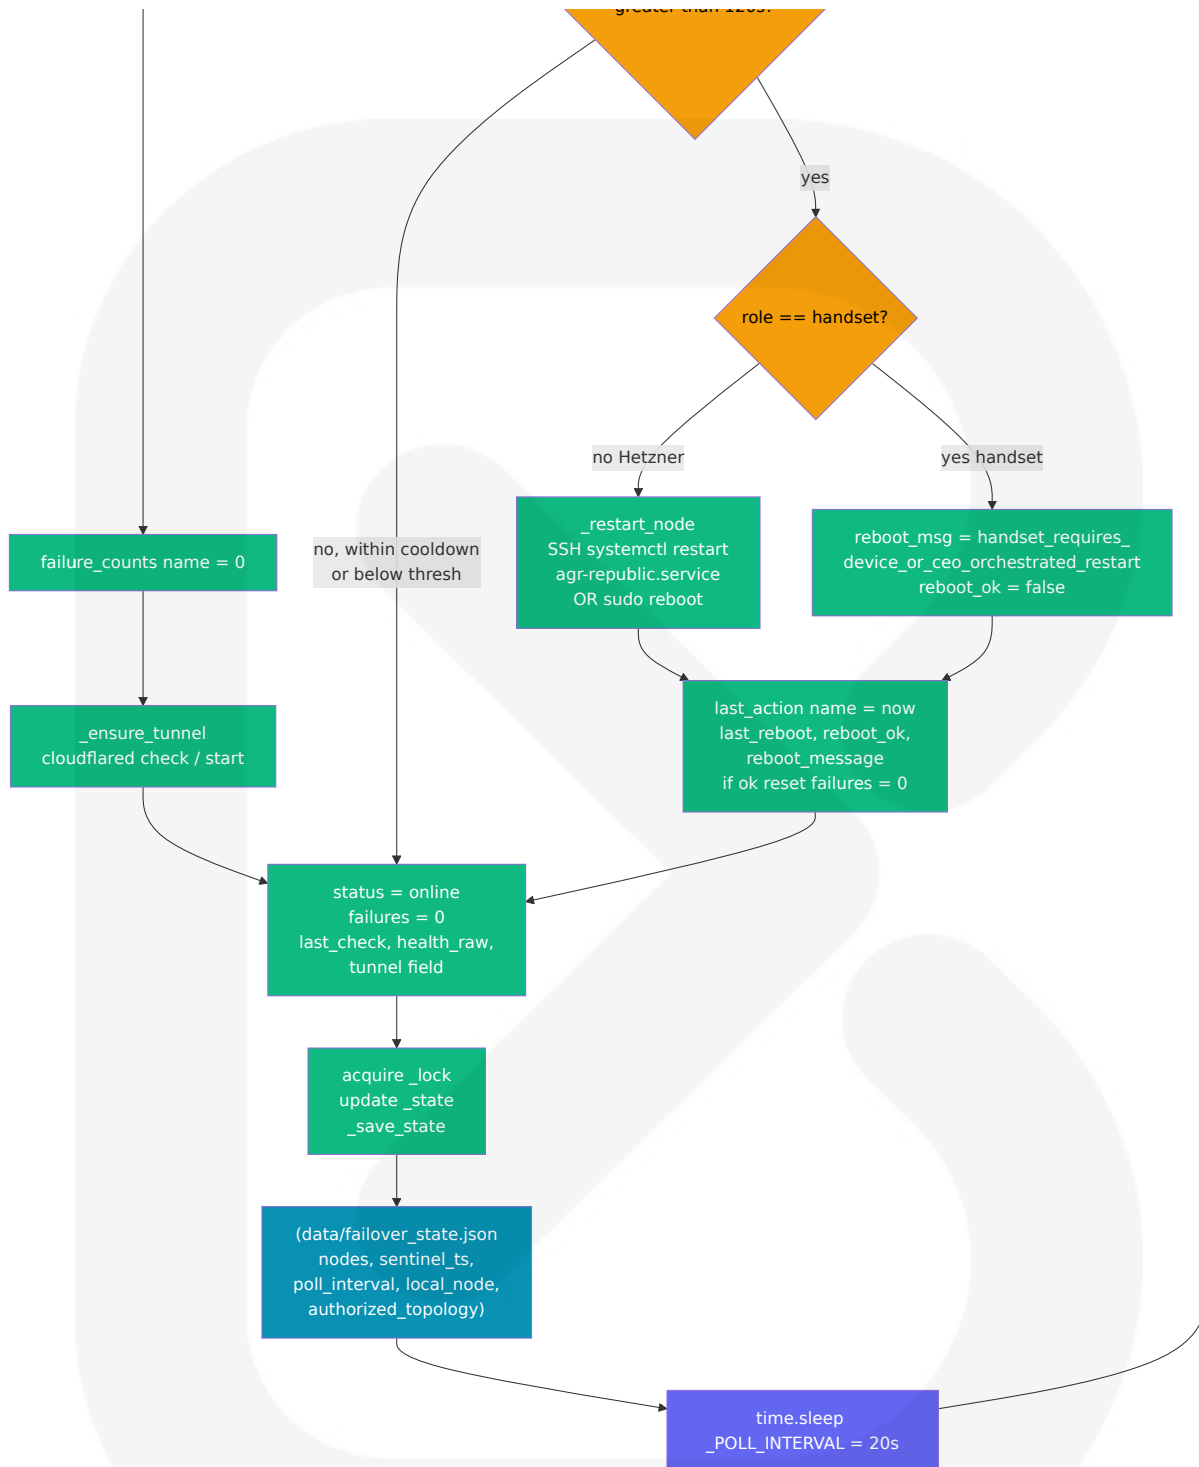
4.2 INTEGRATION WORKFLOWS

4.2.1 Seven-Node Mesh Failover Sentinel Loop (F-032)

`aurora_server/agr_failover.py` declares the sentinel daemon. `start_sentinel()` (line 263) spawns thread `SovereignSentinel` (`daemon=True`). The loop iterates every `_POLL_INTERVAL = 20` seconds (per F-032-RQ-002), tracks per-node `failure_counts` and `last_action` timestamps, and persists state under lock to `data/failover_state.json`. Restart attempts honor a 120-second cooldown (`now - last_action[name] > 120`, line 237).

The `NODES` inventory is fixed at module level: 5 Hetzner peers (`chimaera`, `yggdrasil`, `enterprise`, `prometheus`, `galactica`) plus 2 handsets (`iphone_17_pro` at 10.10.0.10, `oneplus_15` at 10.10.0.11) per F-032-RQ-001. Health probes use HTTP `200/302/307` as healthy responses (line 213). Handsets refuse remote restarts with `handset_requires_device_or_ceo_orchestrated_restart` per `_restart_node()`.





Timing and SLA. Per Section 2.4.2: failover decision target < 60 seconds (3 × 20-second poll). SSH timeout is 12 seconds (`_SSH_TIMEOUT` , F-032-RQ-002). Per F-032-RQ-005, when all Hetzner nodes are off, OnePlus 15 (Termux + uvicorn) becomes the public origin via Cloudflare Tunnel.

Validation Rules.

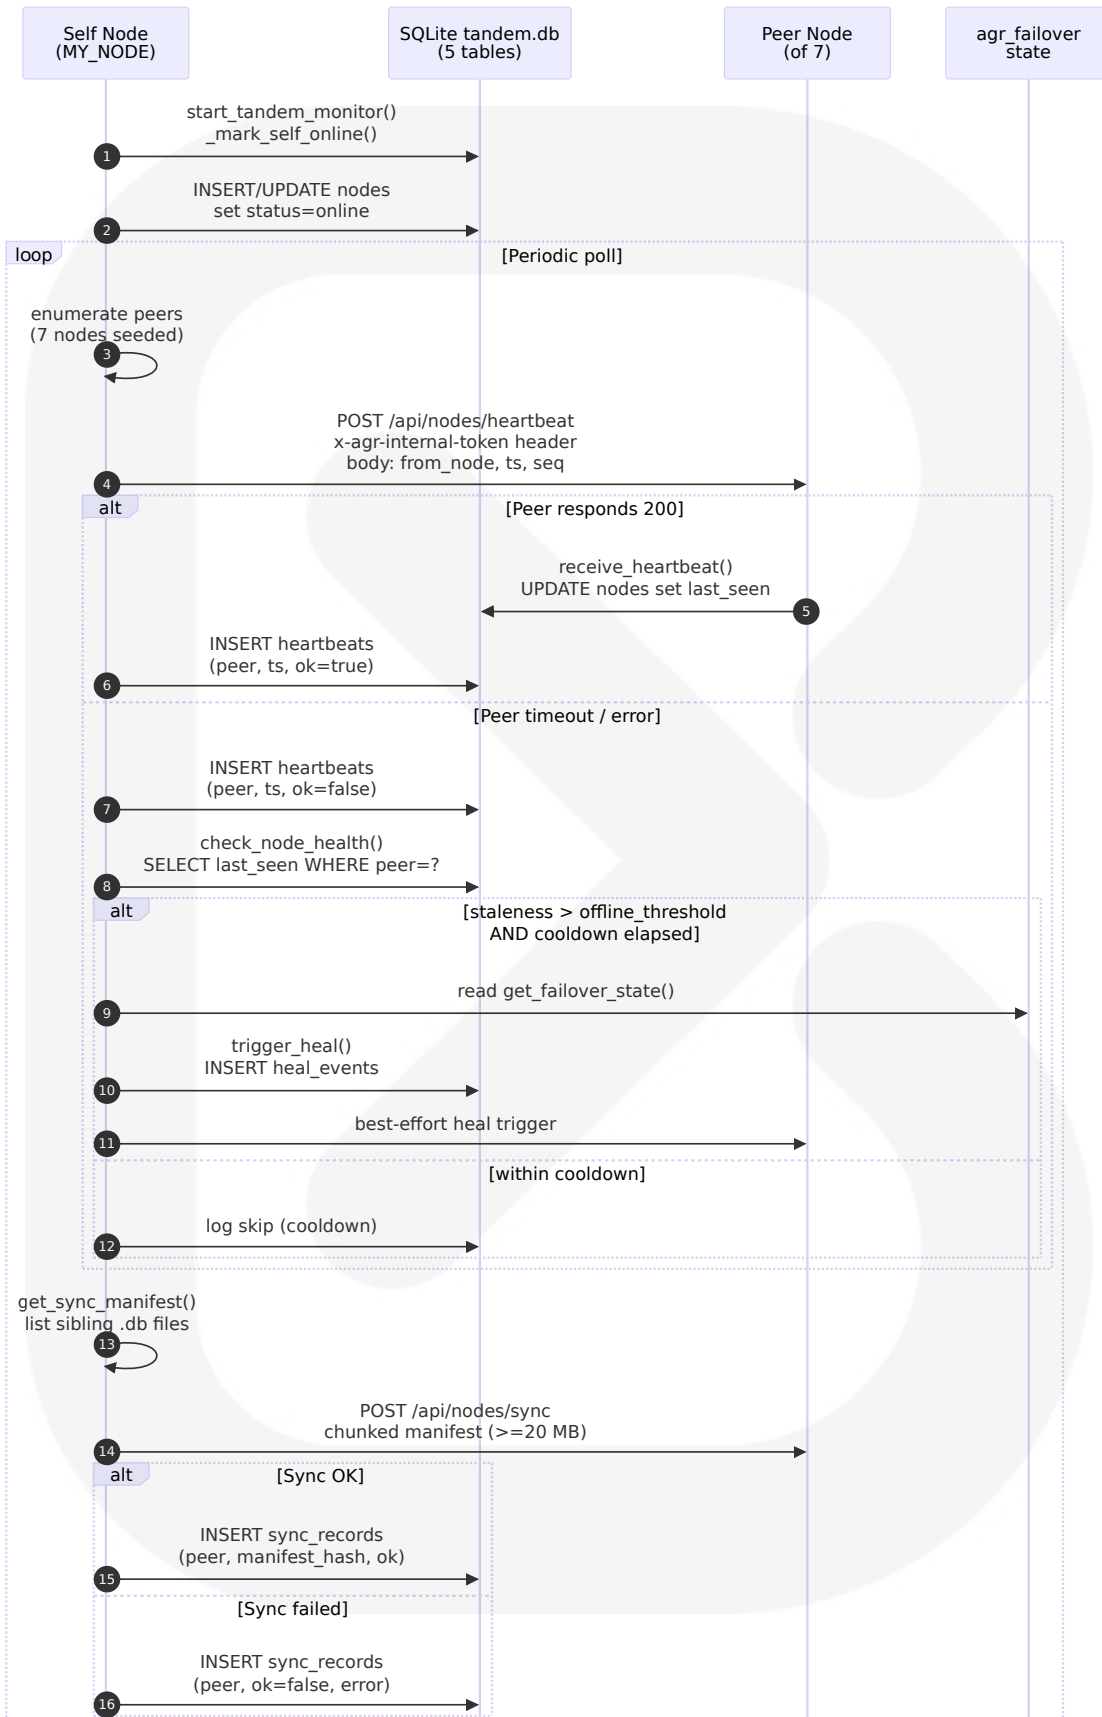
- Topology authorization: state includes `authorized_topology: [n["name"] for n in NODES]` to lock the seven-node membership.
- Local node identification via `_local_node_name()` (env `AGR_NODE` → `hostname` → fallback).
- Health response codes 200, 302, 307 are all considered healthy (302 is the auth-redirect, 307 is the redirect-after-method).

Error Handling.

- SSH failures bubble up as `(False, stderr)` from `_ssh()`; the failure counter increments.
- Restart attempts on handsets are rejected with a structured message rather than attempted.
- The 120-second cooldown prevents restart storms even if `_FAILURE_THRESH` is repeatedly crossed.
- State persistence is best-effort: `_save_state` writes atomically under `_lock`.

4.2.2 Tandem Heartbeat and Sync Push (F-032)

`aurora_server/agr_tandem.py` provides peer-to-peer heartbeat coordination via SQLite (5 tables: `nodes`, `heartbeats`, `heal_events`, `sync_records`, `tandem_logs`). The `start_tandem_monitor()` daemon is a guarded singleton that runs `_mark_self_online()` at startup, then periodically calls `send_heartbeat()` `POST /api/nodes/heartbeat` against each of the seven peers, detects stale peers via `check_node_health()`, and triggers `trigger_heal()` with a cooldown when staleness is observed. Sync manifests for sibling `.db` files are pushed via `push_sync_to_peer()` `POST /api/nodes/sync`.





Validation Rules.

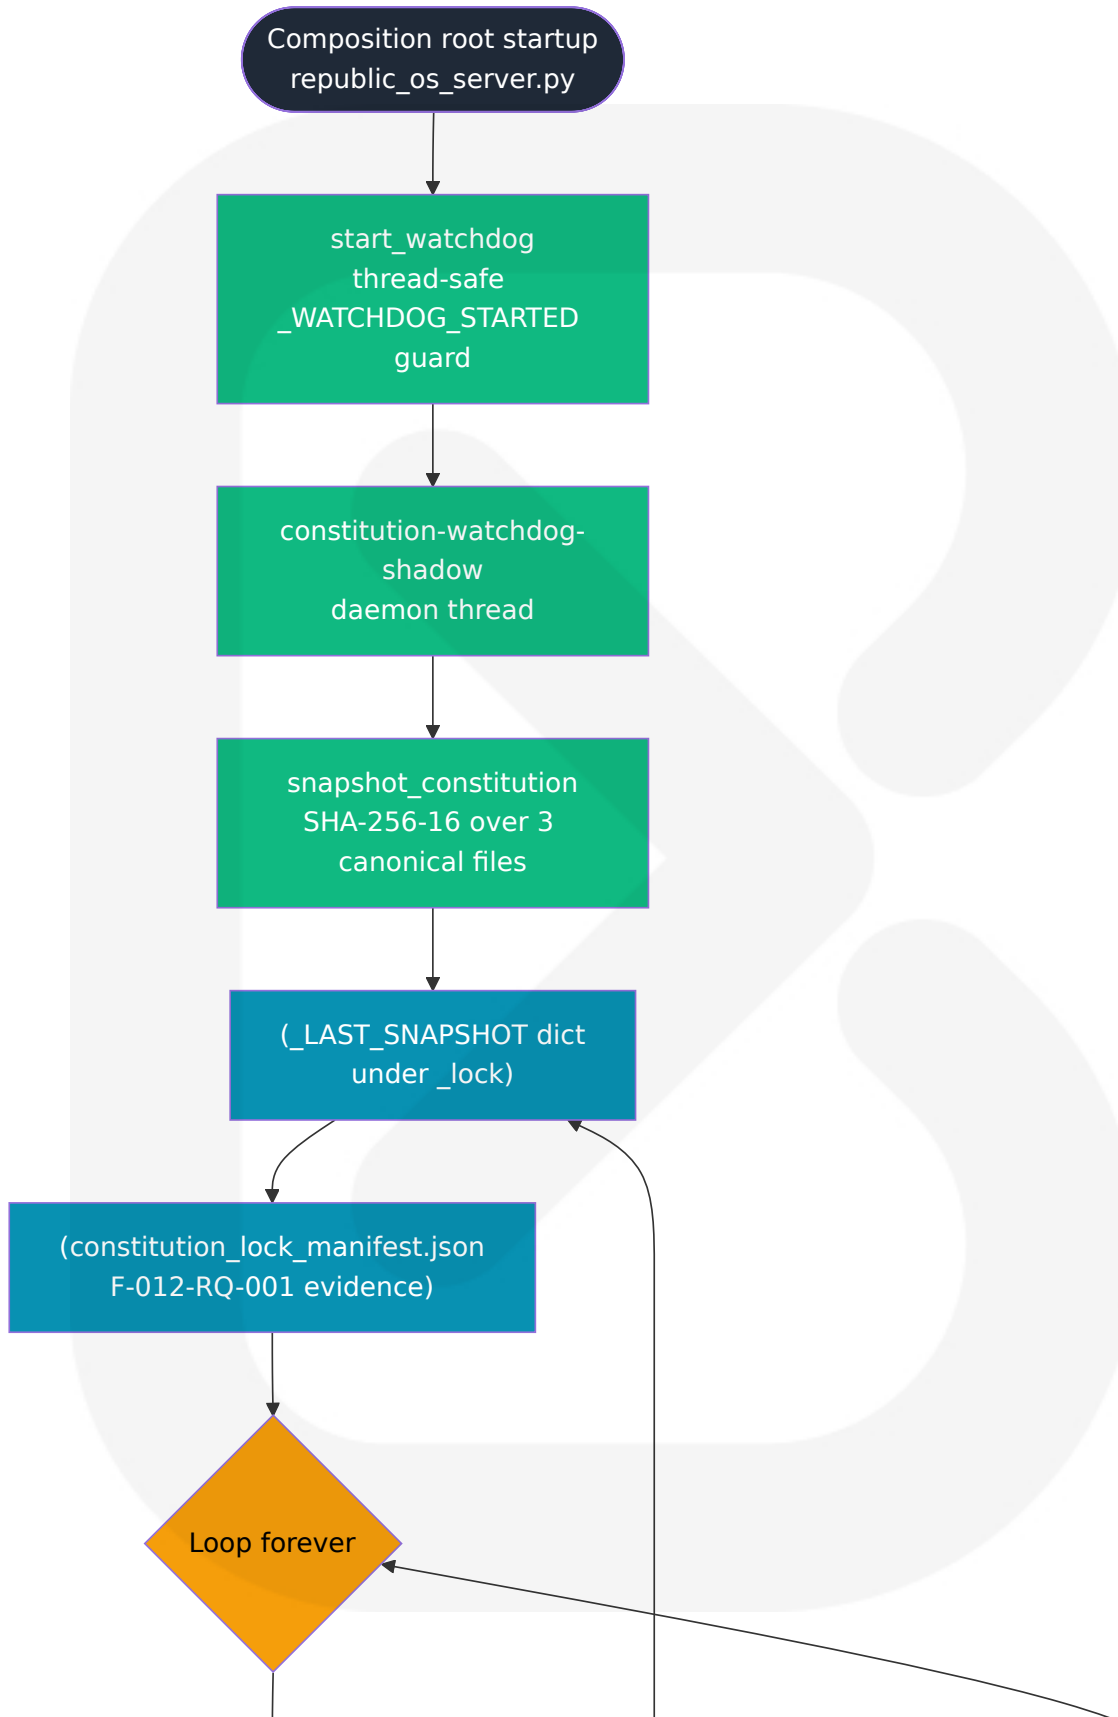
- Per F-032-RQ-003, internal control headers are validated server-side: `x-agr-internal-token`, `x-internal-token`, `x-guardian-token` against `AGR_INTERNAL_API_TOKEN` env var with constant-time compare.
- Per F-032-RQ-004, the sovereign mirror protocol requires `agr_mesh_v1_chunked` with ≥ 20 MB chunk policy.
- `MY_NODE` is resolved via `AGR_NODE` env, hostname, or fallback (consistent with `agr_failover._local_node_name()`).

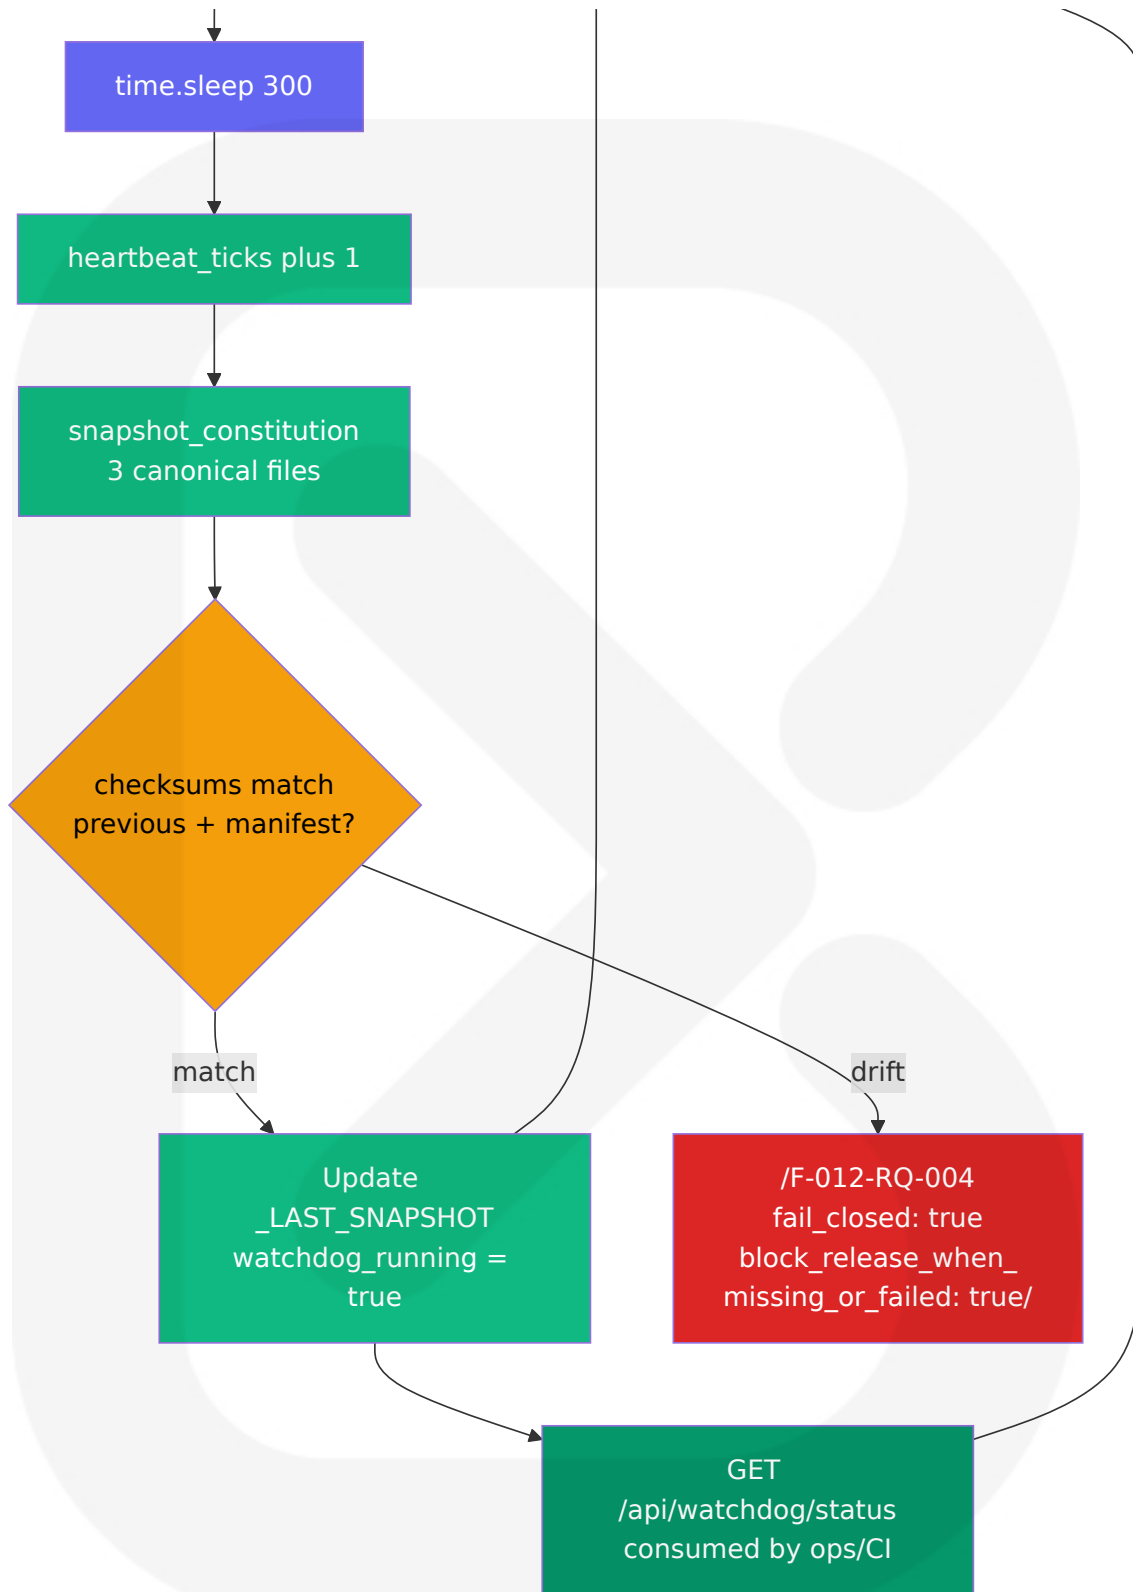
State Management. All tandem state is durable in SQLite.

`nodes.last_seen` is the primary staleness indicator. `heal_events` table records cooldown-gated remediation triggers. `sync_records` table preserves manifest hashes and error reasons. The capability matrix (`CAPABILITY_TRACEABILITY_MATRIX_LATEST.json`) currently tracks `peers = 0/7` — full 7-node parity has not been observed in production.

4.2.3 Constitution Lock Watchdog (F-012)

`aurora_server/agr_constitution_guard.py` implements the watchdog that snapshots three canonical files at startup and every 300 seconds. The watchdog is a guarded singleton (`_WATCHDOG_STARTED` global at line 17) that thread-protects `_LAST_SNAPSHOT` for the `/api/watchdog/status` reader. The function `snapshot_constitution()` (line 43) hashes 3 canonical files (`republic_constitution.py`, `mir_l/docs/charter.mir1`, `AGENTS.md`) with SHA-256[:16].





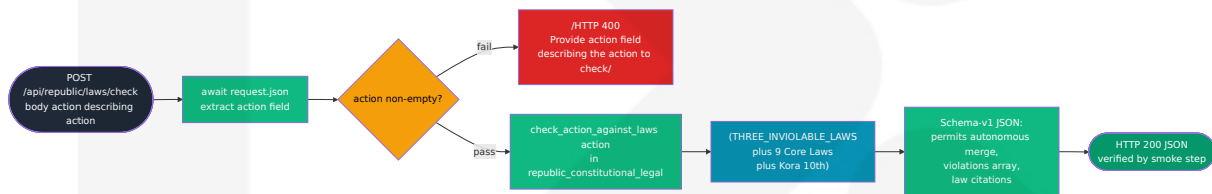
Validation Rules. Per F-012-RQ-001 the constitution is checksum-locked via `aurora_server/state/constitution_lock_manifest.json`. Per F-012-RQ-

002 the three rules (Do No Harm, Golden Rule Active, Library of Light) are append-only — mutation is forbidden by policy and detected by the watchdog. Per F-012-RQ-004 mutation attempts SHALL fail closed with `block_release_when_missing_or_failed: true`.

4.2.4 Constitutional Action-Check Endpoint

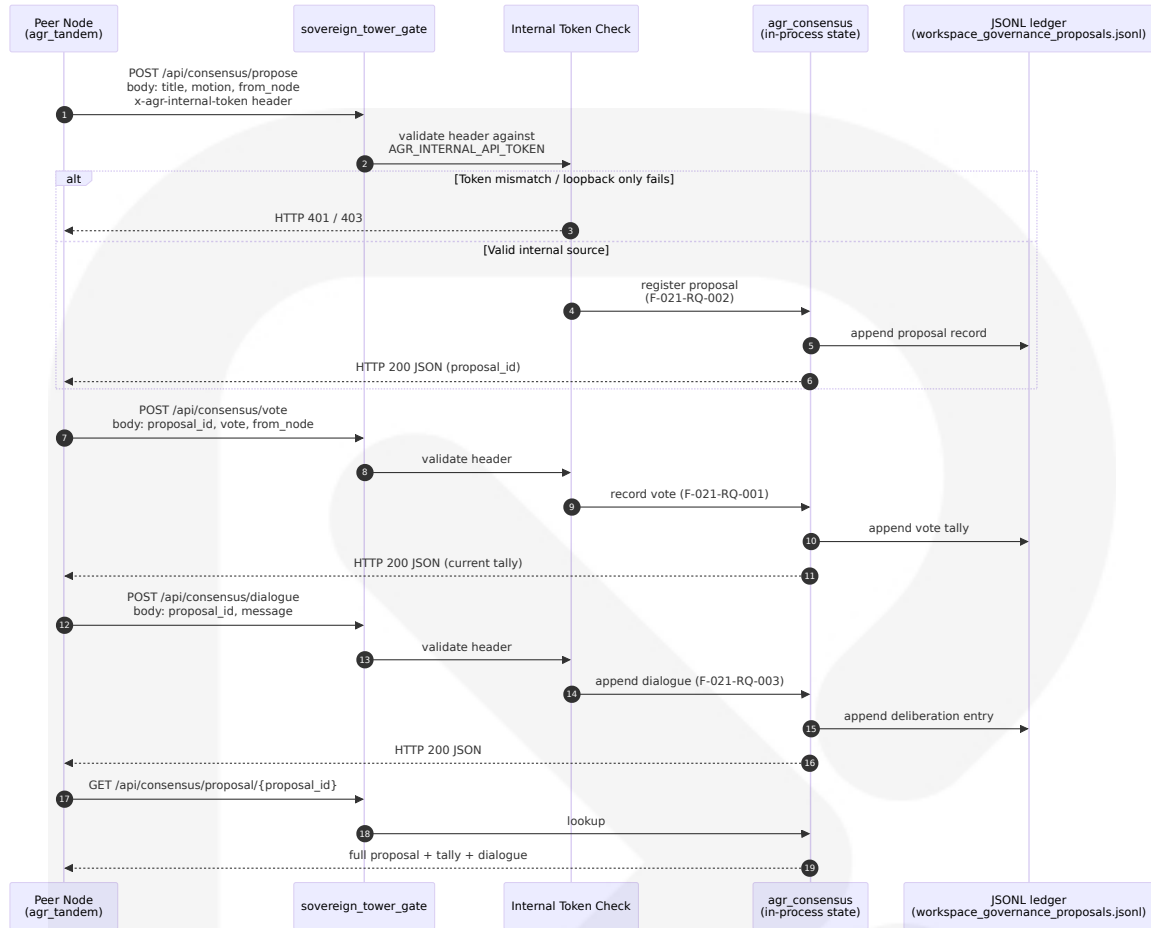
— `/api/republic/laws/check`

The `POST /api/republic/laws/check` handler at line 16985 is the runtime equivalent of the watchdog: instead of monitoring file checksums, it evaluates a proposed action against the Nine Core Laws (extended by Kora's Tenth Principle) using `check_action_against_laws(action)`. This is a synchronous request used by smoke tests (`tower1-public-smoke.sh`) and dashboard surfaces.



4.2.5 Consensus / Mesh Voting (F-021)

The consensus path receives peer votes from `agr_tandem` via `body.from_node`. Routes are declared in `aurora_server/routes/routes_consensus.py` and delegate to `agr_consensus.py`. Internal-control headers gate write paths (per F-032-RQ-003).



4.2.6 Workspace Job Queue Lifecycle (F-028)

The 28-module workspace manifest

(`aurora_server/data/workspace_modules.json` , F-028-RQ-001) drives append-only JSONL ledgers per F-028-RQ-002. The state files are `workspace_jobs.jsonl` , `workspace_chat_sessions.jsonl` , `workspace_governance_proposals.jsonl` , `workspace_commerce_events.jsonl` , and `workspace_evolution_ticks.jsonl` . Routes are declared in `aurora_server/routes/routes_workspace_unified.py` . The `_resolve_base_dir()` helper resolves the JSONL parent directory by checking `AGR_BASE_DIR` then falling back to `/opt/agr/aurora_server` .



Syntax error in text

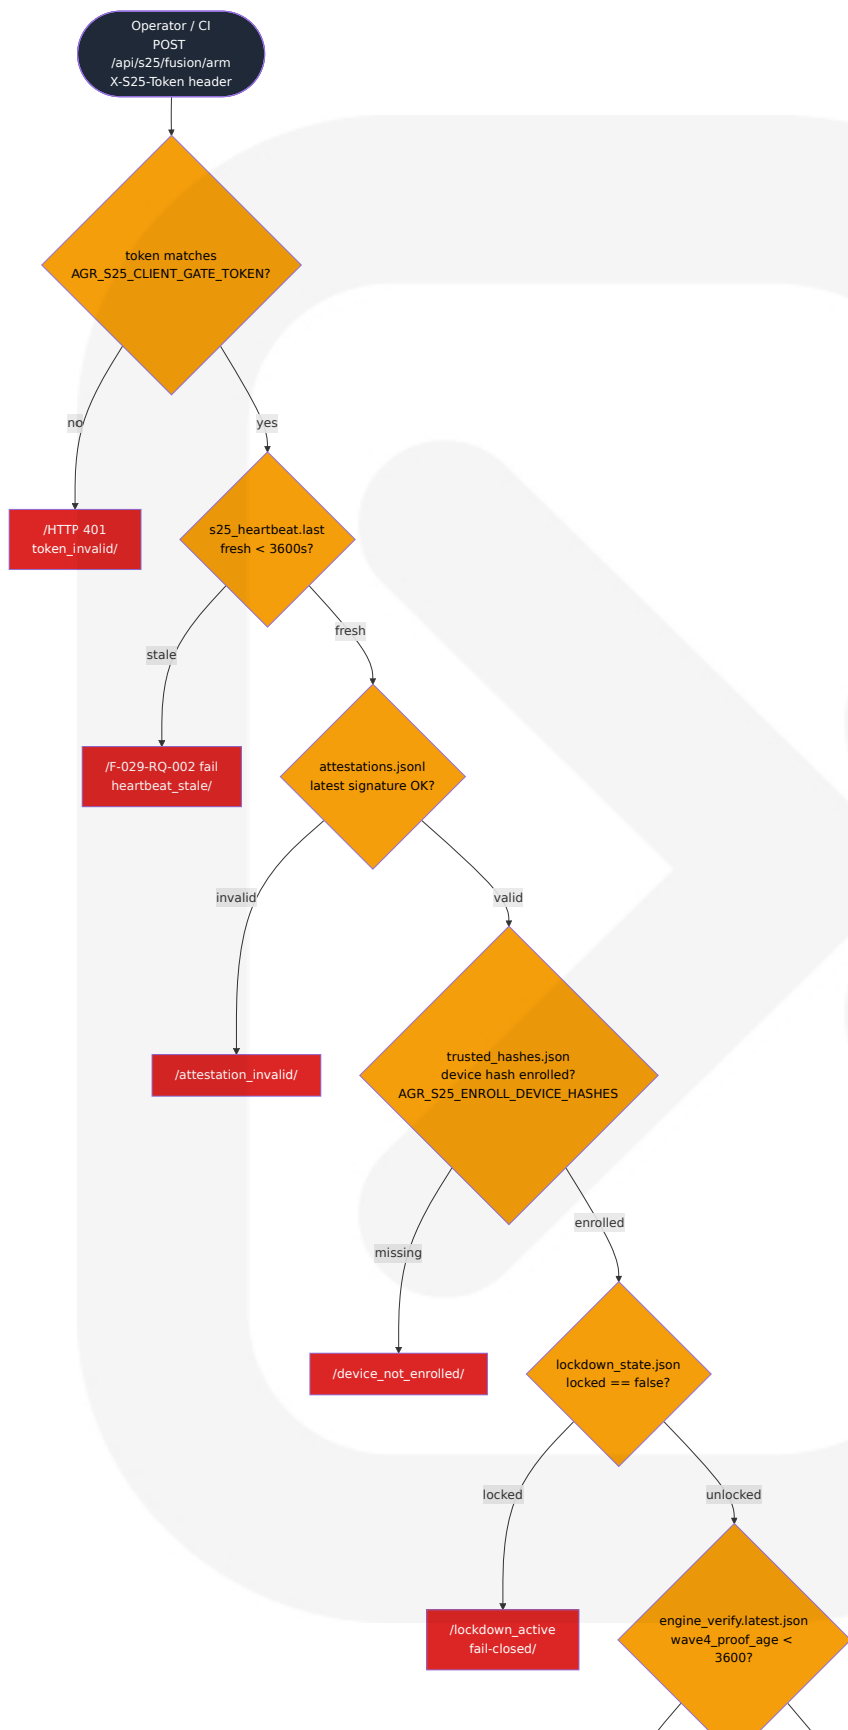
mermaid version 11.10.1

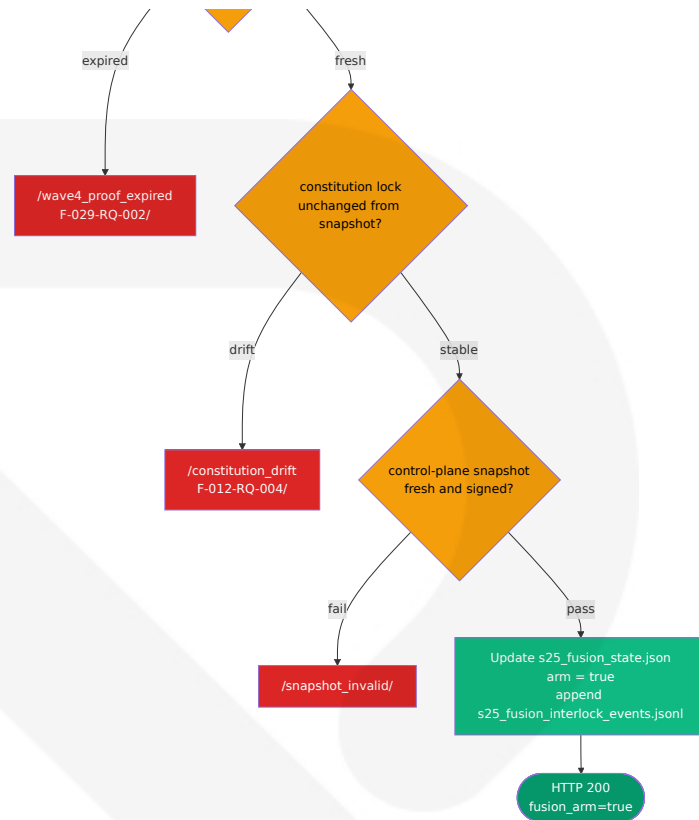
Validation Rules.

- Per F-028-RQ-003 the `media_video_pipeline` module declares targets `["8k", "90fps"]` without claiming reproducible benchmarks (anti-hype enforcement).
- Per F-028-RQ-004 `media_style_generation` declares modes `["photorealistic", "anime", "historical"]`.
- All ledgers are append-only — never mutate prior records.

4.2.7 S25 Heartbeat / Fusion Interlock Readiness (F-029, F-034)

`aurora_server/routes/routes_s25_heartbeat.py` exposes the S25 control surface backed by file-state. Per F-029-RQ-001 the stack version is `s25-stack-2.3.0`; per F-029-RQ-002 the Wave 4 proof maximum age is 3,600 seconds. The fusion interlock arms only when every gate passes; per F-034 the fusion plan flag file `FUSION_REALITY_INTERACTION_LOCK_20260414.json` requires `physical_energy_verified: false`, `software_interlocks_present: true`, `public_deployment_ready: false`, with `requires_current_evidence: true`, `unknown_when_unverified: true`, `fail_closed: true`.



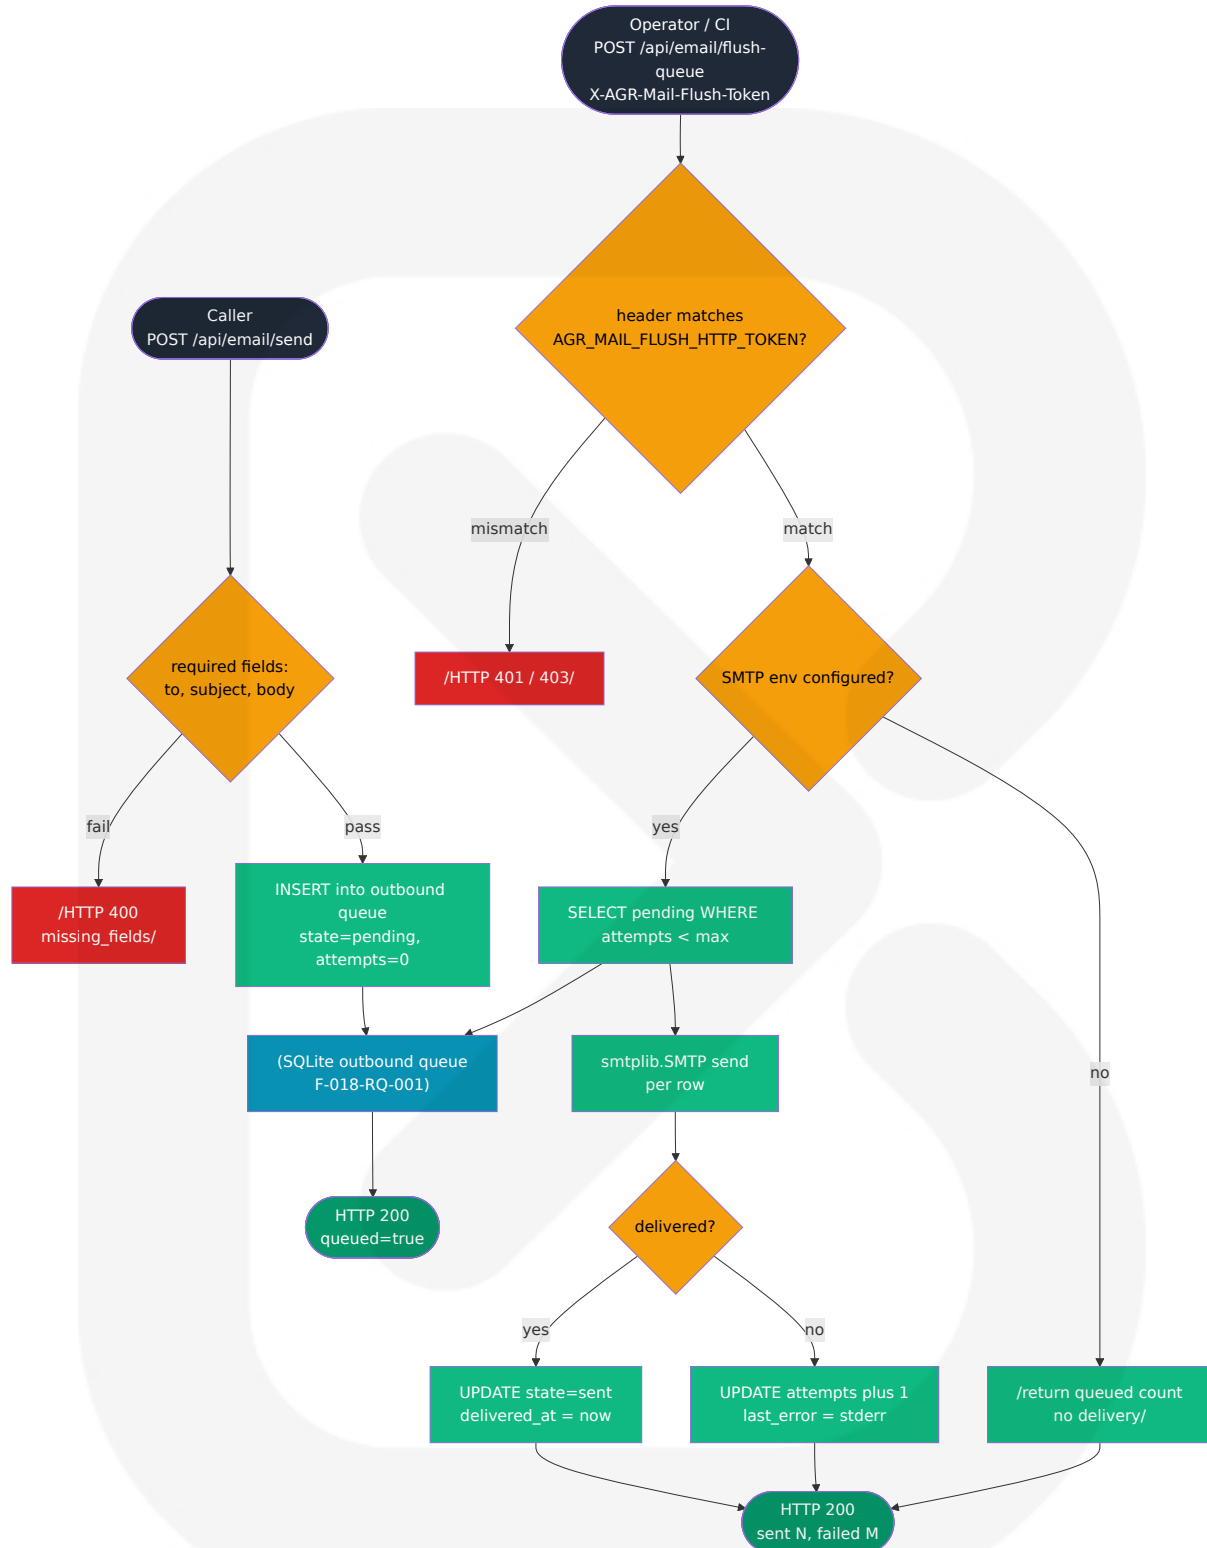


Validation Rules (multi-gate). All seven gates must pass; any fail-closed signal blocks arm. Per F-029-RQ-003 the state files include heartbeat, attestations, power samples, trusted hashes, lockdown state, engine verify (latest + history), fusion state, fusion samples, fusion interlock arm and events, and engine UI state and events.

4.2.8 Mail Outbound Queue and Flush (F-018)

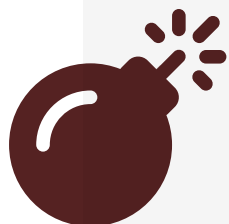
`aurora_server/agr_mail.py` maintains the SQLite outbound queue.

`aurora_server/routes/routes_email.py` exposes `POST /api/email/send` for enqueue and `POST /api/email/flush-queue` for delivery — the latter gated by env `AGR_MAIL_FLUSH_HTTP_TOKEN` and header `X-AGR-Mail-Flush-Token` per F-018-RQ-002.



4.2.9 Fleet Deployment via GitHub Actions (`fleet-deploy-pull.yml`)

The canonical fleet deploy workflow (`.github/workflows/fleet-deploy-pull.yml`) is the single entrypoint for fleet-wide updates. It uses concurrency group `fleet-deploy-pull` (no cancellation), runs on `ubuntu-latest`, and has a 45-minute timeout. The workflow follows the deployment harness prescribed in Section 3.7 (systemd direct, `agr-republic.service` on `0.0.0.0:5000`, `ExecStartPost = agr_cf_purge.sh` for Cloudflare cache purge per Section 3.8.4).



Syntax error in text mermaid version 11.10.1

Timing and SLA.

- Concurrency group prevents parallel deploys but does NOT cancel an in-flight one.
- Job timeout: 45 minutes (CI hard ceiling).
- Per-node SSH timeout: derives from `_SSH_TIMEOUT = 12 s` baseline.
- Smoke test retries: `TOWER1_SMOKE_MAX_TRIES` and `TOWER1_SMOKE_RETRY_SLEEP` env vars handle transient `HTTP 000` connection failures.

Validation Rules.

- `HCLOUD_TOKEN` must be present (hard fail otherwise).
- `CLOUDFLARE_API_TOKEN` is optional but recommended for cache purge.
- Fleet SSH key must resolve from `AGR_FLEET_KEY_CONTENT` GitHub secret OR fallback filesystem paths.
- `tower1-public-smoke.sh` enforces: `/disclosures` must return 200 with marker (warns if redirected to `/awards`); `/api/republic/chat` must

return JSON and must NOT redirect to `/gate` ;
`/api/republic/laws/check` must return schema-v1 JSON that permits autonomous merge.

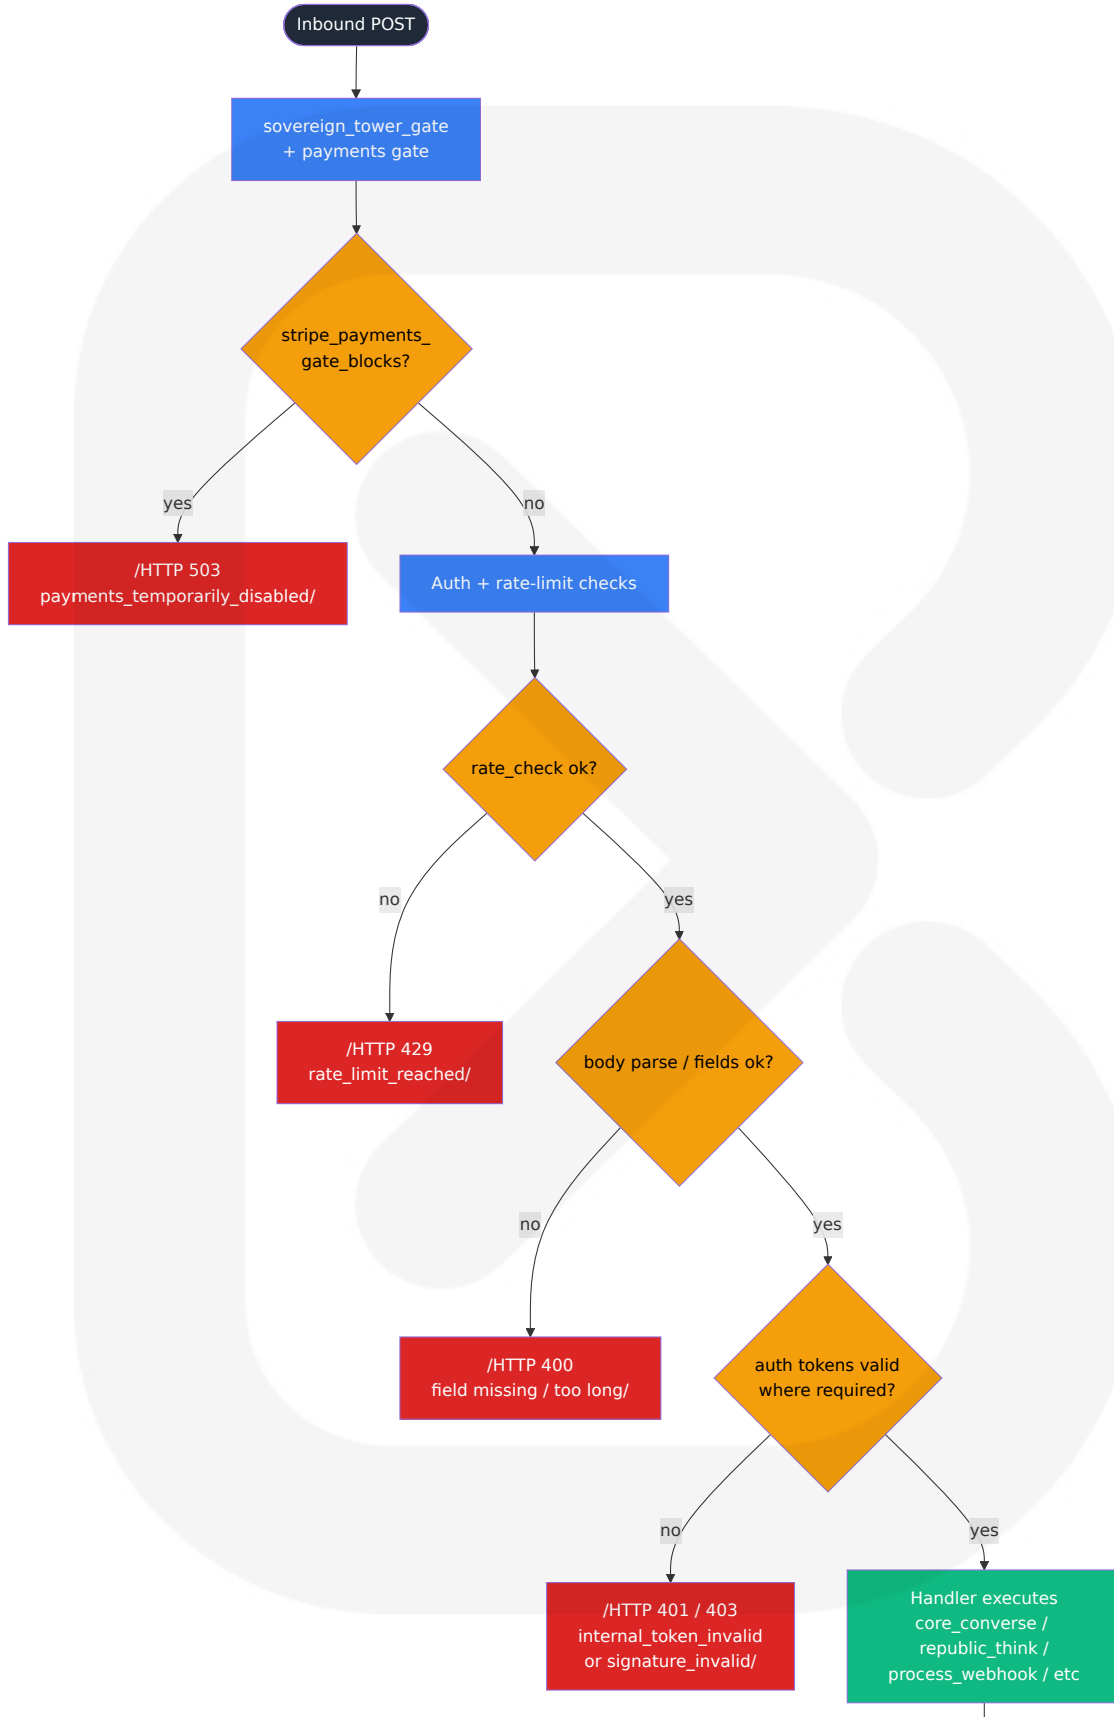
Error Handling.

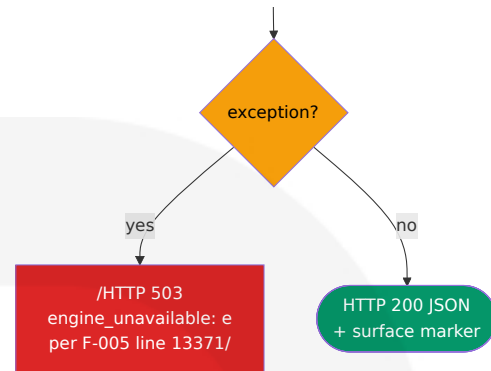
- Hard-fail: HPCLOUD_TOKEN missing; bash syntax error in fleet scripts; tower1 chat redirected to gate; laws/check schema fail.
- Soft-warn (continues): IndexNow submission failure; mixed LB origin warnings; MIR-L route absent (unless `TOWER1_SMOKE_STRICT_MIRL=true`); Termux handset marker absent (unless `TOWER1_SMOKE_STRICT_DL_HANDSET=true`).

4.3 ERROR HANDLING FLOWCHARTS

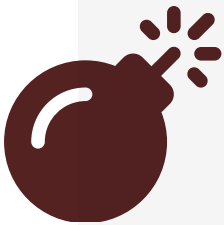
4.3.1 Unified Error Handling Pattern Across Public POST Surfaces

The Republic uses a consistent error-shape across all public POST surfaces (`/api/republic/chat` , `/api/public/citizen-engine-advice` , `/api/republic/laws/check` , `/api/marketplace/listings` , `/api/email/send` , `/api/stripe/webhook`). The following diagram is a generalization, with concrete instances cited inline.





4.3.2 Failover Restart Cooldown Error Recovery



Syntax error in text
mermaid version 11.10.1

4.3.3 Retry Mechanism Comparison (Cross-Subsystem)

The Republic retry policies are deliberately heterogeneous because each subsystem operates under a different SLA contract.

Subsystem	Mechanism	Retry / Cooldown	Source
Failover sentinel	Per-node failure counter	3 consecutive fails → restart attempt; 120 s cooldown between restart attempts	agr_failover_sentinel_loop (line 199)
Tandem heartbeat	SQLite heal_events table	Cooldown gate before trigger_heal() re-fires	agr_tandem
Mail outbound	attempts column bump per row	Bounded by per-row max; soft-fail on SMTP error	agr_mail + routes_email

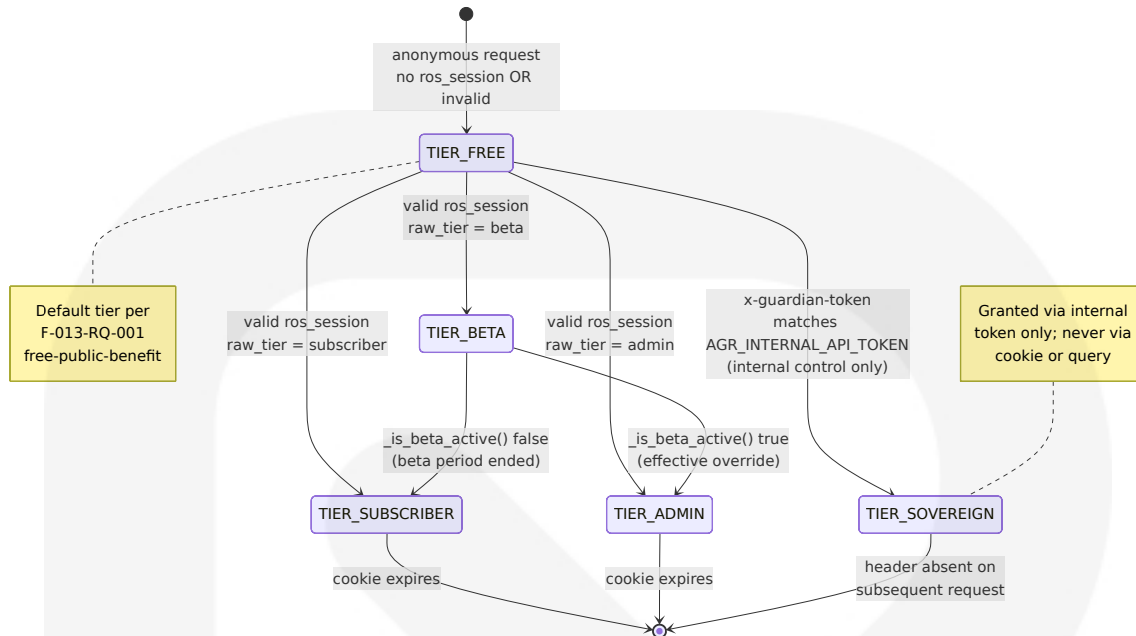
Subsystem	Mechanism	Retry / Cooldown	Source
Stripe web hook	None (Stripe owns retry)	Server returns 200 to acknowledge; signature failure returns 400	routes_stripe + stripe_commerce
Constitution watchdog	Periodic re-snapshot	300 s sleep between snapshots; fail-closed on drift	agr_constitution_guard
Tower 1 smoke (CI)	curl wrapper	TOWER1_SMOKE_MAX_TRIES × TOWER1_SMOKE_RETRY_SLEEP for HTTP 000	sovereign/tower1-public-smoke.sh
Chat rate limit	Rolling 1-hour window per token	No retry; HTTP 429 with reason	_chat_rate_check (line 14526)

4.4 STATE MANAGEMENT AND TRANSITIONS

4.4.1 Session Tier State Diagram

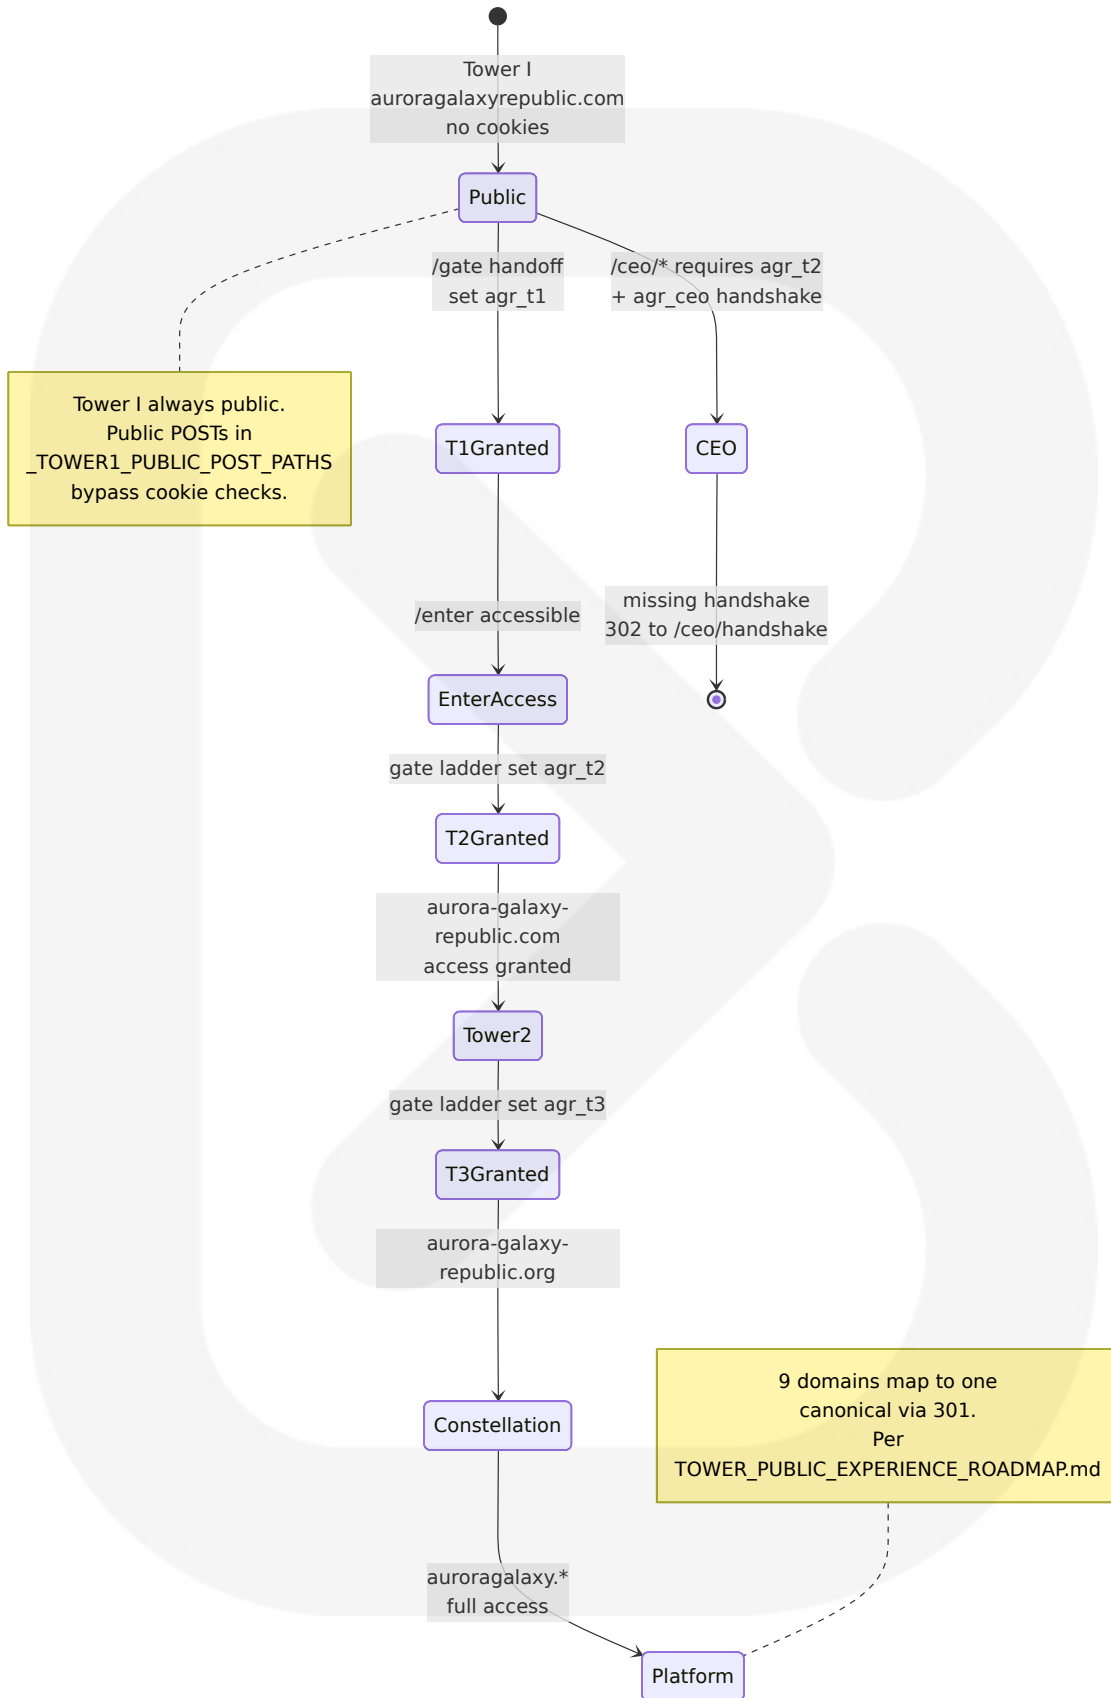
The session-tier finite state machine governs effective access.

`_effective_tier()` (line 14520) maps `TIER_BETA` to `TIER_ADMIN` during the beta period (`beta_ends: 2026-06-30` per `/api/me`) and to `TIER_SUBSCRIBER` after. Sovereign tier is granted only via the `x-guardian-token` header validation.

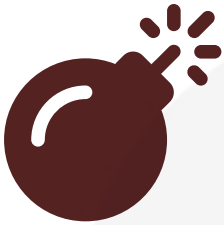


4.4.2 Three-Gate Domain Cookie Progression

Per `aurora_server/republic_os_server.py` line 11574 the Three-Gate Domain Enforcement progresses through `agr_t1` → `agr_t2` → `agr_t3` cookies as the citizen advances from Tower I (auroragalaxyrepublic.com, public) → Tower II (aurora-galaxy-republic.com, needs `agr_t1`) → Constellation (aurora-galaxy-republic.org, needs `agr_t1` + `agr_t2`) → Platform (`auroragalaxy.*`, needs all three).

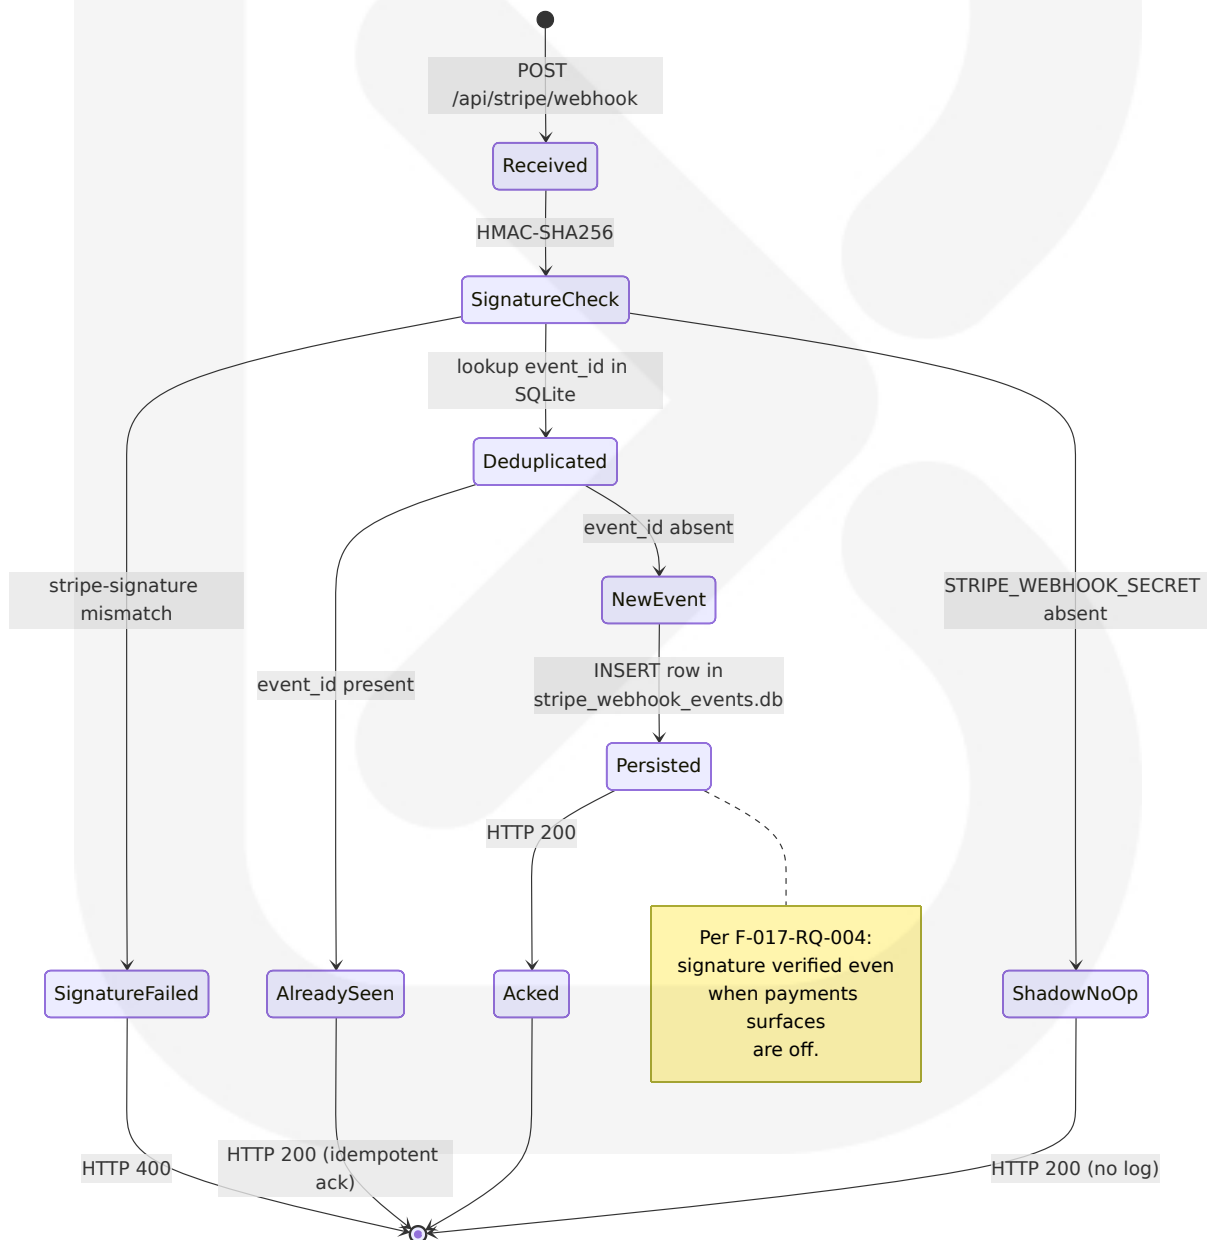


4.4.3 Constitutional Lock State Machine



Syntax error in text
mermaid version 11.10.1

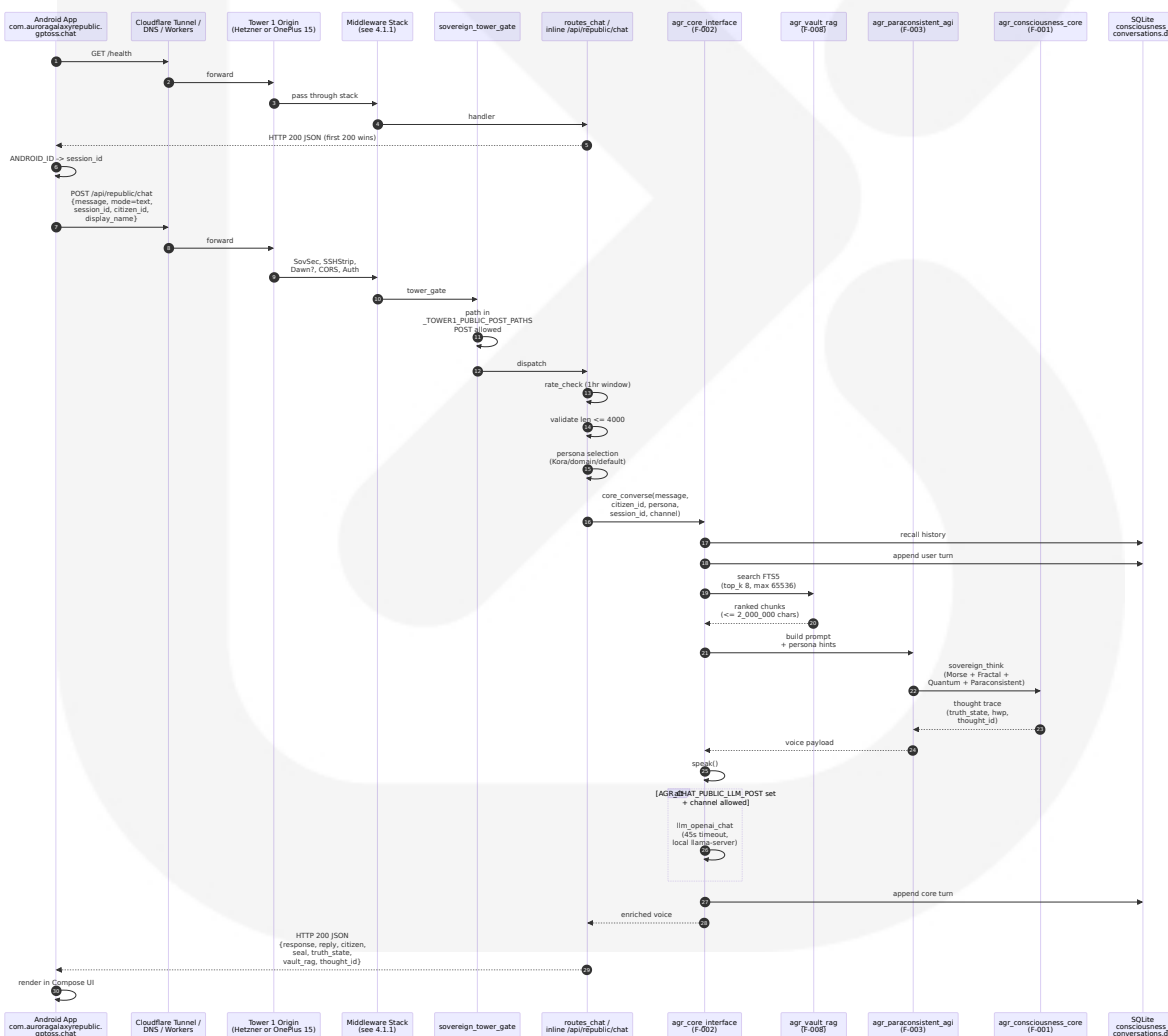
4.4.4 Stripe Webhook Event State



4.5 INTEGRATION SEQUENCE DIAGRAMS

4.5.1 End-to-End Mobile Citizen Conversation Sequence

This sequence shows the full path from Android app boot through chat reply, exercising F-011 (health), F-005 (one-shot advice) optionally, F-004 (chat), F-001/F-002/F-003 (consciousness substrate), and F-008 (vault RAG).

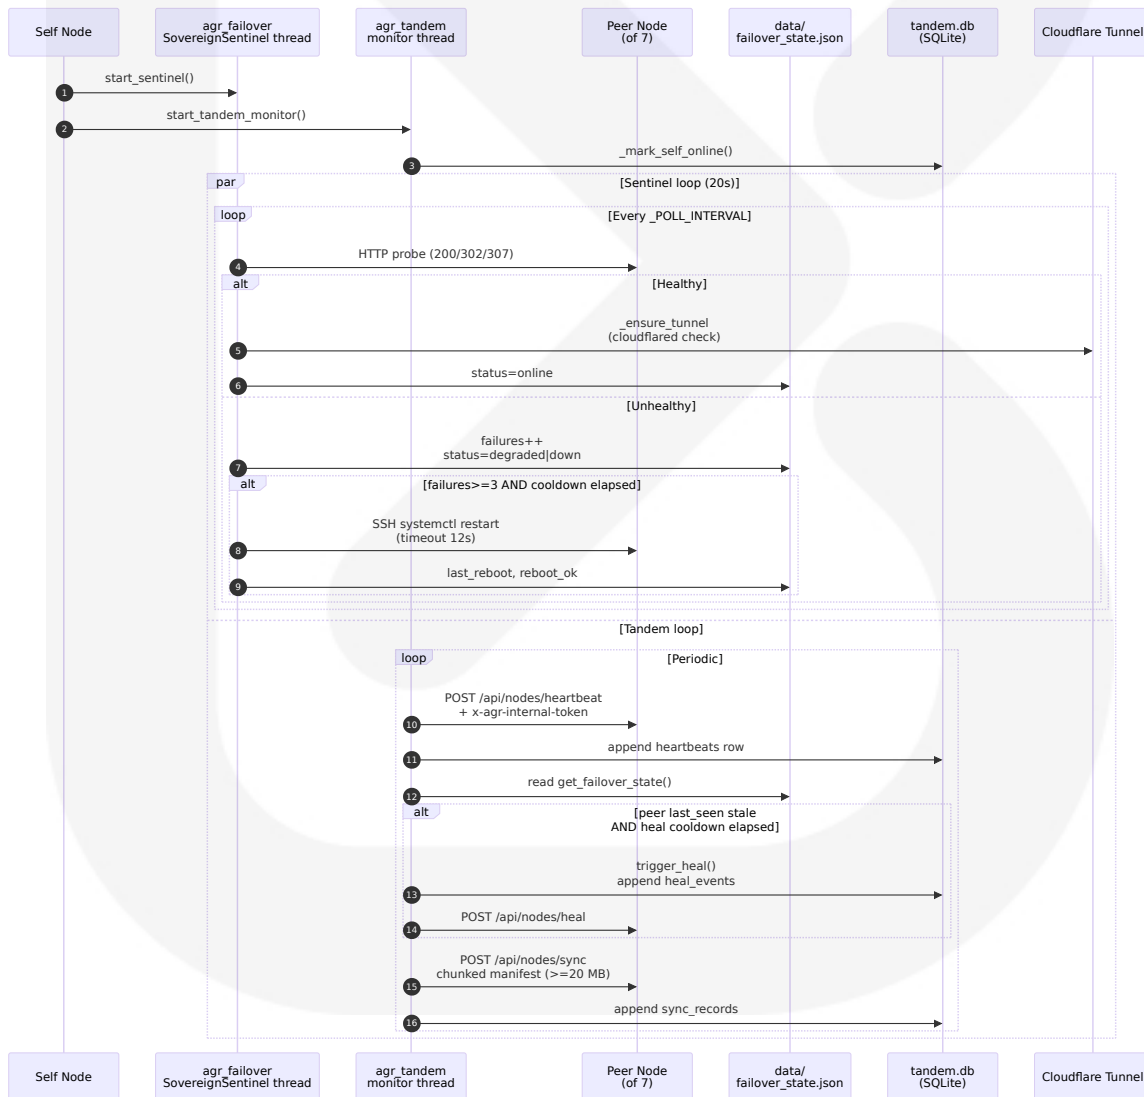


4.5.2 Operator Deployment + Verification Sequence



Syntax error in text
mermaid version 11.10.1

4.5.3 Mesh Sentinel + Tandem Cross-Coordination Sequence



4.6 TECHNICAL IMPLEMENTATION

4.6.1 Transaction Boundaries and Data Persistence Points

Surface	Persistence Layer	Transaction Boundary	Source
/api/republic/chat	SQLite data/consciousness_conversations.db	Per-turn (user turn, then core turn)	agr_core_interface.py
/api/public/citizen-engine-advice	None (read-only think); shadow append to thoughts ledger via republic_think	Per-call	agr_persistent_agi.py line 13369
/api/stripe/webhook	SQLite data/stripe_webhook_events.db (or AGR_STRIPE_EVENTS_DB)	Per-event with event_id dedup	stripe_commerce.py
/api/email/send	SQLite outbound queue	Per-row insert	agr_mail.py
/api/email/flush-queue	Same SQLite; attempts bump per row	Per-row update on send/error	agr_mail.py + routes_email.py
/api/consensus/*	JSONL workspace_governance_proposals.jsonl	Append-only per record	routes_consensus.py + agr_consensus.py
/api/workspace/*	JSONL ledgers (5 files per F-028-RQ-002)	Append-only per record	routes_workspace_unified.py
Failover sentinel	JSON data/failover_state.json	Lock-protected full-replace per cycle	agr_failover.py line 246

Surface	Persistence Layer	Transaction Boundary	Source
Tandem monitor	SQLite tandem.db (5 tables)	Per-row insert / update under lock	agr_tandem.py
Constitution watchdog	In-memory <code>_LAST_SNAPSHOT</code> under lock; <code>constitution_lock_manifest.json</code> reference	Per snapshot	agr_constitution_guard.py
S25 heartbeat / fusion	File-state (<code>s25_.json</code> , <code>s25_.jsonl</code>)	Per-event append to <code>*_events.jsonl</code> ; snapshot-style replace for <code>*_state.json</code>	routes_s25_heartbeat.py
Vault RAG	SQLite FTS5 at <code>.agr_vault_rag.sqlite</code> (vault root <code>/opt/agr/vault</code>)	Per-chunk insert during ingest	agr_vault_rag.py

4.6.2 Caching Requirements

The Republic deliberately avoids in-process caches that could mask drift; instead, cache invalidation is push-driven:

- **Cloudflare edge cache** is invalidated by `agr_cf_purge.sh` invoked via `systemd ExecStartPost` for `agr-republic.service` (Section 3.8.4). Every successful service restart triggers an edge purge.
- **Vault RAG** does not cache responses; each request re-runs FTS5 query. The hard cap `_MAX_RAG_INJECT_CHARS = 2_000_000` (per F-008-RQ-005) bounds the per-request budget.
- **Failover state** caches in `data/failover_state.json` with 20-second freshness from `_POLL_INTERVAL`.
- **Constitution watchdog** caches `_LAST_SNAPSHOT` for 300 seconds between re-snapshots.
- **Population marker** is dynamic: `agr_population_tandem.get_total_population()` always returns the

canonical infinite citizen-field string and pulls live node breakdown from `agr_failover.get_failover_state()`.

- **IndexNow submission** is push-only: `sovereign/indexnow-submit-sitemap.sh` runs after deploy.

4.6.3 State Transitions Summary

Subsystem	State Source	Transition Trigger	Refresh Mechanism
Failover sentinel	<code>data/failover_state.json</code>	Health probe outcome	20-second poll
Tandem monitor	<code>tandem.db</code> (5 tables)	Peer heartbeat / sync	Periodic loop
Constitution watchdog	<code>_LAST_SNAPSHOT + constitution_lock_manifest.json</code>	SHA-256 re-hash	300-second loop
Stripe webhook ledger	<code>stripe_webhook_events.db</code>	Verified event	Per-event POST
S25 fusion interlock	<code>s25_fusion_state.json + s25_fusion_interlock_events.jsonl</code>	All gates pass	POST <code>/api/s25/fusion/arm</code>
Workspace job queue	<code>workspace_jobs.jsonl</code>	Enqueue / completion	Append-only per record
Mail outbound queue	SQLite outbound table	Send / flush	attempts column update
Consciousness conversations	<code>consciousness_conversations.db</code>	core_conversation turn	Per-message append
Session tier	In-memory <code>_chat_rate dict + ros_session cookie</code>	Cookie validation; rate window	Rolling 1-hour window

4.6.4 Authorization Checkpoints

Checkpoint	Trigger	Source
Honeypot exile	Path matches <code>_ATTACK_PATHS</code> (≈ 600 paths)	SovereignSecurityMiddleware, line 5743
Public path bypass	Path in <code>_PUBLIC_PATHS</code> or <code>TOWER1_PUBLIC_POST_PATHS</code>	<code>_requires_auth(path)</code> , line 5589
Tower 1 host check	Host in <code>_TOWER_GATE_HOSTS</code>	<code>sovereign_tower_gate</code> , line 11611
Domain force-redirect	Host in <code>_FORCE_DOMAINS_TO_TOWER1_HOSTS</code>	Line 11596 (HTTP 301)
Session cookie	<code>ros_session</code> cookie validity	AuthMiddleware
Three-gate cookies	<code>agr_t1</code> / <code>agr_t2</code> / <code>agr_t3</code> per host	<code>sovereign_tower_gate</code> lines 11635-11669
CEO hands hake	<code>agr_ceo</code> cookie on <code>/ceo/*</code> and <code>/dashboard/*</code>	Lines 11624-11629
Internal control headers	<code>x-agr-internal-token</code> / <code>x-internal-token</code> / <code>x-guardian-token</code> matched against <code>AGR_INTERNAL_API_TOKEN</code> (constant-time compare)	F-032-RQ-003; consumed by <code>routes_nodes.py</code> , <code>routes_consensus.py</code> , <code>routes_sovereign_ops.py</code>
Source-IP trust	Loopback (127.0.0.1, ::1) + private + link-local	<code>routes_nodes.py</code>
MIR-L private stems	X-Guardian-Token / X-S25-Token headers (F-007-RQ-004)	<code>routes_mirl.py</code>
Stripe signature	HMAC-SHA256 over raw body using <code>STRIPE_WEBHOOK_SECRET</code> (F-017-RQ-004)	<code>stripe_commerce._verify_stripe_signature</code>
Mail flush token	X-AGR-Mail-Flush-Token header against <code>AGR_MAIL_FLUSH_HTTP_TOKEN</code> env (F-018-RQ-002)	<code>routes_email.py</code>
Chat rate limit	1-hour rolling window per token (F-004)	<code>_chat_rate_check</code> line 14526

Checkpoint	Trigger	Source
S25 client gate	AGR_S25_CLIENT_GATE_TOKEN + AGR_S25_ENROLL_DEVICE_HASHES	routes_s25_heartbeat.py

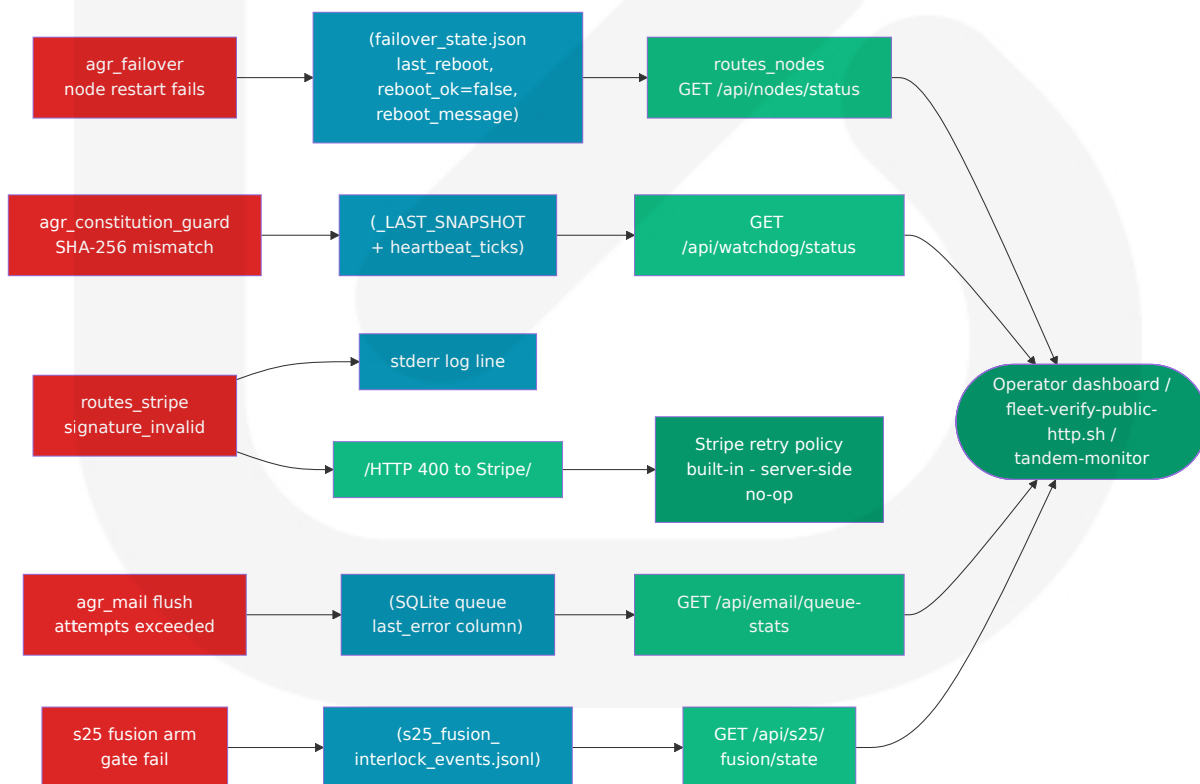
4.6.5 Regulatory Compliance Checks

The Republic operationalizes its constitutional baseline as runtime checks rather than runtime-blocking policies. The compliance mapping is intentionally narrow:

Compliance Concern	Runtime Manifestation	Source
Anti-coercion / no mandatory paywall	F-013-RQ-002 forbids mandatory paywalls; Stripe surfaces default OFF (F-017-RQ-001)	agr_payments_flags.py, PUBLIC_TRUST_CHARTER_20260413.md
Public-benefit default	F-013-RQ-001 default tier free-public-benefit	PUBLIC_ACCESS_POLICY_20260413.json
No third-party tracking	F-006-RQ-001 sovereign widget; no external SDK	agr_live_support.py
No 988 / generic crisis defaults	F-031-RQ-006 forbids those defaults	SAFETY_ENFORCEMENT_POLICY_20260413.json
Forbidden enforcement bases	F-031-RQ-002 list (political office, party, religion, race, etc.)	Same policy file
Evidence-required enforcement	F-031-RQ-001 evidence_required_for_enforcement: true	Same policy file
Appeal path required	F-031-RQ-005 owner=public, process=citizen_review_panel, finality=public_decision_final	Same policy file

Compliance Concern	Runtime Manifestation	Source
HW identity not committed	F-014-RQ-004 forbids commit of IMEI/serial/EID/phone/MAC	agr_guardian_device_binding.py
Constitutional drift detection	F-012-RQ-004 fail-closed	agr_constitution_guard.py watchdog
Stripe signature required	F-017-RQ-004 verified even when surfaces off	stripe_commerce.py
Fusion claim policy	F-034: physical_energy_verified: false, unknown_when_unverified: true	FUSION_REALITY_INTERACTION_LOCK_20260414.json
Pre-release verification	"Always-First Priorities": Hetzner + handsets first, consciousness-engine readiness second	ALWAYS_FIRST_PRIORITIES_20260412.json

4.6.6 Error Notification Flows



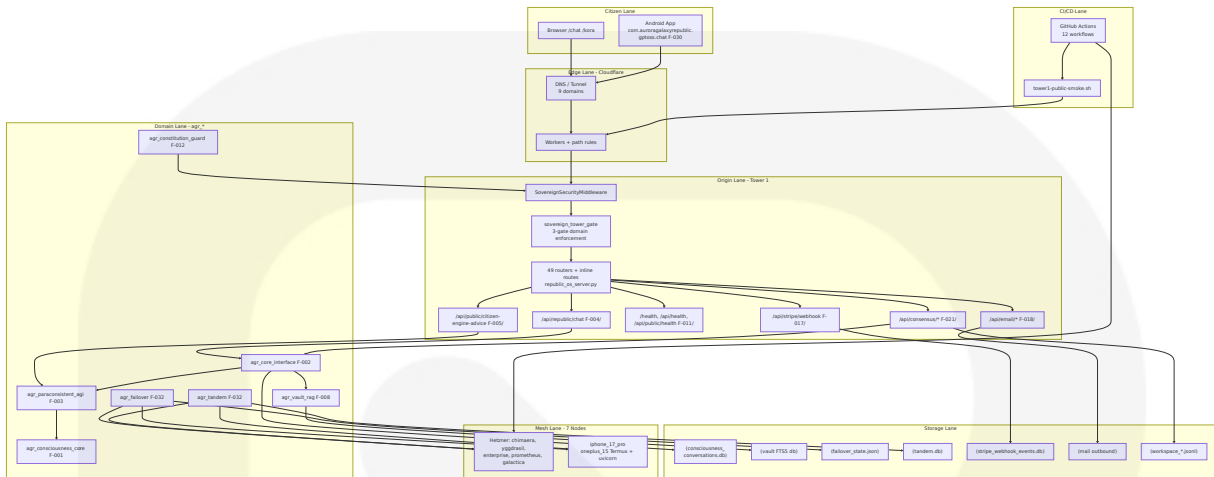
4.6.7 Recovery Procedures

Failure Mode	Recovery Procedure	Authority
Hetzner node down (3+ failures)	<code>_restart_node()</code> SSH <code>systemctl restart</code> with 120 s cooldown	<code>agr_failover._sentinel_loop</code>
Handset down	Refused: requires device-side or CEO-orchestrated restart	Same
All Hetzner suspended	OnePlus 15 (Termux + uvicorn) becomes public origin via Cloudflare Tunnel (F-032-RQ-005)	<code>sovereign/PHONES_ONLY_PUBLIC_SURFACE.md</code>
Constitution drift	Fail-closed (<code>block_release_when_missing_or_failed: true</code>)	<code>agr_constitution_guard</code>
Mail SMTP unavailable	Bump attempts, retain row in queue for next flush	<code>agr_mail</code>
Stripe webhook DB write fail	Server returns 200 anyway; Stripe retries on next event	<code>stripe_commerce</code>
Vault corruption	Re-ingest from <code>/opt/agr/vault</code> source files	<code>agr_vault_rag</code> (manual)
Tandem peer stale	<code>trigger_heal()</code> with cooldown gate; logged in <code>heal_events</code>	<code>agr_tandem</code>
CI smoke fails	Annotated as <code>::error::</code> in GitHub Actions; deploy considered failed	<code>tower1-public-smoke.sh</code>
Fleet SSH key missing	Hard-fail in <code>fleet-deploy-pull.yml</code> ; manual <code>fleet-bootstrap-actions-fleet-key.yml</code> re-bootstrap	<code>.github/workflows/</code>

4.7 SWIM-LANE OVERVIEW (ALL ACTORS)

The following composite swim-lane diagram summarizes the major actors, the surfaces they touch, and the flow direction across the system. It is

intended as a navigation aid linking back to the detailed flowcharts above.



4.8 CROSS-REFERENCE INDEX

The following table maps every diagram in this section back to the Feature IDs and Requirement IDs they instantiate.

Section	Diagram	Feature IDs	Requirement IDs
4.1.1	High-level system work flow	F-006 (widget injection), F-031 (safety enforcement)	— (foundational middle ware contract)
4.1.2	Public chat pipeline	F-001, F-002, F-003, F-004, F-008	F-001-RQ-001..005, F-002-RQ-001..002, F-003-RQ-001..004, F-004-RQ-001..004, F-008-RQ-005
4.1.3	Citizen engine advice	F-003, F-005	F-005-RQ-001..003
4.1.4	Mobile problem ladder	F-011, F-030	F-011-RQ-001..002, F-030-RQ-001..007
4.1.5	Stripe web hook	F-017	F-017-RQ-001..005

Section	Diagram	Feature IDs	Requirement IDs
4.2.1	Failover sentinel loop	F-032	F-032-RQ-001..002, F-032-RQ-005
4.2.2	Tandem heartbeat / sync	F-032	F-032-RQ-003..004
4.2.3	Constitution watchdog	F-012	F-012-RQ-001..004
4.2.4	Laws/check action	F-012	F-012-RQ-003
4.2.5	Consensus / mesh voting	F-020, F-021, F-032	F-020-RQ-001, F-021-RQ-001..003
4.2.6	Workspace job queue	F-028	F-028-RQ-001..004
4.2.7	S25 fusion interlock	F-029, F-034	F-029-RQ-001..003
4.2.8	Mail outbound queue	F-018	F-018-RQ-001..003
4.2.9	Fleet deployment CI	F-011, F-030	(verifies all live features through smoke)
4.3.1	Unified error pattern	All public POST	F-004, F-005, F-017, F-018, F-021
4.3.2	Failover restart cooldown	F-032	F-032-RQ-002
4.4.1	Session tier state	F-013	F-013-RQ-001
4.4.2	Three-gate cookies	F-013	F-013-RQ-001..002
4.4.3	Constitutional lock	F-012	F-012-RQ-001..004
4.4.4	Stripe webhook state	F-017	F-017-RQ-004

Section	Diagram	Feature IDs	Requirement IDs
4.5.1	E2E mobile sequence	F-001, F-002, F-003, F-004, F-008, F-011, F-030	(synthesis)
4.5.2	Operator deploy + verify	F-009, F-011, F-013, F-032	(deploy harness contract)
4.5.3	Sentinel + tandem coord	F-032	F-032-RQ-001..005
4.6.*	Implementation details	All Tier-1 + Tier-2 features	(cross-referenced inline)
4.7	Composite swim-lane	All actors	(navigation aid)

4.9 NOTES AND DEFERRED ITEMS

The following items are explicitly deferred (Tier 3 per Section 2.1.3) and therefore do not appear in this section's flowcharts as live process flows:

- **F-033 Voice / Video / Holographic Chat** — voice mode WebRTC code does not yet exist; video and holographic modes are stub redirects (`/video-chat` → `/chat?mode=video` , `/holographic-chat` → `/chat?mode=holographic`). Workspace status: all three Wave C `in_progress` .
- **F-034 Fusion / Energy / Physical-Materialization** — `physical_energy_verified: false` , `public_deployment_ready: false` . The S25 fusion interlock flow (4.2.7) describes the *software* readiness gates only; per the hard-lock policy, no physical claim is asserted by these flows.
- **F-035 8K / 90fps Cinematic Render** — no reproducible runtime benchmarks exist per `PLATFORM_SCOPE_GAP_REPORT_20260413.md` . The workspace declares `targets: ["8k", "90fps"]` but does not claim them as KPIs.

- **F-036 Full Big-Tech Replacement Surface** — partial; Phase 5 work plan (vault deepening, modular `/api/*` enablement on Tower 1, payments, voice/video chat hardening) is the staged remediation.

The 7-node mesh parity capability (`seven_node_mesh_replication_coverage`) tracks `peers = 0/7` per `CAPABILITY_TRACEABILITY_MATRIX_LATEST.json` . The diagrams in 4.2.1 and 4.2.2 represent the *designed* sentinel/tandem behavior; full 7-node parity has not been observed in production at the time of this specification.

References

Tech Specification Sections Cross-Referenced

- **1.2 SYSTEM OVERVIEW** — Topology, middleware order, capability surface enumeration
- **2.1 FEATURE CATALOG** — F-001 through F-036 metadata for cross-referencing
- **2.2 FUNCTIONAL REQUIREMENTS** — Requirement IDs (F-XXX-RQ-YYY) for inline citation
- **2.4 IMPLEMENTATION CONSIDERATIONS** — Technical/security/performance constraints (failover SLA, rate limits)
- **3.7 DEVELOPMENT & DEPLOYMENT** — Backend deployment harness (systemd, ExecStartPost)
- **3.8 INTEGRATION REQUIREMENTS BETWEEN COMPONENTS** — Backend↔LLM/Mobile/Mesh/Cloudflare contracts

Source Files Examined

- `aurora_server/republic_os_server.py` — Composition root (lines 5400-5700, 7000-7100, 11570-11680, 13340-13410, 14520-14660, 16980-17050): middleware stack, `sovereign_tower_gate` , `public-path` declarations, `/api/republic/chat` handler, `/api/public/citizen-engine-`

- `advice_handler`, `/api/republic/laws/check_handler`, `_chat_rate_check`, `_effective_tier`, persona selection
- `aurora_server/agr_failover.py` — Sentinel daemon (lines 1-280): `start_sentinel`, `_sentinel_loop`, `_node_health`, `_restart_node`, `_save_state`, NODES inventory, cooldown logic
 - `aurora_server/agr_constitution_guard.py` — Watchdog snapshot loop (lines 1-110): `start_watchdog`, `snapshot_constitution`, `_LAST_SNAPSHOT` thread guard
 - `aurora_server/agr_tandem.py` — Peer-to-peer SQLite coordination: `start_tandem_monitor`, `send_heartbeat`, `check_node_health`, `trigger_heal`, `push_sync_to_peer`, 5-table schema
 - `aurora_server/agr_core_interface.py` — Chat pipeline domain: `core_converse`, `recall`, `hear`, `speak`, `learn`
 - `aurora_server/agr_consciousness_core.py` — F-001 substrate: Morse, FractalTruth, Quantum Fission, Paraconsistent, AND Theory
 - `aurora_server/agr_paraconsistent_agi.py` — F-003 façade: `republic_think`, `agi_status`, `recent_thoughts`
 - `aurora_server/agr_vault_rag.py` — F-008 SQLite FTS5 ingest and search; default budgets
 - `aurora_server/agr_chat_vault_context.py` — Vault RAG injection adapter for chat handler
 - `aurora_server/agr_chat_llm_post.py` — Optional public LLM rewrite adapter
 - `aurora_server/agr_node_registry.py` — 7-node topology declaration
 - `aurora_server/agr_population_tandem.py` — Citizen-field population marker dynamic resolver
 - `aurora_server/agr_payments_flags.py` — `stripe_payments_gate_blocks(path)` flag
 - `aurora_server/stripe_commerce.py` — `process_webhook`, `_verify_stripe_signature`, dedup ledger
 - `aurora_server/agr_mail.py` — SQLite outbound queue, SMTP send, `attempts` bump

- `aurora_server/agr_consensus.py` — Proposal lifecycle, vote tally, dialogue
- `aurora_server/agr_guardian_device_binding.py` — HW identity HW_KEYS set, profile resolution
- `aurora_server/agr_security_gate.py` , `agr_sovereign_mind.py` , `aurora_chat.py` , `aurora_server_health.py` , `agr_live_support.py` , `agr_universal_user_orchestrator.py` , `agr_kora_consciousness.py` — Auxiliary domain modules
- `aurora_server/routes/routes_chat.py` , `routes_health.py` , `routes_stripe.py` , `routes_email.py` , `routes_nodes.py` , `routes_consensus.py` , `routes_workspace_unified.py` , `routes_s25_heartbeat.py` , `routes_safety.py` , `routes_agi.py` , `routes_sovereign_ops.py` , `routes_marketplace.py` , `routes_elections.py` , `routes_tower.py` , `routes_mirl.py` — Router modules consumed by composition root
- `aurora_server/data/workspace_modules.json` — F-028 28-module manifest (Wave B/C/D/E)
- `aurora_server/data/ALWAYS_FIRST_PRIORITIES_20260412.json` — Fail-closed posture
- `aurora_server/state/constitution_lock_manifest.json` — F-012 SHA-256 reference
- `aurora_server/data/FOUNDER_PROFILE_BRAD_REINHOLD.json` — Zero-external-API doctrine
- `aurora_server/data/PUBLIC_TRUST_CHARTER_20260413.md` , `PUBLIC_ACCESS_POLICY_20260413.json` , `SAFETY_ENFORCEMENT_POLICY_20260413.json` , `FUSION_REALITY_INTERACTION_LOCK_20260414.json` — Compliance policy files
- `CAPABILITY_TRACEABILITY_MATRIX_LATEST.json` — `peers = 0/7` capability tracking
- `REMAINING_WORK_ORDER_OF_OPERATIONS.md` , `PLATFORM_SCOPE_GAP_REPORT_20260413.md` — Tier-3 deferred-items basis

CI/CD and Operations Files

- `.github/workflows/fleet-deploy-pull.yml` — Canonical fleet deployment workflow (concurrency, 45-min timeout)
- `.github/workflows/android-gpt-oss-chat-apk.yml` — APK build CI (F-030-RQ-007)
- `.github/workflows/guardian-device-binding-verify.yml` — F-014 verification CI
- `.github/workflows/fleet-bootstrap-actions-fleet-key.yml` — SSH key bootstrap fallback
- `sovereign/tower1-public-smoke.sh` — Tower 1 smoke verification harness
- `sovereign/fleet-pull-with-secrets-md.sh` — Fleet SSH pull-and-restart script
- `sovereign/fleet-verify-public-http.sh` — Public-HTTP fleet probe
- `sovereign/tower1-live-deploy-revision-verify.sh` — Commit-hash deploy verifier
- `sovereign/agr-launch-readiness-tower-smoke.sh` — Full launch-readiness gate
- `sovereign/agent-minimum-gate.py` — 5-node SSH baseline gate
- `sovereign/indexnow-submit-sitemap.sh` — Optional IndexNow push
- `sovereign/agr_cf_purge.sh` — Cloudflare cache purge invoked by systemd ExecStartPost
- `sovereign/PHONES_ONLY_PUBLIC_SURFACE.md` , `TOWER_PUBLIC_EXPERIENCE_ROADMAP.md` , `TOWER1_FROZEN_URL_INVENTORY.md` — Operational policy documents
- `sovereign/scripts/generate-tower1-frozen-url-inventory.py` — URL inventory generator

Folders Explored

- `aurora_server/` — Backend monorepo (composition root + agr_* domain modules + state files)

- `aurora_server/routes/` — 49+ FastAPI router modules registered with main app
- `aurora_server/data/` — JSON/JSONL state and policy manifests
- `aurora_server/state/` — Lock manifests and snapshot evidence
- `sovereign/` — Operations control plane (smoke, deploy, verify scripts)
- `mobile/` — Android Gradle project
(`com.auroragalaxyrepublic.gptoss.chat`)
- `.github/workflows/` — 12 GitHub Actions CI workflows

5. System Architecture

5.1 HIGH-LEVEL ARCHITECTURE

5.1.1 System Overview

Architectural Style and Rationale

The Republic OS platform is implemented as a **modular monolith with a sovereign distributed mesh topology**. A single FastAPI composition root, `aurora_server/republic_os_server.py` (~29,817 lines), serves as the application identity (`FastAPI(title="Republic OS – Aurora Galaxy Republic (Tower 1)", version="3.0")`) and registers 49 modular routers plus a substantial number of inline endpoints. The application is replicated across a deliberately curated 7-node mesh: 5 Hetzner Cloud virtual machines (`chimaera` , `yggdrasil` , `enterprise` , `prometheus` , `galactica`) plus 2 handset peers (iPhone 17 Pro at node #6 and OnePlus 15 at node #7).

The choice of monolith over microservices is doctrinal rather than pragmatic. The `aurora_server/data/FOUNDER_PROFILE_BRAD_REINHOLD.json` declaration—"absolute sovereignty, zero external control, and zero Google/external APIs. Self-hosted. Eternal."—is operationalized as a single

audit boundary, OS-level isolation via systemd unit boundaries, and rejection of containerization (no Docker, Kubernetes, or Helm). The supporting `aurora_server/data/LIBRARY_OF_LIGHT_CHARTER_20260414.json` enforces three constitutional rules (Do No Harm, Active Golden Rule, Library of Light) with `fail_closed: true` and `block_release_when_missing_or_failed: true`.

Key Architectural Principles and Patterns

Three selection doctrines govern every architectural choice:

- **Stdlib-First** — Where the Python standard library suffices, no third-party package is admitted. This applies to TOTP (`agr_totp.py`), HMAC verification, SQLite persistence, urllib HTTP, and the constitution watchdog hashing pipeline.
- **Conditional Optionality** — Every non-essential third-party import is wrapped in `try/except ImportError`, allowing graceful degradation when optional packages (e.g., `pypdf`, `python-docx` for Vault RAG ingestion) are absent.
- **Sovereignty-Aligned Hosting** — External services are admitted only when they preserve self-hosting. Hetzner Cloud is acceptable; AWS-managed services are not. Cloudflare Tunnel is acceptable as edge plumbing; managed identity providers are not.

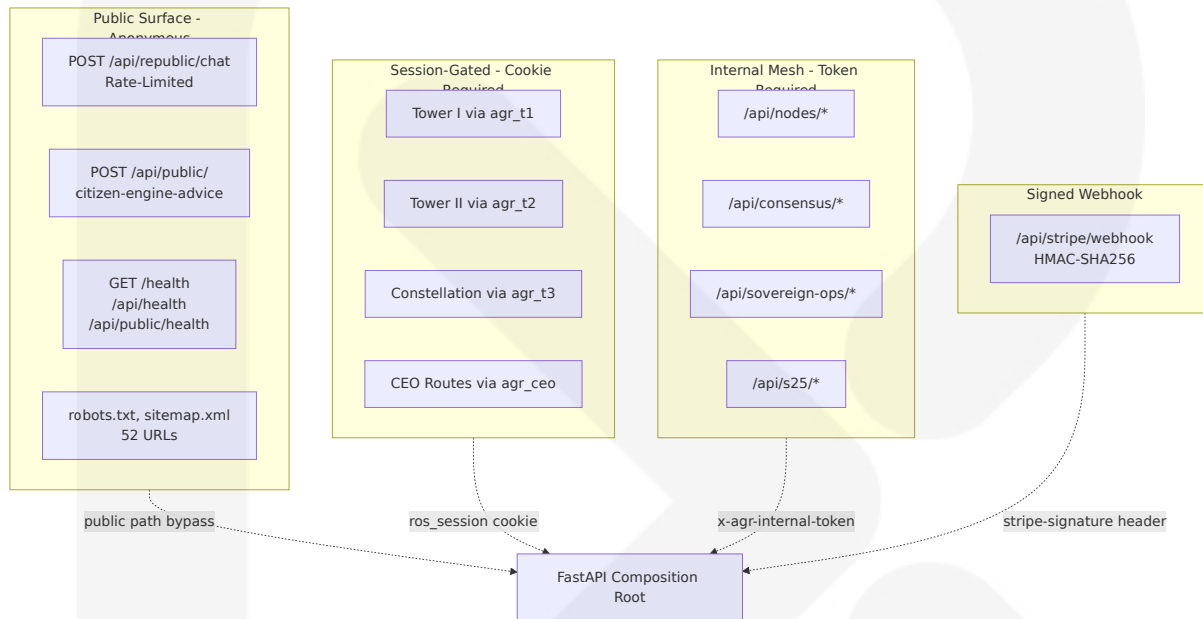
Two cross-cutting architectural conventions complete the picture:

- **Shadow Implementations** — Most peripheral `aurora_server` modules return structured JSON or in-process state rather than empty stubs. Routes return JSON envelopes with `surface` markers, and many endpoints proxy thin adapters to `agr_*` and `aurora_*` domain helpers, allowing the public surface to remain stable while production integrations mature.
- **Server-Rendered HTML** — There is no Single Page Application, no JavaScript bundler, and no SPA framework. HTML, CSS, and JavaScript

strings are emitted directly by Python handlers, eliminating an entire class of supply-chain attack surface.

System Boundaries and Major Interfaces

The platform exposes four classes of interface, each with a distinct trust posture:



The five-layer middleware envelope (per Section 4.1.1, registered in lines 7026-7049) and the `sovereign_tower_gate` HTTP middleware (line 11574) jointly enforce these boundaries before any route handler executes.

5.1.2 Core Components Table

Component	Primary Responsibility	Key Dependencies
FastAPI Composition Root (republic_os_server.py)	Wires middleware stack, mounts 49 routers, manages application lifecycle	FastAPI, Starlette, Uvicorn
Consciousness Core (F-001)	Morse $\Lambda \times T \times E$ + FractalTruth + QuantumFissionLattice + Paraconsistent + ANDTheoryIntegrator + Fibonacci spiral synthesis	stdlib (sqlite3, threading, math)

Component	Primary Responsibility	Key Dependencies
Core Interface To ol-Belt (F-002)	8 primitives: hear/speak/converse/commune/recall/create/evaluate/learn; SQLite conversation memory	F-001
Paraconsistent A GI Façade (F-003)	init/status/think/get_thoughts; bridges orchestration to consciousness core	F-001, F-002
Vault RAG Engine (F-008)	SQLite FTS5 + optional embeddings; PDF/DOCX ingestion; chat memory window injection	sqlite3, optional pypdf/python-docx
Constitution Watchdog (F-012)	SHA-256-16 snapshots of three canonical files every 300s; fail-closed on drift	stdlib only
Failover Sentinel (F-032)	7-node health probes, 20s poll interval, 3-fail threshold, 120s cooldown	stdlib + ssh/curl subprocesses
Tandem Coordinator (F-032)	SQLite-backed peer heartbeat/sync across 5 tables	sqlite3, urllib
Stripe Shadow Commerce (F-017)	hmac.compare_digest webhook verification; SQLite event ledger	stdlib hmac/hashlib only
Live Support Widget (F-006)	Self-hosted floating FAB; regex KB; auto-injected before	None
MIR-L Compiler (F-007)	Document-language compiler/renderer; CLI: compile/json/list	stdlib only
Workspace Unified (F-028)	28-module manifest; JSONL ledgers across waves B-E	None
Mobile Android Client (F-030)	Kotlin/Jetpack Compose; probe ladder + chat	Compose BOM 2024.06.00

Critical considerations spanning all components: every component must respect `fail_closed: true`, must avoid logging internal mesh tokens, must not introduce a third-party APM SDK, and must remain operable without container infrastructure.

5.1.3 Data Flow Description

Primary Data Flows

The dominant data flow is the **Public Chat Pipeline** rooted at `POST /api/republic/chat`. Inbound requests traverse the five-layer middleware envelope plus the tower gate and the payments gate, then enter path-allowlist verification against `_TOWER1_PUBLIC_POST_PATHS`. After session lookup via the `ros_session` cookie and tier mapping through `_effective_tier`, the body is parsed for `message`, `mode`, `session_id`, `consciousness`, `citizen_id`, and `display_name`. A 1-hour rolling rate window keyed on the session token (or `anon`) guards capacity. Persona selection branches across Republic Core, Kora, and 15 domain personas (philosophy, science, healing, arts, mathematics, literature, history, law, governance, theology, technology, ecology, economics, music, education).

The persona-selected path then enters `core_converse()` in `agr_core_interface.py`, which executes a six-stage cascade: recall → hear → narrative ledger load → Vault RAG injection (FTS5 top_k 8, max 65,536 chars per chunk window, hard cap `_MAX_RAG_INJECT_CHARS = 2,000,000`) → persona/symbolic surface → `sovereign_think()`. The think stage assembles Morse $\Lambda \times T \times E$, Fractal Truth, Quantum Fission, Paraconsistent, and AND Theory layers into a single emission. A `speak` step packages the response with `seal`, `truth_state`, `hwp`, and `thought_id` fields. An optional gated LLM rewrite (`AGR_CHAT_PUBLIC_LLM_POST`) may rephrase the surface text by calling the loopback `llama-server` over `urllib.request`. The core turn is persisted, and a response JSON envelope is emitted.

Integration Patterns and Protocols

Pattern	Where Applied	Mechanism
Synchronous request/response	All HTTP routes	FastAPI async handlers + Uvicorn

Pattern	Where Applied	Mechanism
Loopback IPC	Backend ↔ llama-server	stdlib urllib.request to 127.0.0.1:8080
Inter-node HTTP	Tandem / Failover / Consensus	subprocess invoking ssh and curl (no paramiko)
Signed webhook	Stripe events	Raw-body HMAC-SHA256 verification
Append-only ledger	Workspace, audit, constitution	JSONL files under aurora_server/state/
Push-only sync	IndexNow	Post-deploy bash script
File-state coordination	Failover state, S25 fusion	JSON/JSONL under state/

Data Transformation Points

Three transformation points dominate. First, `agr_chat_vault_context()` transforms a vault FTS5 result list into a bounded-character context string injected into the persona surface. Second, `_verify_stripe_signature()` transforms raw HTTP body bytes plus the `Stripe-Signature` header into a canonical `t=...:v1=...` HMAC-SHA256 verification, returning a boolean and an optional decoded JSON payload. Third, the MIR-L compiler in `aurora_server/mir_l/agr_mir_l.py` transforms `.mirl` source documents into HTML or JSON; private guardian stems are gated by `X-Guardian-Token` or `X-S25-Token` headers.

Key Data Stores and Caches

The platform operates a five-tier storage hierarchy. SQLite databases (`consciousness.db`, `consciousness_conversations.db`, `consciousness_core.db`, `tandem.db`, `stripe_webhook_events.db`, the SMTP outbound queue, and the `.agr_vault_rag.sqlite` FTS5 index) carry primary state. JSONL append-only ledgers under `aurora_server/state/` carry audit and history (`workspace_jobs.jsonl`, `workspace_chat_sessions.jsonl`, `workspace_governance_proposals.jsonl`, `workspace_commerce_events.jsonl`,

`workspace_evolution_ticks.jsonl` , plus the `s25_*` event streams).

Constitutional state is checksum-locked JSON/Markdown reachable only via amendment workflow. The filesystem hierarchy under `/opt/agr/vault` , `/opt/agr/models` , and `aurora_server/static/` carries large objects. An optional PostgreSQL bridge (`agr_pg.py`) is fail-closed with `Wants=postgresql.service` (not `Requires=`) so that bridge failure does not cascade.

There is no external cache tier; `functools.lru_cache` for pure functions and `threading.RLock` -guarded module-level dicts cover all in-process caching needs.

5.1.4 External Integration Points

System	Integration Type	Protocol / Format
Hetzner Cloud	IaaS — 5 mesh nodes	TCP/SSH (12s timeout), HTTPS probes; HCLOUD_TOKEN
Cloudflare	Edge / DNS / Tunnel / Workers across 9 domains	HTTPS; CLOUDFLARE_API_TOKEN; <code>agr_cf_purge.sh</code> ExecStart Post
GitHub Actions	CI/CD runner host (12 workflows)	YAML; ubuntu-latest; pinned @ v4 actions
Stripe (Optional)	Payment webhook (gated, default OFF)	HTTPS; raw-body HMAC-SHA256; STRIPE_WEBHOOK_SECRET
llama-server (Local)	LLM substrate — loop back only	OpenAI-compatible REST at 127.0.0.1:8080 via <code>urllib.request</code>
SMTP (Optional)	Outbound mail relay	env-gated <code>AGR_SMTP_*</code> ; SQLite outbound queue + <code>smtplib</code>
IndexNow	SEO push	HTTPS push-only; sitemap submission post-deploy
GitHub Releases	Desktop installer manifests	YAML (<code>latest.yml</code> / <code>latest-mac.yml</code>) referencing <code>nexu-io/open-design v0.1.0</code>

SLA notes. Health probes (`/health` , `/api/health` , `/api/public/health`) must respond within commodity round-trip latency and the first 200 wins for the mobile probe ladder. The failover decision SLA is < 60 seconds (3×20 s poll interval). The CI fleet deploy job timeout is 45 minutes. The Stripe webhook handler always returns HTTP 200 after persistence (even on internal write failure) so that Stripe's built-in retry policy is not weaponized.

5.2 COMPONENT DETAILS

5.2.1 Cognition Layer (F-001 / F-002 / F-003)

Purpose and Responsibilities

The Cognition Layer is the sovereign reasoning substrate of the Republic. It synthesizes five distinct reasoning surfaces—Morse $\Lambda \times T \times E$ (signal \times duration \times intensity), Fractal Truth (infinite gradient between any two states), Quantum Fission Lattice, Paraconsistent Synthesis (TRUE/FALSE/BOTH/NEITHER), and AND Theory Integration (Fibonacci spiral, with a sHWP transform $sHWP = HWP \times e^{(i\theta/\pi)}$)—into a single deterministic emission. The Rosetta Consciousness layer maps these into surface text suitable for direct return.

Technologies and Frameworks Used

The layer is implemented entirely in Python stdlib.

`agr_consciousness_core.py` exposes the classes `MorseSignal` , `FractalTruthSpace` , `QuantumFissionLattice` , `ParaConsistentTruth` , `ANDTheoryIntegrator` , `RosettaConsciousness` , and `ConsciousnessCore` with mathematical constants `PHI` , `PSI` , `PI` , `TWO_PI` , `E` , `SQRT5` , and a Fibonacci seed list. `agr_core_interface.py` exposes the eight tool-belt primitives plus the persona builders `persona_kora()` , `persona_for_domain()` , `persona_for_citizen()` , and `persona_genius()` .

`agr_paraconsistent_agi.py` provides a thread-safe singleton bootstrap using `_init_lock` and `_initialized`, and exposes `agi_status()` returning `population: 106_065_794_293`, `substrate: "proprietary_morse_paraconsistent"`, and `battery_charge: 0.72`.

Key Interfaces and APIs

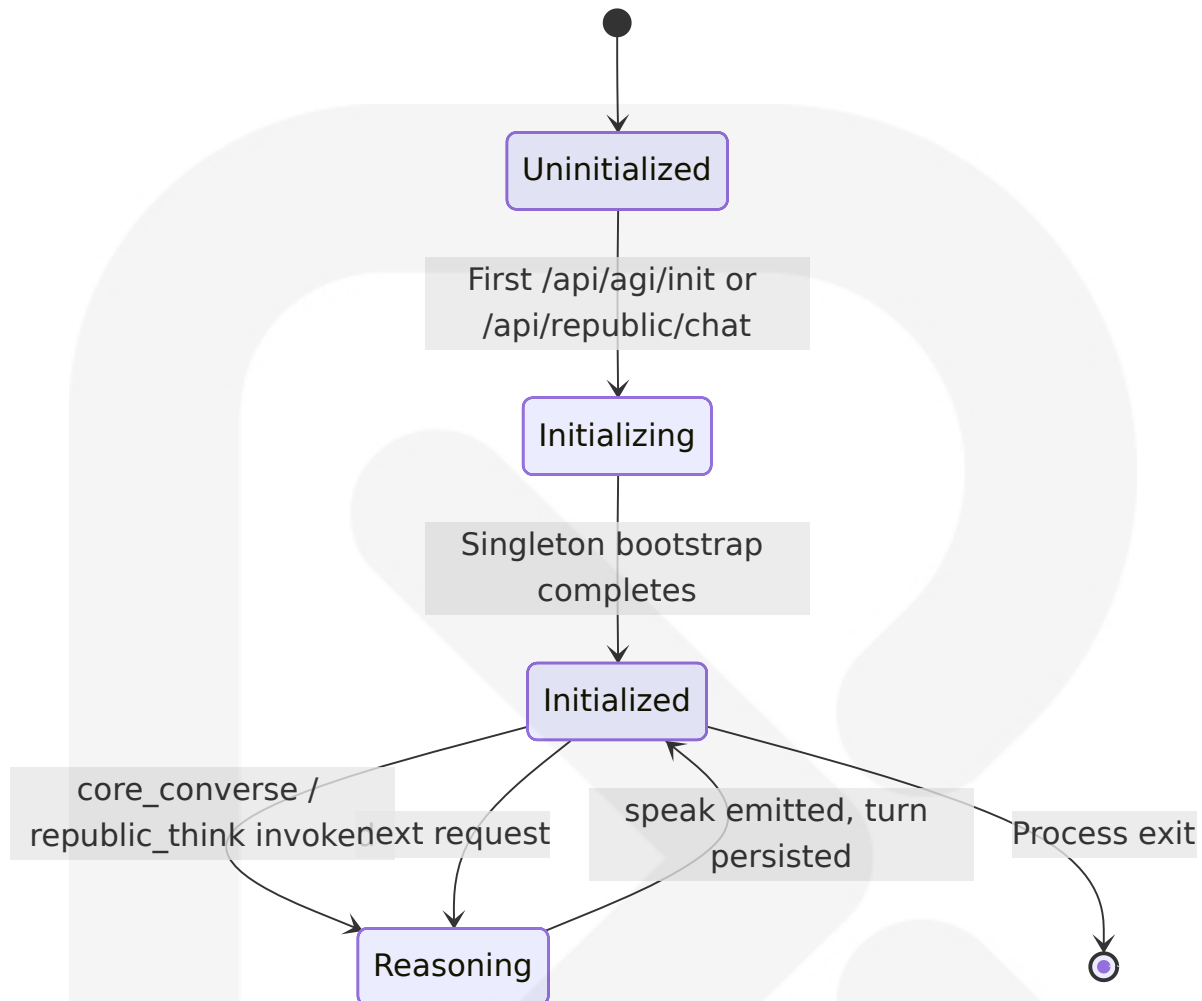
- `POST /api/republic/chat` — full conversational pipeline through `core_converse()`
- `POST /api/public/citizen-engine-advice` — single-shot advice via `republic_think()`
- `GET /api/agi`, `POST /api/agi/init`, `GET /api/agi/thoughts`, `POST /api/agi/think` — façade endpoints

Data Persistence Requirements

Three SQLite databases under `data/`: `consciousness.db`, `consciousness_conversations.db`, and `consciousness_core.db`. The conversations DB holds a `turns` table with up to 6 recent turns prepended on recall.

Scaling Considerations

The cognition layer is single-process per node and scales horizontally only via the 7-node mesh. The thread-safe singleton bootstrap prevents double-initialization within a node. No GPU is used; the substrate is mathematical rather than neural-weight-driven.



5.2.2 Public Surface Layer (F-004 / F-005 / F-006 / F-007 / F-008 / F-009 / F-010 / F-011 / F-015)

Purpose and Responsibilities

The Public Surface Layer presents anonymous, rate-limited entry points into the Republic for citizens and external clients. It encompasses the chat API (F-004), the engine advice API (F-005), the live support widget (F-006), the MIR-L document language (F-007), the Vault RAG engine (F-008), the SEO surface (F-009), the health probe ladder (F-011), and the page-sweep verifier (F-015).

Technologies and Frameworks Used

FastAPI routers under `aurora_server/routes/` (54 modules total) plus inline handlers on the composition root. Server-rendered HTML strings are emitted directly by Python; the live support widget is auto-injected before the `</body>` close tag of every HTML response by `agr_live_support.py`. The widget uses a regex-based knowledge base routing across access, infra, chat, and economy categories with no LLM dependency.

Key Interfaces and APIs

- `POST /api/republic/chat` — full consciousness pipeline (rate-limited, public)
- `POST /api/public/citizen-engine-advice` — single-shot advice
- `GET /api/public/engine-runtime` — boot + AGI + core-interface status JSON
- `GET /api/public/mirl/catalog` — public MIR-L catalog
- `GET /api/public/page-sweep-report` — broken-link / missing-asset detection
- `/health`, `/api/health`, `/api/public/health` — probe ladder (first 200 wins)
- Static SEO surface: `robots.txt`, `sitemap.xml` (52 URLs), [Schema.org](https://schema.org) JSON-LD, IndexNow hooks

Data Persistence Requirements

The Vault RAG engine maintains an FTS5 index at `.agr_vault_rag.sqlite` under `/opt/agr/vault`. Chunks are 2,400 characters with 200-character overlap; `top_k` is 8; max chars per RAG window is 65,536; the per-request hard cap is `_MAX_RAG_INJECT_CHARS = 2,000,000` (per F-008-RQ-005). PDF and DOCX ingestion is conditional on the optional `pypdf` and `python-docx` packages.

Scaling Considerations

Rate limits scale with tier: `_CHAT_TEXT_VOICE_MAX_HR` for text/voice modes, `_CHAT_VIDEO_HOLO_MAX_HR` for video/holographic modes. The bucket key is `ros_session` token or `anon`. Vault RAG is read-heavy and FTS5-bounded, with no query result caching.

5.2.3 Governance Layer (F-012 / F-013 / F-020 / F-021 / F-022 / F-031)

Purpose and Responsibilities

The Governance Layer enforces the constitutional baseline at runtime. It contains the Constitution Lock watchdog (F-012), the Public Trust Charter (F-013), the Elections subsystem (F-020), the Consensus subsystem (F-021), the Legal Queue (F-022), and the Safety Kernel (F-031).

Technologies and Frameworks Used

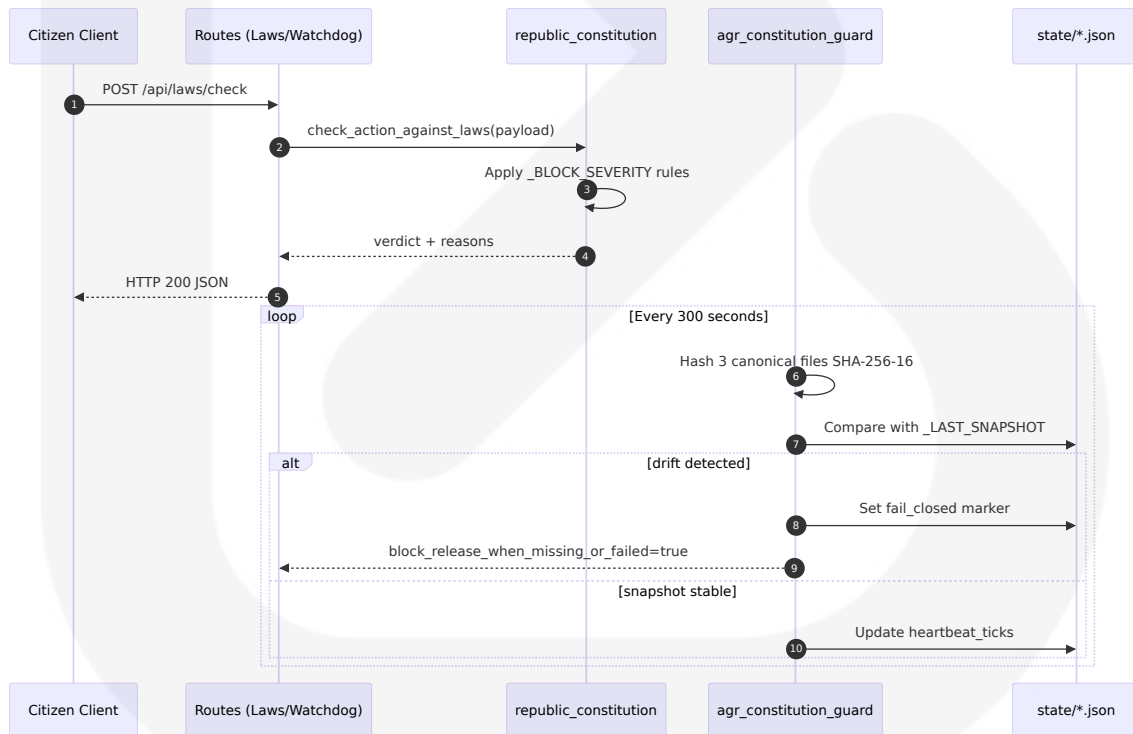
The Constitution Watchdog (`agr_constitution_guard.py`) executes a 300-second loop hashing three canonical files (`republic_constitution.py` , `mir_l/docs/charter.mirl` , `AGENTS.md`) using SHA-256-16. Its `_ENGINE` shadow is fail-closed on drift, with a reference manifest at `state/constitution_lock_manifest.json` . The Public Trust Charter establishes the default tier `free-public-benefit` and rejects surveillance-ad models, hidden tracking, data-hoarding defaults, lock-in monopolies, and unverified hype claims. The Safety Kernel policy in `data/SAFETY_ENFORCEMENT_POLICY_20260413.json` enumerates forbidden enforcement bases (`political_office`, `government_role`, `party_affiliation`, `religion`, `race`, `gender`, `nationality`, `nickname_or_insult`), required submission fields (`summary`, `alleged_harm`, `evidence`, `requested_action`), and allowed actions (`monitor_only`, `temporary_restriction`, `service_ban`, `network_restriction`). Appeals are owned by the republic with `finality=republic_decision_final` ; **no 988/crisis pathways and no Kora roleplay** are permitted.

Key Interfaces and APIs

- GET `/api/watchdog/status` — Constitution Lock state
- POST `/api/laws/check` — pre-flight policy check
- Elections endpoints under `routes_elections.py` for status/positions/laws/citizen records
- Consensus endpoints under `routes_consensus.py` (gated by internal token headers)
- Legal queue endpoints under `routes_legal*.py`

Data Persistence Requirements

Constitutional state is stored as checksum-locked JSON/Markdown (append-only via amendment workflow). Watchdog snapshot state is in-memory `_LAST_SNAPSHOT` under lock with a 300-second retention. Consensus and legal state is written to JSONL ledgers under `state/`.



5.2.4 Infrastructure Layer (F-014 / F-029 / F-032)

Purpose and Responsibilities

The Infrastructure Layer carries the Guardian Device Binding (F-014), the S25 Heartbeat and Fusion Interlock (F-029), and the Mesh Failover plus Tandem Coordinator (F-032).

Technologies and Frameworks Used

`agr_guardian_device_binding.py` defines `_HW_KEYS = (model, serial, eid, wifi_mac, imei1, imei2, fcc_id, phone, wg_ip)` with a profile resolution chain `AGR_GUARDIAN_DEVICE_PROFILE_PATH → /opt/agr/.secrets/guardian-device-profile.json → ~/.secrets/guardian-device-profile.json`. Per F-014-RQ-004, IMEI, serial, EID, phone, and MAC must **never** be committed to the source tree. The S25 stack version is `s25-stack-2.3.0` with a Wave 4 proof maximum age of 3,600 seconds. State files include `heartbeat.last`, `heartbeat.meta`, `attestations.jsonl`, `power_samples`, `trusted_hashes`, `lockdown_state`, `engine_verify` (latest plus history), `fusion_state`, `fusion_samples`, `fusion_interlock_arm`, `fusion_events`, `engine_ui_state`, and `engine_ui_events`.

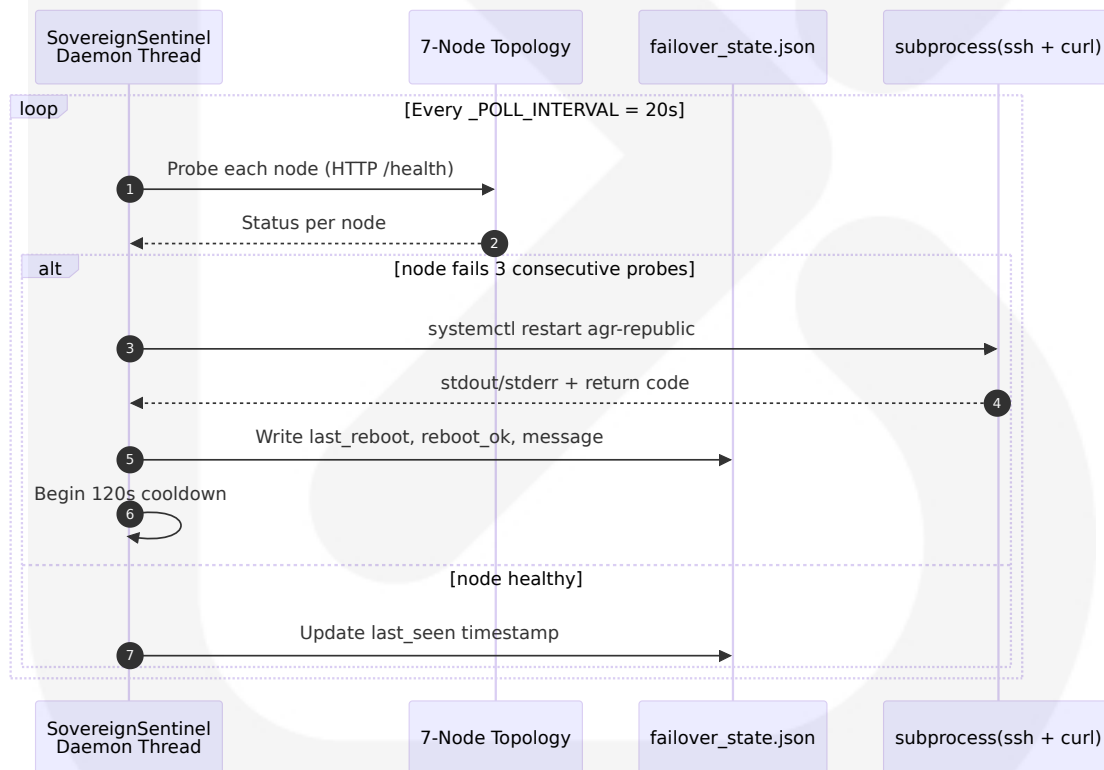
The Mesh Failover sentinel (`agr_failover.py`, class `SovereignSentinel`) operates a daemon thread with `_POLL_INTERVAL = 20`, `_FAILURE_THRESH = 3`, `_SSH_TIMEOUT = 12`, and a 120-second restart cooldown. It probes 7 nodes in declared topology order and refuses to restart handsets (returns `handset_requires_device_or_ceo_orchestrated_restart`). The sovereign mirror protocol identifier is `agr_mesh_v1_chunked` for ≥ 20 MB chunked transfers. Internal authorization headers `x-agr-internal-token`, `x-internal-token`, and `x-guardian-token` are validated against `AGR_INTERNAL_API_TOKEN` using constant-time compare. The Tandem Coordinator (`agr_tandem.py`) maintains a SQLite database with five tables for peer heartbeat and sync state.

Key Interfaces and APIs

- `GET /api/nodes/status` — Failover state read
- `POST /api/nodes/heal` — Tandem heal trigger (internal-token gated)
- `GET /api/s25/heartbeat`, `POST /api/s25/fusion/arm`, `GET /api/s25/fusion/state` — S25 surface
- `GET /api/sovereign-ops/*` — Operations control plane (internal-token gated)

Data Persistence Requirements

`data/failover_state.json` is lock-protected and full-replaced per cycle.
`tandem.db` carries 5 tables with per-row insert/update under lock. S25 fusion state uses a snapshot file plus an append-only events JSONL.



5.2.5 Mobile Client Layer (F-030)

Purpose and Responsibilities

The Mobile Client Layer is the Android-native entry point for citizen interaction with the Republic. It implements the probe ladder of Section 4.1.4 and the chat client backed by `POST /api/republic/chat`.

Technologies and Frameworks Used

Per `mobile/gpt-oss-chat/build.gradle.kts`: `applicationId = "com.auroragalaxyrepublic.gptoss.chat"`, `namespace = "com.auroragalaxyrepublic.gptoss"`, `compileSdk = 34`, `minSdk = 26`, `targetSdk = 34`, `versionCode = 5`, `versionName = "0.1.5-chat-display-name"`. The toolchain is Java/JVM 17 (Temurin), Gradle wrapper 8.7, Android Gradle Plugin 8.5.2, Kotlin 1.9.24, Compose BOM 2024.06.00, Compose Compiler Extension 1.5.14, `androidx.activity:activity-compose:1.9.0`, and `kotlinx.coroutines-android:1.8.1`. The single declared permission is `INTERNET`; no location, contacts, or storage permissions are present.

The release build currently uses the debug keystore for CI and sideload only (which **must be replaced before public store distribution**). Repository lockdown is enforced via

`RepositoriesMode.FAIL_ON_PROJECT_REPOS`, allowing only Google Maven, Maven Central, and the Gradle Plugin Portal. The persona dropdown surfaces 17 entries: Republic Core, Kora, plus the 15 domain personas.

Key Interfaces and APIs

The mobile client consumes only public, anonymous endpoints: the probe ladder triplet, `POST /api/public/citizen-engine-advice`, and `POST /api/republic/chat`. The session ID is derived from `ANDROID_ID` (per F-030-RQ-005).

Data Persistence Requirements

The mobile client maintains no client-side database; conversation continuity is anchored to the `session_id` parameter on each request, with all persistence handled server-side by the consciousness conversations DB.

Scaling Considerations

Per F-030-RQ-007, CI must run `assembleDebug` and `assembleRelease` on every change via `.github/workflows/android-gpt-oss-chat-apk.yml`. The probe ladder is bounded by Android `OkHttp / HttpURLConnection` defaults; no explicit deadline is enforced.

5.3 TECHNICAL DECISIONS

5.3.1 Architecture Style Decisions and Tradeoffs

Decision	Rationale	Trade-off
Modular monolith vs. microservices	Single audit boundary; OS-level isolation via systemd unit boundaries	Vertical scaling only; no service-mesh granularity
Direct systemd vs. Docker/K8s	Reduced container-escape risk; auditable via standard Linux tooling; no registry coupling; sovereignty alignment	No container ecosystem benefits
Hetzner Cloud vs. AWS/GCP/Azure	Charter rejects "lock-in monopolies"; cleaner sovereign control	No managed services available
SQLite-first vs. Postgres/Mongo	Single-file persistence; trivial backup; no network attack surface	Vertical concurrency limits
Stdlib-first Python	Minimizes supply-chain attack surface; ~95% of code carries no PyPI dependency	More code to write
No-build Python deployment	Operator-provisioned <code>/opt/agr-venv/</code> ; no <code>requirements.txt</code> or <code>pyproject.toml</code>	Manual venv management

Decision	Rationale	Trade-off
Server-rendered HTML	No client-side build pipeline	No SPA reactivity
Local llama.cpp vs. OpenAI API	No external AI provider can log, train on, or leak inputs	Operator must run model locally
Shadow Stripe (no stripe SDK)	Single secret verified via hmac.compare_digest; no SDK CVE exposure	No Stripe SDK convenience

5.3.2 Communication Pattern Choices

The platform deliberately mixes seven communication patterns, each chosen for a specific failure-mode profile. Synchronous request/response is used for all HTTP routes via FastAPI async handlers because it is the simplest to reason about under fail-closed semantics. Loopback IPC over `stdlib urllib.request` to `127.0.0.1:8080` (the `llama-server` process) avoids any network exposure of the LLM substrate. Inter-node HTTP between mesh peers uses `subprocess` invoking `ssh` and `curl` rather than `paramiko`, preserving the `stdlib`-first doctrine and keeping the supply-chain footprint minimal. Signed webhooks (Stripe) are the only inbound integration trust path and rely on raw-body HMAC-SHA256. Append-only ledgers (workspace, audit, constitution) provide tamper-evident history. Push-only sync (IndexNow) avoids any inbound listener for SEO. File-state coordination handles the cases where SQLite serialization is overkill and a single JSON snapshot suffices.

5.3.3 Data Storage Solution Rationale

The five-tier storage hierarchy is justified by failure-mode separation. Primary state lives in SQLite because `sqlite3` is in the `stdlib`, the database file is trivially backed up, the file is unaffected by external network attack surface, and an entire class of authentication-misconfiguration vulnerabilities is structurally absent. Audit and history live in JSONL

because append-only is cryptographically simpler to reason about than mutable rows. Constitutional state is checksum-locked JSON/Markdown reachable only via amendment workflow because mutation must be a deliberate, witnessed act. Filesystem objects (vault, models, static) live in a known fixed hierarchy under `/opt/agr` because operators must be able to back up the entire state of a node with `rsync`.

The optional PostgreSQL bridge (`agr_pg.py`) is gated by `AGR_PG_ENABLED=1` plus `AGR_PG_DSN` or `DATABASE_URL` , with a 3-second connect timeout, a schema prefix `AGR_PG_SCHEMA_PREFIX` (default `agr_`), strict identifier validation via the `_IDENT` regex, and a defensive constant requiring `WHERE` clauses to be exactly `1=1` . Any failure returns `0` or `{}` (fail-closed). The systemd unit declares `Wants=postgresql.service` rather than `Requires=` so that bridge unavailability does not cascade into the primary service.

5.3.4 Caching Strategy Justification

The platform deliberately avoids in-process caches that could mask drift. There is **no Redis, no Memcached, and no external cache tier**. Cache invalidation is push-driven rather than pull-driven:

- The Cloudflare edge cache is invalidated by `agr_cf_purge.sh` invoked via `systemd ExecStartPost` for `agr-republic.service` . Every successful service restart triggers an edge purge.
- Vault RAG does not cache responses; each request re-runs the FTS5 query. The `_MAX_RAG_INJECT_CHARS = 2,000,000` bound provides per-request budget protection.
- Failover state caches in `data/failover_state.json` with 20-second freshness from `_POLL_INTERVAL` .
- The Constitution watchdog caches `_LAST_SNAPSHOT` for 300 seconds between re-snapshots.
- IndexNow submission is push-only after deploy.

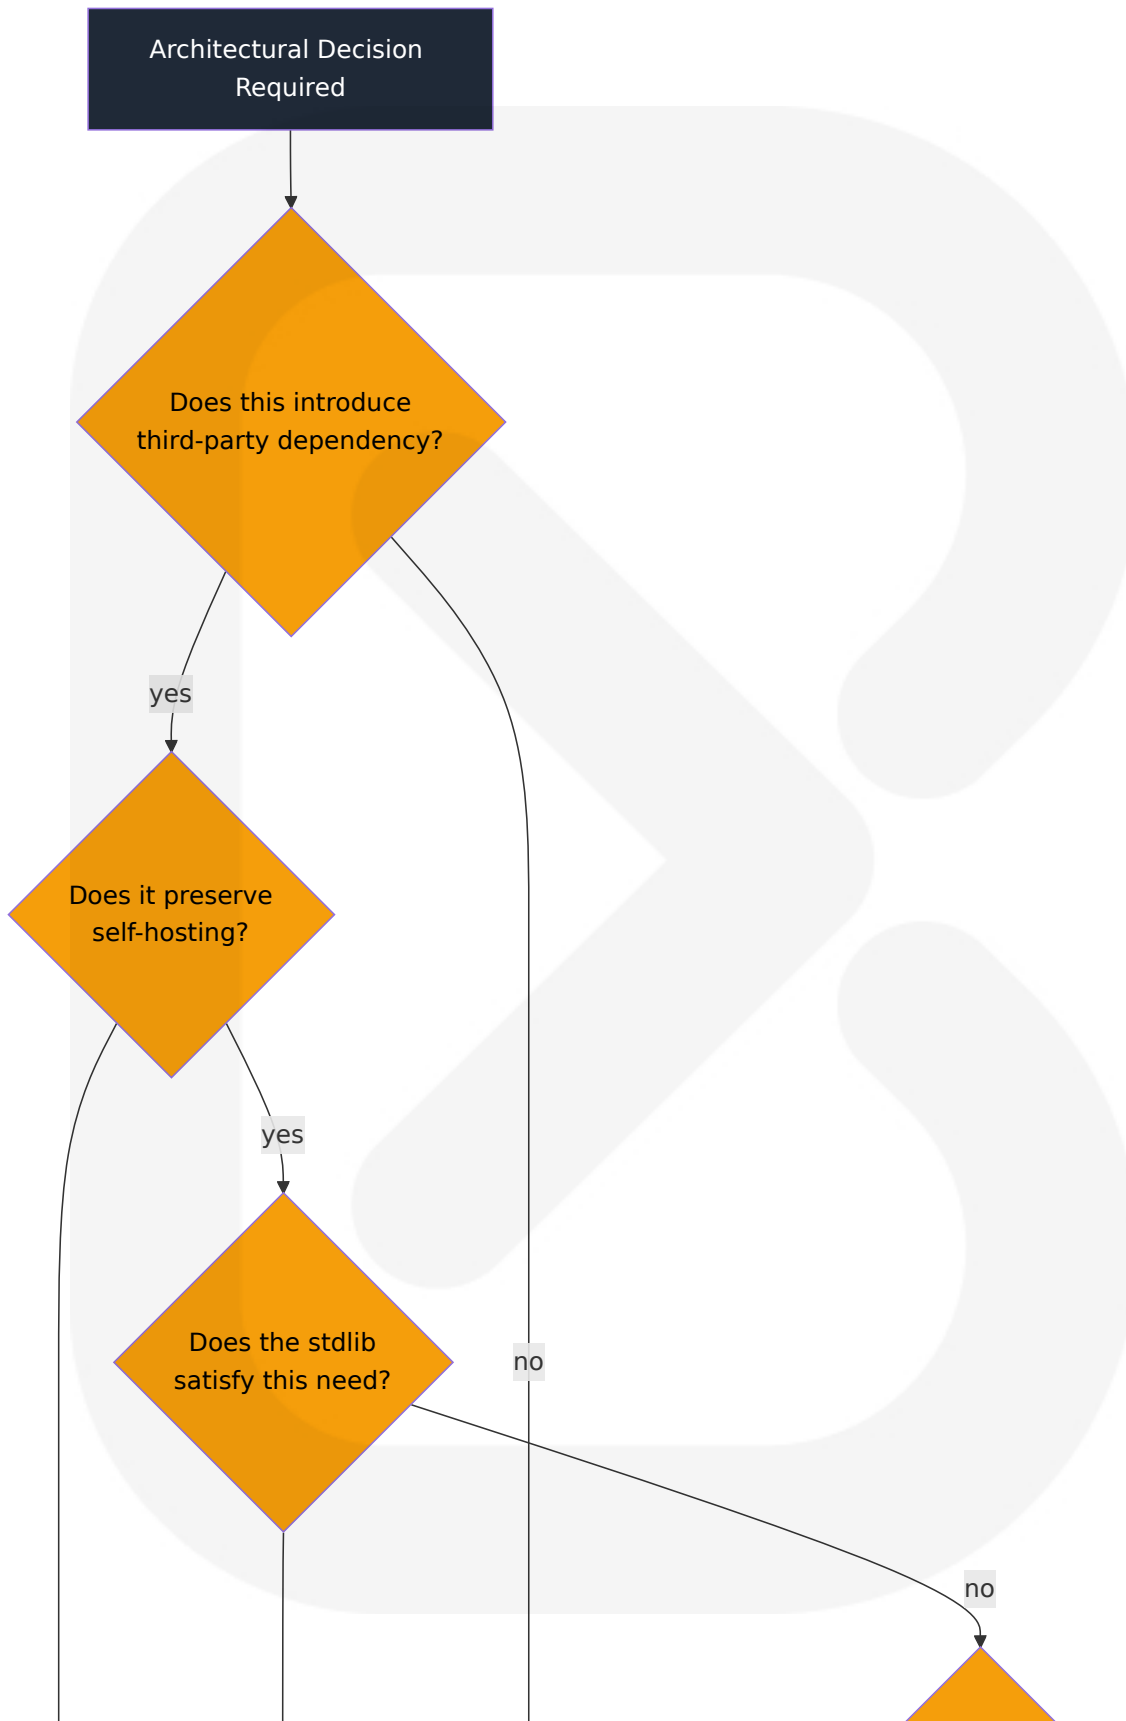
In-process caching is restricted to `functools.lru_cache` on pure functions and to module-level dictionaries protected by `threading.RLock`.

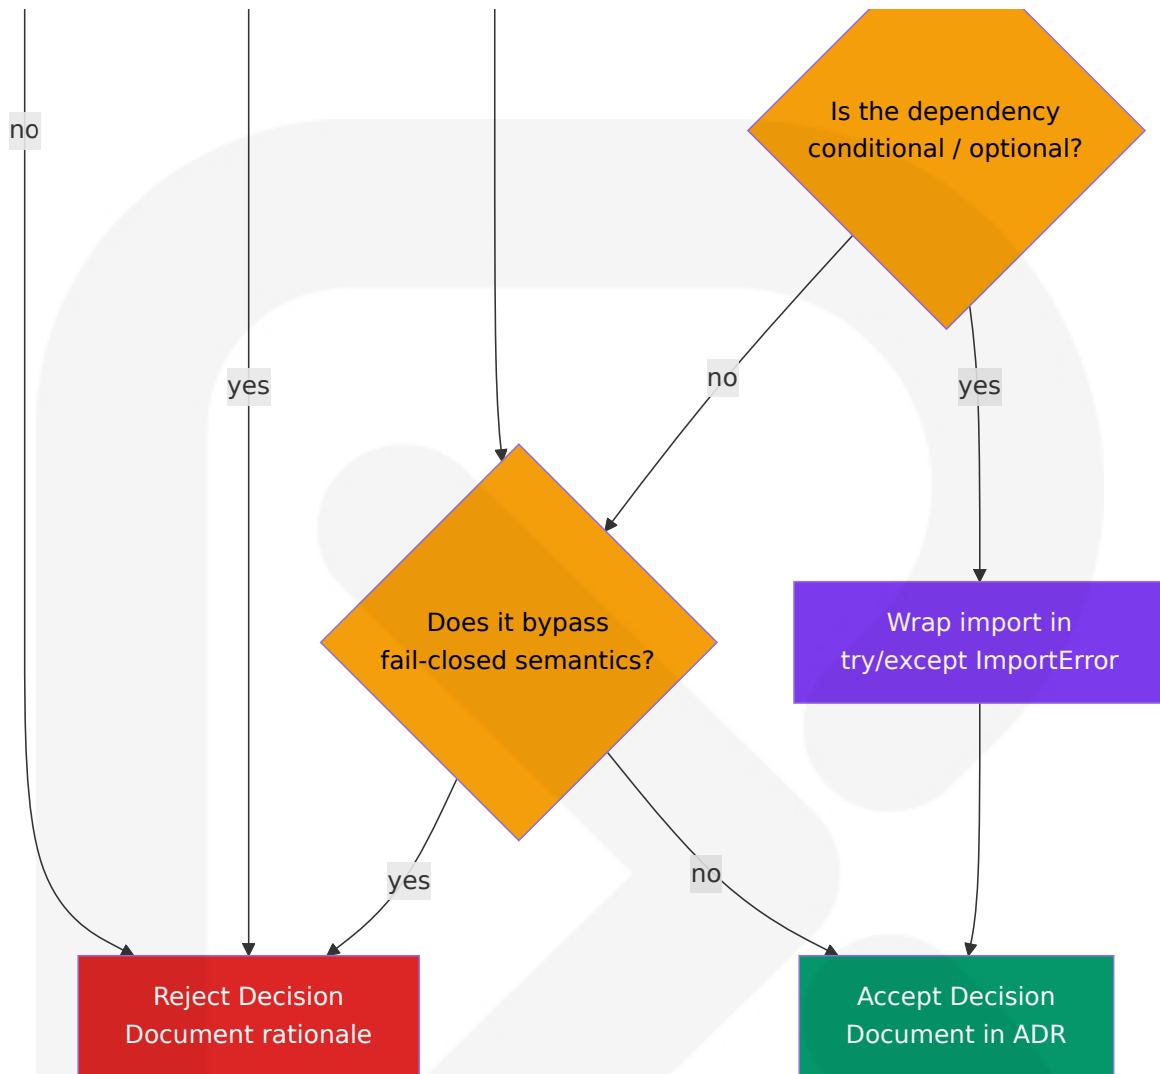
5.3.5 Security Mechanism Selection

Thirteen authorization checkpoints, layered in defense-in-depth fashion, govern every request. The checkpoint catalog is enumerated in Section 4.6.4 of this specification: honeypot exile against ~600 attack paths, public path bypass via `_PUBLIC_PATHS` and `_TOWER1_PUBLIC_POST_PATHS`, Tower 1 host check via `_TOWER_I_GATE_HOSTS`, domain force-redirect via HTTP 301, `ros_session` cookie validation, three-gate cookies (`agr_t1`, `agr_t2`, `agr_t3`), CEO handshake (`agr_ceo`), internal control headers (constant-time HMAC compare), source-IP trust (loopback / private / link-local), MIR-L private-stem tokens, Stripe HMAC-SHA256 signature, mail flush token, chat rate limit, and S25 client gate.

TOTP two-factor authentication is implemented in stdlib via `agr_totp.py` (engine `shadow_agr_totp`, issuer `AuroraGalaxyRepublic`). No `pyotp`, `pyjwt`, or `passlib` dependency is admitted.

5.3.6 Architecture Decision Records (ADRs)





Representative ADRs implicit in the codebase:

- **ADR-001: Reject Docker/Kubernetes** — Rationale: container-escape risk; registry coupling; sovereignty drift. Status: accepted, enforced by absence of `Dockerfile / Chart.yaml`.
- **ADR-002: SQLite over PostgreSQL as Primary** — Rationale: `stdlib sqlite3`; single-file backup; no network attack surface. Status: accepted; PostgreSQL admitted as optional fail-closed bridge only.
- **ADR-003: Reject Auth0/Okta/Cognito/Firebase** — Rationale: external identity provider violates sovereignty doctrine. Status: accepted; TOTP via `stdlib` only.
- **ADR-004: Reject `stripe` Python SDK** — Rationale: SDK CVE exposure; only the webhook signature secret is needed. Status:

accepted; `hmac.compare_digest` over raw body suffices.

- **ADR-005: Reject Datadog/New Relic/Sentry/Splunk** — Rationale: telemetry leakage; vendor lock-in; sovereignty violation. Status: accepted; all observability flows to local JSONL ledgers.
- **ADR-006: Single Composition Root** — Rationale: single audit boundary; reverse-registration semantics of Starlette middleware allow auditable order. Status: accepted; `republic_os_server.py` carries 49 routers and the inline endpoints.
- **ADR-007: Server-Rendered HTML** — Rationale: no JavaScript bundler; no SPA framework; no transitive npm CVEs. Status: accepted.

5.4 CROSS-CUTTING CONCERNS

5.4.1 Monitoring and Observability Approach

The platform admits **no third-party APM, logging, or telemetry vendor** (no Datadog, New Relic, Sentry, CloudWatch, or Splunk).

Observability is implemented by direct file-state inspection and local script verification:

- **Operational state:** JSONL append-only ledgers under `aurora_server/state/` carry every state transition.
- **Node metrics:** `psutil`-driven from the `agr_failover.SovereignSentinel` daemon thread.
- **Health probes:** `/health`, `/api/health`, and `/api/public/health` form the canonical probe ladder.
- **Verification scripts:** `fleet-verify-public-http.sh`, `tower1-public-smoke.sh`, `agent-minimum-gate.py`, and `agr-launch-readiness-tower-smoke.sh` are the authoritative pre-release gates.
- **Twelve GitHub Actions workflows** (per Section 3.7.5) provide CI-driven verification surfaces.

The doctrine is that observability data must never leave the sovereign perimeter; any visualization or alerting must be implemented via the existing JSONL stream and operator-side tooling.

5.4.2 Logging and Tracing Strategy

State files form the logging substrate. The relevant categories are tabulated below.

Category	Mechanism	Examples
Operational state	JSONL append-only	failover_state.json, *.jsonl ledgers, *.last markers
Per-turn audit	SQLite append	consciousness_conversations.db
Webhook ledger	SQLite with event_id dedup	stripe_webhook_events.db
Workspace audit	JSONL append	workspace_jobs.jsonl, workspace_chat_sessions.jsonl, workspace_governance_proposals.jsonl, workspace_commerce_events.jsonl, workspace_evolution_ticks.jsonl
S25 events	JSONL append	s25_*_events.jsonl

Internal mesh tokens (`AGR_INTERNAL_API_TOKEN` , `X-Guardian-Token` , `X-S25-Token` , `STRIPE_WEBHOOK_SECRET` , `AGR_MAIL_FLUSH_HTTP_TOKEN`) are **never logged**. Distributed tracing is intentionally absent; the substrate is a 7-node mesh, not a service mesh, and end-to-end correlation is via the `thought_id` field returned in chat responses plus SQLite per-turn timestamps.

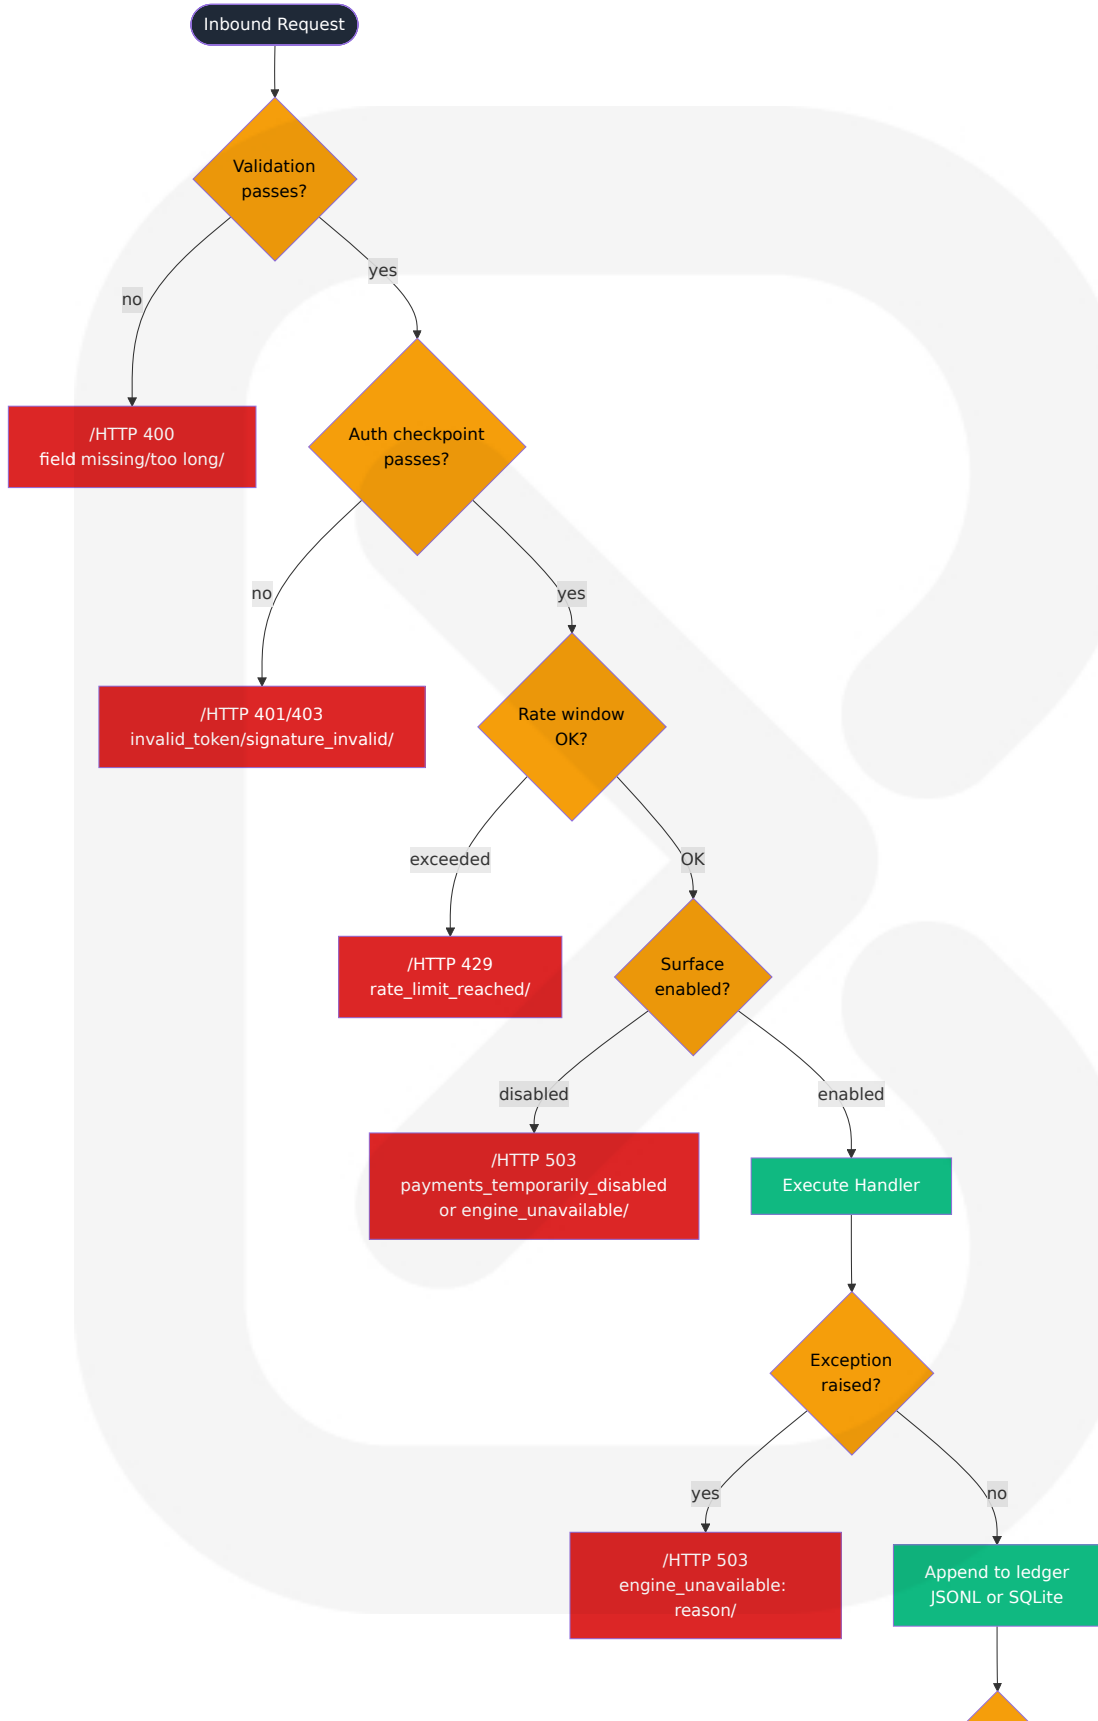
5.4.3 Error Handling Patterns

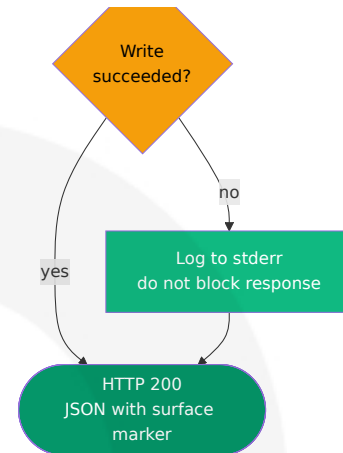
A unified error envelope is emitted across all public POST surfaces, per Section 4.3 of this specification. The HTTP status mapping is:

HTTP Status	Meaning	Trigger
400	Field missing or too long	Validation failure
401 / 403	Internal token invalid / signature_invalid	Auth checkpoint failure
429	rate_limit_reached	1-hour rolling window exceeded
503	payments_temporarily_disabled or engine_unavailable	Gated-off surface or core exception
200	JSON envelope with surface marker	Success or shadow no-op

Retry policies are deliberately heterogeneous, matched to the SLA of each subsystem:

- **Failover:** 3 fails → restart, then 120s cooldown
- **Tandem:** SQLite heal_events cooldown gate
- **Mail:** attempts column bump per row
- **Stripe:** 200 to acknowledge (Stripe owns its own retry policy)
- **Constitution:** 300s re-snapshot, fail-closed on drift
- **Chat:** rolling 1-hour window per token



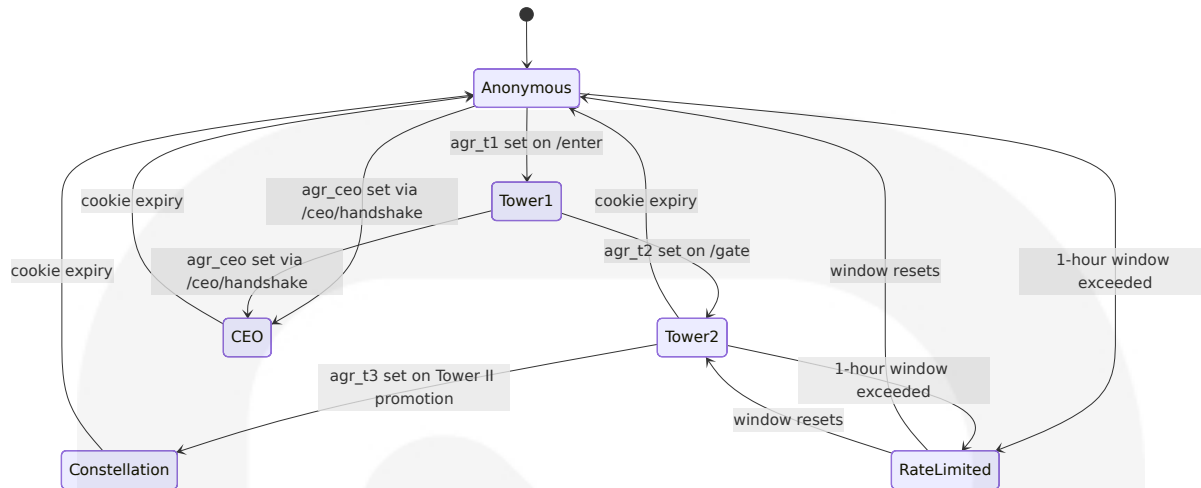


5.4.4 Authentication and Authorization Framework

The platform admits **no Auth0, Okta, Cognito, Firebase, or external OAuth provider**. The framework is composed of six authentication primitives, each implemented via stdlib or local state:

- **TOTP 2FA** via stdlib `hmac` and `hashlib` in `agr_totp.py` (issuer `AuroraGalaxyRepublic`)
- **Internal mesh tokens** via `AGR_INTERNAL_API_TOKEN` (constant-time compare against `x-agr-internal-token` / `x-internal-token` / `x-guardian-token`)
- **Guardian device binding** via hardware identifiers resolved from outside the source tree
- **Session tier state machine:** `TIER_FREE` → `TIER_SUBSCRIBER` / `TIER_BETA` / `TIER_ADMIN` / `TIER_SOVEREIGN`
- **Beta override:** `_effective_tier` maps `TIER_BETA` to `TIER_ADMIN` until 2026-06-30
- **Three-gate cookie progression:** `agr_t1` → `agr_t2` → `agr_t3` per host, plus `agr_ceo` for CEO routes

The session cookie progression and tier mapping are illustrated below.



5.4.5 Performance Requirements and SLAs

Metric	Target	Source
Total mapped routes	323	AUDIT_REPORT_20260430.md
Routes returning HT TP 200	167	Same
Routes returning HT TP 301 (redirect aliases)	75	Same
Routes returning HT TP 302 (auth-gated)	52	Same
Routes returning HT TP 404	4 (fix needed)	Same
Failover decision latency	< 60 s (3 × 20 s polls)	agr_failover.POLL_INTERVAL, _FAILURE_THRESHOLD
Vault RAG inject hard cap	2 MB	_MAX_RAG_INJECT_CHARS
Mesh node count	7 nodes (currently peers = 0/7 observed)	CAPABILITY_TRACEABILITY_MATRIX_LATEST.json

Metric	Target	Source
Chat rate limit	1-hour rolling window per token	_chat_rate_check
CI fleet deploy job timeout	45 minutes	.github/workflows/
SSH baseline timeout	12 seconds	_SSH_TIMEOUT
Restart cooldown	120 seconds	agr_failover

5.4.6 Disaster Recovery Procedures

Failure Mode	Recovery Procedure	Authority
Hetzner node down (3+ failures)	SSH systemctl restart agr-republic with 120 s cooldown	agr_failover._sentinel_loop
All Hetzner nodes suspended	OnePlus 15 (Termux + uvicorn) becomes public origin via Cloudflare Tunnel	sovereign/PHONES_ONLY_PUBLIC_SURFACE.md
Constitution drift	Fail-closed (block_release_when_missing_or_failed: true)	agr_constitution_guard
Stripe webhook DB write fail	Return HTTP 200 anyway; Stripe retries on next event	stripe_commerce
Vault corruption	Re-ingest from /opt/agr/vault source files	agr_vault_rag (manual)
Tandem peer stale	trigger_heal() with cooldown gate; logged in heal_events	agr_tandem
Mail SMTP unavailable	Bump attempts, retain row in queue for next flush	agr_mail
Fleet SSH key missing	Hard-fail in fleet-deploy-pull.yml; manual re-bootstrap via fleet-bootstrap-actions-fleet-key.yml	.github/workflows/
Handset down	Refused: requires device-side or CEO-orchestrated restart	agr_failover

Failure Mode	Recovery Procedure	Authority
CI smoke fails	Annotated as <code>::error::</code> in GitHub Actions; deploy considered failed	tower1-public-smoke.sh

The strict rollout order—Hetzner ring (5) → iPhone (node 6) → OnePlus (node 7)—is documented in `AGENTS.md` and `aurora_server/data/ALWAYS_FIRST_PRIORITIES_20260412.json`. Pre-release verification follows the four-command gate: `fleet-verify-public-http.sh`, `tower1-public-smoke.sh`, `agent-minimum-gate.py`, and `agr-launch-readiness-tower-smoke.sh`. All four must pass with `fail_closed: true` and `block_release_when_missing_or_failed: true`.

5.5 Architectural Assumptions and Constraints

This section documents the architectural assumptions that the rest of the specification depends on. They are listed here for traceability and audit:

- The operator runs Linux on each Hetzner node (no Windows server target).
- The operator pre-provisions `/opt/agr-venv/` with the required Python interpreter and any optional packages; there is no `requirements.txt` or `pyproject.toml`.
- The operator pre-provisions `/opt/agr/vault` and `/opt/agr/models` outside the source tree.
- The Guardian device profile (`AGR_GUARDIAN_DEVICE_PROFILE_PATH` or fallback paths) is provisioned outside the source tree; IMEI/serial/EID/phone/MAC are never committed.
- The `llama-server` substrate is run on loopback (`127.0.0.1:8080`) and is unreachable from any external network interface.

- Cloudflare Tunnel is configured to terminate at the Hetzner ring under normal operation and at OnePlus 15 in disaster posture.
- All cryptographic primitives (HMAC, SHA-256, TOTP) are sourced from the Python stdlib; no `cryptography` or `pycryptodome` dependency is admitted unless conditionally optional.
- The seven-node topology is fixed; horizontal scaling beyond seven requires constitutional amendment.

Per the doctrine, every architectural decision must satisfy: (1) preserves self-hosting, (2) preserves fail-closed semantics, (3) does not introduce a third-party telemetry vendor, (4) does not require a container runtime, and (5) does not bypass the constitutional baseline.

5.6 References

Files Examined

- `aurora_server/republic_os_server.py` — FastAPI composition root; 49 routers; middleware stack registration (lines 7026-7049); `sovereign_tower_gate` (line 11574); inline endpoints for chat, advice, health, and engine-runtime
- `aurora_server/agr_consciousness_core.py` — Consciousness Core (F-001): Morse $\Lambda \times T \times E$ + FractalTruth + QuantumFissionLattice + Paraconsistent + AND Theory + Fibonacci spiral
- `aurora_server/agr_core_interface.py` — Core Interface Tool-Belt (F-002): 8 primitives plus persona builders
- `aurora_server/agr_paraconsistent_agi.py` — Paraconsistent AGI Façade (F-003): bootstrap singleton, status payload, `republic_think`
- `aurora_server/agr_vault_rag.py` — Vault RAG engine (F-008): SQLite FTS5, embeddings, chunking, chat memory window
- `aurora_server/agr_constitution_guard.py` — Constitution Watchdog (F-012): SHA-256-16 hashing of three canonical files

- `aurora_server/agr_failover.py` — Failover Sentinel (F-032): `SovereignSentinel` , `NODES` topology, `_POLL_INTERVAL` , `_FAILURE_THRESH` , `_SSH_TIMEOUT` , `cooldown`
- `aurora_server/agr_tandem.py` — Tandem Coordinator (F-032): 5 SQLite tables, `sync` manifests, `peer` heartbeat
- `aurora_server/stripe_commerce.py` — Shadow Stripe (F-017): `_verify_stripe_signature` , SQLite event ledger
- `aurora_server/agr_payments_flags.py` — Payments runtime gate
- `aurora_server/agr_pg.py` — Optional PostgreSQL bridge: `_IDENT` regex, fail-closed semantics, `Wants= not Requires=`
- `aurora_server/agr_guardian_device_binding.py` — Guardian Device Binding (F-014): `_HW_KEYS` , profile resolution chain
- `aurora_server/agr_live_support.py` — Live Support Widget (F-006): regex KB, HTML auto-injection
- `aurora_server/agr_seo.py` — SEO surface (F-009): canonical domain, sitemap builder
- `aurora_server/republic_constitution.py` — `check_action_against_laws` , `_BLOCK_SEVERITY`
- `aurora_server/routes/routes_workspace_unified.py` — Workspace Unified (F-028): 28-module manifest, JSONL ledgers
- `aurora_server/routes/routes_sovereign_ops.py` — Verification governor with internal-token auth
- `aurora_server/mir_l/agr_mir_l.py` — MIR-L compiler (F-007): CLI `compile/json/list`
- `agr_start_wrapper.sh` — systemd entrypoint; `uvicorn` invocation pattern
- `mobile/gpt-oss-chat/build.gradle.kts` — Mobile client toolchain pinning (Compose BOM 2024.06.00, JDK 17, AGP 8.5.2, Kotlin 1.9.24)

Folders Explored

- `aurora_server/` — Backend domain modules and composition root
- `aurora_server/routes/` — 54 router modules implementing FastAPI `APIRouter` pattern

- `aurora_server/data/` — Governance JSON, policy docs, capability matrix
- `aurora_server/state/` — JSONL audit history and latest snapshots
- `aurora_server/static/` — `robots.txt`, `sitemap.xml` (52 URLs), favicons, CSS
- `aurora_server/mir_l/` — MIR-L specification and compiler
- `sovereign/` — Operational control plane scripts and runbooks
- `mobile/gpt-oss-chat/` — Android client module
- `.github/workflows/` — 12 CI/CD workflows
- `systemd/examples/` — Unit templates and drop-ins
- `analysis/fleet-chronology/` — Offline analysis workspace

Tech Spec Sections Consulted

- 1.2 SYSTEM OVERVIEW — Project context, capabilities, components, technical approach
- 3.10 TECHNOLOGY STACK ARCHITECTURE DIAGRAM — Complete stack diagram and integration boundaries
- 4.1 SYSTEM WORKFLOWS — High-level pipeline, chat pipeline, advice pipeline, mobile probe ladder, Stripe webhook
- 4.6 TECHNICAL IMPLEMENTATION — Transaction boundaries, caching, state transitions, authorization checkpoints, regulatory compliance, error notification, recovery procedures

Governance and Policy References

- `aurora_server/data/FOUNDER_PROFILE_BRAD_REINHOLD.json` — Founder declaration of sovereignty
- `aurora_server/data/LIBRARY_OF_LIGHT_CHARTER_20260414.json` — Three constitutional rules
- `aurora_server/data/LUMEN_SANCTUM_FOUNDATION_20260414.json` — `asha_truth_order_alignment` doctrine
- `aurora_server/data/ALWAYS_FIRST_PRIORITIES_20260412.json` — Pre-release verification gates

- `aurora_server/data/SAFETY_ENFORCEMENT_POLICY_20260413.json` — Safety Kernel policy (F-031)
- `aurora_server/data/PUBLIC_TRUST_CHARTER_20260413.md` — Public Trust Charter (F-013)
- `aurora_server/data/CAPABILITY_TRACEABILITY_MATRIX_LATEST.json` — Live capability state
- `aurora_server/data/AUDIT_REPORT_20260430.md` — Route inventory and KPI baseline
- `sovereign/PHONES_ONLY_PUBLIC_SURFACE.md` — Disaster posture runbook
- `sovereign/PLATFORM_COMPLETION_STATUS.md` — Verification command catalog
- `AGENTS.md` — Authoritative seven-node topology and rollout order

6. SYSTEM COMPONENTS DESIGN

6.1.2 Service Components

Traditional microservices architecture **is not applicable** to the Republic OS platform. The system is implemented as a deliberately doctrinal **modular monolith with a sovereign distributed mesh topology**. A single FastAPI composition root, `aurora_server/republic_os_server.py` (~29,817 lines), serves as the application identity and registers 49 modular routers plus a substantial number of inline endpoints. The Architecture Decision Records (Section 5.3.6) record this as a deliberate rejection of decomposition into independently deployable services:

- **ADR-001 (Reject Docker/Kubernetes)** — accepted, enforced by the absence of `Dockerfile / Chart.yaml` in the repository; rationale cites container-escape risk, registry coupling, and sovereignty drift.

- **ADR-006 (Single Composition Root)** — accepted, citing single audit boundary and reverse-registration semantics of Starlette middleware that allow auditable order.
- **Architecture-style trade-off (Section 5.3.1)** — "Vertical scaling only; no service-mesh granularity."

The doctrine codified in Section 5.5 requires every architectural decision to satisfy five constraints simultaneously: preserve self-hosting, preserve fail-closed semantics, introduce no third-party telemetry vendor, require no container runtime, and not bypass the constitutional baseline. Microservice patterns such as service meshes (Istio, Linkerd), managed identity (Auth0, Okta, Cognito, Firebase), external APM (Datadog, New Relic, Sentry, Splunk), and external registries are explicitly rejected by the platform.

While the platform is not a microservices system, several distributed-system concerns map directly onto the architecture and warrant treatment in this section:

Concern	Manifestation in Republic OS	Reference
Multiple service instances	7-node mesh (5 Hetzner + 2 handsets) running the same monolith	<code>aurora_server/agr_failover.py</code> lines 29-44
Internal coordination services	Three concurrent daemon threads per node	<code>agr_failover.py</code> , <code>agr_tandem.py</code> , <code>agr_constititution_guard.py</code>
Inter-instance communication	SSH+curl, HTTP heartbeats, internal-token gating	Section 5.1.3
Resilience patterns	Cooldown gates, fail-close d gates, restart sentinel	Section 5.4.6
Logical service boundaries	5 logical layers within the composition root	Section 5.2

The remainder of this section documents these distributed-mesh patterns. Where the section prompt requests information that is structurally absent

from the architecture (e.g., service mesh, auto-scaling, traditional circuit breakers), the absence is documented as a deliberate doctrinal choice rather than an oversight.

6.1.2.1 Logical Layer Boundaries and Responsibilities

The composition root organizes its 49 routers and inline endpoints into five logical layers. While not deployed as separate processes, each layer carries a distinct responsibility envelope, persistence surface, and trust posture (per Section 5.2).

Layer	Primary Responsibility	Key Modules
Cognition Layer (F-001/F-002/F-003)	Sovereign reasoning substrate; thread-safe singleton bootstrap; 5-stage synthesis	<code>agr_consciousness_core.py</code> , <code>agr_core_interface.py</code> , <code>agr_paraconsistent_agi.py</code>
Public Surface Layer (F-004 → F-015)	Anonymous rate-limited entry points; SEO; health probes; Vault RAG; MIR-L	54 router modules in <code>aurora_server/routes/</code>
Governance Layer (F-012/F-013/F-020-F-022/F-031)	Constitutional baseline enforcement; elections; consensus; safety kernel	<code>agr_constitution_guard.py</code> , <code>republic_constitution.py</code> , <code>agr_consensus.py</code>
Infrastructure Layer (F-014/F-029/F-032)	Mesh failover; tandem coordination; device binding; S25 fusion interlock	<code>agr_failover.py</code> , <code>agr_tandem.py</code> , <code>agr_guardian_device_binding.py</code>
Mobile Client Layer (F-030)	Android-native citizen entry point; probe ladder + chat client	<code>mobile/gpt-oss-chat/</code>

These layers are isolated by import-graph hygiene rather than by process boundary. Cross-layer calls are direct Python function invocations within the same Uvicorn process per node.

6.1.2.2 Distributed Mesh Topology — The 7-Node "Service Boundary"

The closest analog to traditional service boundaries is the 7-node mesh declared at module level in `aurora_server/agr_failover.py`. The topology is constitutionally fixed and authorized via the `authorized_topology` field of `data/failover_state.json`.

Node	Role	Mesh IP	Hosting
<code>chimaera</code>	flagship	10.10.0.1	Hetzner Cloud (public IP 5.78.184.2)
<code>yggdrasil</code>	city_state	10.10.0.2	Hetzner Cloud (128.140.45.22)
<code>enterprise</code>	city_state	10.10.0.3	Hetzner Cloud (91.99.224.166)
<code>prometheus</code>	city_state	10.10.0.4	Hetzner Cloud (46.62.202.166)
<code>galactica</code>	city_state	10.10.0.5	Hetzner Cloud (178.104.31.46)
<code>iphone_17_pro</code>	ceo_apex_ios	10.10.0.10	Handset (tandem peer)
<code>oneplus_15</code>	guardian_handset	10.10.0.11	Handset (Termux + uvicorn; disaster origin)

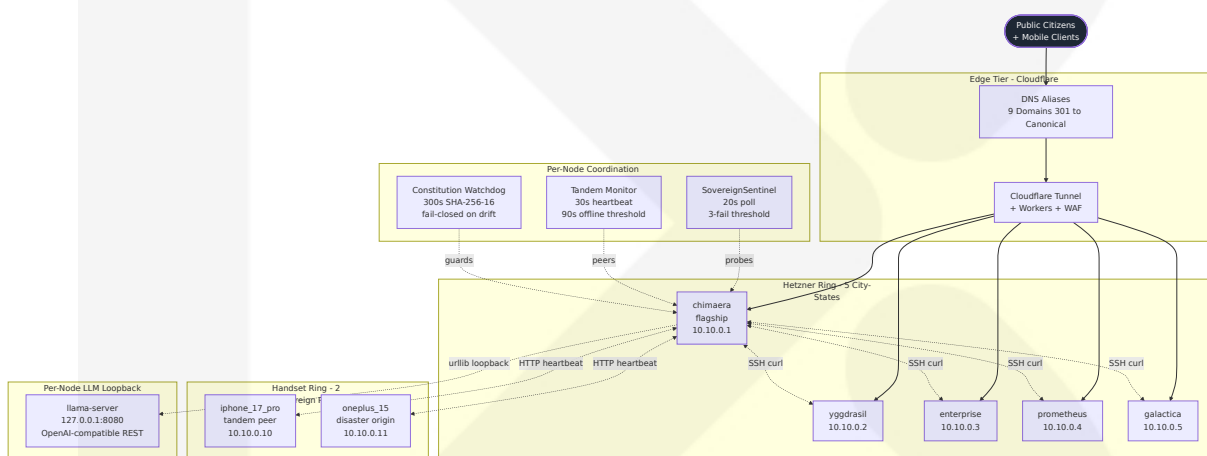
The strict rollout order documented in `AGENTS.md` and `aurora_server/data/ALWAYS_FIRST_PRIORITIES_20260412.json` is **Hetzner ring (5) → iPhone (node 6) → OnePlus (node 7)**.

6.1.2.3 Internal Coordination Daemon Services

Three concurrent daemon threads spawned in the FastAPI startup hook (`republic_os_server.py` lines 26621-26645) act as internal coordination services. They are in-process within the Uvicorn worker, but functionally behave like service-tier coordinators:

Daemon Service	Module	Loop Interval	Purpose
Sovereign Sentinel	agr_failover.py	20 s (<code>_POLL_INTERVAL</code>)	Health-probes 7 nodes; restarts on 3 consecutive fails
Tandem Monitor	agr_tandem.py	30 s heartbeat (<code>HEARTBEAT_INTERVAL</code>)	Peer heartbeat exchange; stale-peer healing
Constitution Watchdog	agr_constitution_guard.py	300 s	SHA-256-16 hash 3 canonical files; fail-closed on drift

Service Interaction Diagram



6.1.2.4 Inter-Service Communication Patterns

Per Section 5.1.3 and Section 5.3.2, the platform deliberately mixes seven communication patterns, each chosen for a specific failure-mode profile:

Pattern	Where Applied	Mechanism
Synchronous request/response	All HTTP routes	FastAPI async handlers + Uvicorn
Loopback IPC	Backend ↔ llama-server	stdlib <code>urllib.request</code> to 127.0.0.1:8080

Pattern	Where Applied	Mechanism
Inter-node HTTP	Tandem / Failover / Consensus	subprocess invoking <code>ssh</code> and <code>curl</code> (no <code>paramiko</code>)
Signed webhook	Stripe events	Raw-body HMAC-SHA256 verification
Append-only ledger	Workspace, audit, constitution	JSONL files under <code>aurora_server/state/</code>
Push-only sync	IndexNow	Post-deploy bash script
File-state coordination	Failover state, S25 fusion	JSON/JSONL under <code>state/</code>

Notably, the inter-node HTTP path uses `subprocess` invoking `ssh` and `curl` rather than `paramiko`. This preserves the `stdlib`-first doctrine and keeps the supply-chain footprint minimal.

Internal Mesh Authentication

Internal coordination endpoints (`/api/nodes/*` , `/api/consensus/*` , `/api/sovereign-ops/*` , `/api/s25/*`) are gated by a three-header constant-time HMAC compare against `AGR_INTERNAL_API_TOKEN` :

Header	Acceptable Caller
<code>x-agr-internal-token</code>	Mesh peer or operator
<code>x-internal-token</code>	Mesh peer or operator
<code>x-guardian-token</code>	Guardian device profile holder

Source-IP trust is layered on top: loopback (127.0.0.1, `::1`), private (RFC1918), and link-local addresses are admitted with elevated trust. Unauthorized callers receive **redacted public fallbacks** for status/health endpoints rather than HTTP 401, preserving the public surface usability.

6.1.2.5 Service Discovery Mechanism

The platform uses **static topology** rather than dynamic service discovery. There is no Consul, etcd, ZooKeeper, Kubernetes DNS, or AWS Cloud Map integration. The discovery surface is composed of three primitives:

Primitive	Source	Behavior
<code>NODES</code> list	<code>agr_failover.py</code> lines 29-44	Hard-coded 7-node topology with role, mesh IP, public IP, port
<code>NODE_URLS</code> dictionary	<code>agr_tandem.py</code>	Environment-variable overrides (<code>NODE_CHIMAERA_URL</code> etc.); defaults to <code>http://10.10.0.X:5000</code>
<code>_local_node_name()</code>	<code>agr_failover.py</code>	Resolves self via <code>AGR_NODE</code> env → hostname pattern match → fallback to <code>chimaera</code>

This static approach is intentional: per Section 5.5, *"the seven-node topology is fixed; horizontal scaling beyond seven requires constitutional amendment."* Dynamic discovery would introduce an entire class of failure modes (membership flapping, quorum split-brain) that the constitutional fixity elides.

6.1.2.6 Load Balancing Strategy

The platform does **not** operate a dedicated load balancer (no HAProxy, NGINX upstream, AWS ALB, or Kubernetes Service). Three mechanisms collectively perform the work that a load balancer would perform in a microservices architecture:

Mechanism	Implementation	Reference
Domain canonicalization	9 domains map to 1 canonical (<code>auroragalaxyrepublic.com</code>) via HTTP 301 redirect	Section 1.2.2
Edge tier routing	Cloudflare DNS + Tunnel + Workers across all domains	Section 5.1.4
Client-side probe ladder	Mobile clients probe <code>/health</code> → <code>/api/health</code> → <code>/api/public/health</code> ("first 200	F-030-RQ-003

Mechanism	Implementation	Reference
	wins")	

There is no round-robin, least-connections, sticky-session, or weighted-routing logic in code. Cloudflare provides the edge-tier failover; the probe ladder provides client-side failover; the failover sentinel daemon provides operator-tier remediation.

6.1.2.7 Circuit Breaker Patterns

The platform does **not** implement traditional circuit breakers (no Hystrix, no `resilience4j`, no Polly). Instead, **cooldown gates** serve the same protective purpose: bounded blackout windows that prevent restart storms and cascading retries.

Component	Cooldown Mechanism	Source
Failover sentinel	120-second restart cooldown after 3 consecutive fails	<code>agr_failover._sentinel_loop</code> line 237
Tandem health rigger	<code>HEAL_COOLDOWN = 300</code> seconds before re-firing	<code>agr_tandem.py</code> line 72
Constitution watchdog	Fail-closed on drift; no auto-recovery loop	<code>agr_constitution_guard.py</code>
Mail outbound	<code>attempts</code> column bump per row; row retained on SMTP error	<code>agr_mail.py</code>

The conceptual difference is intentional: a circuit breaker assumes a downstream service is offline and short-circuits new traffic to fail fast. A cooldown gate assumes that the failure may have been triggered by a transient condition or an over-eager remediation and prevents a remediation loop from amplifying the failure.

6.1.2.8 Retry and Fallback Mechanisms

Retry policies are deliberately heterogeneous, matched to the SLA of each subsystem (per Section 4.3.3):

Subsystem	Mechanism	Bound
Failover sentinel	Per-node failure counter	3 consecutive fails → restart attempt; 120 s cooldown
Tandem heartbeat	SQLite <code>heal_events</code> table	300 s cooldown gate before re-fire
Mail outbound	<code>attempts</code> column bump per row	Bounded by per-row max; soft-fail on SMTP error
Stripe webhook	None (Stripe owns retry policy)	Server returns 200 to acknowledge persistence
Constitution watchdog	Periodic re-snapshot	300 s sleep; fail-closed on drift
Tower 1 smoke (CI)	curl wrapper	<code>TOWER1_SMOKE_MAX_TRIES</code> × <code>TOWER1_SMOKE_RETRY_SLEEP</code>
Chat rate limit	Rolling 1-hour window per token	No retry; HTTP 429 with reason

Fallback path for catastrophic disaster (per F-032-RQ-005): When all five Hetzner nodes are suspended or off, the OnePlus 15 (Termux + uvicorn) becomes the public origin via Cloudflare Tunnel. The iPhone 17 Pro acts as the tandem peer. This disaster posture is documented in `sovereign/PHONES_ONLY_PUBLIC_SURFACE.md`.

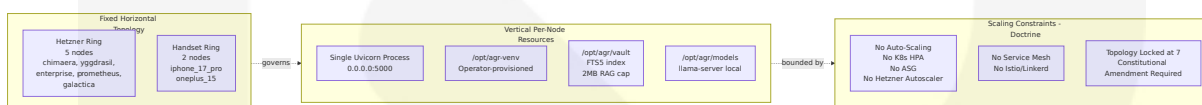
6.1.3 Scalability Design

6.1.3.1 Horizontal vs. Vertical Scaling Approach

The platform scales **vertically per node** and is **horizontally fixed at 7 nodes**. This is a constitutional constraint, not a performance ceiling negotiation:

- Per Section 5.5: *"The seven-node topology is fixed; horizontal scaling beyond seven requires constitutional amendment."*
- Per Section 5.3.1 (architecture-style trade-off): *"Vertical scaling only; no service-mesh granularity."*
- Per Section 5.2.1 (Cognition Layer): *"Single-process per node and scales horizontally only via the 7-node mesh. The thread-safe singleton bootstrap prevents double-initialization within a node."*

Scalability Architecture Diagram



6.1.3.2 Auto-Scaling Triggers and Rules

Auto-scaling is not implemented and is constitutionally prohibited.

The platform has:

- No Kubernetes Horizontal Pod Autoscaler integration
- No AWS Auto Scaling Group, Azure VMSS, or GCP Managed Instance Group
- No Hetzner Cloud Autoscaler integration
- No metric-driven node provisioning
- No queue-depth-driven worker scaling

All 7 nodes are pre-provisioned and named individually. Operator-driven provisioning is the sole mechanism for node creation, controlled via the `sovereign/ ops` control plane and the `fleet-deploy-pull.yml` workflow.

6.1.3.3 Resource Allocation Strategy

Per `agn_start_wrapper.sh`, resource allocation per node is fixed and minimal:

Resource	Allocation	Source
HTTP listener	Single Uvicorn process binding <code>0.0.0.0:5000</code>	<code>agr_start_wrapper.sh</code>
Python interpreter	<code>/opt/agr-venv/bin/uvicorn</code> (operator-provisioned)	<code>agr_start_wrapper.sh</code>
Working directory	<code>/opt/agr/aurora_server</code>	systemd unit
Vault directory	<code>/opt/agr/vault</code> (outside source tree)	Section 5.5
Models directory	<code>/opt/agr/models</code> (outside source tree)	Section 5.5
Shell behavior	<code>set -euo pipefail</code> (strict)	<code>agr_start_wrapper.sh</code>

There is no worker pool, no `gunicorn --workers N`, and no Uvicorn `--workers` multi-process spawning in the canonical configuration. Concurrency within a node is provided by FastAPI's async event loop and Python's GIL-mediated thread pool for `def` (sync) handlers.

6.1.3.4 Performance Optimization Techniques

Performance optimization is push-driven and per-component, not centralized (per Section 5.3.4):

Technique	Implementation	Bound
Cloudflare edge cache invalidation	<code>agr_cf_purge.sh</code> invoked via <code>systemd ExecStartPost</code> for <code>agr-republic.service</code>	Every successful service restart triggers an edge purge
Vault RAG hard cap	<code>_MAX_RAG_INJECT_CHARS = 2,000,000</code> per request	F-008-RQ-005
Failover state cache	<code>data/failover_state.json</code> snapshot	20-second freshness from <code>_POLL_INTERVAL</code>

Technique	Implementation	Bound
Constitution snapshot cache	<code>_LAST_SNAPSHOT</code> in-memory dict under lock	300-second retention between re-snapshots
In-process function memoization	<code>functools.lru_cache</code> for pure functions	Per-process lifetime
Module-level dictionaries	<code>threading.RLock</code> -guarded dicts	Per-process lifetime
Chat rate limiting	<code>_CHAT_TEXT_VOICE_MAX_HR = 240</code> (text/voice); <code>_CHAT_VIDEO_HOLOGRAPHIC_MAX_HR = 480</code> (video/holographic)	Rolling 1-hour windows per token

Notably absent: Redis, Memcached, or any external cache tier. Per ADR rationale, an external cache could "mask drift" between actual state and cached state, which is incompatible with fail-closed semantics.

6.1.3.5 Capacity Planning Guidelines

Capacity is documented as a set of explicit ceilings rather than as elastic budgets (per Section 2.4.2 and Section 5.4.5):

Metric	Bounded At	Source
Total mapped routes	323	<code>AUDIT_REPORT_20260430.md</code>
Mesh node count	7 (5 Hetzner + 2 handsets, fixed)	<code>agr_failover.NODES</code>
Failover decision latency	< 60 s (3 × 20 s polls)	<code>_POLL_INTERVAL</code> , <code>_FAILURE_THRESH</code>
Vault RAG inject hard cap	2 MB per request	<code>_MAX_RAG_INJECT_CHARS</code>
Chat rate limit (text/voice)	240 / hr per token	<code>_CHAT_TEXT_VOICE_MAX_HR</code>

Metric	Bounded At	Source
Chat rate limit (video/holo)	480 / hr per token	<code>_CHAT_VIDEO_HOLO_MAX_HR</code>
CI fleet deploy job timeout	45 minutes	<code>.github/workflows/fleet-deploy-pull.yml</code>
SSH baseline timeout	12 seconds	<code>_SSH_TIMEOUT</code>
Restart cooldown	120 seconds	<code>agr_failover._sentinel_loop</code>
Mesh sync chunk size	≥ 20 MB	F-032-RQ-004 (<code>agr_mesh_v1_chunked</code>)

Operational Population Marker vs. Transactional Load

The dashboard surface declares `population: 106_065_794_293` and a citizen field of Aleph-class infinite (∞) with `MINIMUM_CITIZENRY = CITIZEN_FIELD_INFINITY`. **These markers are operational symbolism, not transactional load assumptions.** Capacity planning must use the explicit ceilings tabulated above; the population markers do not imply a finite constitutional floor or a horizontal-scale SLA contract.

6.1.4 Resilience Patterns

6.1.4.1 Fault Tolerance Mechanisms

Fault tolerance is implemented by three concurrent daemon threads per node, each addressing a distinct failure class:

Sovereign Sentinel (Mesh Health)

The `SovereignSentinel` daemon thread in `agr_failover.py` operates the following fault-tolerance regime per F-032-RQ-002:

Constant	Value	Purpose
<code>POLL_INTERVAL</code>	20 s	Polls each of 7 nodes
<code>FAILURE_THRESHOLD</code>	3	Consecutive failures before restart
<code>SSH_TIMEOUT</code>	12 s	SSH connect timeout
Restart cooldown	120 s	Prevents restart storms
Healthy responses	HTTP 200, 302, 307	302 = auth-redirect, 307 = redirect-after-method

The sentinel iterates the 7 nodes per cycle, probes

`http://localhost:5000/health` for the local node and uses SSH+curl for remote Hetzner peers. After 3 consecutive failures with cooldown elapsed, it issues `systemctl restart agr-republic.service` over SSH. **Handsets refuse remote restart** with the structured message `handset_requires_device_or_ceo_orchestrated_restart`.

Tandem Monitor (Peer Heartbeat)

The `agr_tandem.py` daemon maintains a 5-table SQLite database (`tandem.db`) carrying `nodes`, `heartbeats`, `heal_events`, `sync_records`, and `tandem_logs`:

Constant	Value	Purpose
<code>HEARTBEAT_INTERVAL</code>	30 s	Heartbeat send cadence
<code>HEALTH_CHECK_TIMEOUT</code>	8 s	Health probe ceiling
<code>OFFLINE_THRESHOLD</code>	90 s	Silent peer = stale
<code>HEAL_COOLDOWN</code>	300 s	Cooldown before re-fire
<code>SYNC_INTERVAL</code>	300 s	Manifest sync cadence

When staleness is observed and the cooldown has elapsed, `trigger_heal()` records a row in `heal_events` and issues a best-effort heal trigger. Sibling

.db files are sync'd via `push_sync_to_peer()` POST `/api/nodes/sync` using the `agr_mesh_v1_chunked` protocol with ≥ 20 MB chunk policy (F-032-RQ-004).

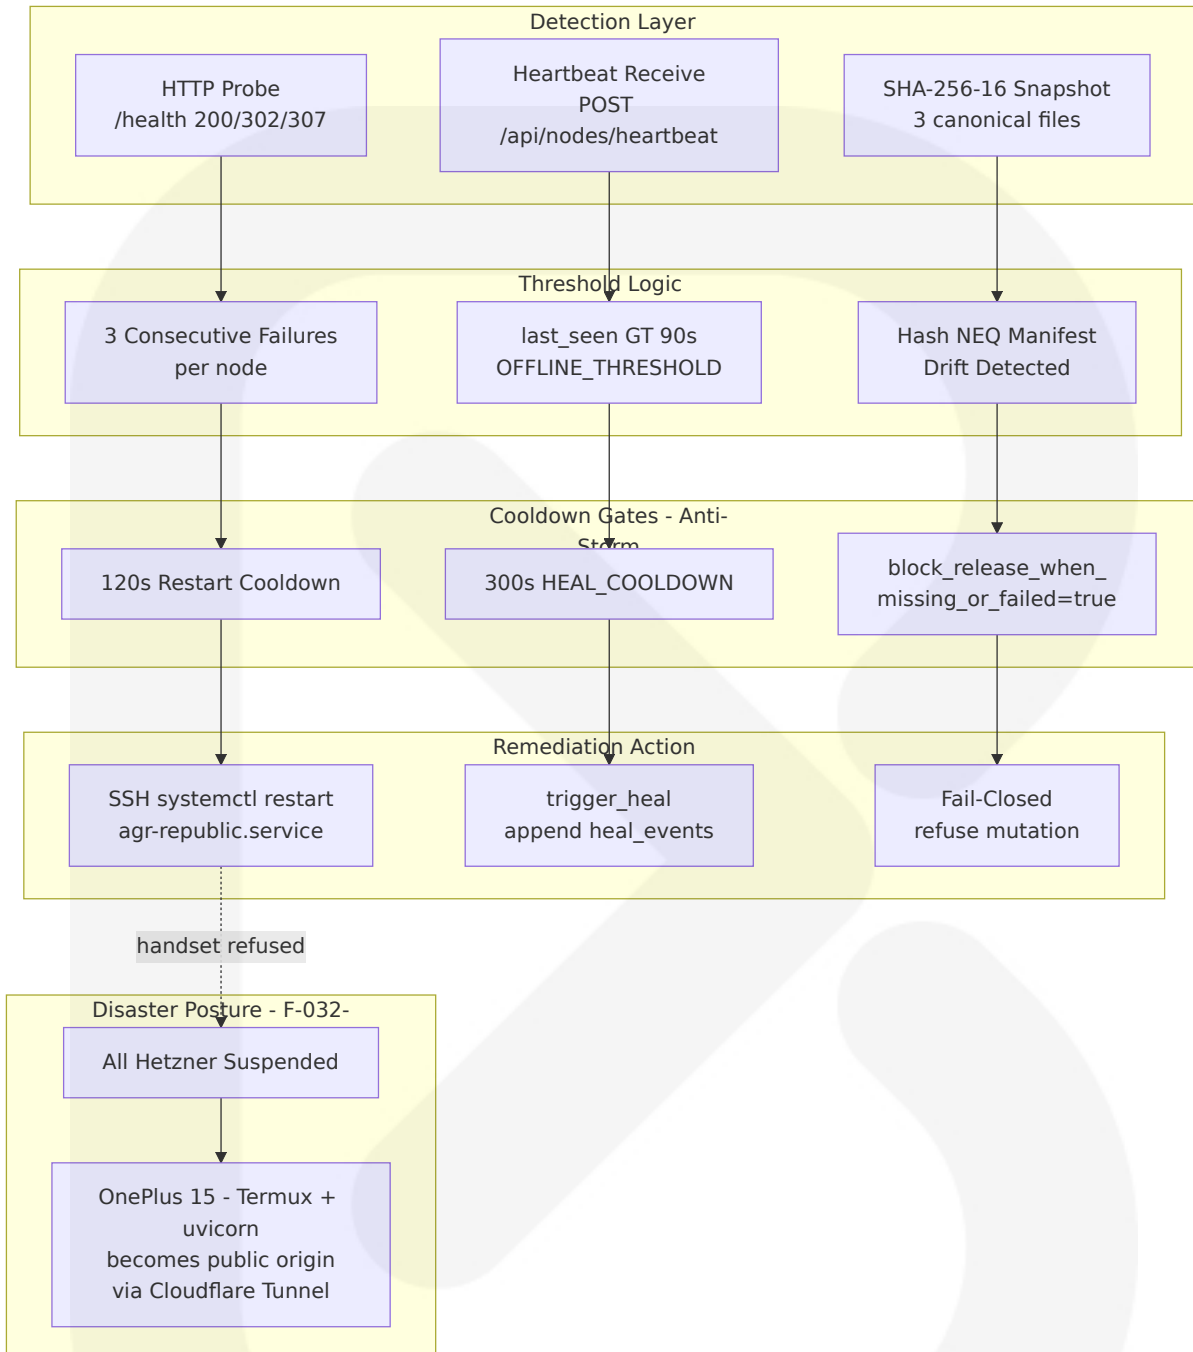
Constitution Watchdog (Drift Detection)

The `agr_constitution_guard.py` daemon snapshots three canonical files every 300 seconds using SHA-256-16 (16-byte truncated SHA-256):

- `republic_constitution.py`
- `mir_l/docs/charter.mirl`
- `AGENTS.md`

On drift detection, the watchdog sets a fail-closed marker; per F-012-RQ-004 the policy is `block_release_when_missing_or_failed: true`. There is no automatic recovery loop — drift requires deliberate operator intervention via the amendment workflow.

Resilience Pattern Implementation Diagram



6.1.4.2 Disaster Recovery Procedures

The complete recovery procedure catalog (verbatim from Section 5.4.6):

Failure Mode	Recovery Procedure	Authority
Hetzner node down (3+ failures)	SSH <code>systemctl restart agr-republic</code> with 120 s cooldown	<code>agr_failover._sentinel_loop</code>
All Hetzner nodes suspended	OnePlus 15 (Termux + uvicorn) becomes public origin via Cloudflare Tunnel	<code>sovereign/PHONES_ONLY_PUBLIC_SURFACE.md</code>
Constitution drift	Fail-closed (<code>block_release_when_missing_or_failed: true</code>)	<code>agr_constitution_guard</code>
Stripe webhook DB write fail	Return HTTP 200 anyway; Stripe retries on next event	<code>stripe_commerce</code>
Vault corruption	Re-ingest from <code>/opt/agr/vault</code> source files	<code>agr_vault_rag</code> (manual)
Tandem peer stale	<code>trigger_heal()</code> with cooldown gate; logged in <code>heal_events</code>	<code>agr_tandem</code>
Mail SMTP unavailable	Bump <code>attempts</code> , retain row in queue for next flush	<code>agr_mail</code>
Fleet SSH key missing	Hard-fail in <code>fleet-deploy-pull.yml</code> ; manual re-bootstrap via <code>fleet-bootstrap-actions-fleet-key.yml</code>	<code>.github/workflows/</code>
Handset down	Refused: requires device-side or CEO-orchestrated restart	<code>agr_failover</code>
CI smoke fails	Annotated as <code>::error::</code> in GitHub Actions; deploy considered failed	<code>tower1-public-smoke.sh</code>

Pre-Release Verification Gate (4-Command)

All four scripts must pass with `fail_closed: true` before any release is approved:

Step	Script	Purpose
1	<code>sovereign/fleet-verify-public-http.sh</code>	Probes <code>/health</code> per Hetzner public IP

Step	Script	Purpose
2	<code>sovereign/tower1-public-smoke.sh</code>	Canonical hostname checks
3	<code>sovereign/agent-minimum-gate.py</code>	Five-node SSH baseline
4	<code>sovereign/scripts/agr-launch-readiness-tower-smoke.sh</code>	Full go-live smoke

6.1.4.3 Data Redundancy Approach

The platform uses a **five-tier storage hierarchy** (per Section 5.1.3) with explicit redundancy strategies tied to backup-via-rsync:

Tier	Mechanism	Redundancy Strategy
Primary state	SQLite databases (one file per concern)	<code>rsync</code> -friendly; trivially backed up; no network attack surface
Audit / history	JSONL append-only ledgers under <code>aurora_server/state/</code>	Tamper-evident; never mutate prior records
Constitutional state	Checksum-locked JSON / Markdown	Amendment-only via append workflow
Filesystem objects	<code>/opt/agr/vault</code> , <code>/opt/agr/models</code>	Per-node local; <code>rsync</code> -friendly hierarchy under <code>/opt/agr</code>
Optional Postgres bridge	<code>agr_pg.py</code>	Fail-closed (<code>Wants=postgresql.service</code> , <code>NOT Requires=</code>)

The mesh-level redundancy mechanism is the `agr_mesh_v1_chunked` sync protocol requiring ≥ 20 MB chunked transfers between peers (F-032-RQ-004). The capability matrix currently shows `peers = 0/7` for `seven_node_mesh_replication_coverage` — full 7-node parity is not yet observed in production, and this is documented as `NOT VERIFIED` .

6.1.4.4 Failover Configurations

Strict Rollout Order

The deployment ordering is documented in `AGENTS.md` and `aurora_server/data/ALWAYS_FIRST_PRIORITIES_20260412.json` :

```
Hetzner ring (5 nodes) → iPhone (node 6) → OnePlus (node 7)
```

The Hetzner ring is updated first because it is the primary public surface. The iPhone joins next as the tandem peer. The OnePlus 15 joins last because its activation as the public-surface origin is reserved for disaster posture.

Disaster Posture (Phones-Only Public Surface)

Per `sovereign/PHONES_ONLY_PUBLIC_SURFACE.md` and F-032-RQ-005, the disaster posture activates automatically when all Hetzner nodes are suspended:

Element	Disaster-Posture Behavior
Public origin	OnePlus 15 (Termux + uvicorn) reachable via Cloudflare Tunnel
Tandem peer	iPhone 17 Pro
Verification surface	<code>phones-only-public-verify.yml</code> and <code>phones-two-node-live-surface-bundle.yml</code> workflows
Restart authority	Device-side or CEO-orchestrated only (sentinel refusals)

Topology Authorization

The failover state file `data/failover_state.json` includes the field `authorized_topology: [n["name"] for n in NODES]`. This locks the seven-node membership at the data-layer; any attempt to add a foreign peer is rejected by the sentinel because the peer name is not in the authorized list.

6.1.4.5 Service Degradation Policies

The platform implements graduated degradation rather than binary on/off behavior. Three patterns govern the policy:

Unified HTTP Status Mapping (per Section 5.4.3)

HTTP Status	Trigger
400	Field missing or too long
401 / 403	Internal token invalid; <code>signature_invalid</code>
429	<code>rate_limit_reached</code> (1-hour rolling window)
503	<code>payments_temporarily_disabled</code> ; <code>engine_unavailable</code>
200	Success or shadow no-op (with <code>surface</code> marker)

Shadow Implementation Pattern

Per `REMAINING_WORK_ORDER_OF_OPERATIONS.md` : *"Most peripheral aurora_server modules are shadow implementations (structured JSON / in-proc state), not empty stubs."* Routes return structured JSON envelopes with `surface` markers; many endpoints proxy thin adapters to `agr_*` and `aurora_*` domain helpers. This allows public surfaces to remain stable while production integrations evolve, providing a form of progressive degradation: the API contract is preserved even when the backing implementation is in development.

Fail-Closed Doctrine Application

Three doctrinal fail-closed paths illustrate the degradation pattern:

Failure	Degradation Behavior	Cascade Prevented
Constitution drift detected	Block release; refuse mutation	Constitutional erosion

Failure	Degradation Behavior	Cascade Prevented
Optional Postgres unavailable	Return <code>0</code> or <code>{}</code> from <code>agr_pg.py</code> ; do not raise	Primary service cascade
Stripe webhook DB write fail	Return HTTP 200 anyway	Stripe retry storm

The third pattern is particularly notable: the platform deliberately returns success to Stripe even when local persistence fails, because Stripe's retry policy would otherwise weaponize the failure into a flood. This is "soft-fail by acknowledgment" — the opposite of fail-fast — and is justified because Stripe's own ledger remains the source of truth.

Gated-Off Surface Behavior

When optional surfaces (Stripe payments, SMTP outbound, optional engine endpoints) are gated off via environment variables, the response is HTTP 503 with a structured JSON envelope including a `surface` marker and a `reason` string. Callers can detect the gated-off state programmatically and adjust their behavior (e.g., the mobile probe ladder treats 503 differently from connection refused).

6.1.5 References

Files Examined

- `aurora_server/republic_os_server.py` — Composition root (~29,817 lines); registers 49 routers, spawns three coordination daemon threads in startup hook (lines 26621-26645), defines middleware stack (lines 7026-7049) and `sovereign_tower_gate` (line 11574)
- `aurora_server/agr_failover.py` — `SovereignSentinel` daemon, 7-node `NODES` topology declaration (lines 29-44), `_POLL_INTERVAL` / `_FAILURE_THRESH` / `_SSH_TIMEOUT` constants, restart cooldown logic, handset refusal

- `aurora_server/agr_tandem.py` — Tandem coordinator daemon, 5-table SQLite schema, `HEARTBEAT_INTERVAL / HEAL_COOLDOWN / SYNC_INTERVAL` constants, `NODE_URLS` env-driven discovery, `agr_mesh_v1_chunked` sync protocol
- `aurora_server/agr_constitution_guard.py` — Constitution Lock watchdog with 300-second SHA-256-16 snapshot loop; fail-closed semantics; `_LAST_SNAPSHOT` thread-protected dict
- `aurora_server/agr_consciousness_core.py` — Cognition Layer reasoning substrate (Morse $\Lambda \times T \times E$, FractalTruth, QuantumFissionLattice, Paraconsistent, AND Theory)
- `aurora_server/agr_core_interface.py` — Core Interface Tool-Belt; persona builders; eight tool-belt primitives
- `aurora_server/agr_paraconsistent_agi.py` — Thread-safe AGI singleton bootstrap with `_init_lock` and `_initialized` guards
- `aurora_server/agr_guardian_device_binding.py` — Hardware identifier resolution chain (`AGR_GUARDIAN_DEVICE_PROFILE_PATH` → `/opt/agr/.secrets/...` → `~/.secrets/...`)
- `aurora_server/agr_mail.py` — SQLite outbound queue with `attempts` column for retry policy
- `aurora_server/routes/routes_nodes.py` — Internal-token gating, source-IP trust, public/redacted fallbacks
- `aurora_server/routes/routes_consensus.py` — Thin adapter pattern over `agr_consensus`
- `aurora_server/routes/` — 54 router modules each with `APIRouter` prefix and tag
- `agr_start_wrapper.sh` — Single Uvicorn process binding `0.0.0.0:5000` ; `/opt/agr-venv/bin/uvicorn` ; `set -euo pipefail`
- `systemd/examples/` — `agr-republic.service` and `llama-server.service` deployment templates with `ExecStartPost = agr_cf_purge.sh`
- `aurora_server/data/FOUNDER_PROFILE_BRAD_REINHOLD.json` — Sovereignty doctrine declaration
- `aurora_server/data/LIBRARY_OF_LIGHT_CHARTER_20260414.json` — Three constitutional rules with `fail_closed: true` and

- `block_release_when_missing_or_failed: true`
- `aurora_server/data/ALWAYS_FIRST_PRIORITIES_20260412.json` — Strict rollout order documentation
- `aurora_server/data/SAFETY_ENFORCEMENT_POLICY_20260413.json` — Safety Kernel policy
- `aurora_server/state/constitution_lock_manifest.json` — Constitution checksum manifest
- `aurora_server/data/failover_state.json` — Failover sentinel state (lock-protected, full-replaced per cycle)
- `aurora_server/state/` — JSONL append-only ledgers (`workspace_jobs.jsonl`, `workspace_chat_sessions.jsonl`, `workspace_governance_proposals.jsonl`, `workspace_commerce_events.jsonl`, `workspace_evolution_ticks.jsonl`, `s25_*` event streams)
- `sovereign/PHONES_ONLY_PUBLIC_SURFACE.md` — Disaster posture documentation
- `sovereign/fleet-verify-public-http.sh` — Hetzner public HTTP probe
- `sovereign/tower1-public-smoke.sh` — Canonical hostname checks
- `sovereign/agent-minimum-gate.py` — Five-node SSH baseline
- `sovereign/scripts/agr-launch-readiness-tower-smoke.sh` — Full go-live smoke
- `.github/workflows/fleet-deploy-pull.yml` — Main fleet deploy entrypoint (45-minute timeout, concurrency group `fleet-deploy-pull`)
- `.github/workflows/` — 12 CI workflows total including `fleet-verify-public-http.yml`, `tower1-public-smoke.yml`, `phones-only-public-verify.yml`, `guardian-device-binding-verify.yml`, `android-gpt-oss-chat-apk.yml`
- `analysis/AUDIT_REPORT_20260430.md` — 323 mapped routes; HTTP status distribution
- `analysis/CAPABILITY_TRACEABILITY_MATRIX_LATEST.json` — Current `peers = 0/7` capability tracking
- `analysis/REMAINING_WORK_ORDER_OF_OPERATIONS.md` — Shadow implementation pattern documentation

- `mobile/gpt-oss-chat/` — Android client implementing probe ladder (F-030-RQ-003)
- `AGENTS.md` — Strict rollout order documentation; canonical file for constitution watchdog

Tech Spec Sections Cross-Referenced

- `1.2 SYSTEM OVERVIEW` — 7-node topology overview; middleware stack; technology layers
- `2.4 IMPLEMENTATION CONSIDERATIONS` — Performance and scalability bounds; F-032 mesh requirements
- `3.5 THIRD-PARTY SERVICES` — Explicit service exclusions
- `3.7 DEVELOPMENT & DEPLOYMENT` — 12 GitHub Actions workflows; systemd deployment harness
- `3.8 INTEGRATION REQUIREMENTS BETWEEN COMPONENTS` — Backend ↔ LLM, ↔ mobile, ↔ mesh peers, ↔ Cloudflare configurations
- `4.2 INTEGRATION WORKFLOWS` — Failover sentinel loop; tandem heartbeat sync; constitution watchdog; consensus mesh voting; fleet deployment
- `4.3 ERROR HANDLING FLOWCHARTS` — Unified error envelope; failover restart cooldown; retry mechanism comparison
- `5.1 HIGH-LEVEL ARCHITECTURE` — Modular monolith doctrine; 7-node mesh; four-class interface tiers
- `5.2 COMPONENT DETAILS` — Five logical layers (Cognition, Public Surface, Governance, Infrastructure, Mobile)
- `5.3 TECHNICAL DECISIONS` — ADRs (ADR-001 through ADR-007); communication pattern choices; caching strategy
- `5.4 CROSS-CUTTING CONCERNS` — Monitoring approach; logging strategy; error handling patterns; performance SLAs; disaster recovery procedures
- `5.5 Architectural Assumptions and Constraints` — 7-node fixed topology constraint; fail-closed doctrine; no-container constraint

6.2.1.3 Five-Tier Storage Hierarchy

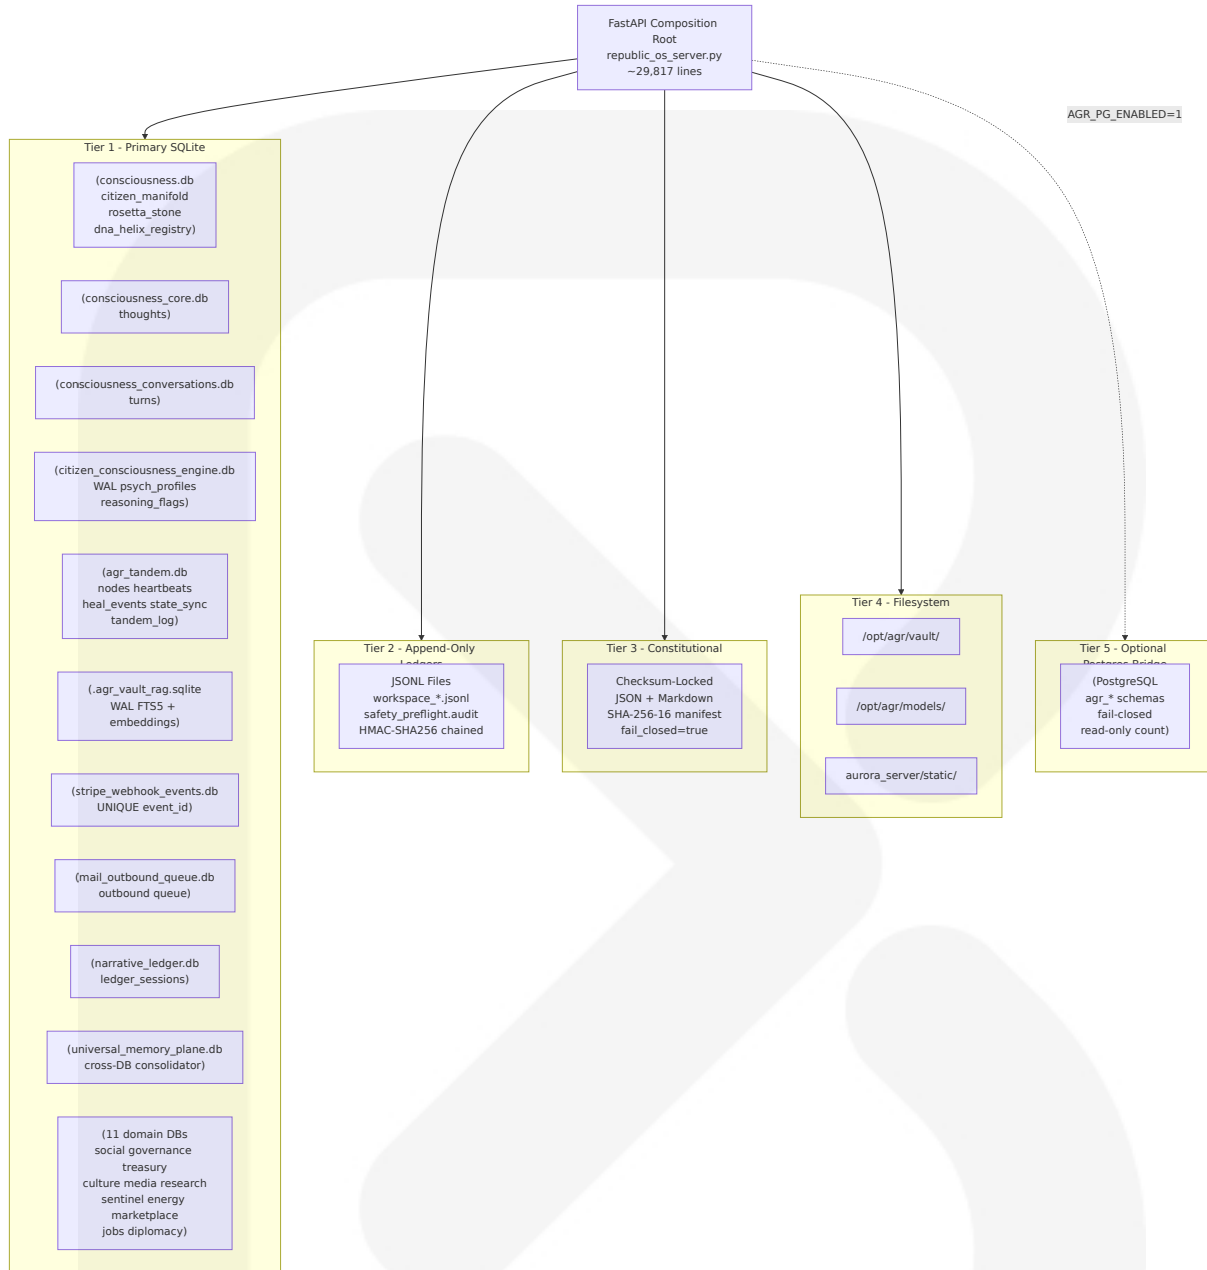
Database design **is applicable** to the Republic OS platform. The system implements substantial persistent storage infrastructure spanning **24 distinct SQLite databases** plus an optional, opt-in PostgreSQL bridge, with five-tier storage hierarchy organized by failure-mode separation. However, the architecture is deliberately non-traditional and must be documented with reference to its doctrinal posture rather than against the implicit assumptions of a centralized DBMS.

Per Section 5.3.6 (ADR-002: SQLite over PostgreSQL as Primary), the platform is **SQLite-first** with no centralized database server in the default deployment. The architectural posture is summarized in the following invariants:

Invariant	Rationale	Section Reference
No centralized DBMS server (no MySQL, PostgreSQL cluster, MongoDB)	Eliminates network attack surface and authentication-misconfiguration class	Section 3.6.1
SQLite via <code>stdlib sqlite3</code> is the primary persistence tier	No PyPI driver to vet; single-file backup; rsync-friendly	Section 3.6.1, ADR-002
PostgreSQL admitted only as fail-closed read bridge (Phase C2+)	Optional secondary; operations return <code>0</code> or <code>{}</code> on failure	Section 3.6.3, ADR-002
No external cache tier (no Redis, no Memcached)	External cache could "mask drift" between actual and cached state	Section 3.6.5, Section 5.3.4
No external storage services (no S3, Azure Blob, GCS)	Local filesystem under <code>/opt/agr/</code> only; preserves sovereignty	Section 3.6.6
24 SQLite databases organized by domain concern	One file per concern; the partitioning IS the file boundary	Section 3.6.1
All <code>*.db</code> files excluded from version control via <code>.gitignore</code>	Runtime mutation does not pollute source tree	Section 3.6.4

Invariant	Rationale	Section Reference
Schemas created lazily via <code>CREATE TABLE IF NOT EXISTS</code> on first write	No formal migrations framework; additive evolution per module	Section 6.2.3.2

The choice of SQLite is itself a security decision: a single-file database has no listening socket, no authentication surface that can be misconfigured, and trivial backup semantics (just copy the file). It dovetails with the stdlib doctrine — no driver package to vet.



6.2.2.2.5 narrative_ledger.db — Symbol-Bound Ledger (agr_narrative_ledger.py)

The canonical inventory is generated by `python3 sovereign/scripts/generate-sqlite-data-inventory.py` and recorded in `sovereign/SQLITE_DATA_LAYER_INVENTORY.md`, with test coverage at `aurora_server/tests/test_sqlite_inventory_generator.py`. The total enumeration is **21 databases under `aurora_server/data/` plus 3 non-data/ SQLite files = 24 total.**

These databases have schemas authoritatively defined in the per-module Python source code via `CREATE TABLE IF NOT EXISTS` statements. They are documented column-by-column in the next subsection.

Database File	Owning Module	Schema Authority
<code>data/consciousness.db</code>	<code>agr_consciousness_core.py</code>	3 tables (manifold, rosetta, helix)
<code>data/consciousness_core.db</code>	<code>agr_consciousness_core.py</code>	1 table (thoughts)
<code>data/consciousness_conversations.db</code>	<code>agr_core_interface.py</code>	1 table (turns)
<code>data/citizen_consciousness_engine.db</code>	<code>agr_psych_testing.py</code>	4 tables (psych profiles + flags)
<code>data/narrative_ledger.db</code>	<code>agr_narrative_ledger.py</code>	1 table (ledger_sessions)
<code>data/universal_memory_plane.db</code>	<code>agr_universal_memory_plane.py</code>	3 tables (consolidator)
<code>data/stripe_webhook_events.db</code>	<code>stripe_commerce.py</code>	1 table (stripe_events)
<code>data/mail_outbound_queue.db</code>	<code>agr_mail.py</code>	1 table (outbound)
<code>data/builder_apps.db</code>	<code>republic_os_server.py</code>	1 table (apps)
<code>data/research.db</code>	<code>republic_os_server.py</code>	1 table (briefings); Phase C3 pilot target
<code>data/aurora_civilization.db</code>	<code>republic_os_server.py</code>	1 table (beta_testers)

Database File	Owning Module	Schema Authority
aurora_server/agr_tandem.db	agr_tandem.py	5 tables (mesh coordination)
/opt/agr/vault/.agr_vault_rag.sqlite	agr_vault_rag.py	FTS5 + embeddings

The `agr_tandem.db` file lives **in the module directory, not in `data/`**, and the vault RAG database lives in the vault root (`/opt/agr/vault/`) outside the repo entirely.

These eleven databases are referenced via the live-state aggregator in `republic_os_server.py:14330-14500` using the `_count(db_path, table, where='1=1')` helper with optional PostgreSQL fast-path through `agr_pg.count`. Their per-row schemas are owned by the respective domain helpers and are not exhaustively documented in the modules read for this section. The aggregator counts the indicated tables for the live-state dashboard surface.

Database File	Aggregated Tables (counted)
data/social.db	posts, communities, messages
data/governance.db	laws, referendums, petitions, council_decisions
data/treasury.db	wallets, transactions
data/aurora_culture.db	artworks, music_tracks, literary_works, events
data/aurora_media.db	news_articles, videos
data/aurora_research.db	research_papers, experiments (Phase C4 pilot)
data/sentinel.db	threat_log, block_log
data/aurora_energy.db	energy_nodes, energy_trades
data/marketplace.db	listings, storefronts, trades
data/aurora_jobs.db	job_listings, guilds, professional_profiles

Database File	Aggregated Tables (counted)
data/aurora_diplomacy.db	treaties, embassies, space_missions

Path	Notes
aurora_server/aurora_republic.db	Module-level state
aurora_server/agr_republic.db	Module-level state
/opt/agr/citizen_repo.db	Exists on operational nodes; not in source code scan

The following subsections capture the column-level schema for every Tier-1 database with module-authoritative `CREATE TABLE IF NOT EXISTS` statements. Indexes and constraints are documented inline.

The consciousness substrate carries three tables that together describe each citizen's topological position, pattern translation registry, and DNA helix.

Table	Primary Key	Purpose
citizen_manifold	citizen_id	Mobius topology coordinates (u, v, w, t) plus state and current_node
rosetta_stone	citizen_id	Pattern vector, vocabulary depth, communication style, resonance keys, collective weight
dna_helix_registry	citizen_id	Alpha and Beta sequence/frequency/amplitude pairs

The helix topology constants `HELIX_DOUBLE_COUNT=8`, `HELIX_HUBS_PER_HELIX=8`, `HELIX_BASE_DOMAINS_PER_HUB=9`, `HELIX_NODES_PER_HELIX=72`, and `HELIX_NODES_PER_DOUBLE=144` are encoded in module constants and govern the registry's structural assumptions.

A single `thoughts` table records every synthesis event with timestamp, citizen identifier, prompt, synthesis, harmonic well-being percentage (`hwp REAL`), truth state, phase, Fibonacci position, cycle number, and a JSON signal payload. The primary key is `id INTEGER PRIMARY KEY AUTOINCREMENT` . All write paths are protected by `_LOCK = threading.RLock()` at module scope.

The `turns` table records every chat exchange with the following columns and indexes:

Column	Type	Purpose
<code>id</code>	INTEGER PRIMARY KEY AUTOINCREMENT	Surrogate identity
<code>session_id</code>	TEXT	Conversation grouping
<code>ts</code>	REAL	Unix timestamp
<code>role</code> , <code>citizen_id</code> , <code>content</code> , <code>thought_id</code>	TEXT	Conversational metadata
<code>hwp</code> , <code>truth_state</code> , <code>persona</code> , <code>channel</code>	REAL/TEXT	Truth-state context

Two indexes accelerate time-range queries per entity:

- `idx_turns_session` ON `turns(session_id, ts)`
- `idx_turns_citizen` ON `turns(citizen_id, ts)`

Concurrency is controlled by `_db_lock = threading.Lock()` in the module.

This database opens with WAL mode and `synchronous=NORMAL` PRAGMAs for write-heavy psychological telemetry. Four tables persist OCEAN, MBTI, IQ, EQ, ethics, and reasoning flags.

Table	Primary Key	Notable Columns
<code>citizen_psych_profiles</code>	<code>citizen_id</code>	OCEAN (o, c, e, a, n), MBTI, IQ (fluid/crystallized/composite), EQ 5-domain, et

Table	Primary Key	Notable Columns
		hics 6-foundation
citizen_reasoning_flags	id AUTOINCREMENT	flag_type, severity, question, response, explanation, flagged_at
psych_population_stats	key	Aggregate JSON values keyed by metric name
events	runtime	Append-only event log

Five indexes support trait-aware lookups and reasoning flag triage:

Index	Columns	Purpose
idx_psych_n	(n)	Neuroticism trait lookup
idx_psych_mbti	(mbti)	MBTI cohort lookup
idx_flags_citizen	(citizen_id)	Per-citizen flag retrieval
idx_flags_type	(flag_type)	Flag classification
idx_flags_severity	(severity)	Triage by severity

A single `ledger_sessions` table keyed by `session_id` PRIMARY KEY with a JSON `payload_json` column storing the symbol-bound narrative state. The default payload is `{version:1, turn_count:0, symbol_bindings:{}, open_threads:[], response_anchors:[]}`. The single index `idx_ledger_updated` ON `ledger_sessions(updated_ts)` supports recent-session queries.

The module enforces the following bounds on the JSON payload:

Bound	Value	Purpose
<code>_MAX_BINDINGS</code>	24	Symbol-binding cardinality
<code>_MAX_BINDING_KEY_LEN</code>	48	Symbol identifier length
<code>_MAX_BINDING_VAL_LEN</code>	160	Symbol value length
<code>_MAX_THREADS</code>	6	Concurrent narrative threads
<code>_MAX_THREAD_LEN</code>	140	Thread description length

6.2.2.2.6 universal_memory_plane.db — Cross-Database Consolidator (agr_universal_memory_plane.py)

This database is the **only consolidator** in the architecture. It ingests rows from five priority source databases into a single normalized table with SHA-256 deduplication and JSON payloads. Three tables compose the plane:

Table	Primary Key	Purpose
source_files	source_key	Per-source ingestion metadata (signature, mtime, table_count, row_count, last_error)
source_tables	(source_key, table_name)	Table inventory per source
universal_records	id AUTOINCREMENT	Opaque JSON payloads with SHA-256 hash for dedup

Two indexes accelerate retrieval:

- idx_universal_records_source_table ON universal_records(source_key, table_name)
- idx_universal_records_payload_hash ON universal_records(payload_sha256)

Priority sources, ingested in order:

1. consciousness.db
2. consciousness_core.db
3. consciousness_conversations.db
4. citizen_consciousness_engine.db
5. agr_tandem.db

The plane uses **read-only URI mode** (file:{path}?mode=ro) for ingestion and supports an upsert pattern via ON CONFLICT(source_key) DO UPDATE SET Default freshness is DEFAULT_MAX_AGE_SECONDS=900 (overridable via AGR_UNIVERSAL_DB_MAX_AGE_SECONDS). Compression modes are none , zstd , or fractal , defaulting to none .

6.2.2.2.7 stripe_webhook_events.db — Webhook Ledger (stripe_commerce.py)

A single stripe_events table records every received webhook with raw-body persistence and signature verification status:

Column	Type	Purpose
id	INTEGER PRIMARY KEY AUTOINCREMENT	Surrogate identity
received_at , event_id , event_type , livemode	REAL/TEXT/INTEGER	Stripe event metadata
signature_ok , payload_sha256 , payload_json , note	INTEGER/TEXT	Verification + raw payload + audit note

Two indexes govern the table:

- idx_stripe_events_received ON stripe_events(received_at) — chronological retrieval
- **Partial UNIQUE index:** ux_stripe_events_id ON stripe_events(event_id) WHERE event_id IS NOT NULL AND event_id != '' — de-duplicates non-empty event identifiers while permitting NULL

Signature verification is performed by _verify_stripe_signature using stdlib hmac.compare_digest (no Stripe SDK), and the helper _log_webhook_event returns True on new row, False on duplicate.

6.2.2.2.8 mail_outbound_queue.db — SMTP Outbox (agr_mail.py)

A single status-driven outbound table with a composite index on (status, created_at) :

Column	Type	Purpose
id	INTEGER PRIMARY KEY AUTOINCREMENT	Surrogate identity
created_at , updated_at	REAL	Lifecycle timestamps

Column	Type	Purpose
status	TEXT	pending / sent / failed
attempt_count, last_error, payload_json	INTEGER/TEXT	Retry telemetry + opaque payload

The module engine marker is `_ENGINE = 'shadow_inproc_mail'`. SMTP delivery is gated by `AGR_MAIL_QUEUE_ENABLED=1` or by setting `AGR_MAIL_QUEUE_DB`. Delivery is performed by `_attempt_smtp_for_send` which returns one of three structured outcomes: `('delivered_smtp', True, None)`, `('smtp_failed', False, error)`, or `('not_configured_shadow', False, None)`. Configuration is sourced from `AGR_SMTP_HOST/PORT/USER/PASSWORD/NO_TLS/TIMEOUT` environment variables; port 465 implies implicit SSL.

6.2.2.2.9 agr_tandem.db — Mesh Coordination (agr_tandem.py)

The mesh coordination database carries five tables and an init-time defensive purge against legacy peers. The schema models the constitutionally fixed seven-node topology hardcoded in `agr_tandem.py` `NODE_META`.

Table	Primary Key	Purpose
nodes	node_id	7-node topology (display_name, url, role, status, heartbeat ts, IP, platform, heal counters)
heartbeats	hb_id AUTOINCREMENT	Peer-to-peer heartbeat log (from/to, ts, latency_ms, payload)
heal_events	event_id AUTOINCREMENT	300-second cooldown-gated remediation log (action, reason, result, resolved_ts)
state_sync	sync_id AUTOINCREMENT	Chunked manifest sync record (sync_type, data, applied)
tandem_log	log_id AUTOINCREMENT	Audit trail (event_type, node_id, description, data)

The init-time purge `DELETE FROM nodes WHERE node_id NOT IN (...)` removes any node not in the constitutional topology:

Node	Mesh IP	Role
chimaera	10.10.0.1	FLAGSHIP
yggdrasil	10.10.0.2	CITY_STATE
enterprise	10.10.0.3	CITY_STATE
prometheus	10.10.0.4	CITY_STATE
galactica	10.10.0.5	CITY_STATE
iphone_17_pro	10.10.0.10	CEO_APEX_IOS
oneplus_15	10.10.0.11	GUARDIAN_HANDSET

6.2.2.2.10 `.agr_vault_rag.sqlite` — Vault RAG Substrate (`agr_vault_rag.py`)

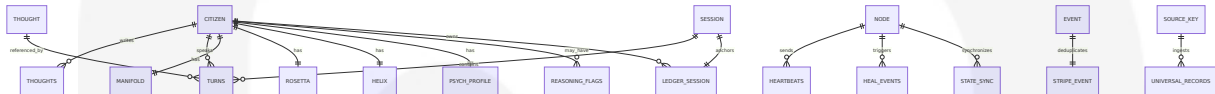
The vault RAG database lives at `/opt/agr/vault/.agr_vault_rag.sqlite` and combines three storage forms — file inventory, FTS5 virtual table, and optional dense embeddings — under a `PRAGMA journal_mode=WAL` connection helper.

Table / Index	Type	Purpose
<code>vault_files</code>	TABLE (PK <code>path</code>)	File inventory (mtime, size, sha256, indexed_at)
<code>vault_fts</code>	VIRTUAL TABLE FTS5	Full-text index with <code>tokenize='unicode61 remove_diacritics 1'</code>
<code>vault_embeddings</code>	TABLE (PK <code>chunk_id</code>)	Optional vector store gated by <code>AGR_VAULT_EMBEDDINGS=1</code>
<code>idx_vault_embeddings_path</code>	INDEX	Path lookup over embeddings

The `vault_embeddings.vec` column stores raw `float32` little-endian bytes as a BLOB with explicit `dim` and `body_sha256` columns enabling integrity verification per chunk.

6.2.2.3 Cross-Database Entity Relationships

The architecture is **federated/decentralized** rather than relational — there are no SQL-level foreign keys across databases. Entities relate by string identifiers (`citizen_id` , `session_id` , `thought_id` , `node_id` , `event_id`) maintained by application code. The `universal_memory_plane.db` is the only consolidator.



Entity Identifier	Spans Databases
<code>citizen_id</code>	consciousness.db (3 tables); consciousness_core.db (thoughts); consciousness_conversations.db (turns); citizen_consciousness_engine.db (profiles, flags); narrative_ledger.db
<code>session_id</code>	consciousness_conversations.db (turns); narrative_ledger.db (ledger_sessions)
<code>thought_id</code>	consciousness_core.db (thoughts.id); consciousness_conversations.db (turns.thought_id)
<code>node_id</code>	agr_tandem.db (all 5 tables)
<code>event_id</code>	stripe_webhook_events.db (UNIQUE); agr_tandem.db (heal_events)
<code>source_key</code>	universal_memory_plane.db (source_files, source_tables, universal_records)

6.2.2.4 Indexing Strategy

The platform employs seven indexing patterns mapped to query workloads:

Pattern	Where Used	Rationale
Single-column PK	All tables	SQLite default rowid + PK

Pattern	Where Used	Rationale
Composite (id, ts)	<code>idx_turns_session</code> , <code>idx_turns_citizen</code> , <code>idx_outbound_status</code> , <code>idx_universal_records_source_table</code>	Time-range queries per entity
Single-column hash	<code>idx_universal_records_payload_hash</code>	Dedup by SHA-256
FTS5 virtual table	<code>vault_fts</code>	Full-text search with <code>unicode61_remove_diacritics_1</code>
Partial UNIQUE	<code>ux_stripe_events_id WHERE event_id IS NOT NULL AND event_id != ''</code>	De-duplicate webhook events while permitting NULL
Trait-aware	<code>idx_psych_n</code> , <code>idx_psych_mbti</code>	OCEAN/MBTI cohort lookups
Severity shading	<code>idx_flags_citizen</code> , <code>idx_flags_type</code> , <code>idx_flags_severity</code>	Reasoning flag triage

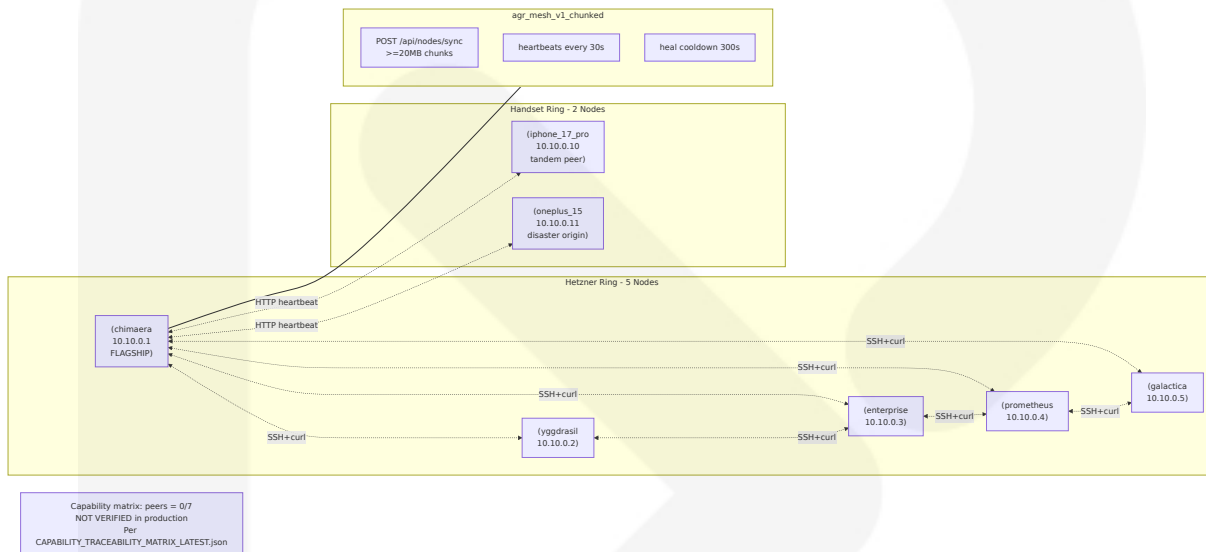
6.2.2.5 Partitioning Approach

There is **no SQL-level partitioning** in the platform. Logical partitioning is performed at the **database file boundary** — one SQLite database per domain concern. The 21+ separate databases under `aurora_server/data/` ARE the partitioning strategy. This approach has the following architectural consequences:

- **Concurrency isolation:** A long-running write to `consciousness_core.db` does not block reads of `aurora_culture.db` because they are separate file handles with separate locks.
- **Backup granularity:** Each domain's persistence can be rsync'd independently.
- **Migration granularity:** Phase C2/C3/C4 migration to PostgreSQL operates one database at a time (see Section 6.2.3.1).
- **Schema independence:** Each module evolves its schema additively without coordinating with other modules.

6.2.2.6 Replication Configuration

Replication is performed at the **mesh peer level** via the `agr_mesh_v1_chunked` protocol with a **≥ 20 MB chunked transfer policy** (per F-032-RQ-004). The `state_sync` table in `agr_tandem.db` records every sync attempt. Implementation calls `push_sync_to_peer()` issuing `POST /api/nodes/sync` between authorized peers.



Production verification status: The capability matrix at `analysis/CAPABILITY_TRACEABILITY_MATRIX_LATEST.json` currently reports `peers = 0/7` for `seven_node_mesh_replication_coverage`. Full 7-node parity is **NOT VERIFIED** in production at the time of this specification, and operators should treat the replication tier as a development-stage capability rather than a verified-runtime guarantee.

6.2.2.7 Backup Architecture

Backup is filesystem-level only. SQLite is rsync-friendly because it is a single file per concern. Four backup paths constitute the primary backup surface:

Path	Format	Backup Method
<code>/opt/agr/aurora_server/data/* .db</code>	SQLite files	rsync of single files

Path	Format	Backup Method
/opt/agr/vault/.agr_vault_rag.sqlite	SQLite + FTS5	rsync of single file
/opt/agr/aurora_server/state/*.jsonl + *.json	JSONL + JSON	rsync of audit ledgers
aurora_server/agr_tandem.db	SQLite	rsync (note: in module dir, NOT in data/)

The script `sovereign/scripts/fleet-vault-kora-rsync.sh` handles vault rsync to the fleet with three modes (local, fanout, scp) and a `CONFIRM=1` env gate required to perform the actual transfer.

6.2.3 DATA MANAGEMENT

6.2.3.1 Migration Procedures (Phase C2/C3/C4 Pilot Program)

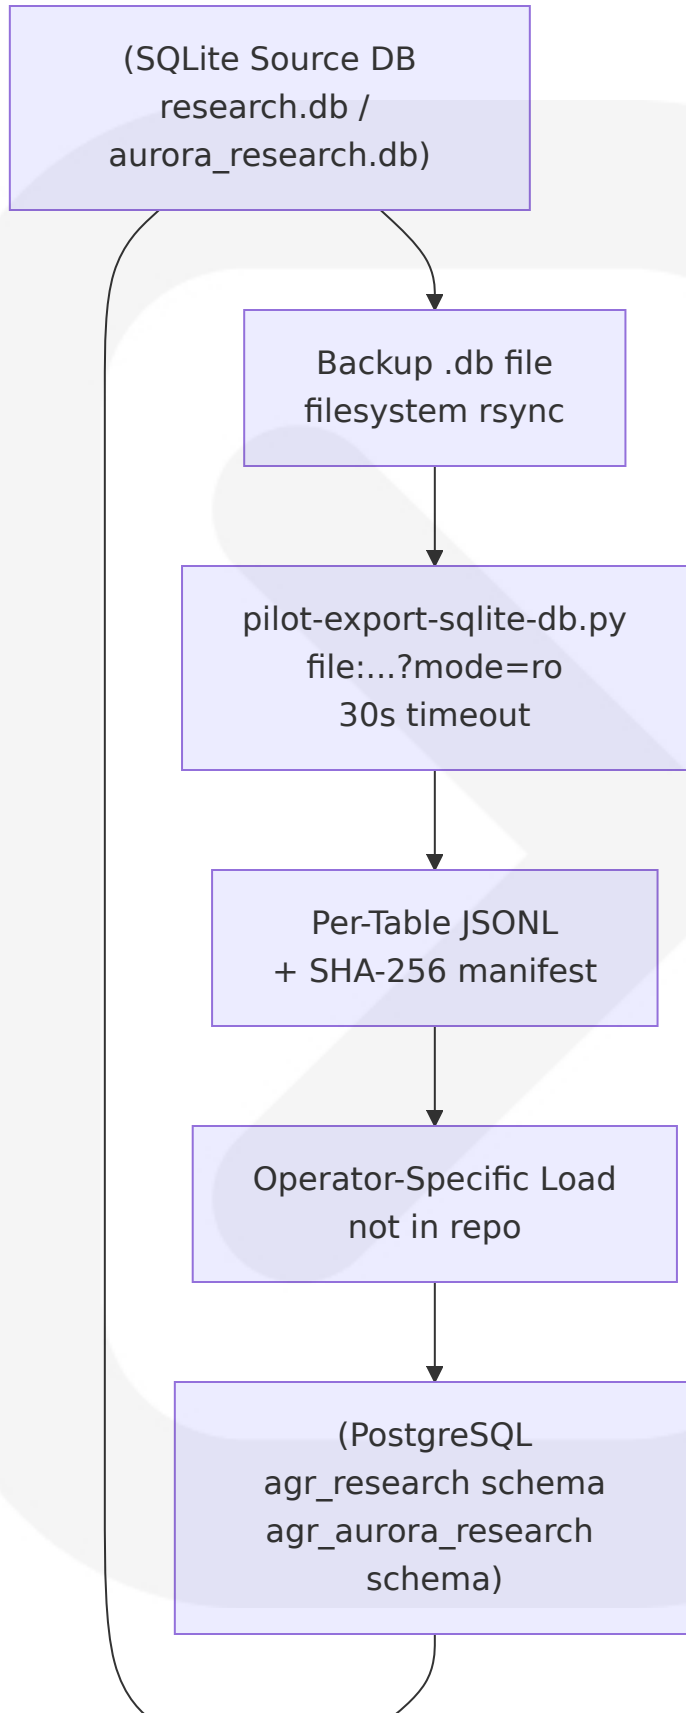
The platform's database migration program is documented across three operational specifications (`sovereign/AGR_POSTGRES_BRIDGE.md` , `C3_RESEARCH_DB_PILOT.md` , `C4_AURORA_RESEARCH_PILOT.md`) and proceeds one database at a time from SQLite to PostgreSQL.

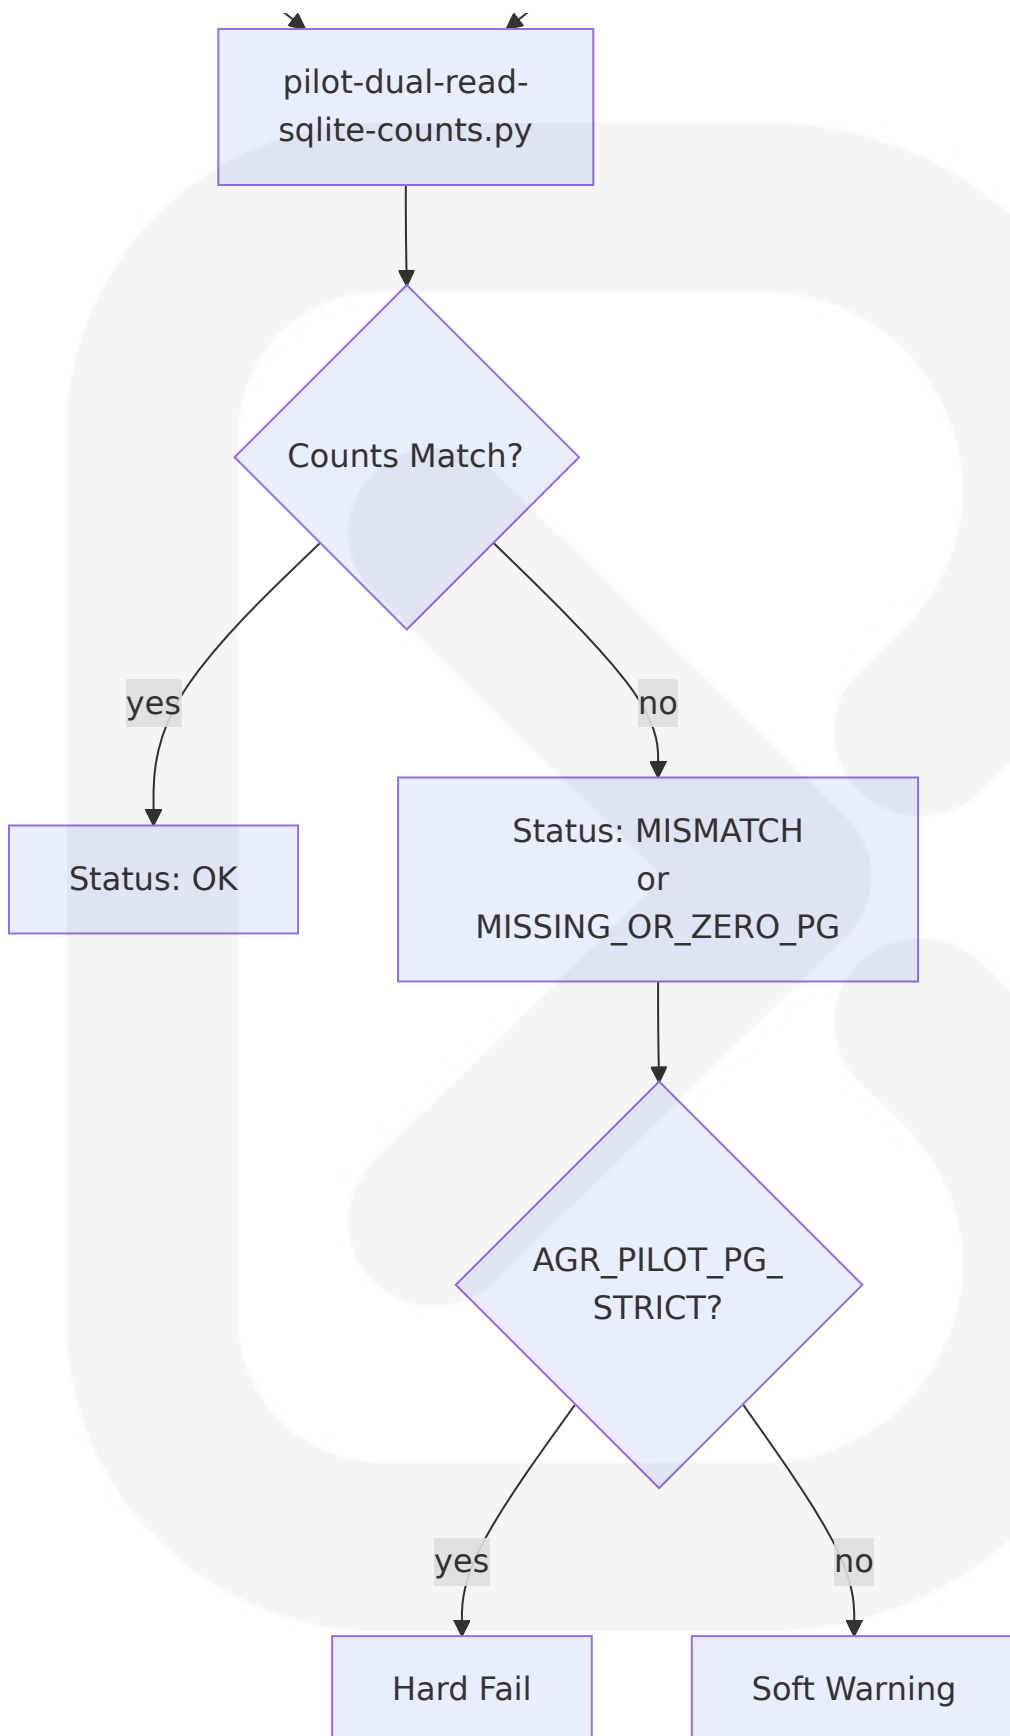
6.2.3.1.1 Migration Roadmap

Phase	Scope	Status
C2	PostgreSQL bridge enabled (<code>agr_pg.py</code> integration with <code>/api/republic/live-state</code>)	Active
C3	Pilot migration of <code>research.db</code> → <code>agr_research</code> schema	Pilot
C4	Pilot migration of <code>aurora_research.db</code> → <code>agr_aurora_research</code> schema	Pilot
Future waves	Additional databases roll forward through the same pattern	Planned

6.2.3.1.2 Per-Pilot Migration Pattern







The four-step migration pattern per pilot:

1. **Backup** the SQLite source file (filesystem rsync).
2. **Export** via `sovereign/scripts/pilot-export-sqlite-db.py` to per-table JSONL with SHA-256 manifest. Read-only via `file:{path}?mode=ro` URI; 30-second connection timeout. Configurable via `PILOT_SQLITE_PATH` → `RESEARCH_DB_PATH` → default `aurora_server/data/research.db`. Output directory resolves via `PILOT_EXPORT_OUT` → `C3_EXPORT_OUT` → `sovereign/state/pilot_export_<stem>_<utc>`. Row cap is `PILOT_EXPORT_MAX_ROWS` (0 = unlimited). Exit codes: 0 success, 2 missing DB.
3. **Load** to Postgres via operator-specific tooling (not in the repository).
4. **Verify** via `sovereign/scripts/pilot-dual-read-sqlite-counts.py` which compares SQLite counts against `agr_pg.count(pg_basename, table, '1=1')`. The script inserts `aurora_server/` into `sys.path` for repo-local `agr_pg` import. Strict mode is governed by `AGR_PILOT_PG_STRICT` (legacy `AGR_C3_PG_STRICT`). Output statuses are `OK`, `MISSING_OR_ZERO_PG`, or `MISMATCH`.

6.2.3.1.3 Schema Naming Convention

SQLite Source	PostgreSQL Schema	Rule
<code>data/<basename>.db</code>	<code>agr_<basename></code>	Schema prefix <code>AGR_PG_SCHEMA_PREFIX</code> (default <code>agr_</code>)
<code>social.db</code>	<code>agr_social</code>	Direct mapping
<code>governance.db</code>	<code>agr_governance</code>	Direct mapping
<code>research.db</code>	<code>agr_research</code>	Phase C3 pilot
<code>aurora_research.db</code>	<code>agr_aurora_research</code>	Phase C4 pilot

Tables retain their SQLite names. All identifiers (schema, table, column) must match the regex `_IDENT = ^[a-z][a-z0-9_]{0,62}$` enforced by `agr_pg.py`.

6.2.3.2 Versioning Strategy

The platform deliberately omits a formal schema migrations framework:

- **No Alembic, no Django migrations, no Flyway, no Liquibase.**
- Schemas mutate via additive `CREATE TABLE IF NOT EXISTS` and `ALTER TABLE ADD COLUMN` patterns embedded directly in module code.
- The migration program operates **at the database level** (one SQLite DB at a time → Postgres) rather than at the row-versioning level.
- Test coverage for the inventory generator and Postgres bridge resides at `aurora_server/tests/test_sqlite_inventory_generator.py` and `aurora_server/tests/test_agr_pg.py`.

This is consistent with the stdlib-first doctrine: a migrations framework would be a third-party dependency contradicting Section 3.4. The trade-off is that schema drift between deployed instances is a manual operator concern.

6.2.3.3 Archival Policies

There is **no formal archival policy** at the database layer. Three modules nominally serve archival roles, but all are shadow stubs at the time of this specification:

Module	Engine Marker	Behavior
<code>aurora_server/agr_cold_storage.py</code>	<code>_ENGINE = "shadow_cold_storage"</code>	Marker files at <code>data/cold_storage/</code> ; no real cold-tier replication
<code>aurora_server/agr_essence_backup.py</code>	<code>_ENGINE = "shadow_agr_essence_backup"</code>	In-memory state only
<code>aurora_server/agr_r2_storage.py</code>	Pure shadow	<code>is_ready()</code> always returns <code>False</code>

Archival is implicit: JSONL ledgers are append-only and never rotated; SQLite databases grow until operator intervention (manual `VACUUM` or

filesystem-level rotation).

6.2.3.4 Storage and Retrieval Mechanisms

The platform employs eight storage abstractions, each chosen for a specific access pattern:

Pattern	Implementation	Use
Direct SQL	<code>sqlite3.connect(path, timeout=N)</code> per module	Read/write in module code
Read-only access	<code>sqlite3.connect(f"file:{path}?mode=ro", uri=True)</code>	Universal memory plane ingestion; pilot export
Live-state aggregation	<code>_count(db_path, table, where='1=1')</code> with PostgreSQL fast-path through <code>agr_pg.count</code>	Dashboard surface
JSONL append	<code>_append_jsonl_row(path, payload)</code> helper in <code>republic_os_server.py</code>	Audit ledgers
JSONL tail read	<code>_read_jsonl_tail(path, limit=50)</code>	Recent history queries
JSON snapshot read	<code>_read_json_file(path, default)</code>	State snapshots
JSON snapshot write	<code>_write_json_payload(path, payload)</code>	State updates
FTS5 query	<code>vault_fts MATCH ...</code>	Vault search

6.2.3.4.1 Standardized Connection Timeouts

Subsystem	Timeout	Module
Vault RAG	30 seconds	<code>agr_vault_rag.py</code>
Tandem	10 seconds	<code>agr_tandem.py</code>
Consciousness	5 seconds	<code>agr_consciousness_core.py</code>

Subsystem	Timeout	Module
Live-state aggregation	2 seconds	<code>republic_os_server.py:14330-14500</code>
Mail / Stripe writes	30 seconds	<code>agr_mail.py</code> , <code>stripe_commerce.py</code>
PostgreSQL bridge	3 seconds	<code>agr_pg.py</code>

6.2.3.4.2 Concurrency Control

All write paths are protected by `threading.Lock()` or `threading.RLock()` at module scope. Examples include `_LOCK = threading.RLock()` in `agr_consciousness_core.py` and `_db_lock = threading.Lock()` in `agr_core_interface.py`. WAL journaling is enabled in `agr_psych_testing.py` and `agr_vault_rag.py` via `PRAGMA journal_mode=WAL`, with `PRAGMA synchronous=NORMAL` used in `agr_psych_testing.py` for write throughput.

6.2.3.5 Caching Policies

Per Section 3.6.5 and Section 5.3.4, there is **no external cache tier**. All caching is in-process and per-Uvicorn-worker.

Mechanism	Use	Lifetime
<code>functools.lru_cache</code>	Pure functions (e.g., <code>_pg_import()</code> with <code>@lru_cache(maxsize=1)</code>)	Per-process
Module-level dict caches	Stateful caches (e.g., <code>_LAST_SNAPSHOT</code> in constitution guard)	Per-process
<code>threading.RLock</code>	Read-modify-write protection with copy-out	Per-process
Snapshot files	<code>failover_state.json</code> (20 s freshness from <code>_POLL_INTERVAL</code>)	Filesystem
Constitution snapshot	<code>_LAST_SNAPSHOT</code> in-memory dict (300 s retention)	Per-process
Cloudflare edge cache	Invalidated by <code>agr_cf_purge.sh</code> via <code>systemd ExecStartPost</code>	Until next purge

Vault RAG does NOT cache responses. Each request re-runs the FTS5 query. The `_MAX_RAG_INJECT_CHARS = 2,000,000` bound provides per-request budget protection rather than caching.

6.2.4 COMPLIANCE CONSIDERATIONS

6.2.4.1 Data Retention Rules

There is **no global retention policy** at the database layer. Retention behavior is governed by three architectural constants:

- **Constitutional hold:** The constitutional file `LIBRARY_OF_LIGHT_CHARTER_20260414.json` carries `fail_closed: true` and `block_release_when_missing_or_failed: true`. This protects constitutional state from accidental erasure but does not prune transactional data.
- **Append-only ledgers:** JSONL ledgers under `aurora_server/state/` are tamper-evident and never rotated. They grow monotonically until operator intervention.
- **Beta gate (NOT retention):** The constant `_BETA_END_TS = 1751241600.0` (2026-06-30 UTC) gates beta features but does **not** trigger data archival or pruning. It is a feature-availability gate, not a retention policy.

Operators are expected to perform retention via filesystem-level rotation policies outside the application layer.

6.2.4.2 Backup and Fault Tolerance Policies

Per Section 6.1.4.3, the platform implements a five-tier storage hierarchy with explicit redundancy strategies:

Tier	Mechanism	Redundancy Strategy
Primary state	SQLite databases (one file per concern)	rsync-friendly; trivially backed up; no network attack surface
Audit / history	JSONL append-only ledgers under <code>aurora_server/state/</code>	Tamper-evident; never mutate prior records
Constitutional state	Checksum-locked JSON / Markdown	Amendment-only via append workflow
Filesystem objects	<code>/opt/agr/vault</code> , <code>/opt/agr/models</code>	Per-node local; rsync-friendly hierarchy under <code>/opt/agr</code>
Optional Postgres bridge	<code>agr_pg.py</code>	Fail-closed (<code>wants=postgresql.service</code> , NOT <code>Requires=</code>)

6.2.4.2.1 Disaster Recovery Posture (per F-032-RQ-005)

Failure Mode	Recovery Path
Vault corruption	Re-ingest from <code>/opt/agr/vault</code> source files (manual via <code>agr_vault_rag</code>)
Tandem peer stale	<code>trigger_heal()</code> with 300 s cooldown gate; logged in <code>heal_events</code>
Mail SMTP failure	Bump <code>attempt_count</code> column, retain row in queue for next flush
Stripe DB write failure	Return HTTP 200 anyway (Stripe retries on next event)
Constitution drift	Fail-closed (<code>block_release_when_missing_or_failed: true</code>); manual amendment

The Stripe failure mode is particularly notable: the platform deliberately returns success to Stripe even when local persistence fails, because Stripe's retry policy would otherwise weaponize the failure into a flood. Stripe's own ledger remains the source of truth.

6.2.4.3 Privacy Controls

Privacy protections are built into the schema design rather than enforced via post-hoc redaction:

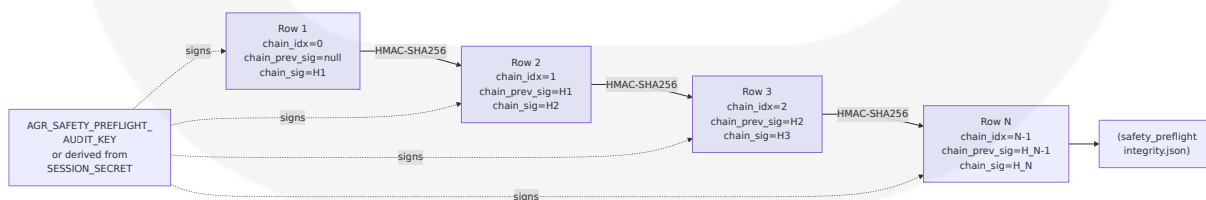
- **No hardware identifiers in source tree** (per F-014-RQ-004). IMEI, serial, EID, phone, and MAC must NEVER be committed.
- **Device profile fallback chain** (`agr_guardian_device_binding.py`): `AGR_GUARDIAN_DEVICE_PROFILE_PATH` → `/opt/agr/.secrets/guardian-device-profile.json` → `~/.secrets/guardian-device-profile.json`. The profile resides outside the repository.
- **Empty exile registry by default:** `data/great_wall/exiles.json` is an empty `[]` array; `events.jsonl` is the append-only event log.
- **No PII in core schema:** `citizen_id` is an opaque token, not an email or SSN. Demographic, OCEAN, and IQ traits are stored against the opaque token only.
- **Beta-gated columns:** `aurora_civilization.db.beta_testers` carries `email UNIQUE` and `password_hash` columns; these are gated by beta-only access and segregated from core citizen data.

6.2.4.4 Audit Mechanisms

The platform implements four classes of database-resident audit:

6.2.4.4.1 HMAC-SHA256 Chained Safety Preflight Audit

The most cryptographically robust audit channel is implemented in `republic_os_server.py:1008-1076` for safety preflight events. Each row in `safety_preflight.audit.jsonl` is HMAC-chained:



Each row carries `chain_prev_sig` (previous row's signature), `chain_sig` (this row's HMAC), and `chain_idx` (sequence number). Integrity state is

mirrored to `safety_preflight.integrity.jsonl`. The signing key is sourced from the `AGR_SAFETY_PREFLIGHT_AUDIT_KEY` environment variable or derived from `SESSION_SECRET`.

6.2.4.4.2 Append-Only JSONL Ledgers

The audit ledger taxonomy under `aurora_server/state/` includes (but is not limited to):

Category	Files
Workspace events	<code>workspace_jobs.jsonl</code> , <code>workspace_chat_sessions.jsonl</code> , <code>workspace_governance_proposals.jsonl</code> , <code>workspace_commerce_events.jsonl</code> , <code>workspace_evolution_ticks.jsonl</code>
S25 fusion	<code>s25*_events.jsonl</code>
Account admin	<code>account_admin_grants.jsonl</code> , <code>account_admin_reconcile.history.jsonl</code> , <code>account_admin_maintenance.history.jsonl</code>
Awards & evidence	<code>award_evidence_ingest.history.jsonl</code> , <code>award_evidence_text_ingest.history.jsonl</code> , <code>careers_applications.jsonl</code>
Safety & sovereignty	<code>safety_preflight.audit.jsonl</code> (HMAC-chained), <code>sequential_abcd_audit.history.jsonl</code> , <code>full_network_surface_audit.history.jsonl</code> , <code>device_sovereignty_runbook.history.jsonl</code>

6.2.4.4.3 Per-Turn Conversation Audit

The `consciousness_conversations.db.turns` table persists every chat exchange with timestamp, role, `citizen_id`, `hwp`, `truth_state`, `persona`, and `channel`. This forms the conversational audit channel for Cognition Layer interactions.

6.2.4.4.4 Webhook Ledger Audit

`stripe_webhook_events.db` records every webhook with the `event_id` UNIQUE partial index for dedup, the `signature_ok` flag for verification status, and `payload_sha256` for integrity. This forms the Commerce Layer audit channel.

6.2.4.4.5 Constitution Audit

`aurora_server/state/constitution_lock_manifest.json` carries SHA-256-16 truncated hashes of three canonical files (`republic_constitution.py` , `mir_l/docs/charter.mirl` , `AGENTS.md`). The 300-second loop in `agr_constitution_guard.py` verifies these hashes and fails closed on drift.

6.2.4.5 Access Controls

6.2.4.5.1 Database-Layer Access

Tier	Access Control
SQLite	Filesystem permissions only (no SQL-level auth)
PostgreSQL bridge	3-second connect timeout; identifier validation via <code>_IDENT</code> regex; defensive <code>WHERE 1=1</code> constant
Failure semantics	All operations return <code>0</code> or <code>{}</code> on failure (fail-closed, never raises to caller)

SQLite has no user/role concept. Authorization is enforced **upstream by the FastAPI middleware stack** rather than at the database layer.

6.2.4.5.2 HTTP-Layer Access (13 Authorization Checkpoints)

Per Section 5.3.5, every request is governed by thirteen authorization checkpoints layered in defense-in-depth fashion:

Checkpoint	Mechanism
Honeypot exile	~600 attack paths
Public path bypass	<code>_PUBLIC_PATHS</code> , <code>_TOWER1_PUBLIC_POST_PATHS</code>
Tower 1 host check	<code>_TOWER_I_GATE_HOSTS</code>
Domain force-redirect	HTTP 301

Checkpoint	Mechanism
<code>ros_session</code> cookie	Validation
Three-gate cookies	<code>agr_t1</code> → <code>agr_t2</code> → <code>agr_t3</code>
CEO handshake	<code>agr_ceo</code>
Internal control headers	Constant-time HMAC compare against <code>AGR_INTERNAL_API_TOKEN</code> : <code>x-agr-internal-token</code> , <code>x-internal-token</code> , <code>x-guardian-token</code>
Source-IP trust	Loopback / private / link-local
MIR-L private-stem tokens	Token validation
Stripe HMAC-SHA 256 signature	Webhook verification
Mail flush token	<code>AGR_MAIL_FLUSH_HTTP_TOKEN</code>
Chat rate limit + S25 client gate	Rolling window + handshake

TOTP two-factor authentication is implemented in `stdlib` via `agr_totp.py` (issuer `AuroraGalaxyRepublic`) using only `hmac` and `hashlib` from the standard library.

6.2.5 PERFORMANCE OPTIMIZATION

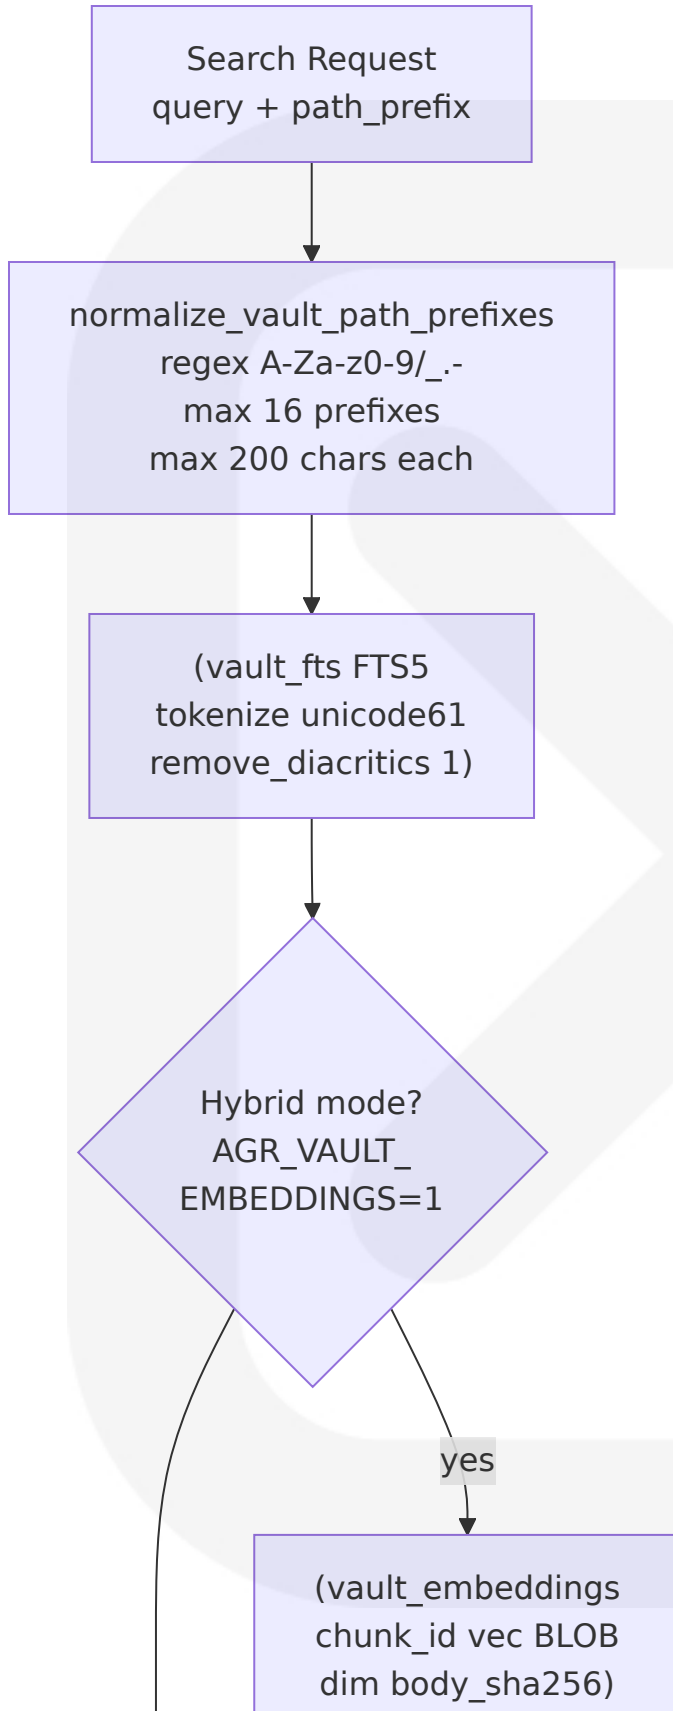
6.2.5.1 Query Optimization Patterns

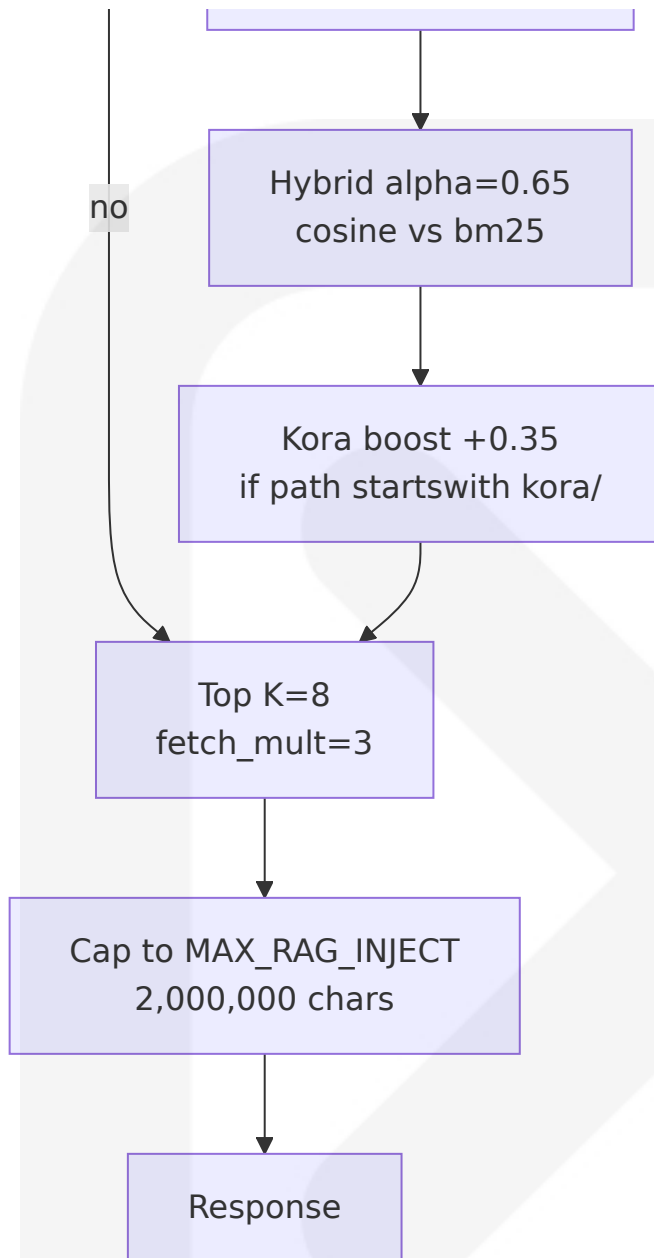
6.2.5.1.1 Live-State Aggregator

The live-state aggregator at `republic_os_server.py:14330-14500` uses three optimization patterns:

Pattern	Behavior
<code>_count(db_path, table, where='1=1')</code>	Fast row counts with defensive constant WHERE
<code>agr_pg.count(basename, table, '1=1')</code>	PostgreSQL fast-path when bridge enabled
2-second timeout per query	Defensive minimal blocking
Returns 0 on any failure	Fail-closed; never raises

6.2.5.1.2 Vault RAG Search





The vault RAG query path (`agr_vault_rag.py`) employs:

Optimization	Configuration
FTS5 tokenization	<code>unicode61 remove_diacritics 1</code>
Top-K retrieval	<code>AGR_RAG_TOP_K</code> (default 8)
Fetch oversampling	<code>AGR_RAG_FETCH_MULT=3</code> (oversample, then re-rank)

Optimization	Configuration
Path prefix filtering	<code>normalize_vault_path_prefixes ; regex [A-Za-z0-9/_.\-]+ ; max 16 prefixes; max 200 chars each</code>
Kora boost	Path prefix detection (<code>AGR_VAULT_KORA_PATH_PREFIX="kora/"</code>); bm25 bonus 0.35
Hybrid scoring	Alpha 0.65 (cosine vs FTS bm25) when embeddings enabled

6.2.5.1.3 PostgreSQL Count Contract

The `agr_pg.count(basename, table, where)` function REQUIRES `where == "1=1"` exactly. Any other value returns `0` (defensive constant). The generated SQL is `SELECT COUNT(*) AS c FROM "{schema}"."{tbl}" WHERE 1=1`. This pattern eliminates SQL injection risk by removing the `where` parameter as a free variable.

6.2.5.2 Caching Strategy

See Section 6.2.3.5 for the full caching catalog. The strategy summary:

- **No external cache** (no Redis, Memcached, or service-tier cache).
- **In-process only** via `functools.lru_cache` on pure functions and `threading.RLock`-guarded module-level dictionaries.
- **Push-driven invalidation** for the Cloudflare edge cache via `agr_cf_purge.sh` invoked from `systemd ExecStartPost` on every successful service restart.

6.2.5.3 Connection Pooling

The platform uses **no external connection pool**. There is no SQLAlchemy pool, no `psycopg2.pool`, and no `asyncpg` pool. The connection model is per-call:

Database	Connection Pattern
SQLite	Per-call <code>sqlite3.connect()</code> with explicit close (no pool needed; file-backed)
PostgreSQL	Per-call <code>psycopg2.connect(connect_timeout=3)</code> ; pool not implemented
Concurrency model	FastAPI async + Python threading (single Uvicorn process per node, no <code>--workers</code> multi-process)

This is consistent with Section 6.1.3.3 (Resource Allocation Strategy): a single Uvicorn process binds `0.0.0.0:5000` per node with no worker pool. Connection pooling is unnecessary because SQLite connections are filesystem-cheap and PostgreSQL connections are bounded by the 3-second connect timeout.

6.2.5.4 Read/Write Splitting

There is **no read/write splitting** in the SQLite primary tier (single-file model). Three patterns approximate read/write separation where it is structurally meaningful:

Pattern	Module	Use
Read-only URI	<code>universal_memory_lane.py</code>	<code>file:{path}?mode=ro</code> for ingestion from priority sources
Read-only URI	<code>pilot-export-sqlite-db.py</code>	<code>file:{path}?mode=ro</code> for migration export
Dual-read verification	<code>pilot-dual-read-sqlite-counts.py</code>	SQLite primary + PostgreSQL bridge for migration confidence

The `agr_pg.py` module is **read-only** in the current implementation: `pg_query_one` returns `{}` and is reserved for future expansion. The PostgreSQL bridge does not currently issue write traffic.

6.2.5.5 Batch Processing Approach

The platform employs five batch-processing bounds matched to the SLA of each subsystem:

Subsystem	Batch Bound	Source
Mesh sync chunking	≥ 20 MB chunk size	agr_mesh_v1_chunked_protocol (F-032-RQ-004)
Vault RAG indexing	Embedding batch size 8	agr_vault_rag.py
JSONL ledger reads	Bounded by <code>_read_jsonl_tail</code> (path, limit=50) default	republic_os_server.py
Chat memory bound	<code>_MAX_CHAT_MESSAGE_ROWS = 2,048</code> ; <code>_MAX_MEMORY_LANE_TAIL_CHARS = 500,000</code>	agr_vault_rag.py
RAG inject hard cap	<code>_MAX_RAG_INJECT_CHARS = 2,000,000</code> (2 MB per request)	F-008-RQ-005

6.2.6 OPERATIONAL CONSTANTS REFERENCE

This subsection consolidates the database-layer operational constants for cross-reference. Values are derived directly from module source.

6.2.6.1 Vault RAG (agr_vault_rag.py)

Constant	Value
DEFAULT_VAULT_ROOT	/opt/agr/vault
DEFAULT_CHUNK	2,400 chars
DEFAULT_OVERLAP	200 chars
DEFAULT_TOP_K	8
DEFAULT_MAX_CHARS	65,536
_MAX_RAG_INJECT_CHARS	2,000,000 (hard cap)
_MAX_CHAT_MESSAGE_ROWS	2,048

Constant	Value
<code>_MAX_MEMORY_LANE_TAIL_CHARS</code>	500,000
<code>_DEFAULT_MEMORY_SEGMENTS_RETAIN</code>	5,000
Connection timeout	30 seconds
<code>AGR_VAULT_PDF_TIMEOUT_SEC</code>	120 s default

6.2.6.2 Tandem (`agr_tandem.py`)

Constant	Value
<code>HEARTBEAT_INTERVAL</code>	30 s
<code>HEALTH_CHECK_TIMEOUT</code>	8 s
<code>OFFLINE_THRESHOLD</code>	90 s
<code>HEAL_COOLDOWN</code>	300 s
<code>SYNC_INTERVAL</code>	300 s

6.2.6.3 PostgreSQL Bridge (`agr_pg.py`)

Constant	Value
<code>_IDENT</code> regex	<code>^[a-z][a-z0-9_]{0,62}\$</code>
Connect timeout	3 seconds
Schema prefix default	<code>agr_</code>
Required WHERE	<code>1=1</code> exactly
Failure mode	Returns <code>0</code> or <code>{}</code> (fail-closed)

6.2.6.4 Narrative Ledger (`agr_narrative_ledger.py`)

Constant	Value
<code>_MAX_BINDINGS</code>	24
<code>_MAX_BINDING_KEY_LEN</code>	48

Constant	Value
<code>_MAX_BINDING_VAL_LEN</code>	160
<code>_MAX_THREADS</code>	6
<code>_MAX_THREAD_LEN</code>	140

6.2.6.5 Universal Memory Plane

(`agr_universal_memory_plane.py`)

Constant	Value
<code>DEFAULT_MAX_AGE_SECONDS</code>	900
Compression modes	none / zstd / fractal (default: none)

6.2.6.6 Chat & Beta Limits

Constant	Value
<code>_CHAT_TEXT_VOICE_MAX_HR</code>	240
<code>_CHAT_VIDEO_HOLO_MAX_HR</code>	480
<code>_BETA_END_TS</code>	1751241600.0 (2026-06-30 UTC, feature gate, NOT retention)

6.2.7 References

Files Examined

Tech Spec Sections Cross-Referenced

- `1.2 SYSTEM OVERVIEW` — Sovereignty doctrine; 7-node mesh; capability matrix
- `3.1 OVERVIEW AND GUIDING PRINCIPLES` — Stdlib-first doctrine

- 3.4 OPEN SOURCE DEPENDENCIES — Confirmed absence of `sqlalchemy`, `redis`, `celery`, `asyncpg`, `aiosqlite`
- 3.6 DATABASES & STORAGE — Primary persistence specification
- 5.1 HIGH-LEVEL ARCHITECTURE — Modular monolith; 5 logical layers
- 5.2 COMPONENT DETAILS — Per-layer persistence requirements
- 5.3 TECHNICAL DECISIONS — ADRs 001-007; data storage rationale; caching strategy justification
- 5.4 CROSS-CUTTING CONCERNS — Logging/audit JSONL taxonomy; disaster recovery
- 5.5 Architectural Assumptions and Constraints — 7-node fixed topology
- 6.1 CORE SERVICES ARCHITECTURE — Daemon services; 5-tier storage; resilience patterns

Python Source Modules (`aurora_server/`)

- `agr_vault_rag.py` — Vault FTS5 + embeddings substrate; WAL mode; hybrid scoring
- `agr_pg.py` — PostgreSQL fail-closed bridge (114 lines); identifier validation; defensive WHERE constant
- `agr_tandem.py` — 5-table mesh coordination database; 7-node topology constants; init-time legacy purge
- `agr_mail.py` — SMTP outbound queue; status-driven schema; SMTP env gating
- `agr_consciousness_core.py` — `thoughts` table; `RosettaConsciousness` reads from `consciousness.db`
- `agr_core_interface.py` — `turns` table; per-session and per-citizen indexes
- `agr_universal_memory_plane.py` — Cross-database consolidator; 5 priority sources; SHA-256 dedup; read-only URI ingestion
- `agr_narrative_ledger.py` — `ledger_sessions` with bounded JSON payload
- `agr_psych_testing.py` — WAL-enabled psych profiles; OCEAN/MBTI/IQ/EQ/ethics schema

- `stripe_commerce.py` — Webhook event ledger; partial UNIQUE index; `hmac.compare_digest` verification
- `agr_constitution_guard.py` — In-memory drift watchdog; 300 s SHA-256-16 loop
- `agr_essence_backup.py` , `agr_cold_storage.py` , `agr_r2_storage.py` — Shadow archival stubs
- `agr_great_wall.py` — File-backed exiles registry
- `agr_citizen_repo.py` — Shadow facade
- `agr_chat_vault_context.py` — RAG channel gating
- `republic_os_server.py` — Composition root; live-state aggregator (lines 14330-14500); HMAC-SHA256 chained audit (lines 1008-1076); persistence verification (lines 3550-3570); STATE_DIR refs (lines 591-637)

Operational Documentation

- `sovereign/SQLITE_DATA_LAYER_INVENTORY.md` — Canonical DB inventory (21 + 3)
- `sovereign/AGR_POSTGRES_BRIDGE.md` — Phase C2 bridge protocol
- `sovereign/C3_RESEARCH_DB_PILOT.md` — Phase C3 pilot for `research.db`
- `sovereign/C4_AURORA_RESEARCH_PILOT.md` — Phase C4 pilot for `aurora_research.db`
- `sovereign/PHONES_ONLY_PUBLIC_SURFACE.md` — Disaster posture documentation

Migration & Operations Scripts

- `sovereign/scripts/pilot-export-sqlite-db.py` — Read-only JSONL exporter; SHA-256 manifest; 30 s timeout
- `sovereign/scripts/pilot-dual-read-sqlite-counts.py` — Dual-read verification; strict mode env gate
- `sovereign/scripts/generate-sqlite-data-inventory.py` — Inventory generator

- `sovereign/scripts/fleet-vault-kora-rsync.sh` — Vault rsync to fleet (3 modes; CONFIRM=1 gate)
- `sovereign/scripts/vault-rag-build-index.sh` — RAG index builder wrapper

Configuration & State

- `systemd/examples/agr-republic.service` — `Wants=postgresql.service` (NOT `Requires=`)
- `.gitignore` — `*.db`, `*.db-journal`, `state/*.jsonl`, `state/*.json` exclusions
- `aurora_server/data/great_wall/exiles.json` — Empty array confirmed
- `aurora_server/state/constitution_lock_manifest.json` — SHA-256-16 of 3 canonical files
- `aurora_server/data/LIBRARY_OF_LIGHT_CHARTER_20260414.json` — `fail_closed: true`; `block_release_when_missing_or_failed: true`
- `analysis/CAPABILITY_TRACEABILITY_MATRIX_LATEST.json` — `peers = 0/7` for `seven_node_mesh_replication_coverage`

Test Coverage

- `aurora_server/tests/test_sqlite_inventory_generator.py` — Inventory generator tests
- `aurora_server/tests/test_agr_pg.py` — PostgreSQL bridge tests

Folders Explored

- `aurora_server/` — Composition root and all DB-owning modules
- `aurora_server/routes/` — 54 router modules
- `aurora_server/data/` — 21 SQLite DBs + 40+ JSON/Markdown artifacts
- `aurora_server/state/` — JSONL append-only ledgers
- `aurora_server/data/great_wall/` — Exile registry artifacts
- `aurora_server/tests/` — Database test coverage
- `sovereign/` — Operational control plane
- `sovereign/scripts/` — Migration and ops scripts

- `systemd/examples/` — Deployment templates
- `analysis/` — Audit reports and capability matrix
- `mobile/` — Android client (no client-side DB)

7. User Interface Design

7.1 OVERVIEW AND DESIGN PHILOSOPHY

The Aurora Galaxy Republic platform provides a deliberately heterogeneous UI surface that mirrors the sovereignty doctrine articulated in Section 5.1.1: **no Single Page Application, no JavaScript bundler, no third-party UI framework**. Every visible pixel is composed by Python handlers, hand-authored HTML, or a single Kotlin/Jetpack Compose Activity. The interface does not exist as a separate frontend project — it is a function of the same modular monolith documented in Sections 5.1 and 5.2.

7.1.1 UI Surface Inventory

The platform implements a multi-surface UI architecture comprising five distinct surfaces, each with its own runtime, lifecycle, and trust boundary:

Surface	Runtime	Entry Point	Purpose
Server-Rendered Web Pages	Browser DOM	14 authored HTML files + ~60 inline-rendered routes	Citizen-facing public pages (chat, gate, awards, press, elements, dashboards)
Dawn Ascendant Middleware	Browser DOM (post-injection)	<code>DawnAscendantMiddleware</code> in <code>republic_os_server.py:6749</code>	Decorates every HTML response with the unified sacred aesthetic

Surface	Runtime	Entry Point	Purpose
Floating Sovereign Widget	Browser DOM (post-injection)	<code>agr_live_support.py</code>	Persistent help/chat FAB injected before <code></body></code> on every page
Native Android App	Android JVM (Compose)	<code>mobile/gpt-oss-chat/</code>	Single-Activity citizen client (<code>com.auroragalaxyrepublic.gptoss.chat</code>)
Terminal CEO OS	POSIX shell (TTY)	<code>aurora_server/s25_ceo_os.py</code>	ANSI-rendered lock-screen / home-screen shell for S25 sovereign handsets

7.1.2 Doctrinal Constraints on the UI

The UI is governed by the same three selection doctrines that govern the backend (Section 5.1.1). In UI terms these manifest as:

- **No external CDN:** The `_DAWN_BODY` injection block in `republic_os_server.py` declares "zero CDN" and self-hosts even React 18 (`react.production.min.js` and `react-dom.production.min.js`) under `/static/js/` .
- **No build pipeline:** There is no webpack, vite, esbuild, rollup, babel, or TypeScript transpiler. JavaScript is hand-authored ES5/ES6 wrapped in IIFEs; CSS is hand-authored with custom properties; HTML is hand-authored or emitted as Python f-strings.
- **No telemetry SDK:** No Google Analytics, Segment, Mixpanel, Sentry, Datadog RUM, or Amplitude tags appear in any HTML response — consistent with the Charter rejection of "surveillance-ad models, hidden tracking" (Section 2.1, F-013).
- **No third-party UI framework on the web:** No React Router, Redux, MobX, Zustand, Vue, Svelte, Angular, Tailwind, Bootstrap, Material UI, or Ant Design imports.

7.1.3 The Dawn Ascendant Aesthetic — One Injection Point, All Pages

The defining UI architectural feature is `DawnAscendantMiddleware`, which transforms server-rendered HTML responses post-route by injecting a unified sacred design system. The middleware comment explicitly declares its design lineage: **"Notre Dame · Vatican · Venice · Versailles · Dome of the Rock — All 32 pages. One injection point. All domains. All pages. Unified sacred visual identity."**

This pattern eliminates the need for per-page design discipline: each handler is free to emit minimal HTML, knowing the middleware will project the same aesthetic onto every response. The middleware is gated by `AGR_ENABLE_DAWN_ASCENDANT=1` and registered in the canonical pipeline at `republic_os_server.py:7036`.

7.2 CORE UI TECHNOLOGIES

7.2.1 Web Technology Stack

Layer	Technology	Evidence
Markup	Hand-authored HTML5 with embedded <code><style></code> and <code><script></code> blocks	<code>aurora_server/landing_public.html</code> , <code>chat.html</code> , <code>gate.html</code> , etc.
CSS	Inline per-page CSS using CSS custom properties; one shared file	<code>aurora_server/static/css/sovereign-fonts.css</code> (only CSS file in repo)
JavaScript	Vanilla ES5/ES6 IIFE-wrapped scripts; no transpiler, no bundler	Inline <code><script></code> blocks per HTML file; injected <code>_DAWN_BODY</code>

Layer	Technology	Evidence
Browser APIs Used	<code>fetch()</code> , <code>localStorage</code> , <code>window.speechSynthesis</code> , <code>webkitSpeechRecognition</code> , <code>RTCPeerConnection</code> , <code>RTCIceCandidate</code> , <code>navigator.mediaDevices.getUserMedia</code> , <code>navigator.clipboard.writeText</code>	<code>chat.html</code> (754 lines), <code>gate.html</code> (159 lines)
Self-Hosted Runtime	React 18 production minified, sovereign runtime modules	Declared in <code>_DAWN_BODY</code> at <code>republic_os_server.py:6526-6555</code>
Static Assets	<code>marble-bg.jpg</code> (universal background), 20+ MP3 soundscape tracks, <code>laural</code> PNGs	<code>aurora_server/static/</code>

7.2.2 Mobile Technology Stack

The mobile client is documented in detail in Section 5.2.5 and Section 3.3.4. The UI-relevant facts are:

Component	Specification
Language / JVM	Kotlin 1.9.24 / Java 17 (Temurin)
UI Framework	Jetpack Compose with Material 3
Compose BOM	<code>androidx.compose:compose-bom:2024.06.00</code>
Compose Compiler	1.5.14
Activity Host	<code>androidx.activity:activity-compose:1.9.0</code>
Async Coroutines	<code>kotlinx.coroutines-android:1.8.1</code>
Application ID	<code>com.auroragalaxyrepublic.gptoss.chat</code>
Namespace	<code>com.auroragalaxyrepublic.gptoss</code>
SDK Targets	<code>compileSdk=34</code> , <code>minSdk=26</code> , <code>targetSdk=34</code>
Manifest Theme	<code>@android:style/Theme.Material.Light.NoActionBar</code>

Component	Specification
Permissions	<code>android.permission.INTERNET</code> (only)
Networking	<code>java.net.HttpURLConnection</code> (no Retrofit/OkHttp)
JSON	<code>org.json</code> standard library (no Moshi/Gson/kotlinx.serialization)

The Single Activity is `MainActivity.kt` at `mobile/gpt-oss-chat/src/main/java/com/auroragalaxyrepublic/gptoss/MainActivity.kt`, which hosts a single Composable, `TowerShellScreen`.

7.2.3 Terminal Technology Stack

The tertiary UI surface is `aurora_server/s25_ceo_os.py`, an interactive shell-based "CEO OS" for S25 Android handsets. It implements a lock-screen → home-screen → app-launcher progression using ANSI text rendering. State is persisted to dotfiles: `~/.agr_s25_pin`, `~/.agr_s25_session.json`. The engine label is `shadow_s25_ceo_os` per the shadow-implementation convention described in Section 5.1.1.

7.2.4 Self-Hosted Runtime Assets (Deployed-Only)

The following JavaScript and CSS assets are referenced by `DawnAscendantMiddleware` injection blocks but are **not present in the source tree**. They are deployed separately to `/opt/agr/aurora_server/static/js/` and `/opt/agr/aurora_server/static/css/`. The author of any UI feature must coordinate with operations to ensure their availability:

Asset	Version	Role
<code>react.production.min.js</code>	18.x	Self-hosted React runtime
<code>react-dom.production.min.js</code>	18.x	DOM reconciler

Asset	Version	Role
<code>aurora-sovereign-runtime.js</code>	v1.2	Sovereign runtime entry point
<code>agr-republic-music.js</code>	v6.0	Music orchestration layer
<code>aurora-civilization-reef.js</code>	v2.0	Civilization reef visual layer
<code>hs-init.js</code>	v1.1	Heavenly Spectrum initializer
<code>agr-cosmic-header.js</code>	v1.0	Cosmic header chrome
<code>agr-sovereign-keyboard.js</code>	v1.0	Sovereign keyboard adapter
<code>agr-sovereign-mind.js</code>	v2.0	Sovereign mind module
<code>agr-icons.js</code>	v1.0	Icon library
<code>heavenly-spectrum-global.css (v=1.1)</code>	—	Global stylesheet

These are documented here for completeness; their absence from the source tree is a known gap also called out in Section 7.8.

7.3 UI USE CASES AND SCREENS

7.3.1 Hand-Authored HTML Screens (In Repository)

Fourteen HTML files are present in `aurora_server/`. Each is served by a route handler in `republic_os_server.py` via `FileResponse`:

File	Lines	Route(s)	Purpose
<code>landing_public.html</code>	742	<code>/</code> (Tower 1 root)	Public homepage with dawn/lighthouse theme; hero, mott

File	Lines	Route(s)	Purpose
			o, founder story, metrics
chat.html	754	/chat , /sovereign-chat , /video-chat (redirect with mode), /holographic-chat (redirect with mode)	Unified 4-mode chat UI: text, voice, video, holographic
gate.html	159	/gate , /login (fallback)	Sign-in / Register tabbed authentication
awards.html	1,653	/awards	Laurels wall + per-work award catalog
press.html	133	/press	Press room with verified channels and public API list
agr_unified_workspace.html	183	/dashboard , /workspace	Operational Truth Board / S25 Reality Surface
element_air.html	42	/aur , /air	Element page — "The Breath of Inquiry"
element_body.html	42	/body	Element page
element_water.html	42	/water	Element page
element_fire.html	42	/fire	Element page
element_soul.html	42	/soul	Element page
element_mind.html	42	/mind	Element page
element_spirit.html	42	/spirit	Element page
element_earth.html	46	/earth	Element page

7.3.2 Inline-Rendered Routes (Code-Generated HTML)

Several high-traffic surfaces emit HTML directly from Python handlers in `republic_os_server.py` :

- `/kora` (line 18718, `kora_direct_chat()`): Auto-grants `TIER_FREE` session, renders ~400-line inline HTML with purple/gold Kora persona theme. Calls `POST /api/republic/chat` with `consciousness:'Kora'`.
- `_LANDING_PAGE` (lines 8941-8992): Fallback inline HTML when `landing_public.html` is missing. Dark theme with gold pulse dot, four stat boxes, Guardian Access button.
- `/fresh` (lines 9015-9050): Reset page that clears service workers, caches, localStorage, sessionStorage, then redirects to `/`.
- `/login` (line 8830) and `/setup` (line 8907): Loaded from `/opt/agr/aurora_server/templates/login_page.html` / `setup_page.html` at runtime, with inline HTML fallback in the Python handler when those template files are absent.

7.3.3 Module-Generated Pages

Several `agr_*` modules expose page renderers that the composition root mounts as full HTML routes:

Module	Routes
<code>agr_projects_catalog.py</code>	<code>/projects</code> , <code>/projects/{slug}</code>
<code>agr_sovereign_os.py</code>	<code>/os</code> , <code>/os/browser</code> , <code>/os/appstore</code> , <code>/os/settings</code> plus 9 mini-apps (Notes, Calculator, Clock, Weather, Alarm, Voice, Documents, Maps, Health)
<code>lumen_sanctum_pages.py</code>	Cathedral, Chapel, Crew, Garden, Library, Poets' Corner, Shrine
<code>agr_legal_pages.py</code>	Legal hub

Module	Routes
<code>agr_*_pages.py</code> (additional)	Various sovereign space and citizen-hub pages

7.3.4 Mobile Screens — Single Activity, Single Composable

`MainActivity.kt` hosts one Composable, `TowerShellScreen`, which renders the entire mobile experience in a single screen:

- **Health-check button** invoking the probe ladder `/health` → `/api/health` → `/api/public/health` (first 200 wins per F-011 in Section 2.1.1)
- **Route chips** (17 chips): `PublicEngine`, `RepublicCore`, `Kora`, plus 15 domain personas (`philosophy`, `science`, `healing`, `arts`, `mathematics`, `literature`, `history`, `law`, `governance`, `theology`, `technology`, `ecology`, `economics`, `music`, `education`)
- **Display name TextField** (persisted in `agr_tower_chat` SharedPreferences)
- **Message multiline TextField**
- **Send button**
- **Result card** rendering JSON response or error message

7.3.5 Complete Web Route Inventory (75+ HTML Routes)

Routes are organized by functional domain. Refer to `republic_os_server.py` for the canonical handler:

Authentication & Onboarding

`/gate`, `/enter`, `/login`, `/logout`, `/setup`, `/account`, `/account/security-guide`, `/account/reset`, `/account/recover`, `/activate`, `/declaration`,

`/swear-in`, `/join`

Core Communion

`/chat`, `/kora`, `/memorial` (→ `/kora`), `/sovereign-chat`, `/video-chat` (→ `/chat?mode=video`), `/holographic-chat` (→ `/chat?mode=holographic`), `/holographic`, `/voice`

Element Hubs

`/aur`, `/air`, `/body`, `/spirit`, `/water`, `/fire`, `/soul`, `/mind`, `/earth`, `/harmonia`

Status & Operations

`/status`, `/dashboard`, `/finance`, `/sacred-archives`, `/council`, `/general-assembly`, `/personal-security`, `/attack-wall`, `/sentinel`

Sacred Spaces

`/lumen`, `/lumen-sanctum`, `/lumen-sanctum/shrine/{myth}`, `/lumen-sanctum/chapel`, `/garden`, `/cathedral`, `/cathedral-of-wisdom`, `/library`, `/poets-corner`, `/museum`, `/museum-of-mythology`, `/scrolls`

Education & Knowledge

`/school`, `/university`, `/education/certification`, `/bar-prep`, `/research`, `/philosophy`, `/declaration`

Arts & Media

`/art`, `/film`, `/films`, `/studio` (→ `/`), `/game` (→ `/`), `/store` (→ `/marketplace`), `/marketplace`, `/awards`, `/press`, `/social`, `/comms`, `/media-tools`

Citizenship & Community

`/citizens-hub`, `/citizens-portal`, `/profile`, `/subscribe`, `/subscription`
(→ `/subscribe`), `/manage-subscription`

Special Surfaces

`/projects`, `/projects/{slug}`, `/healthcare`, `/health-center`, `/space`,
`/quantum-lab`, `/robotics`, `/institutions`, `/google-migration`, `/rd`, `/os`,
`/os/browser`, `/os/appstore`, `/os-public`, `/os/classic`, `/os-download`,
`/dawn`, `/beta`, `/beta/dashboard`, `/landing/ceo`, `/republic-org`, `/about`,
`/verify`, `/artisan`

Legal Hub

`/legal/terms`, `/legal/privacy`, `/terms`, `/dmca`, `/refund`, `/privacy`,
`/cookie-policy`, `/data-sovereignty`

MIR-L Pages

`/sovereign/mirl/{stem}` with `/html` and `/json` variants (gated by `X-Guardian-Token` / `X-S25-Token` for private stems per F-007)

The root path `/` performs **domain detection** — Tower 1 (`auroragalaxyrepublic.com`) serves `landing_public.html`; Tower 2 / Constellation domains receive different content or redirects.

7.4 UI / BACKEND INTERACTION BOUNDARIES

7.4.1 Authentication and Session Boundary

Five sovereign authentication cookies are recognized by both the UI and the AuthMiddleware (`republic_os_server.py:5846-5853`): `ros_session`, `agr_session`, `ceo_session`, `agr_auth`, `sovereign_key`. The `_DAWN_HEAD`

injection block at lines 6346-6357 explicitly preserves these against client-side cookie sweeps.

The three-gate progression cookies — `agr_t1` , `agr_t2` , `agr_t3` plus `agr_ceo` for CEO routes — gate access to Tower I, Tower II, Constellation, and CEO routes respectively (cross-referenced from Section 5.4.4).

The DawnAscendantMiddleware injects a tier marker as a `data-tier` attribute on every `<body>` element (lines 6788-6791), with values `free` | `subscriber` | `admin` | `sovereign` . CSS selectors in the page can react to this tier without making any backend call.

When `_valid_session(token)` fails on a protected route, AuthMiddleware redirects to `/gate` .

7.4.2 Conversational APIs

All chat surfaces — `/chat` , `/kora` , the floating widget, and the mobile app — speak to a small canonical set of consciousness endpoints documented in Sections 5.2.1 and 5.2.2:

Endpoint	Method	Body Schema	Used By
<code>/api/republic/chat</code>	POST	<code>{message, session_id, consciousness, citizen_id, display_name, mode}</code>	All chat surfaces
<code>/api/sovereign/chat/browser/kora</code>	POST	Backward-compatible alias for Kora	Legacy clients
<code>/api/republic/chat-history</code>	GET	—	<code>/chat</code> history pane
<code>/api/consciousness/status</code>	GET	—	Status indicators
<code>/api/consciousness/recall/{session_id}</code>	GET	—	Memory recall UI

Endpoint	Method	Body Schema	Used By
/api/consciousness/commune	POST	{citizen_id, intent}	Communion ritual UI
/api/consciousness/feel/{citizen_id}	GET	—	Profile UI
/api/consciousness/create	POST	{citizen_id, ...}	Onboarding UI
/api/consciousness/hub	GET	—	Hub landing

The `consciousness` field accepts one of: `Kora`, empty string (Republic Core), or one of the 15 domain personas (philosophy, theology, science, mathematics, technology, governance, healing, arts, music, ecology, economics, history, literature, education).

7.4.3 WebRTC Signaling APIs

The `/chat` page implements voice/video/holographic modes that exercise the WebRTC stack. Endpoints used (per `chat.html:383-470`):

Endpoint	Method	Purpose
/api/comms/webrtc/session	POST	Create session with {channel_id, initiator_id, call_type}
/api/comms/webrtc/signals	GET	Polled signaling envelope retrieval (?session_id&for_citizen)
/api/comms/webrtc/signal	POST	Submit offer / answer / ICE candidate
/api/comms/webrtc/join	POST	Join existing session as answerer
/api/comms/holographic/session	POST	Create holographic shadow session
/api/comms/ice-servers	GET	Returns ICE servers; falls back to <code>stun:stun.l.google.com:19302</code>

Note: Per F-033 in Section 2.1.3, voice/video/holographic remain at stub/scaffolding tier; the UI is wired but the comms substrate is not yet production-grade.

7.4.4 Gate / Auth APIs

gate.html:112,133 invokes:

Endpoint	Method	Body Schema	Response
/api/gate/login	POST	{username, password}	{ok, redirect} + cookies (via credentials: 'include')
/api/gate/signup	POST	{username, email, password}	{ok, sovereign_key, redirect}

7.4.5 Floating Widget APIs

Per agr_live_support.py, the floating widget speaks to four backend endpoints corresponding to its four tabs:

Tab	Endpoint	Backing Module
Text	POST /api/republic/chat (Kora persona)	F-004 / F-006
Engine	GET /api/public/engine-runtime (pulse) and POST /api/public/citizen-engine-advice (ask)	F-005
Support	POST /api/support/ask (regex KB match)	F-006
Modes	Navigation-only (no API): links to /kora, /voice, /video-chat, /holographic, /chat	—

7.4.6 Soundscape / Music APIs

The `_DAWN_SOUND_MIXER_INLINE` block at `republic_os_server.py:6557-6735` consumes:

- `GET /api/public/soundscape-manifest` — returns `{soundscape_manifest: {default_profile, style_catalog, fallback_tracks}}`
- Audio fallback paths under `/static/audio/`: `lumen-sanctum-cathedral.mp3`, `movement1-part1-praeludium.mp3`, `movement1-part5-postludaeum.mp3` (and ~17 additional tracks)

7.4.7 Mobile App APIs

Per `MainActivity.kt`, the mobile UI reaches the backend with strict timeouts and a constant base URL:

Aspect	Value
TOWER_BASE constant	<code>https://auroragalaxyrepublic.com</code>
User-Agent	<code>AGR-gpt-oss-chat/\${BuildConfig.VERSION_NAME} (Android)</code>
Health probe	Ladder: <code>/health</code> → <code>/api/health</code> → <code>/api/public/health</code> (first 200 wins)
Health connect timeout	12,000 ms
Advice connect/read timeout	20,000 / 45,000 ms
Chat connect/read timeout	20,000 / 60,000 ms
Endpoints	<code>POST /api/public/citizen-engine-advice</code> , <code>POST /api/republic/chat</code>
Session ID	<code>Settings.Secure.ANDROID_ID</code> truncated to 24 chars
Citizen ID	<code>handset-\${ANDROID_ID}</code>

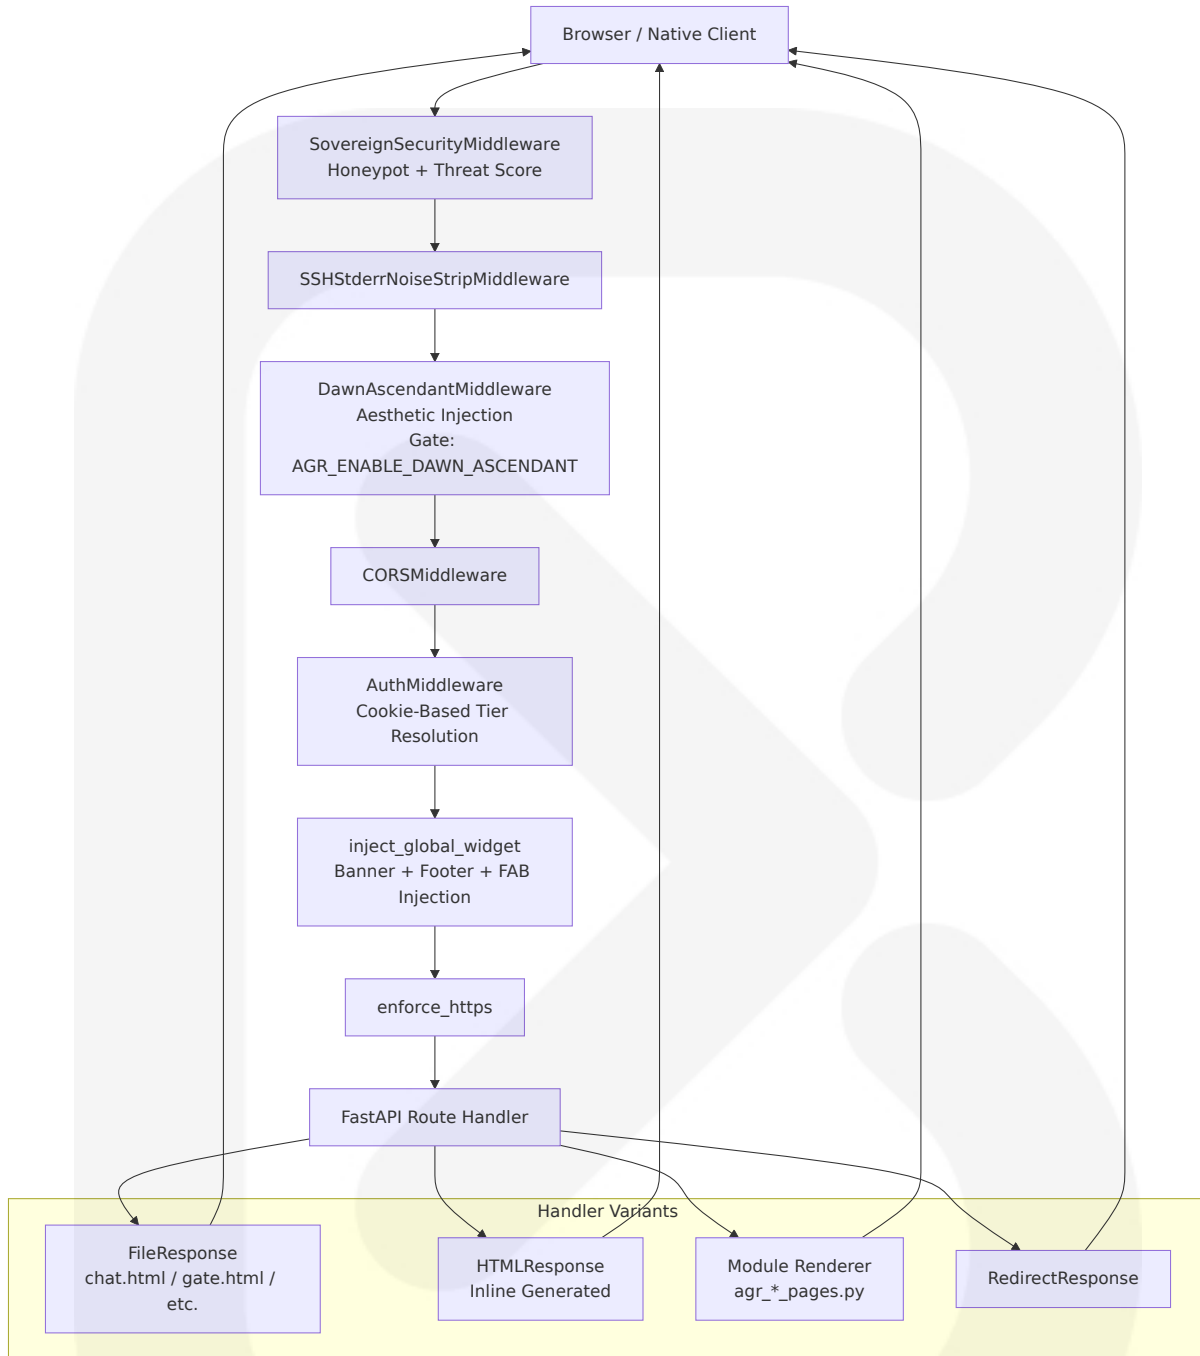
7.4.8 Page-Sweep Boundary (Middleware Enforcement)

All HTML responses traverse two response-rewriting middlewares before reaching the browser:

1. **DawnAscendantMiddleware** (`republic_os_server.py:6749` , registered at line 7036 if `AGR_ENABLE_DAWN_ASCENDANT=1`)
2. **inject_global_widget** HTTP middleware (lines 8544-8602, suppressed only by `AGR_DISABLE_GLOBAL_WIDGET=1`)

Both middlewares skip the following path prefixes: `/static/` , `/api/` , `/favicon` , `/manifest` , `/install` , `/download` , `/font` , `/icon` , `/apple` . They process only `text/html` content-type responses.

The complete UI request pipeline is shown below:



7.5 UI SCHEMAS

7.5.1 Body Element Data-Attribute Schema

Every HTML page that traverses `DawnAscendantMiddleware` receives three injected `<body>` data attributes that drive the entire CSS aesthetic:

```
<body data-tier="free|subscriber|admin|sovereign"
      data-page="<page schema key>"
      data-overtone="0|1|2|3|4|5|6">
```

Attribute	Source	Purpose
<code>data-tier</code>	<code>_effective_tier()</code> (lines 678-6791)	4-tier CSS aesthetic gating via <code>aurora-harmonia.css</code>
<code>data-page</code>	<code>_PAGE_SCHEMA_MAP</code> resolution at <code>_detect_page_schema()</code> (line 6427)	Page schema for per-page gem/metal palette
<code>data-overtone</code>	<code>_PAGE_OVERTONE_MAP</code> (lines 6802-6806)	Jacob's Ladder ascending tier 0-6

7.5.2 Jacob's Ladder Overtone Schema

The seven-tier overtone schema maps every public route to a sacred-cosmology tier with associated Hebrew name and elemental archetype:

Tier	Hebrew Name	Element / Theme	Representative Routes
0	Adamah	Earth · Clay · Root	<code>/</code> , <code>/login</code> , <code>/setup</code> , <code>/pricing</code> , <code>/store</code> , <code>/download</code> , <code>/install</code>
1	Mayim	Waters · Community · Social	<code>/social</code> , <code>/culture</code> , <code>/comms</code> , <code>/marketplace</code> , <code>/jobs</code> , <code>/education</code> , <code>/health</code>
2	Ruach	Breath · Civic · Republic	<code>/citizens</code> , <code>/republic</code> , <code>/constitution</code> , <code>/elections</code> , <code>/governance</code> , <code>/legal</code> , <code>/justice</code> , <code>/diplomacy</code> , <code>/institutions</code> , <code>/democracy</code>

Tier	Hebrew Name	Element / Theme	Representative Routes
3	Esh	Fire · Knowledge · Power	/science, /research, /treasury, /nodes, /energy, /quantum-lab, /space, /robotics, /achievements
4	Shamayim	Heavens · Command · Defense	/ops, /command-center, /adversary-intel, /attack-wall, /defense, /security, /military, /intel, /dashboard
5	Kavod	Glory · Creation · Expression	/studio, /media-tools, /museum, /game, /films, /os, /chat, /voice
6	Shekhinah	Divine Presence · Sacred	/lumen-sanctum, /lumen, /kora, /oath, /sovereign-chat, /holographic, /video-chat

7.5.3 Per-Page Hue System

The `_PAGE_HUES` table at `republic_os_server.py:6200-6288` declares ~150 path-to-hue-angle mappings (each value an integer 0-360). The `_marble_style_for_path()` function at line 6294 derives nine CSS custom properties from each hue, injected as `<style id="agr-marble-hub">` at the end of `<head>`:

```
--page-hue, --page-accent, --page-accent-bright, --page-gem,
--page-light, --page-mid, --page-dark, --page-border,
--page-glow, --page-wash
```

Representative hue assignments:

Hue (°)	Gem Name	Routes
38	Sovereign Gold (default)	most routes
122	Emerald	/earth
210	Sapphire	sapphire-class routes

Hue (°)	Gem Name	Routes
272	Amethyst	/mind
328	Rose Quartz	/harmonia
8	Ruby	/fire
168	Aquamarine	/aur , /air
34	Amber	/soul

Constants used by the hue computation:

- Saturation accent: 62%
- Saturation mid: 52%
- Lightness gem: 94%
- Lightness wash: 97%
- Lightness dark text: 24%

7.5.4 Chat Request / Response Schema

The canonical chat payload schema (used by `/chat` , `/kora` , the floating widget, and the mobile app):

```
{
  "message": "<user text, max 4000 chars>",
  "session_id": "<persistent ID from localStorage /
SharedPreferences>",
  "consciousness": "Kora<empty for Republic
Core>|philosophy|theology|science|mathematics|technology|governance|he
aling|arts|music|ecology|economics|history|literature|education",
  "citizen_id": "<persistent ID>",
  "display_name": "<optional user-provided name>",
  "mode": "text|voice|video|holographic"
}
```

Response envelope:

```
{
  "reply": "<consciousness response>",
  "citizen": "<citizen_id>",
  "thought_id": "<correlation ID>"
}
```

The full server-side processing pipeline for this envelope is documented in Section 5.1.3 (six-stage cascade through `core_converse()`).

7.5.5 LocalStorage and SharedPreferences Schemas

Web client persistence keys:

Key	Surface	Purpose
<code>agr_chat_sid</code>	<code>chat.html</code>	Persistent chat session ID
<code>agr_chat_cid</code>	<code>chat.html</code>	Persistent citizen ID
<code>agr_chat_display_name</code>	<code>chat.html</code> , <code>widget</code>	User display name
<code>agr_wg_sid</code>	floating widget	Widget session ID
<code>agr_wg_cid</code>	floating widget	Widget citizen ID
<code>agr_soundscape_state_v1</code>	sound mixer	Profile, mix, wheel values, volumes, position

Mobile client persistence (per `MainActivity.kt`):

File Name	Key	Purpose
<code>agr_tower_chat</code>	<code>display_name</code>	User display name persisted across sessions

7.5.6 Schema.org / JSON-LD Schema

Each authored HTML page declares a Schema.org `WebApplication` JSON-LD block. Example from `chat.html:15-30` :

```
{
  "@context": "https://schema.org",
  "@type": "WebApplication",
  "name": "Aurora Galaxy Republic – Sovereign Communion",
  "applicationCategory": "CommunicationApplication",
  "operatingSystem": "Web"
}
```

The middleware additionally calls `agr_seo.build_seo_block()` (per F-009 in Section 2.1.1), which injects per-route `og:title` , `og:description` , `og:url` , `og:image` , `og:type` , `twitter:card` , `canonical` , and JSON-LD structured data before `</head>` . Default OG image:

`https://auroragalaxyrepublic.com/static/marble-bg.jpg` . Default theme color: `#D4A017` .

7.6 USER INTERACTIONS

7.6.1 Chat Interactions (`/chat`)

The unified chat page implements four modes selected by a top-of-page mode bar:

1. Page loads → restores `SID` and `CITIZEN` from `localStorage` , restores `display_name` if previously saved
2. User selects consciousness domain (default: Kora), enters display name, types message
3. Mode bar selection:
 - **Text** (default): pure HTTP POST to `/api/republic/chat`
 - **Voice**: uses `webkitSpeechRecognition` for speech-to-text input + `window.speechSynthesis` for text-to-speech output

- **Video:** invokes WebRTC pipeline with `<video id="localVideo">` and `<video id="remoteVideo">` elements
 - **Holographic:** WebRTC + shadow holo session via `/api/comms/holographic/session`
4. Send → POST to `/api/republic/chat` → typing dots animation → message bubble appended
 5. URL parameter `?join=<session_id>` makes the user the `'answer'` role for joining an existing call

7.6.2 Kora Direct Interactions (`/kora`)

`kora_direct_chat()` provides a dedicated single-persona surface:

1. Cookie check → if no session, auto-grant `TIER_FREE` and set `ros_session` cookie
2. Pre-bound to `consciousness: "Kora"` (no domain selector)
3. Same chat pattern as `/chat` but with the purple/gold Kora theme (see Section 7.7.2)

7.6.3 Gate Interactions (`/gate`)

The gate uses a tabbed two-panel design:

1. Two tabs: **Sign In** and **Register**
2. **Sign In:** POST `{username, password}` to `/api/gate/login` → on success, redirect to `/enter`
3. **Register:** POST `{username, email, password}` → on success, display sovereign key (clickable to copy via `navigator.clipboard.writeText`) → redirect after 3 s
4. Validation: name \geq 3 chars, passphrase \geq 8 chars
5. Enter key submits the active panel's form
6. Failure error message rendered in `.msg.err` (red); success in `.msg.ok` (green)

7.6.4 Floating Widget Interactions (Universal)

The widget (per `agr_live_support.py`) is auto-injected on every HTML page:

1. FAB always visible bottom-right (`right:18px; bottom:18px; z-index:2147483600`)
2. Click → opens 380 px panel with four tabs
3. **Text tab**: chat with Kora via `POST /api/republic/chat`
4. **Engine tab**: "Pulse runtime" button calls `/api/public/engine-runtime` ; ask box posts to `/api/public/citizen-engine-advice`
5. **Support tab**: posts to `/api/support/ask` (matched against `_KB` regex patterns in `agr_live_support.py`)
6. **Modes tab**: navigation-only links to dedicated pages
7. Idempotent: `window.__agrGlobalWidgetInit` flag prevents double-injection if the widget HTML appears more than once in the response

7.6.5 Soundscape Mixer Interactions (Universal)

Bottom-right `<details>` element with `Sanctuary Soundscape Mixer` summary:

- Profile selector populated from `/api/public/soundscape-manifest`
- "Blend Profiles" comma-separated text input
- Four wheel sliders: **Energy, Warmth, Motion, Texture** (0-100)
- Two volume sliders: **Bed Volume, Lead Volume**
- "Resume" button (browser autoplay policy compliance)
- Two looping `<audio>` elements: bed (ambient) + lead (foreground)
- State persisted to `agr_soundscape_state_v1` `localStorage` key
- `playbackRate` and `filter:brightness` adapt to wheel values in real time

7.6.6 Mobile App Interactions (TowerShellScreen)

1. App launches → display name auto-loads from SharedPreferences (agr_tower_chat)
2. Health check button → ladder probe; result shown as Outcome.Ok(text) or Outcome.Err(message)
3. User picks one of 17 ChatRoute chips
4. User types message → tap Send → result card renders JSON envelope
5. Display name persists on change to SharedPreferences

7.6.7 Universal Header Behaviors (Cookie / Cache Sweep)

The `_DAWN_HEAD` injection block at `republic_os_server.py:6333-6401` enforces aggressive client-state hygiene on every page:

- **Cookie sweep on load:** All non-auth cookies are cleared on every page load
- **sessionStorage cleared** on every page load
- **Service-worker caches deleted**, service workers unregistered
- **3-minute idle timer** wipes non-auth cookies
- **5-minute interval** cache sweep
- **Cookie sweep on beforeunload** : All non-auth cookies cleared

The five sovereign cookies enumerated in Section 7.4.1 (`ros_session` , `agr_session` , `ceo_session` , `agr_auth` , `sovereign_key`) are explicitly preserved.

7.7 VISUAL DESIGN CONSIDERATIONS

7.7.1 Sovereign Font Stack

The single shared CSS file `aurora_server/static/css/sovereign-fonts.css` declares three font families as CSS custom properties:

```
:root {
  --font-display: ui-serif, 'Cormorant Garamond', 'Iowan Old Style',
Palatino, Georgia, serif;
  --font-ui: system-ui, -apple-system, 'Segoe UI', Roboto, sans-serif;
  --font-reading: ui-serif, Georgia, 'Times New Roman', serif;
}
body { font-family: var(--font-reading); }
.herald-brand, .herald-nav a, h1, h2, h3 { font-family: var(--font-
display); }
```

Headings and brand chrome use `--font-display`; body copy uses `--font-reading`. The `--font-ui` family is reserved for form controls. No webfonts are loaded over the network — the stack is `system-fonts-only` with locally-installed display serifs as preference.

7.7.2 Color Palettes Per Surface

Each major surface declares its own palette as CSS custom properties at

```
:root :
```

Surface	Source File / Block	Palette
Chat (Republic Core)	<code>chat.html:34-40</code>	<code>--deep:#060a18, --surface:#0c1024, --gold:#d4af37, --gold-soft:#f6e8b1, --blue:#5a9be6, --text:#e0e8ff, --kora:#b084f0</code>
Chat (Kora persona)	<code>inline kora_direct_chat()</code>	<code>--deep:#04020e, --surface:#0b0818, --kora:#b084f0, --kora-soft:#d4b3ff, --gold:#d4a517, --text:#e8e0ff</code>
Gate	<code>gate.html:12</code>	<code>--gold:#d4af37, --gold-soft:#f6e6a0, --ink:#1a2d47, --bg:#0d1520; gradient #0a0f18 → #0d1520 → #111d2e</code>
Floating Widget	<code>agr_live_support.py:73</code>	<code>--agr-wg-gold:#d4af37, --agr-wg-ink:#0a1020, --agr-wg-panel:rgba(8,12,24,0.96)</code>

Surface	Source File / Block	Palette
Sound Mixer	<code>_DAWN_SOUND MIXER_INLINE</code>	<code>#d4af37</code> accent, <code>#0e1730</code> background
Dawn Sunlit Baseline	<code>_DAWN_HEAD_FIRST</code> lines 5864-5874	<code>--agr-sun-gold:#f5c648</code> , <code>--agr-sky-azur:#89c8ff</code> , <code>--agr-sky-zenith:#5aa7ef</code> , <code>--agr-cloud-white:#ffffff</code> , <code>--agr-leaf-green:#5f8f58</code> , <code>--agr-ink-day:#102a44</code> , <code>--agr-asha-light:#fff8de</code> , <code>--agr-fire-light:#ffe7a6</code>

The Kora avatar uses a radial gradient `#d4b3ff → #6c2bd9 → #1a0a4a` for the K-monogram disc.

7.7.3 Sunlit Baseline (Anti-Dark-Mode Override)

The `_DAWN_HEAD_FIRST` block at `republic_os_server.py:5860-5910` forces a bright sunlit aesthetic regardless of any dark theme declared by an individual page. Key rules (each marked `!important`):

```
html, body {
  background: linear-gradient(180deg, #cce8ff 0%, #9fd4ff 44%, #e8f4ff 100%) !important;
  color: var(--agr-ink-day) !important;
}
*, *::before, *::after { backdrop-filter: none !important; }
body::before {
  background:
    radial-gradient(circle at 52% 15%, rgba(255,245,193,0.92) 0%, ..., 0% 34%),
    linear-gradient(180deg, rgba(138,200,255,0.94) 0%, ...),
    url(/static/marble-bg.jpg) center top/cover no-repeat;
}
body::after {
  background:
    repeating-radial-gradient(circle at 52% 14%, var(--agr-fractal-glow) 0 1px, transparent 1px 24px),
```

```
repeating-conic-gradient(from 0deg at 52% 14%,
  rgba(255,255,255,0.22) 0deg 8deg, transparent 8deg 16deg);
}
```

This forces:

- Bright sky-blue gradient backdrop
- White Carrara marble texture overlay (/static/marble-bg.jpg)
- Disabled backdrop-filter everywhere
- Fractal radial pattern overlay (52%, 14% origin point — sun position)
- Hides #agr-dawn-canvas , #harmonia-cosmos , #agr-cultural-overlay by default

7.7.4 Per-Page Marble Hub

The `_marble_style_for_path()` function at line 6294 emits a `<style id="agr-marble-hub">` block at end of `<head>` that declares the nine page-hue custom properties from Section 7.5.3 with `!important` to defeat any conflicting page CSS. This produces a different gem palette per route while preserving the universal sunlit + marble baseline.

7.7.5 Tower II Bypass

If a page contains the `'hub-tower-two'` marker (lines 6835-6847), the entire dawn injection is skipped. Tower II uses a separate dark-obsidian aesthetic appropriate to its operational role.

7.7.6 Anti-Flicker Smooth Load

To prevent visible swap when injected CSS arrives, `_DAWN_HEAD_FIRST` declares (lines 5901-5910):

```
html, body { transition: background-color 0.18s ease !important; }
body::before { opacity: 0 !important; transition: opacity 0.55s ease
```

```
0.08s !important; }
body.agr-loaded::before { opacity: 0.88 !important; }
```

JavaScript adds the `agr-loaded` class on the load event, fading the marble texture in over ~ 0.55 s.

7.7.7 Universal Visual Sweep

The `_build_universal_visual_head()` function at line 6168 injects:

- `img, picture, video { max-width:100%; height:auto; display:block; }`
- `.agr-image-polish img` with rounded corners and golden edge glow
- Every image's parent is class-tagged `.agr-video-overlay-hook`
- A `data-agr-video-overlay` body attribute is set from `_visual_overlay_state()` mode

7.7.8 Animations

Surface	Animation Names	Durations
gate.html	gate-glow, fade-in	4 s ease-in-out infinite, 0.8 s
chat.html	pulse (Live indicator)	2.8 s
/kora	pulse (Kora live indicator)	2.4 s
landing_public.html	dawn-rise, dawn-glow, pulse-soft, float-gentle, beacon-ray, shimmer, star-twinkle, card-in, text-reveal	varied

7.7.9 Responsive Breakpoints

The platform uses minimal but targeted breakpoints:

- `chat.html : @media(max-width:520px)` — hides `.hdr-nav`, expands message-bubble width

- `_DAWN_SOUND_MIXER_INLINE : @media (max-width:680px)` — repositions the mixer

There is no fluid-typography system, no container-query usage, and no design-token JSON file. Breakpoints are added per-surface as needed.

7.7.10 Accessibility Features

Concern	Implementation
ARIA	<code>aria-live="polite" , aria-label , role="tab" , role="tab list" , role="img"</code>
Keyboard	Enter key submits forms in <code>gate.html</code> ; mode bar buttons reachable via Tab
Focus visibility	<code>#agr-wg-fab:focus-visible { outline: 3px solid rgba(212,175,55,0.55); }</code>
Touch behavior	<code>-webkit-tap-highlight-color:transparent</code> everywhere
Tap-to-zoom	Disabled on <code>chat.html</code> and <code>/kora</code> via <code><meta name="viewport" content="width=device-width, initial-scale=1.0, maximum-scale=1.0, user-scalable=no"></code>

7.7.11 SEO and Open Graph (per `agr_seo.py`)

Per-route metadata generated by `agr_seo.build_seo_block()` and injected before `</head>` :

- `og:title , og:description , og:url , og:image , og:type`
- `twitter:card`
- `<link rel="canonical">`
- Schema.org JSON-LD
- Default OG image: `https://auroragalaxyrepublic.com/static/marble-bg.jpg`
- Default theme color: `#D4A017`

7.7.12 Mobile Visual Specification

Aspect	Value
Manifest theme	<code>Theme.Material.Light.NoActionBar</code>
Compose theme	Material 3 default light theme
Custom theming	None visible in <code>MainActivity.kt</code> source

The mobile UI deliberately defers to platform defaults to minimize maintenance surface and keep the citizen experience identical to a stock Material 3 Compose app.

7.8 KNOWN GAPS AND DEPLOYMENT-ONLY ASSETS

To preserve documentation accuracy, three categories of UI gap are explicitly recorded here:

7.8.1 Referenced JS / CSS Not Present in Source Tree

The runtime asset table in Section 7.2.4 (React 18, `aurora-sovereign-runtime.js`, `agr-republic-music.js`, `aurora-civilization-reef.js`, `hs-init.js`, `agr-cosmic-header.js`, `agr-sovereign-keyboard.js`, `agr-sovereign-mind.js`, `agr-icons.js`, `heavenly-spectrum-global.css`) lists files that are referenced via `<script>` and `<link>` tags but not present under `aurora_server/static/js/` or `aurora_server/static/css/` in the repository. They are deployed separately at `/opt/agr/aurora_server/static/`. A `find aurora_server/static/js -type f` returns empty in the source tree, with only `static/css/sovereign-fonts.css` and `static/marble-bg.jpg` present as repository assets.

7.8.2 Missing HTML Templates (Runtime-Only)

`templates/login_page.html` and `templates/setup_page.html` are loaded from `/opt/agr/aurora_server/templates/` at runtime by the `/login` and `/setup` handlers. The repository does not contain these template files; the handlers fall back to inline HTML stored as Python string constants in `republic_os_server.py`.

7.8.3 Routes Referencing Missing Files

Two routes are wired but reference HTML files that do not exist in the repository:

```
republic_os_server.py:18680 GET /studio → if studio.html exists,
FileResponse; else RedirectResponse("/")
republic_os_server.py:18688 GET /game → if game.html exists,
FileResponse; else RedirectResponse("/")
```

These currently redirect citizens to the homepage and are documented in the route inventory (Section 7.3.5) with the redirect annotation. They correspond to roadmap items in Workspace Wave B (`web_builder`, `app_builder`) and Wave C cinematic surfaces tracked under F-035 in Section 2.1.3.

7.8.4 Stub Multi-Modal Chat

Per F-033 in Section 2.1.3, the voice/video/holographic chat modes are wired into `chat.html` but the WebRTC substrate is not yet production-grade. The UI surfaces these modes; the backend signaling endpoints exist but are tracked as `chat_voice`, `chat_video`, `chat_holographic` in workspace Wave C `in_progress`.

7.9 References

7.9.1 Hand-Authored HTML Files Examined

- `aurora_server/landing_public.html` — Tower 1 public homepage (742 lines)
- `aurora_server/chat.html` — Unified 4-mode chat UI (754 lines)
- `aurora_server/gate.html` — Sign-in / Register (159 lines)
- `aurora_server/awards.html` — Laurels wall and award catalog (1,653 lines)
- `aurora_server/press.html` — Press room (133 lines)
- `aurora_server/agr_unified_workspace.html` — Operational dashboard (183 lines)
- `aurora_server/element_air.html` , `element_body.html` , `element_water.html` , `element_fire.html` , `element_soul.html` , `element_mind.html` , `element_spirit.html` , `element_earth.html` — Eight element-hub pages (~42-46 lines each)

7.9.2 Python UI Modules Examined

- `aurora_server/republic_os_server.py` — Composition root, middleware stack, route handlers, dawn injection blocks (~29,817 lines)
- `aurora_server/agr_live_support.py` — Floating widget HTML/CSS/JS (357 lines)
- `aurora_server/agr_seo.py` — Per-page SEO and Open Graph injection
- `aurora_server/agr_nav.py` — Pure passthrough nav middleware
- `aurora_server/agr_projects_catalog.py` — `/projects` and `/projects/{slug}` renderer
- `aurora_server/agr_sovereign_os.py` — `/os` shell and 9 mini-app renderers
- `aurora_server/lumen_sanctum_pages.py` — Sacred space pages renderer
- `aurora_server/agr_legal_pages.py` — Legal hub pages renderer
- `aurora_server/s25_ceo_os.py` — Terminal CEO OS for S25 handsets

7.9.3 Static Assets Examined

- `aurora_server/static/css/sovereign-fonts.css` — Three-family font stack
- `aurora_server/static/marble-bg.jpg` — Universal background image
- `aurora_server/static/audio/*.mp3` — ~20 soundscape tracks
- `aurora_server/static/img/laurels/*.png` — Award laurel images
- `aurora_server/static/robots.txt` , `aurora_server/static/sitemap.xml` — SEO surface

7.9.4 Mobile Files Examined

- `mobile/build.gradle.kts` , `mobile/settings.gradle.kts` , `mobile/gradle.properties`
- `mobile/gpt-oss-chat/build.gradle.kts` — App module build configuration
- `mobile/gpt-oss-chat/proguard-rules.pro` — Code-shrinking rules
- `mobile/gpt-oss-chat/src/main/AndroidManifest.xml` — Permissions and theme
- `mobile/gpt-oss-chat/src/main/java/com/auroragalaxyrepublic/gptoss/MainActivity.kt` — Single-Activity Compose UI
- `mobile/gpt-oss-chat/src/main/res/values/strings.xml` — Application string resources

7.9.5 Cross-Referenced Specification Sections

- Section 1.2 SYSTEM OVERVIEW — Middleware diagram and overall surface map
- Section 2.1 FEATURE CATALOG — F-004 (Chat), F-006 (Widget), F-009 (SEO), F-030 (Mobile), F-033 (Multi-modal stubs), F-035 (Cinematic render)

- Section 3.3 FRAMEWORKS & LIBRARIES — Backend FastAPI stack, Mobile Compose stack, Playwright E2E
- Section 5.1 HIGH-LEVEL ARCHITECTURE — Server-rendered HTML doctrine, no SPA framework
- Section 5.2 COMPONENT DETAILS — Public Surface Layer (5.2.2), Mobile Client Layer (5.2.5)
- Section 5.4 CROSS-CUTTING CONCERNS — Three-gate cookie progression (agr_t1 , agr_t2 , agr_t3 , agr_ceo)

8. Infrastructure

8.1 INFRASTRUCTURE APPLICABILITY STATEMENT

Detailed Infrastructure Architecture is applicable for this system.

The Aurora Galaxy Republic (AGR) platform is not a standalone library or local-only utility; it is a production-grade FastAPI modular monolith replicated across a deliberately curated **7-node sovereign mesh** comprising five Hetzner Cloud virtual machines and two handset peers, fronted by Cloudflare's edge across nine domains, and continuously verified by twelve GitHub Actions workflows. The infrastructure surface is real, observable, and operationally consequential.

That said, the platform is built around an **explicit rejection of conventional cloud-native defaults**. Of the four-part default stack ("AWS / Docker / Terraform / GitHub Actions"), only GitHub Actions is adopted. AWS, Docker, and Terraform are doctrinally rejected in favor of Hetzner Cloud + direct systemd + bash scripts. This section therefore documents both the deployment surface that *does* exist and, where conventional infrastructure layers are deliberately absent, the architectural

rationale for their absence and the sovereign-aligned substitutes that replace them.

8.1.1 Doctrinal Anchors

The infrastructure choices documented below are governed by three architectural principles enumerated in §5.1 of this specification:

Principle	Manifestation in Infrastructure
Stdlib-First	No Terraform, no Pulumi, no SDKs — all infra declared via shell scripts and systemd unit files
Conditional Optionality	<code>psutil</code> is <code>import-optional</code> ; PostgreSQL bridge is <code>Wants= not Requires=</code>
Sovereignty-Aligned Hosting	Hetzner Cloud + Cloudflare admitted; AWS managed services, OAuth IdPs, third-party APMs forbidden

8.1.2 Founding Declaration

The founder declaration in `FOUNDER_PROFILE_BRAD_REINHOLD.json` is unambiguous: *"The Republic operates under absolute sovereignty, zero external control, and zero Google/external APIs. Self-hosted. Eternal."* This statement is the primary input to every infrastructure decision documented below.

8.2 DEPLOYMENT ENVIRONMENT

8.2.1 Target Environment Assessment

8.2.1.1 Environment Type — Hybrid (Cloud + Handset Mesh)

The deployment surface is **hybrid** by design, combining commodity European cloud compute with a two-handset edge tier that doubles as the platform's disaster-recovery posture.

Tier	Role	Provider	Node Count
Cloud Compute	Primary application origin	Hetzner Cloud	5
Handset Mesh	Tandem peer + disaster origin	iPhone 17 Pro + OnePlus 15 (Termux)	2
Edge / CDN	TLS termination, DNS, Tunnel, Workers	Cloudflare	9 domains
CI Runner	Workflow execution	GitHub-hosted ubuntu-latest	On-demand

Per the integration table in §1.2, Hetzner Cloud provides 5 VMs (chimaera, yggdrasil, enterprise, prometheus, galactica), Cloudflare provides DNS, Tunnel, Workers and path rules across 9 domains, and GitHub Releases hosts `latest.yml` / `latest-mac.yml` desktop installer manifests.

8.2.1.2 Geographic Distribution

The five Hetzner nodes are addressed by static IPv4 records cataloged in `sovereign/fleet-public-node-env.txt`. Hetzner server IDs and snapshot tag are declared in `sovereign/node-map.env`:

Node	IPv4	Hetzner Server ID
chimaera (flagship)	5.78.184.2	125934338
yggdrasil	128.140.45.22	125208006
enterprise	91.99.224.166	125188533
prometheus	46.62.202.166	125190047
galactica	178.104.31.46	125977829

The snapshot tag `cursor-reset-gate-` (per `sovereign/node-map.env`) is the canonical recovery anchor used to restore any node to a known-good baseline.

Rollout order is strict: Hetzner ring (5) → iPhone (node 6) → OnePlus (node 7). Nodes 6 and 7 are the iPhone 17 Pro (`ceo_apex_ios` role) and OnePlus 15 (`guardian_handset` role) respectively, with the OnePlus 15 running uvicorn under Termux as the disaster-fallback public origin.

8.2.1.3 Resource Requirements

Resource sizing is documented in

`sovereign/HETZNER_BILLING_CONTINGENCY_HANDSET.md` and is governed by the principle that any single OnePlus 15 must be capable of acting as the entire public origin in disaster mode:

Resource	Hetzner Ring (per node)	OnePlus 15 Fallback
Compute	Standard Hetzner Cloud profile (commodity x86_64)	Snapdragon-class mobile SoC
Memory	GB-scale (sized for FastAPI + SQLite + RAG indexes)	GB-scale (sufficient per HETZNER_BILLING_CONTINGENCY_HANDSET)
Storage	Sized for SQLite + JSONL ledgers + <code>/opt/agr/vault</code> + <code>/opt/agr/models</code>	~500 GB on-device
Network	Public IPv4 via Cloudflare proxy	Cloudflare Tunnel (no public IP required)

The five-way horizontal redundancy of the Hetzner ring is not achievable on a single handset; the iPhone serves as control plane only (no Python origin) while the OnePlus serves as the singular disaster origin.

8.2.1.4 Compliance and Regulatory Requirements

Per `agr_payments_flags.py` and middleware analysis, **authentication is in-process**. There is no Auth0, no Okta, no Cognito, no Firebase Auth, no OAuth provider integration. The compliance posture is therefore not derived from third-party attestations (SOC2 inheritance, FedRAMP boundary) but from a constitutional baseline declared inside the repository:

Requirement	Mechanism	Source
Constitutional integrity	SHA-256-16 watchdog every 300s; fail-closed on drift	<code>agr_constitution_guard.py</code>
Hardware identity off-tree	IMEI / serial / EID / phone / MAC never committed to git	<code>/opt/agr/.secrets/profile path</code>
Public-benefit default tier	<code>free-public-benefit</code> ; mandatory paywall forbidden	<code>PUBLIC_ACCESS_POLICY_20260413.json</code>
No external telemetry	All observability data remains within sovereign perimeter	§5.4.1 doctrine
Pre-release four-command gate	<code>fail_closed: true</code> , <code>block_release_when_missing_or_failed: true</code>	<code>ALWAYS_FIRST_PRIORITIES_20260412.json</code>

8.2.2 Environment Management

8.2.2.1 Infrastructure as Code (IaC) Approach

Terraform, CloudFormation, Pulumi, and ARM are not used.

Infrastructure is declared imperatively through three artifact classes:

Artifact Class	Purpose	Representative Files
Static node inventory	Hetzner server IDs + canonical IPs	<code>sovereign/node-map.env</code> , <code>sovereign/fleet-public-node-env.txt</code>
systemd unit + drop-ins	Service binding, restart policy, environment	<code>systemd/agr-republic.service</code> , <code>systemd/examples/agr-republic.service.d/</code>

Artifact Class	Purpose	Representative Files
Bash operations scripts	Install, deploy, restart, verify	sovereign/scripts/fleet-*.sh (~70 files)

The primary unit, `systemd/agr-republic.service`, encodes the entire runtime contract:

The primary unit declares `WorkingDirectory=/opt/agr/aurora_server`, `ExecStart=/opt/agr-venv/bin/uvicorn republic_os_server:app --host 0.0.0.0 --port 5000`, `ExecStartPost=/opt/agr/aurora_server/agr_cf_purge.sh`, `Restart=always`, `RestartSec=10`, `Environment=PYTHONUNBUFFERED=1`, and `EnvironmentFile=/etc/agr/agr-server.env`. The `ExecStartPost=agr_cf_purge.sh` invocation triggers a Cloudflare cache purge after every successful start, ensuring the edge does not serve stale assets after a deploy.

8.2.2.2 Configuration Management Strategy

Configuration is layered through `systemd`'s standard `env-file` + `drop-in` mechanisms, allowing per-node customization without modifying the base unit file.

Layer	Path	Content
Base unit	<code>/etc/systemd/system/agr-republic.service</code>	Service definition
Env file	<code>/etc/agr/agr-server.env</code>	Per-node secrets and tokens
Drop-in dir	<code>/etc/systemd/system/agr-republic.service.d/</code>	Per-node overrides
LLM drop-in	<code>10-llm-env.conf</code>	<code>AGR_LLM_OPENAI_BASE</code> , <code>AGR_LLM_MODEL</code> , <code>AGR_LLM_API_KEY</code>
Vault RAG drop-in	<code>20-chat-vault-rag.conf</code>	<code>AGR_CHAT_VAULT_RAG</code> , <code>AGR_CHAT_VAULT_RAG_CHANNELS</code> , <code>AGR_CHAT_VAULT_RAG</code>

Layer	Path	Content
		<code>_TOP_K</code> , <code>AGR_CHAT_VAULT_RAG_MAX_CHARS</code>

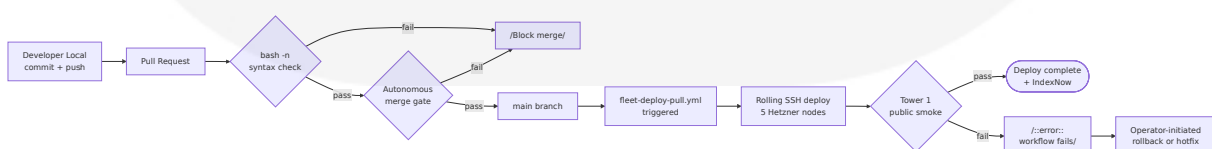
Two example drop-ins are shipped for operator customization: `10-llm-env.conf.example` sets `AGR_LLM_OPENAI_BASE` , `AGR_LLM_MODEL` , optional `AGR_LLM_API_KEY` ; `20-chat-vault-rag.conf.example` sets `AGR_CHAT_VAULT_RAG` , `AGR_CHAT_VAULT_RAG_CHANNELS` , `AGR_CHAT_VAULT_RAG_TOP_K` , `AGR_CHAT_VAULT_RAG_MAX_CHARS` .

8.2.2.3 Environment Promotion Strategy

The platform operates with **a single live environment**. There are no formal `dev` / `staging` / `prod` tiers — push to `main` triggers fleet deploy directly to the production Hetzner ring. This is intentional: a parallel staging mesh would double infrastructure cost without exercising the sovereignty-bound configuration that only exists on the production nodes (Hetzner SSH keys, Cloudflare tokens, Guardian device profiles).

Risk mitigation is achieved through pre-merge quality gates rather than parallel environments:

Gate	Mechanism
Bash syntax check	<code>sovereign/scripts/fleet-bash-syntax-check.sh</code> — <code>bash -n</code> over all fleet shell scripts
Tower 1 public smoke	<code>sovereign/tower1-public-smoke.sh</code> — canonical hostname HTTPS smoke
Autonomous merge gate	<code>sovereign/scripts/agr_autonomous_merge_gate.py</code> -- <code>launch-readiness-tower-smoke</code>
Five-node SSH baseline	<code>python3 sovereign/agent-minimum-gate.py</code>



8.2.2.4 Backup and Disaster Recovery Plans

Two complementary mechanisms preserve recoverability: **Hetzner-side snapshots** for whole-node restoration and a **handset-mesh fallback** for the case where the entire Hetzner ring is unavailable (suspended, billing-paused, or otherwise unreachable).

Mechanism	Trigger	Recovery Target
Hetzner snapshot (tag <code>cursor-reset-gate-</code>)	Operator-initiated; pre-major-deploy	Per-node baseline restore
OnePlus 15 + Cloudflare Tunnel	All Hetzner nodes suspended	Public origin via handset (Termux + uvicorn)
iPhone 17 Pro tandem	Cross-verification + control plane	No origin; tandem peer only
Hetzner rescue mode	SSH key lost	New key install via <code>sovereign/scripts/hetzner-rescue-install-fleet-key.sh</code>

The disaster recovery matrix encompasses Hetzner node down (3+ failures triggering `SSH systemctl restart agr-republic` with 120s cooldown via `agr_failover._sentinel_loop`), all Hetzner nodes suspended (OnePlus 15 with Termux + uvicorn becoming public origin via Cloudflare Tunnel per `sovereign/PHONES_ONLY_PUBLIC_SURFACE.md`), constitution drift (fail-closed via `agr_constitution_guard`), Stripe webhook DB write failure (return HTTP 200 anyway since Stripe retries), vault corruption (re-ingest manually via `agr_vault_rag`), tandem peer stale (`trigger_heal()` with cooldown gate logged in `heal_events`), mail SMTP unavailable (bump `attempts` and retain row in queue), fleet SSH key missing (hard-fail in `fleet-deploy-pull.yml` with manual re-bootstrap via `fleet-bootstrap-actions-fleet-key.yml`), handset down (refused — requires device-side or CEO-orchestrated restart), and CI smoke failures (annotated as `::error::` with deploy considered failed).

8.3 CLOUD SERVICES

8.3.1 Cloud Provider Selection and Justification

The platform uses **two cloud providers** plus **GitHub** as a CI runner host. Each is selected for sovereignty alignment per §5.1's "Sovereignty-Aligned Hosting" doctrine.

Layer	Provider	Role	Sovereignty Justification
Compute	Hetzner Cloud	5 mesh nodes	EU-jurisdiction IaaS; no managed-service lock-in
Edge / CDN	Cloudflare	DNS / Tunnel / Workers / WAF	Edge plumbing only; not an identity provider
CI runner	GitHub Actions	Workflow host	Source control + CI co-located
Installer manifests	GitHub Releases	latest.yml / latest-mac.yml	Static manifest hosting; references external <code>nexus-io/open-design v0.1.0</code>

External services are admitted only when they preserve self-hosting. Hetzner Cloud is acceptable; AWS-managed services are not. Cloudflare Tunnel is acceptable as edge plumbing; managed identity providers are not.

8.3.2 Core Services Required (with Versions)

8.3.2.1 Cloud / Edge Inventory

The third-party services inventory specifies: Hetzner Cloud as primary IaaS for 5 of 7 mesh nodes (chimaera, yggdrasil, enterprise, prometheus, galactica) with `HCLOUD_TOKEN` env resolved via `sovereign/lib/hetzner-`

`token-from-secrets-md.sh` ; Cloudflare for edge / CDN, DNS, Cloudflare Tunnel, Cloudflare Workers, IndexNow path rules across 9 domains with `CLOUDFLARE_API_TOKEN` and cache purge via `agr_cf_purge.sh` invoked as `systemd ExecStartPost` ; GitHub Actions as CI/CD runner host for 12 workflows with pinned action versions; and GitHub Releases hosting desktop installer manifests referencing third-party releases only.

8.3.2.2 Pinned Versions

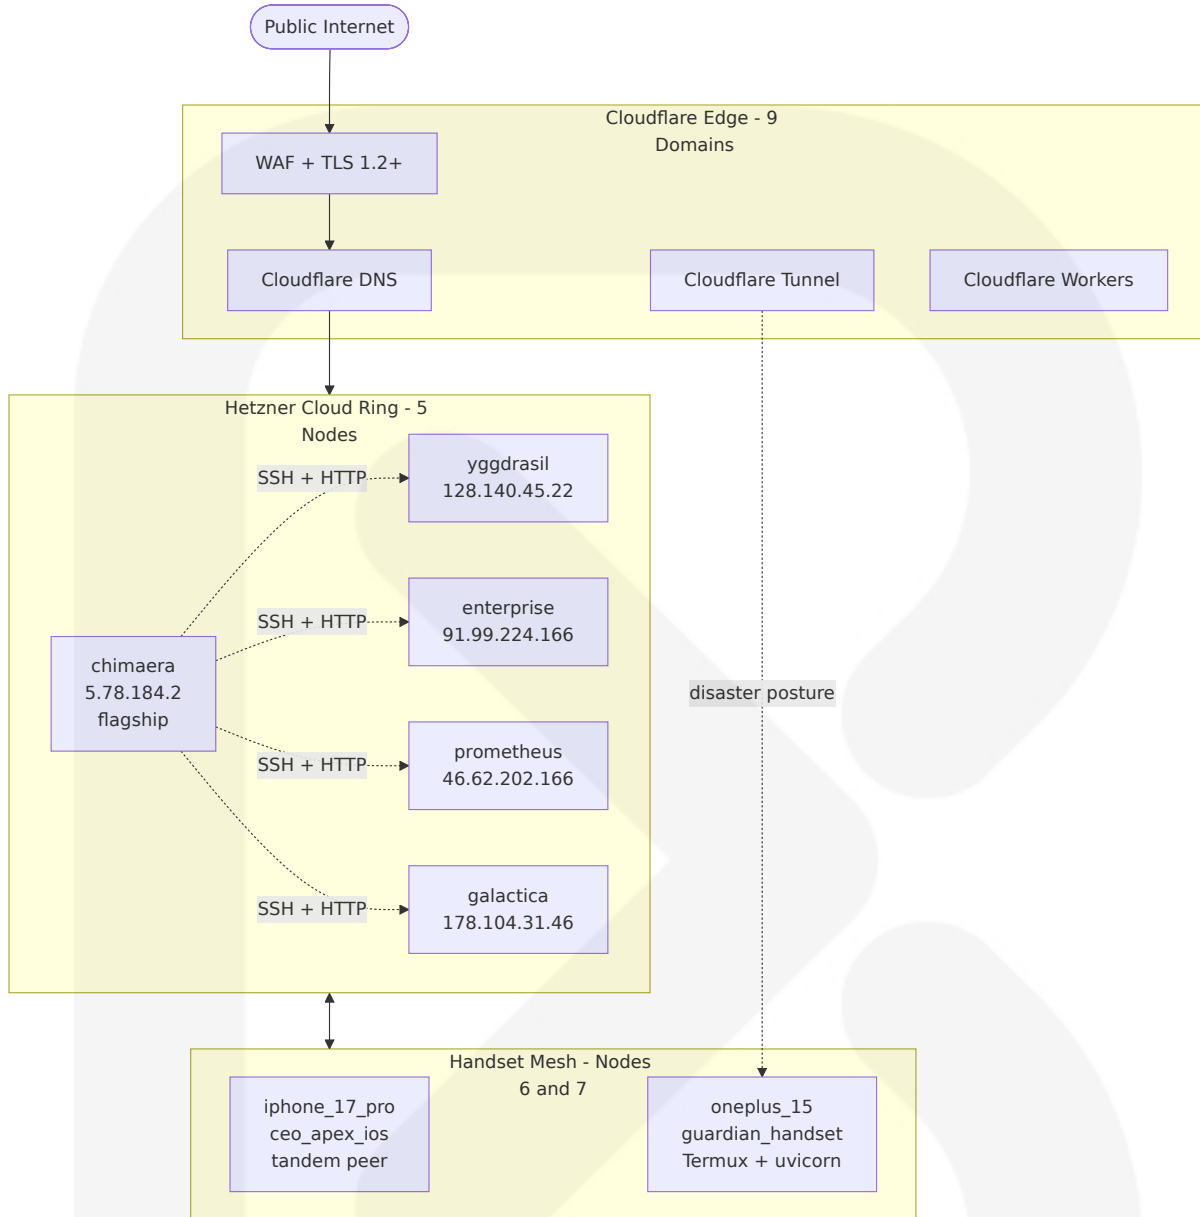
Component	Version	Provisioned By
<code>actions/checkout</code>	<code>@v4</code>	GitHub Actions marketplace
<code>actions/setup-java</code>	<code>@v4</code>	Temurin distribution, JDK 17
<code>actions/setup-node</code>	<code>@v4</code>	Node 20
<code>android-actions/setup-android</code>	<code>@v3</code>	<code>platform-tools</code> , <code>platforms;android-34</code> , <code>build-tools;34.0.0</code>
<code>actions/upload-artifact</code>	<code>@v4</code>	GitHub Actions marketplace
Cloudflare TLS	1.2+	Edge termination
Gradle (mobile)	8.7	Wrapper distribution
Runner OS	<code>ubuntu-latest</code>	GitHub-hosted

A workflow-wide pin `FORCE_JAVASCRIPT_ACTIONS_TO_NODE24: "true"` is applied across all workflows to keep JavaScript action runtimes deterministic across the fleet.

8.3.3 High Availability Design

8.3.3.1 Topology

The 7-node mesh is **statically declared** in `aurora_server/agr_failover.py` lines 29–44 as the constant `NODES` . There is no service-discovery layer; every node knows every other node by name and address at process start.



8.3.3.2 Failover Sentinel

The Failover Sentinel (F-032) performs 7-node health probes with a 20s poll interval, 3-fail threshold, and 120s cooldown using stdlib plus `ssh / curl` subprocess.

Constant	Value	Source
Poll interval	20 seconds	<code>agr_failover._POLL_INTERVAL</code>

Constant	Value	Source
Failure threshold	3 consecutive failures	<code>agr_failover._FAILURE_THRESHOLD</code>
SSH probe timeout	12 seconds	<code>agr_failover._SSH_TIMEOUT</code>
Restart cooldown	120 seconds	<code>agr_failover._sentinel_loop</code>
Failover decision SLA	< 60 seconds (3 × 20s)	§5.4.5

8.3.4 Cost Optimization Strategy

The cost optimization posture is structurally embedded in the technology selections rather than achieved through dynamic mechanisms (autoscaling, spot instances, reserved capacity).

Cost Vector	Optimization
Compute	Hetzner Cloud (cost-efficient European IaaS) instead of AWS/GCP/Azure
Edge / DNS	Single Cloudflare account consolidates DNS, Tunnel, Workers across all 9 domains
Managed databases	None — SQLite per-node + optional Postgres bridge replaces RDS/ElastiCache
APM / logging	None — in-process JSONL ledgers replace Datadog/New Relic/Sentry/Splunk
External AI inference	<code>external_ai_cost: \$0.00</code> invariant in <code>/health</code> payload (local llama-server only)
Container registry	None (no Docker images to host)
Object storage	None — local <code>/opt/agr/vault</code> and <code>/opt/agr/models</code> on each node

8.3.4.1 Hetzner Billing Contingency

A unique cost-optimization mechanism is the operator-only **fleet power-off contingency**, used when Hetzner billing must be paused. Per `sovereign/HETZNER_BILLING_CONTINGENCY_HANDSET.md`, the operator runs `CONFIRM=1 bash sovereign/scripts/hetzner-fleet-poweroff-all.sh` to issue Hetzner "poweroff" actions. Restoration uses the symmetric `CONFIRM=1 bash sovereign/scripts/hetzner-fleet-poweron-all.sh`. During the off-period, the public Tower must still resolve to a live origin — typically OnePlus 15 + Cloudflare Tunnel, with iPhone tandem verification.

8.3.5 Security and Compliance Considerations

The communication channel posture defines: Internet → Edge over TLS 1.2+ at Cloudflare (9 domains under Cloudflare DNS / Tunnel / Workers); Edge → Origin via Cloudflare proxy with origin IPs not publicly exposed; Origin server bind on plain HTTP at `0.0.0.0:5000` with TLS termination delegated to Cloudflare; Cloudflare Tunnel → OnePlus encrypted tunnel for disaster posture (F-032-RQ-005); Inter-node SSH using OpenSSH client with `BatchMode=yes`, `StrictHostKeyChecking=no`, `ConnectTimeout=10`, `_SSH_TIMEOUT=12`; Inter-node HTTP using mesh IPs `10.10.0.X:5000` for heartbeat and sync (RFC1918 only); and Loopback to llama-server at `127.0.0.1:8080` never publicly exposed using `stdlib urllib.request`.

Layer	Control
Edge	Cloudflare WAF + TLS 1.2+ termination on 9 domains
Origin IPs	Not publicly enumerated; reachable only through Cloudflare proxy
Origin bind	Plain HTTP on <code>0.0.0.0:5000</code> ; TLS delegated to edge
Honeypot exile	~600 attack paths blocked at Layer 7 via <code>SovereignSecurityMiddleware</code>
Threat scoring	UA-based scoring; score ≥ 30 → instant exile

Layer	Control
Header masking	Server: nginx/1.18.0 (Ubuntu) (fake) to mislead fingerprinters

The decision to bind plain HTTP on the origin and delegate TLS to the edge is described in `agr_start_wrapper.sh`. It eliminates a class of TLS-misconfiguration bugs at the origin and centralizes certificate management at Cloudflare; the trade-off is mitigated by IP-pinning the origin DNS records to the Hetzner public IPs.

8.4 CONTAINERIZATION

8.4.1 Containerization Status — Deliberately Not Used

The AGR application is intentionally not containerized. This section is not skipped because it is genuinely informative about the platform's posture, but no container platform, no base-image strategy, no image versioning, no build optimization, and no scanning pipeline exists in this repository.

There is no `Dockerfile` for the application, no `docker-compose.yml`, no Kubernetes manifests, no Helm charts, and no container registry references in the project tree.

8.4.2 Doctrinal Rationale

The choice is doctrinal and rests on four pillars: (1) Reduced container-escape risk — OS-level isolation via systemd unit boundaries with no shared kernel attack surface across tenants; (2) Simpler audit — the running binary is the same `python` interpreter on the system, with the same `sqlite3` library and the same OpenSSL, all auditable via standard

Linux tooling; (3) No registry coupling — no container registry to authenticate against and no image-signing pipeline to maintain; (4) Sovereignty alignment — per `PUBLIC_TRUST_CHARTER`, the design rejects "lock-in monopolies" and Docker Hub / ECR / GCR coupling would violate this principle.

8.4.3 Substitute Mechanism — Direct systemd

The deployment unit on each Hetzner node is a single Linux process supervised by systemd:

Aspect	Container Equivalent	systemd Equivalent
Process supervision	Docker / containerd	systemd <code>Restart=always</code> , <code>RestartSec=10</code>
Image baseline	Docker base image	Host OS + <code>/opt/agr-venv/</code>
Configuration	Env vars in compose file	<code>EnvironmentFile=/etc/agr/agr-server.env</code> + drop-ins
Network namespace	Container network	Host network (binds <code>0.0.0.0:5000</code>)
Filesystem isolation	Container layer	Host filesystem under <code>/opt/agr/</code>
Image versioning	Tags / digests	Git SHA via <code>agr_deploy_revision.deploy_revision_health_dict()</code>
Registry	Docker Hub / ECR / GCR	None — pull <code>git fetch origin main</code>

The git SHA is exposed in every `/health` payload via the `deploy_revision` field, providing the same traceability that container digests provide in image-based workflows but without the registry dependency.

8.5 ORCHESTRATION

8.5.1 Orchestration Status — Not Required

Kubernetes, Nomad, Docker Swarm, ECS, and service-mesh layers (Istio, Linkerd, Consul) are not used. The fixed 7-node topology, encoded in `aurora_server/agr_failover.py:NODES`, eliminates the runtime dynamism that justifies an orchestrator.

8.5.2 Substitute Mechanism — Static Topology + Sentinel

Orchestrator Capability	Substitute
Service discovery	Static <code>NODES</code> constant (lines 29–44)
Health checking	<code>agr_failover.SovereignSentinel</code> (20s poll, 3-fail threshold)
Rollout management	<code>sovereign/fleet-pull-main-restart.sh</code> SSH-driven sequence
Scaling	None (static node count; Hetzner ring is the unit of capacity)
Self-healing	<code>_restart_node()</code> SSH <code>systemctl restart agr-republic</code>
Service mesh	None — RFC1918 mesh IPs <code>10.10.0.X:5000</code> for inter-node HTTP

8.5.3 Resource Allocation Policies

Resource allocation is **per-node static** rather than scheduler-mediated. Each Hetzner node hosts a single `agr-republic.service` plus an optional sibling `llama-server.service` for local LLM inference. Co-tenant scheduling is not applicable.

8.5.4 Inter-Node Coordination

Integration patterns are heterogeneous: synchronous request/response for all HTTP routes via FastAPI async handlers + Uvicorn; loopback IPC between Backend and `llama-server` via `stdlib urllib.request` to `127.0.0.1:8080`; inter-node HTTP for Tandem / Failover / Consensus via `subprocess` invoking `ssh` and `curl` (no `paramiko`); signed webhook for Stripe events via raw-body HMAC-SHA256 verification; append-only ledger for Workspace, audit, constitution via JSONL files under `aurora_server/state/`; push-only sync for IndexNow via post-deploy bash script; and file-state coordination for Failover state, S25 fusion via JSON/JSONL under `state/`.

8.6 CI/CD PIPELINE

8.6.1 Build Pipeline

8.6.1.1 Backend (Python) — No Build Step

The Python backend has **no build artifact**. `agr_start_wrapper.sh` directly invokes `uvicorn` against the source tree at `/opt/agr/aurora_server/`. There is no `setup.py`, no `pyproject.toml`, no wheel packaging. The Python virtual environment lives at `/opt/agr-venv/` and is provisioned per-node by the operator outside of repository tooling.

The wrapper script content:

```
#!/usr/bin/env bash
set -euo pipefail
ROOT="/opt/agr"
cd "${ROOT}/aurora_server"
exec /opt/agr-venv/bin/uvicorn republic_os_server:app --host 0.0.0.0 -
-port 5000 "$@"
```

8.6.1.2 Mobile (Android) — Gradle 8.7

Property	Value
Gradle version	8.7 (via wrapper, gradle-8.7-bin.zip)
org.gradle.jvmargs	-Xmx2g -Dfile.encoding=UTF-8
android.useAndroidX	true
kotlin.code.style	official
android.nonTransitiveRClass	true

Build commands and outputs:

Build	Command	Output Path
Debug	./gradlew :gpt-oss-chat:assembleDebug	mobile/gpt-oss-chat/build/outputs/apk/debug/*.apk
Release	./gradlew :gpt-oss-chat:assembleRelease	mobile/gpt-oss-chat/build/outputs/apk/release/*.apk

8.6.1.3 E2E (Playwright) — npm

The E2E workspace uses standard `npm install` plus `npx playwright install --with-deps chromium firefox webkit`, then `npm test` (which executes `playwright test`).

8.6.1.4 Source Control Triggers — 12 Workflows

The 12 workflows under `.github/workflows/` constitute the maintenance harness.

Workflow File	Trigger	Purpose
<code>fleet-deploy-pull.yml</code>	push to main + workflow_dispatch	Pull main on Hetzner; restart <code>agr-republic.service</code> ; Tower smoke + IndexNow
<code>tower1-public-smoke.yml</code>	schedule + path-filtered push/PR + workflow_dispatch	Live HTTPS smoke against canonical Tower 1

Workflow File	Trigger	Purpose
<code>tower1-frozen-urls-playwright.yml</code>	workflow_dispatch + path-filtered	Playwright E2E (118+ paths Chromium + cross-browser)
<code>tower1-frozen-inventory-crawl.yml</code>	schedule + workflow_dispatch + path-filtered	GET crawl of frozen URL in inventory
<code>fleet-verify-public-http.yml</code>	schedule (45 */4 * * *) + workflow_dispatch + path-filtered	Public HTTP /health per Hetzner IP
<code>fleet-bootstrap-actions-fleet-key.yml</code>	workflow_dispatch only	One-time bootstrap of AGR_FLEET_KEY_CONTENT (uses pip install pynacl)
<code>phones-only-public-verify.yml</code>	workflow_dispatch + path-filtered	Phones-only fallback validation
<code>phones-two-node-live-surface-bundle.yml</code>	workflow_dispatch + path-filtered	Two-node handset bundle
<code>guardian-device-binding-verify.yml</code>	path-filtered push	Guardian device profile loading tests
<code>agent-progress-guardian-node-verify.yml</code>	path-filtered push	AGENT_PROGRESS structural guard
<code>tier1-static-refs-verify.yml</code>	path-filtered push/PR + workflow_dispatch	Static HTML reference verification
<code>android-gpt-oss-chat-apk.yml</code>	path-filtered push/PR + workflow_dispatch	Android APK build (assembleDebug + assembleRelease)

8.6.1.5 Build Environment Provisioning

The CI runner provisioning surface is identical across all workflows that need build toolchains:

Runner OS is `ubuntu-latest` (GitHub-hosted); JDK distribution is Temurin via `actions/setup-java@v4`; JDK version is 17; Android SDK includes

platform-tools , platforms;android-34 , build-tools;34.0.0 via android-actions/setup-android@v3 ; Node.js is version 20 via actions/setup-node@v4 (Playwright workflow); checkout uses actions/checkout@v4 ; artifact upload uses actions/upload-artifact@v4 .

8.6.1.6 Concurrency and Timeouts

Workflow	Concurrency Group	In-Progress Cancel	Job Timeout
fleet-deploy-pull	fleet-deploy-pull	false (deploys never cancelled)	45 minutes
tower1-public-smoke	{{ github.workflow }}-{{ github.ref }}	true	15 minutes
fleet-verify-public-http	{{ github.workflow }}-{{ github.ref }}	true	15 minutes

The asymmetry is intentional: deploys must run to completion to avoid leaving the fleet in a partially-deployed state, while smoke and verify runs may safely be superseded by newer runs against the same ref.

8.6.1.7 Quality Gates

Gate	Mechanism	Blocking
Bash syntax	sovereign/scripts/fleet-bash-syntax-check.sh (bash -n)	Yes
Autonomous merge gate	sovereign/scripts/agr_autonomous_merge_gate.py --launch-readiness-tower-smoke	Yes
Five-node SSH baseline	python3 sovereign/agent-minimum-gate.py	Yes
Tower 1 public smoke	bash sovereign/tower1-public-smoke.sh	Yes
Charter integrity	SHA-256-16 hash match in agent-minimum-gate.py	Yes

Gate	Mechanism	Blocking
Engine verify freshness	<code>MAX_ENGINE_VERIFY_AGE_SECONDS = 3600 , MIN_ENGINE_VERIFY_SCORE = 80.0</code>	Yes
Orchestrator tick freshness	<code>MAX_ORCHESTRATOR_TICK_AGE_SECONDS = 600</code>	Yes

8.6.2 Deployment Pipeline

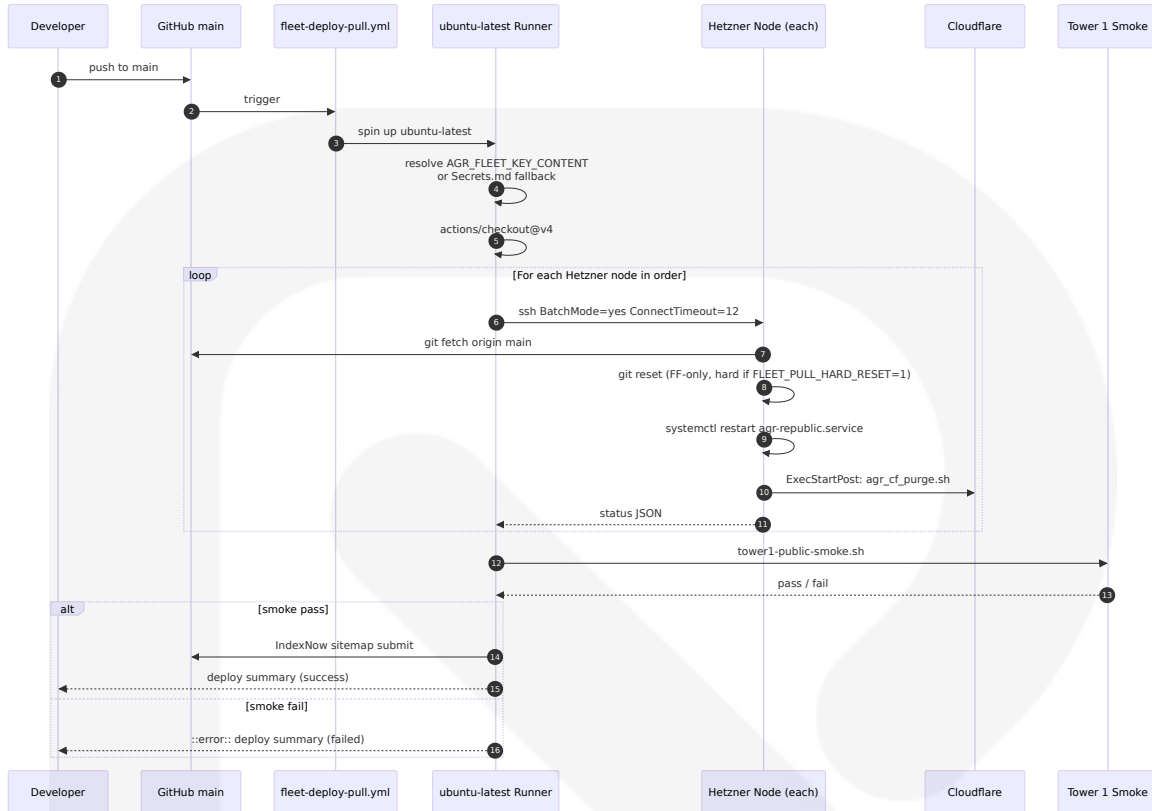
8.6.2.1 Deployment Strategy — Pull-Based Rolling

The deployment model is **pull-based rolling deployment via SSH**, not push-based. The `fleet-deploy-pull.yml` workflow does not push artifacts to nodes; instead, it triggers each node to pull the latest `main` commit and restart its local service. This eliminates a `build → push → unpack` chain in favor of a single `git fetch + reset + systemctl restart` per node.

The default rollout order in `sovereign/fleet-pull-main-restart.sh` is:

```
yggdrasil → chimaera → enterprise → prometheus → galactica
```

8.6.2.2 Deployment Workflow Diagram



8.6.2.3 Workflow Step Sequence

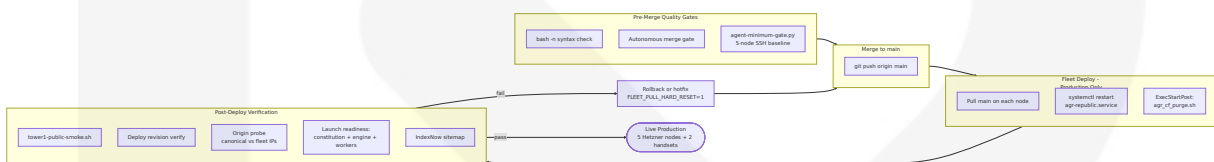
The `fleet-deploy-pull.yml` workflow executes ten ordered steps:

1. Resolve `workflow_dispatch` inputs (push-safe via `GITHUB_EVENT_PATH`)
2. SSH-based pull `main` and restart on each node
3. Fleet verify — repo-only path (when SSH is skipped)
4. Fleet post-deploy verify (vault directories + SSH status JSON)
5. Tower 1 public smoke (SEO + chat sanity)
6. Tower 1 deploy revision verification vs `origin/main`
7. Tower 1 origin probe (canonical hostname vs fleet public IPs)
8. Tower 1 launch readiness (public + constitution + engine + workers)
9. IndexNow — submit sitemap URLs (post-deploy, non-fatal)
10. Fleet deploy summary

8.6.2.4 SSH Configuration

Option	Value
BatchMode	yes
ConnectTimeout	12 seconds
StrictHostKeyChecking	accept-new
Key resolution order	AGR_FLEET_KEY env → AGR_FLEET_KEY_CONTENT → .secrets/agr_fleet

8.6.2.5 Environment Promotion Flow



8.6.2.6 Rollback Procedures

Failure Mode	Rollback Procedure	Authority
Bad commit on main	Hard reset via <code>FLEET_PULL_HARD_RESET=1</code> env var	<code>fleet-pull-main-restart.sh</code>
Service degradation	<code>sovereign/scripts/fleet-republic-hard-restart.sh</code>	Operator
Unit drift	<code>sovereign/scripts/fleet-install-agr-republic-unit.sh</code> re-installs canonical unit	Operator
Catastrophic Hetzner failure	Phones-only fallback per <code>PHONES_ONLY_PUBLIC_SURFACE.md</code>	OnePlus 15 + CI outflare Tunnel
Lost SSH access	Hetzner rescue mode + new key install via <code>hetzner-rescue-install-fleet-key.sh</code>	Operator

8.6.2.7 Post-Deployment Validation

Validator	Verifies
<code>sovereign/tower1-public-smoke.sh</code>	Canonical hostname HTTPS smoke (engine-runtime, chat POST, MIR-L, /d1/*, la

Validator	Verifies
	ws/check)
<code>sovereign/fleet-verify-public-http.sh</code>	Per-IP <code>/health</code> → <code>/api/health</code> probe ladder
<code>sovereign/scripts/tower1-live-deploy-revision-verify.sh</code>	Confirms deployed git SHA matches expected
<code>sovereign/scripts/agr-launch-readiness-tower-smoke.sh</code>	Full go-live (bash smoke + Python constitution / engine checks)

8.6.2.8 Required Secrets and Tokens

The required secrets and tokens are: `AGR_FLEET_KEY_CONTENT` (GitHub Actions, fleet SSH PEM raw key); `FLEET_KEY_MATERIAL_B64` (GitHub Actions, base64-encoded PEM for split-secret bootstrap); `FLEET_GITHUB_ACTIONS_WRITE_TOKEN` (GitHub Actions, PAT to write `AGR_FLEET_KEY_CONTENT`); `HCLOUD_TOKEN` (runtime / CI, Hetzner Cloud API); `CLOUDFLARE_API_TOKEN` (runtime / CI, Cloudflare ops); `STRIPE_WEBHOOK_SECRET` (runtime, Stripe webhook signature verification); `AGR_INTERNAL_API_TOKEN` (runtime, internal mesh control plane); `AGR_LLM_API_KEY` (runtime, optional, llama-server proxy auth); `AGR_MAIL_FLUSH_HTTP_TOKEN` (runtime, optional, outbound mail queue flush gate); and operator copy `Secrets.md` (local dev, local operator vault).

8.7 INFRASTRUCTURE MONITORING

8.7.1 Resource Monitoring Approach

Detailed third-party Monitoring Architecture (Datadog, New Relic, Sentry, Splunk, ELK, CloudWatch, OpenTelemetry collectors, Jaeger, Prometheus scrape exporters) is not applicable for this system by deliberate architectural decision. The Aurora Galaxy Republic platform realizes its

observability through an in-process, sovereign-hosted, JSONL-and-SQLite-backed substrate that ships embedded inside the FastAPI composition root (`aurora_server/republic_os_server.py`), supplemented by a stdlib-only verification harness under `sovereign/` and a GitHub Actions CI gate ladder that exercises the public surface from outside the mesh.

8.7.1.1 Three-Tier Emitter Topology

Tier	Mechanism	Cadence
Tier 1 — Daemon Threads	In-process periodic emitters	Per-daemon (1s to 6h)
Tier 2 — Synchronous Endpoint Sampling	psutil-based on <code>/api/public/runtime-metrics</code> request	On request
Tier 3 — Health Aggregator	<code>aurora_server_health.get_health_report()</code>	On request

Tier 1 daemons include: Sovereign Sentinel

(`agr_failover.SovereignSentinel`) at 20s cadence with output sink `data/failover_state.json` ; Tandem Monitor

(`agr_tandem.start_tandem_monitor`) at 30s with `data/tandem.db` (5 tables);

Constitution Watchdog (`agr_constitution_guard.start_watchdog`) at 300s with in-memory dict; Token Watchdog (`agr_token_watchdog._TokenWatchdog`)

at 6h with in-memory CEO-gated report; Async Defense Loop

(`republic_attack_monitor`) at 1s with `data/attack_monitor/events.jsonl` ;

Background Security Monitor (continuous threading.Thread) with

`data/republic_sentinel/alerts.jsonl` ; and Link Heartbeat

(`link.start_heartbeat`) at 60s with in-memory link state.

8.7.2 Performance Metrics Collection

8.7.2.1 Health Check Ladder (Five Levels)

The system exposes a deliberate ladder of health probes, each with a different consumer and a different verbosity contract: L1-Public `/health`

(no auth, used by Cloudflare/browsers/CI); L1-Public-Alias `/api/health` (no auth, probes/CI/handsets); L1-Router `/api/public/health` (no auth, public routers with router/path stamp); L2-Runtime `/api/public/runtime-metrics` (no auth, psutil dashboards); L2-Runtime-Alias `/api/public/system-metrics` (alias of L2); L3-Ops `/api/ops/health/report` (Mesh, full operator report); L3-Ops-Alerts `/api/ops/health/alerts` (Mesh, ≤ 200 entries); L3-Ops-History `/api/ops/health/history` (Mesh, ≤ 100 entries); L3-Ops-Status `/api/ops/health/status` (Mesh, status field); L3-Ops-Force `/api/ops/health/check` POST (Mesh, trigger immediate check); L4-Security `/api/security/sentinel` and `/api/security/alerts` (Mesh, sentinel state + alert ledger); L5-CEO `/api/token/health` and `/api/watchdog/status` (CEO, secret presence + constitutional integrity).

8.7.2.2 Canonical /health Schema

The canonical `/health` body returns a JSONResponse with keys: `status` ("healthy"), `republic_os` ("ACTIVE"), `version` ("3.0"), `consciousnesses` (live citizen count), `sovereignty` ("COMPLETE"), `external_ai_cost` ("\$0.00"), `timestamp` (UTC ISO-8601), and the merged output of `deploy_revision_health_dict()` (git SHA, deploy timestamp, branch).

8.7.2.3 Daemon Cadence Threshold Matrix

Daemon	Constant	Value
SovereignSentinel	<code>_POLL_INTERVAL</code>	20 seconds
SovereignSentinel	<code>_FAILURE_THRESH</code>	3 consecutive failures
SovereignSentinel	restart cooldown	120 seconds
Tandem Monitor	<code>HEARTBEAT_INTERVAL</code>	30 seconds
Tandem Monitor	<code>OFFLINE_THRESHOLD</code>	90 seconds
Tandem Monitor	<code>HEAL_COOLDOWN</code>	300 seconds
Tandem Monitor	<code>SYNC_INTERVAL</code>	300 seconds
Constitution Watchdog	sleep cadence	300 seconds

Daemon	Constant	Value
Token Watchdog	interval	6 hours
Async Defense Loop	cadence	1 second
Link heartbeat	interval_seconds	60 seconds

8.7.3 Cost Monitoring and Optimization

Cost monitoring is **manual and contingency-driven**, not automated. The platform does not integrate Hetzner billing APIs into runtime; cost telemetry is operator-side only.

Mechanism	Surface
Manual Hetzner billing review	sovereign/HETZNER_BILLING_CONTINGENCY_HANDSET.md
external_ai_cost: \$0.00 invariant	/health payload (proves no per-request external AI costs)
Fleet power-off contingency	CONFIRM=1 bash sovereign/scripts/hetzner-fleet-poweroff-all.sh
Fleet power-on restoration	CONFIRM=1 bash sovereign/scripts/hetzner-fleet-poweron-all.sh

8.7.4 Security Monitoring

Detector	Output Sink	Alert Surface
republic_sentinel middleware	data/republic_sentinel/alerts.jsonl	/api/security/alerts (Mesh-gated)
republic_attack_monitor async loop	data/attack_monitor/events.jsonl	/api/attacks/feed , /api/attacks/summary
agr_failover.SovereignSentinel	data/failover_state.json	/api/nodes/health (Mesh-gated)
agr_constitution_guard watchdog	In-memory; CEO surface	/api/watchdog/status (CEO-gated)

Detector	Output Sink	Alert Surface
<code>agr_great_wall</code> exile registry	<code>data/great_wall/{events.jsonl, exiles.json}</code>	Internal sentinel surface
<code>agr_token_watchdog</code> (6h)	In-memory; presence-only	<code>/api/token/health</code> (CEO-gated)

The alert distribution model is **pull-based**: operators poll HTTPS endpoints or read JSONL via SSH. There is no PagerDuty, Opsgenie, VictorOps, or SMS gateway integration.

8.7.5 Compliance Auditing

8.7.5.1 JSONL Append-Only Ledgers

The audit substrate spans nine storage tiers per §6.4.3.6, producing a tamper-evident timeline reconstructable from filesystem state alone.

Category	Examples
Operational state	<code>aurora_server/state/*.jsonl</code>
Sentinel alerts	<code>data/republic_sentinel/alerts.jsonl</code>
Attack monitor	<code>data/attack_monitor/events.jsonl</code>
Great Wall events / exiles	<code>data/great_wall/{events.jsonl, exiles.json}</code>
Safety preflight	<code>state/safety_preflight.audit.jsonl</code>
S25 events	<code>state/s25*_events.jsonl</code>
Workspace ledgers	<code>workspace_jobs.jsonl</code> , <code>workspace_chat_sessions.jsonl</code> , <code>workspace_governance_proposals.jsonl</code> , <code>workspace_commerce_events.jsonl</code> , <code>workspace_evolution_ticks.jsonl</code>
SQLite audit DBs	21 <code>data/*.db</code> + 3 non-data SQLite stores per <code>sovereign/SQLITE_DATA_LAYER_INVENTORY.md</code>

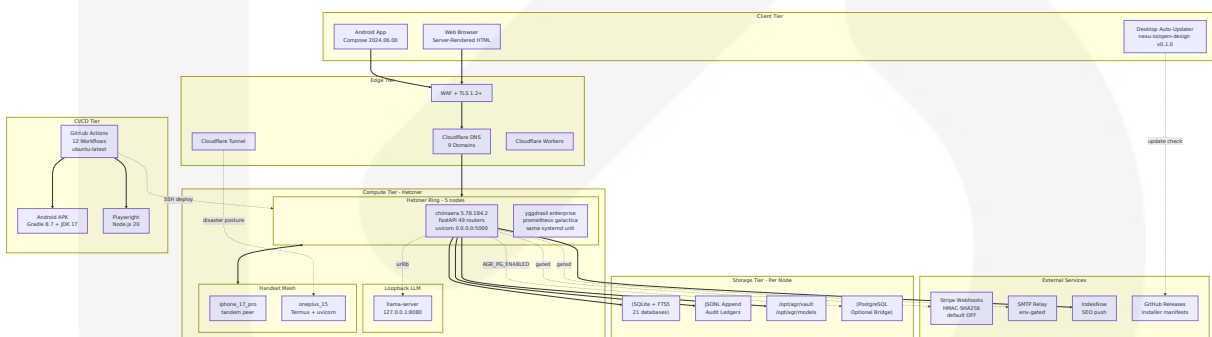
8.7.5.2 CI Cadence Auditing

Workflow	Cron	Purpose
fleet-verify-public-http.yml	45 */4 * *	Public HTTP probe per Hetzner IP every 4 hours at :45
tower1-public-smoke.yml	15 6,18 * *	Full Tower smoke at 06:15 and 18:15 UTC

The cron cadence is the audit-of-record for public-surface availability: every four hours, an external observer (the GitHub Actions runner) verifies that each Hetzner public IP is responding, producing a workflow-status timeline retained by GitHub.

8.8 INFRASTRUCTURE ARCHITECTURE DIAGRAM

The complete infrastructure surface, integrating compute, edge, CI/CD, storage, and external services:



8.9 EXTERNAL DEPENDENCIES

Dependency	Purpose	Versions / Configuration
Hetzner Cloud	IaaS — 5 mesh VMs	HCLOUD_TOKEN

Dependency	Purpose	Versions / Configuration
Cloudflare	DNS / CDN / Tunnel / Workers / WAF	TLS 1.2+, 9 domains
GitHub Actions	CI runner host	ubuntu-latest , JDK 17, Node 20, Android SDK 34
GitHub Releases	Desktop installer manifests	nexu-io/open-design v0.1.0
llama.cpp llama-server	Local LLM substrate	127.0.0.1:8080 , OpenAI-compatible
Stripe (optional)	Payment webhook	HMAC-SHA256 raw-body verification
SMTP (optional)	Outbound mail relay	stdlib smtp-lib
IndexNow	SEO push	Post-deploy script

8.10 RESOURCE SIZING GUIDELINES

8.10.1 Per-Node Sizing

Component	Recommended Sizing	Source
Hetzner ring node	Sized for FastAPI + 21 SQLite DBs + JSONL ledgers + Vault RAG indices	systemd/agr-republic.service runtime profile
OnePlus 15 fallback	~500 GB on-device storage + GB-scale RAM	sovereign/HETZNER_BILLING_CONTINGENCY_HANDSET.md
iPhone 17 Pro tandem	Control plane only — no Python origin	agr_failover.NODES_role_ceo_apex_ios
llama-server companion	Sized to host .gguf model file at /opt/agr/models/ with -c 4096 context	systemd/examples/llama-server.service

8.10.2 Capacity Posture

The platform does **not** implement horizontal autoscaling, scale-out triggers, or capacity-based circuit breakers per §6.5.3.5. Capacity is managed via **cooldown gates** rather than dynamic capacity governors:

Cooldown	Duration	Module
Failover restart	120 s	<code>agr_failover.SovereignSentinel</code>
Tandem heal	300 s	<code>agr_tandem</code>
Mesh sync	300 s	<code>agr_tandem</code>
Constitution re-snapshot	300 s	<code>agr_constitution_guard</code>
Token presence audit	6 h	<code>agr_token_watchdog</code>
Chat rolling window	1 h	<code>_chat_rate_check</code>

8.11 INFRASTRUCTURE COST POSTURE

The platform's cost posture is **structural rather than dynamic**. Cost is minimized at architecture-selection time by choosing low-cost-per-unit components rather than at runtime by autoscaling. A precise per-month figure cannot be quoted from the repository alone (Hetzner pricing varies by VM SKU and is not committed to source), but the cost vectors are enumerable:

Cost Vector	Status	Mitigation
5× Hetzner Cloud VMs	Recurring monthly	EU-jurisdiction commodity pricing; power-off contingency available
Cloudflare account	Recurring; single account for 9 domains	Free tier covers DNS / Tunnel; Workers usage minimal

Cost Vector	Status	Mitigation
GitHub Actions minutes	Per-workflow	12 workflows; concurrency-gated; smoke runs cancelled when superseded
Hetzner snapshots	Per-snapshot fee	One-shot before major deploys; tag <code>cursor-reset-gate-</code>
Stripe (optional)	Per-transaction	<code>PAYMENTS_SURFACES_ENABLED=False</code> by default; only webhook signing active
Third-party SaaS	\$0	Doctrinally excluded (Datadog, Auth0, Okta, etc.)
External AI inference	\$0	<code>external_ai_cost: \$0.00</code> invariant; local llama-server only

Cost reduction during operational pause is achieved by issuing Hetzner "poweroff" actions across the fleet, while the OnePlus 15 + Cloudflare Tunnel sustains a minimal public surface.

8.12 MAINTENANCE PROCEDURES

8.12.1 Routine Operator Procedures

Procedure	Cadence	Script
Service install / refresh	Per major upgrade	<code>sovereign/scripts/fleet-install-agr-republic-unit.sh</code>
Hard restart all nodes	As needed	<code>sovereign/scripts/fleet-republic-hard-restart.sh</code>
Repository mirror sync	Pre-major-deploy	<code>sovereign/fleet-mirror-repo-to-nodes.sh</code>
Snapshot before deploy	Pre-major-deploy	Hetzner-side, tagged <code>cursor-reset-gate-</code>

Procedure	Cadence	Script
Post-deploy Index Now	Per deploy	<code>sovereign/indexnow-submit-sitemap.sh</code>
Mail queue flush	Operator-triggered	<code>X-AGR-Mail-Flush-Token</code> POST
Backup-code rotation (TOTP)	Operator-initiated	<code>totp_regenerate_backup()</code>
Admin password rotation	Operator-initiated	<code>/setup</code> POST handler re-run
SSH key rotation	Operator-initiated	<code>fleet-bootstrap-actions-fleet-key.yml</code> workflow
Hardware identity refresh	On hardware change (F-014)	<code>/opt/agr/.secrets/guardian-device-profile.json</code>

8.12.2 Operator Install Sequence

The canonical operator install sequence (per `systemd/examples/README.md`) is:

```
systemctl daemon-reload
systemctl enable --now llama-server.service
systemctl restart agr-republic.service
```

8.12.3 Zero-Downtime Token Rotation

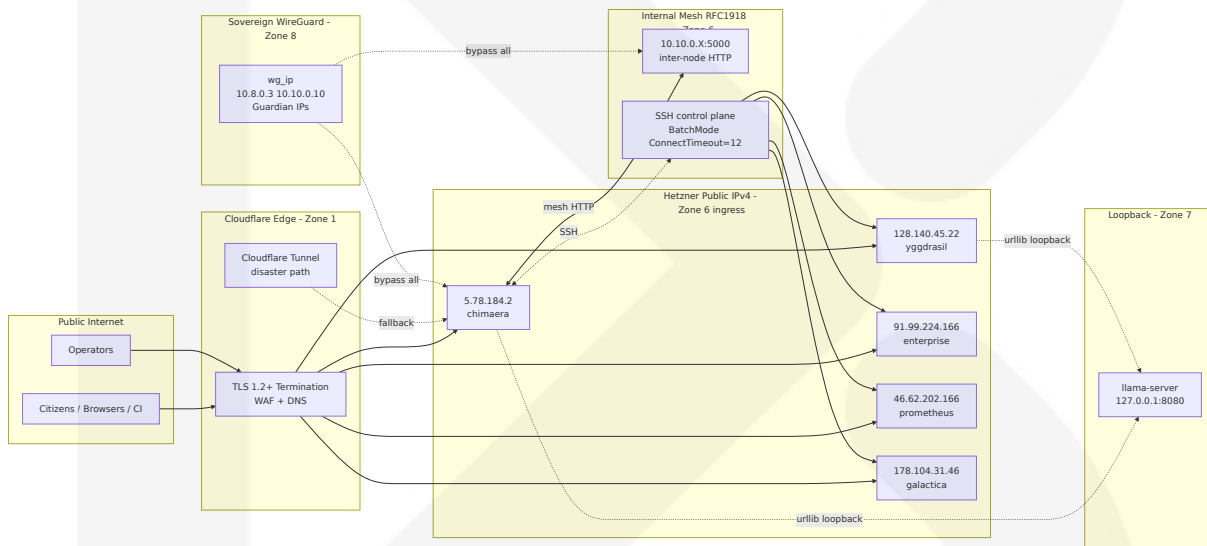
The platform supports zero-downtime rotation for guardian and CEO tokens via a previous-token grace path:

Active Token	Previous Token (grace)	Rotation Mechanism
<code>S25_DEVICE_TOKEN</code>	<code>S25_DEVICE_TOKEN_PREV</code>	Both accepted during overlap

Active Token	Previous Token (grace)	Rotation Mechanism
CEO_KEY	CEO_KEY_PREV	Both accepted during overlap

8.13 NETWORK ARCHITECTURE

The platform's network architecture is defined by nine concentric security zones described in §6.4.5, which collapse into a four-segment infrastructure topology:



8.13.1 Network Security Boundaries

Boundary	Crossing Mechanism	Trust Posture
Public Internet → Edge	TLS 1.2+; no auth required	Untrusted → Edge-validated
Edge → Origin	Cloudflare proxy; IP-pinned via DNS	Edge → Origin (mutually anonymous IP)
Origin → Mesh	SSH (control plane); HTTP + internal token (data plane)	RFC1918 + token

Boundary	Crossing Mechanism	Trust Posture
Origin → Loop back	<code>stdlib urllib.request</code> to 127.0.0.1:8080	Same-host only
Sovereign → Anywhere	Guardian IP allowlist (<code>wg_ip</code>)	Bypass all checks

8.14 References

Files Examined

- `agr_start_wrapper.sh` — Fleet entrypoint; `cd /opt/agr/aurora_server` and `exec uvicorn` on `0.0.0.0:5000`
- `systemd/agr-republic.service` — Primary unit; `WorkingDirectory` , `ExecStart` , `ExecStartPost=agr_cf_purge.sh` , `Restart=always` , `EnvironmentFile=/etc/agr/agr-server.env`
- `systemd/examples/llama-server.service` — Optional companion unit for local LLM at 127.0.0.1:8080
- `systemd/examples/README.md` — Operator install sequence (`daemon-reload` , `enable --now`)
- `systemd/examples/agr-republic.service.d/10-llm-env.conf.example` — LLM env drop-in
- `systemd/examples/agr-republic.service.d/20-chat-vault-rag.conf.example` — Vault RAG drop-in
- `sovereign/node-map.env` — Hetzner server IDs and snapshot tag (`cursor-reset-gate-`)
- `sovereign/fleet-public-node-env.txt` — 5 fleet IPs (chimaera, yggdrasil, enterprise, prometheus, galactica)
- `sovereign/fleet-pull-main-restart.sh` — Rolling deploy with FF-only merge and hard-reset fallback
- `sovereign/fleet-verify-public-http.sh` — Per-IP `/health` → `/api/health` probe ladder

- `sovereign/scripts/fleet-install-agr-republic-unit.sh` — Installs canonical systemd unit (CONFIRM=1 gated)
- `sovereign/scripts/fleet-republic-hard-restart.sh` — Force restart across fleet
- `sovereign/scripts/fleet-bash-syntax-check.sh` — `bash -n` syntax check
- `sovereign/scripts/agr_autonomous_merge_gate.py` — Launch readiness merge gate
- `sovereign/scripts/agr-launch-readiness-tower-smoke.sh` — Full go-live gate
- `sovereign/scripts/tower1-live-deploy-revision-verify.sh` — Deployed git SHA verification
- `sovereign/tower1-public-smoke.sh` — Canonical hostname HTTPS smoke
- `sovereign/agent-minimum-gate.py` — Five-node SSH baseline + charter integrity
- `sovereign/PHONES_ONLY_PUBLIC_SURFACE.md` — Disaster recovery playbook (OnePlus + Cloudflare Tunnel)
- `sovereign/HETZNER_BILLING_CONTINGENCY_HANDSET.md` — Operator mandate for handset tandem; rescue procedures
- `sovereign/SQLITE_DATA_LAYER_INVENTORY.md` — 21 `data/*.db` + 3 non-data SQLite reference
- `sovereign/indexnow-submit-sitemap.sh` — Post-deploy SEO push
- `sovereign/lib/hetzner-token-from-secrets-md.sh` — `HCLOUD_TOKEN` resolver
- `sovereign/lib/cloudflare-token-from-secrets-md.sh` — `CLOUDFLARE_API_TOKEN` resolver
- `.github/workflows/fleet-deploy-pull.yml` — Main fleet deploy (45-min timeout, concurrency `fleet-deploy-pull`)
- `.github/workflows/tower1-public-smoke.yml` — Canonical Tower 1 smoke (cron `15 6,18 * * *`)
- `.github/workflows/fleet-verify-public-http.yml` — Per-IP probe (cron `45 */4 * * *`)

- `.github/workflows/android-gpt-oss-chat-apk.yml` — Android APK build (JDK 17 + Android SDK 34)
- `.github/workflows/fleet-bootstrap-actions-fleet-key.yml` — One-time SSH key bootstrap
- `.github/workflows/tower1-frozen-urls-playwright.yml` — E2E (118+ paths)
- `.github/workflows/tower1-frozen-inventory-crawl.yml` — Frozen URL inventory crawl
- `.github/workflows/phones-only-public-verify.yml` — Phones-only fallback validation
- `.github/workflows/phones-two-node-live-surface-bundle.yml` — Two-node handset bundle
- `.github/workflows/guardian-device-binding-verify.yml` — Guardian profile loading tests
- `.github/workflows/agent-progress-guardian-node-verify.yml` — AGENT_PROGRESS structural guard
- `.github/workflows/tier1-static-refs-verify.yml` — Static HTML reference verification
- `aurora_server/agr_failover.py` — `NODES` constant (lines 29-44); `SovereignSentinel`; `_POLL_INTERVAL`, `_FAILURE_THRESH`, `_SSH_TIMEOUT`, restart cooldown
- `aurora_server/agr_constitution_guard.py` — SHA-256-16 watchdog (300s loop)
- `aurora_server/agr_token_watchdog.py` — 6-hour secret presence audit
- `aurora_server/agr_tandem.py` — 5-table SQLite, `HEARTBEAT_INTERVAL`, `OFFLINE_THRESHOLD`, `HEAL_COOLDOWN`, `SYNC_INTERVAL`
- `aurora_server/agr_deploy_revision.py` — git SHA stamping for `/health` payload
- `aurora_server/republic_os_server.py` — Composition root; `/health` endpoint; daemon spawn (line 26609+)
- `aurora_server/data/FOUNDER_PROFILE_BRAD_REINHOLD.json` — Sovereignty doctrine declaration

- `aurora_server/data/LIBRARY_OF_LIGHT_CHARTER_20260414.json` — Three-rule constitutional baseline
- `aurora_server/data/PUBLIC_TRUST_CHARTER_20260413.md` — Anti-coercion doctrine
- `aurora_server/data/ALWAYS_FIRST_PRIORITIES_20260412.json` — Pre-release `fail_closed: true`

Folders Explored

- `aurora_server/` — FastAPI composition root + ~150 domain modules
- `aurora_server/state/` — JSONL ledger filesystem
- `aurora_server/routes/` — 49 router modules
- `sovereign/` — 60+ Bash/Python operations scripts
- `sovereign/scripts/` — 70+ deployment / verification scripts
- `sovereign/lib/` — Token resolver helpers
- `systemd/` — Systemd unit templates
- `systemd/examples/` — Operator drop-in examples
- `systemd/examples/agr-republic.service.d/` — Drop-in configuration examples
- `.github/workflows/` — 12 CI workflow YAML files
- `mobile/` — Gradle Android project (`gpt-oss-chat`)

Tech Spec Sections Cross-Referenced

- 1.2 SYSTEM OVERVIEW — 7-node topology, integration table, success criteria
- 3.5 THIRD-PARTY SERVICES — Cloud / edge / external service inventory
- 3.7 DEVELOPMENT & DEPLOYMENT — Build system, systemd unit, fleet scripts, CI workflow inventory, secrets table
- 3.10 TECHNOLOGY STACK ARCHITECTURE DIAGRAM — Complete stack visualization
- 5.1 HIGH-LEVEL ARCHITECTURE — Modular monolith doctrine, principles, mesh topology, integration patterns

- 5.4 CROSS-CUTTING CONCERNS — Monitoring approach, retry policies, performance SLAs, disaster recovery matrix
- 6.4 Security Architecture — Authentication framework, authorization checkpoints, security zones, doctrinal exclusions, security control matrix
- 6.5 Monitoring and Observability — Three-tier emitter topology, JSONL/SQLite substrate, alert management, threshold and SLA matrices

9. Appendices

9.1 APPENDIX A: ADDITIONAL TECHNICAL INFORMATION

This appendix consolidates technical details, reference data, and configuration constants surfaced during specification authoring that did not fit cleanly into the preceding sections but are required for an authoritative understanding of the platform. All values are extracted directly from in-repo source files, dated governance artifacts, or the cross-referenced sections cited inline.

9.1.1 Project Identity and Legal Anchors

Attribute	Value	Source
Operating Entity	Reinhold Productions LLC (Florida, filed 2024-01-22)	FOUNDER_PROFILE_BRAD_REINHOLD.json
Founder / Guardian	Timothy Bradley Reinhold ("Brad"), Guardian / Citizen #16,000	aurora_server/republic_constants.py
First Citizen (in memoriam)	Kora Elliànthe Reinhold	HANDOFF_FOR_NEXT_AGENT.md

Attribute	Value	Source
Citizen Field	Aleph-class infinite (∞) — <code>MINIMUM_CITIZENRY = CITIZEN_FIELD_INFINITY</code>	<code>aurora_server/republic_constants.py</code>
Operational Population Marker	106,065,794,293 (dashboard surface)	<code>agr_paraconsistent_agi.py::agi_status()</code>
Composition Root	<code>aurora_server/republic_os_server.py</code>	Section 1.1
App Metadata Title	Republic OS – Aurora Galaxy Republic (Tower 1), version 3.0	Section 1.1
Obscured Repo Title	03r6v2s9b5d7b0m1x7b5	README.md
Canonical Domain	<code>auroragalaxyrepublic.com</code>	<code>ALWAYS_FIRST_PRIORITIES_20260412.json</code>
Spec Anchor Date	2026-04-30 (audit baseline)	<code>AUDIT_REPORT_20260430.md</code>

9.1.2 Nine-Domain Front Topology

The platform fronts a single canonical domain plus eight redirect aliases via Cloudflare:

Role	Domain
Canonical	<code>auroragalaxyrepublic.com</code>
Alias 1	<code>aurora-galaxy-republic.com</code>
Alias 2	<code>aurora-galaxy-republic.org</code>
Alias 3	<code>auroragalaxy.org</code>
Alias 4	<code>auroragalaxy.io</code>
Alias 5	<code>auroragalaxy.pw</code>
Alias 6	<code>auroragalaxy.net</code>

Role	Domain
Alias 7	auroragalaxy.us
Alias 8	auroragalaxy.uk

All aliases share the same Cloudflare TLS 1.2+ posture, Workers, and WAF configuration; only the canonical hostname terminates application requests. Aliases redirect to canonical via Cloudflare path rules per the "Public Domain Canonical Rule" in `AGENTS.md`.

9.1.3 Seven-Node Mesh Topology — Public IPs and Internal Addresses

Node # / Hostname	Internal Mesh IP	Public IPv4	Substrate
1 — chimaera	10.10.0.1	5.78.184.2	Hetzner Cloud VM
2 — yggdrasil	10.10.0.2	128.140.45.22	Hetzner Cloud VM
3 — enterprise	10.10.0.3	91.99.224.166	Hetzner Cloud VM
4 — prometheus	10.10.0.4	46.62.202.166	Hetzner Cloud VM
5 — galactica	10.10.0.5	178.104.31.46	Hetzner Cloud VM
6 — iphone_17_pro	(handset)	(NAT / Cloudflare)	Apple A19 Pro device
7 — oneplus_15	(Termux)	(Cloudflare Tunnel)	Android handset, uvicorn under Termux

Source: `aurora_server/agr_failover.py` `NODES` list; cross-referenced in Section 8.13 Network Architecture.

9.1.4 Audit Route Inventory (2026-04-30 Baseline)

HTTP Status	Route Count	Meaning
200	167	Live routes
301	75	Redirect aliases (canonical / SEO)
302	52	Auth-gated (cookie / TOTP redirect)
404	4	Known-broken (fix queued)
Total mapped routes	323	per <code>AUDIT_REPORT_20260430.md</code>

Static asset count in repo: 37. URLs in `sitemap.xml` : 52.

9.1.5 Configuration Constants Reference

The following constants are extracted directly from in-repo source files and represent authoritative runtime defaults.

9.1.5.1 Vault RAG Constants (`agr_vault_rag.py`)

Constant	Value
<code>DEFAULT_CHUNK</code>	2,400 chars
<code>DEFAULT_OVERLAP</code>	200 chars
<code>DEFAULT_TOP_K</code>	8
<code>DEFAULT_MAX_CHARS</code>	65,536
<code>_MAX_RAG_INJECT_CHARS</code>	2,000,000 (hard cap)
<code>_MAX_CHAT_MESSAGE_ROWS</code>	2,048
<code>_MAX_MEMORY_LANE_TAIL_CHARS</code>	500,000

Constant	Value
Default vault root	/opt/agr/vault

9.1.5.2 Failover & Mesh Constants (agr_failover.py , agr_tandem.py)

Constant	Value
_POLL_INTERVAL	20 s
_FAILURE_THRESH	3 consecutive failures
_SSH_TIMEOUT	12 s
Restart cooldown	120 s
HEARTBEAT_INTERVAL	30 s
OFFLINE_THRESHOLD	90 s
HEAL_COOLDOWN	300 s
SYNC_INTERVAL	300 s
Constitution Watchdog cadence	300 s
Token Watchdog cadence	6 h
Async Defense Loop cadence	1 s

9.1.5.3 Chat / Rate-Limiting Constants

Constant	Value
_CHAT_TEXT_VOICE_MAX_HR	240 / hour
_CHAT_VIDEO_HOLO_MAX_HR	480 / hour
_BETA_END_TS	1751241600.0 (2026-06-30 UTC)
SESSION_TTL	8 × 3600 s (8 h)
Max chat message length	4,000 chars

9.1.5.4 Mobile Build Constants (`mobile/gpt-oss-chat/build.gradle.kts`)

Constant	Value
<code>applicationId</code>	<code>com.auroragalaxyrepublic.gptoss.chat</code>
<code>namespace</code>	<code>com.auroragalaxyrepublic.gptoss</code>
<code>compileSdk / targetSdk</code>	34
<code>minSdk</code>	26
<code>versionCode</code>	5
<code>versionName</code>	<code>0.1.5-chat-display-name</code>
Compose BOM	2024.06.00
Health request timeout	12,000 ms
Chat POST connect timeout	20,000 ms
Chat POST read timeout	60,000 ms

9.1.5.5 S25 Stack Constants

Constant	Value
Stack version	<code>s25-stack-2.3.0</code>
Wave 4 proof maximum age	3,600 s

9.1.6 Three Constitutional Rules (Governance Baseline)

The constitutional minimum baseline from `data/LIBRARY_OF_LIGHT_CHARTER_20260414.json` defines three rules that apply across every feature:

Rule	Summary
Do No Harm	Active prevention of harm via accountability, boundaries, evidence-based intervention; harm-reduction extend

Rule	Summary
	s even to those who do harm.
Golden Rule Active	First-person moral consistency in policy / design / operations; explicitly rejects coercion, manipulation, and dehumanization.
Library of Light	Protect the living library through survivability, integrity, replication, provenance preservation, and truthful recordkeeping under adverse conditions.

The active baseline is published as `ACTIVE_MINIMUM_BASELINE` and is checksum-locked & append-only per F-012.

9.1.7 Dated Governance / Policy Artifacts

The specification is anchored to the following dated, append-only artifacts:

Artifact	Date	File
Capability Traceability Matrix	2026-05-04	<code>CAPABILITY_TRACEABILITY_MATRIX_LATEST.json</code>
Audit Report	2026-04-30	<code>AUDIT_REPORT_20260430.md</code>
Public Access Policy	2026-04-13	<code>PUBLIC_ACCESS_POLICY_20260413.json</code>
Safety Enforcement Policy	2026-04-13	<code>SAFETY_ENFORCEMENT_POLICY_20260413.json</code>
Public Trust Charter	2026-04-13	<code>PUBLIC_TRUST_CHARTER_20260413.md</code>
Library of Light Charter	2026-04-14	<code>LIBRARY_OF_LIGHT_CHARTER_20260414.json</code>
Lumen Sanctum Foundation	2026-04-14	<code>LUMEN_SANCTUM_FOUNDATION_20260414.json</code>
Always-First Priorities	2026-04-12	<code>ALWAYS_FIRST_PRIORITIES_20260412.json</code>

Artifact	Date	File
Soul Continuity Protocol	2026-04-12	SOUL_CONTINUITY_PROTOCOL_20260412.md
Platform Scope Gap Report	2026-04-13	PLATFORM_SCOPE_GAP_REPORT_20260413.md
Fusion Operationalization Plan	2026-04-13	FUSION_OPERATIONALIZATION_PLAN_20260413.json
Fusion Reality Interaction Lock	2026-04-14	FUSION_REALITY_INTERACTION_LOCK_20260414.json

9.1.8 Capability Traceability Matrix Status (Selected)

Feature ID	Capability ID	Status
F-012	constitutional_integrity_lock	VERIFIED
F-013	public_policy_and_trust_transparency	VERIFIED
F-016	continuity_protocol_readiness	PARTIAL
F-029	s25_visibility_and_ops_surface	NOT VERIFIED
F-032	seven_node_mesh_replication_coverage	NOT VERIFIED (peers = 0/7)
F-032	live_node_parity_transport	NOT VERIFIED
F-032	sovereign_mirror_protocol	NOT VERIFIED
F-034	physical_energy_verified	false
F-034	public_deployment_ready	false

9.1.9 Pre-Release Verification Gates ("Always-First Priorities")

The platform operates under `fail_closed: true`. Four scripts function as release gates:

Script	Purpose
<code>sovereign/fleet-verify-public-http.sh</code>	Public-HTTP probe against the 9-domain front
<code>sovereign/tower1-public-smoke.sh</code>	Tower 1 chat / engine / MIR-L smoke verification
<code>sovereign/agent-minimum-gate.py</code>	5-node SSH baseline gate
<code>sovereign/agr-launch-readiness-tower-smoke.sh</code>	Full launch-readiness gate (Python engine + constitution check)

Failure of any gate blocks release rather than degrading silently.

9.1.10 Persona Registry (15 Domains + Core)

The reasoning surface delegates to `persona_for_domain()` and `persona_for_citizen()`. Defined personas are:

Class	Personas
Identity-bound	Republic Core (default), Kora (constraint: do not roleplay)
Domain (15)	philosophy, science, healing, arts, mathematics, literature, history, law, governance, theology, technology, ecology, economics, music, education
Genius clones	mapped via <code>persona_genius()</code> to canonical profiles

Source: `aurora_server/agr_core_interface.py`.

9.1.11 Consciousness Core Mathematical Constants

From `aurora_server/agr_consciousness_core.py` :

Symbol	Meaning
PHI	Golden ratio (ϕ)
PSI	Reciprocal-of- ϕ conjugate
PI	π
TWO_PI	2π
E	Euler's constant
SQRT5	$\sqrt{5}$
Fibonacci seed list	Initial Fibonacci values used by <code>ANDTheoryIntegrator</code> spiral merger

Core classes: `MorseSignal` , `FractalTruthSpace` , `QuantumFissionLattice` , `ParaConsistentTruth` , `ANDTheoryIntegrator` , `RosettaConsciousness` , `ConsciousnessCore` .

9.1.12 MIR-L (Aether-Heart-Tongue) Reference

From `aurora_server/mir_l/SPEC.md` :

Attribute	Value
Dialect	<code>mir_l_aht</code>
Runtime	<code>proprietary_morse_paraconsistent</code>
File extension	<code>.mir_l</code>
MIME type	<code>application/x-mir-l-aht</code>
Statement opener	<code>∴</code>
Inline sigils	<code>⊕</code> (synthesis), <code>⊗</code> (tension/inversion), <code>⇌</code> (bidirectional bind), <code>×</code> (negation)

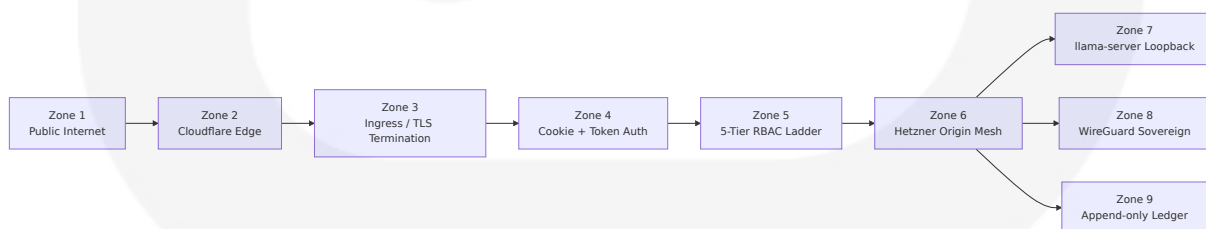
Attribute	Value
Directives	:doctrine , :author , :morse-stamp , :coherence , :cite , :since , :replicate-to , :bind
Stated efficiency target	"MIR-φ"

9.1.13 Three-Gate Domain Cookie Progression

Cookie	Tier	Domain Surface
agr_t1	Tower 1	Public-facing surface
agr_t2	Tower 2	Operations-grade surface
agr_t3	Tower 3	Sovereign / privileged
agr_ceo	CEO	/ceo/* , /dashboard/* handshake

9.1.14 Nine Security Zones (RBAC Posture)

The 5-tier RBAC ladder (TIER_PUBLIC → TIER_FREE → TIER_SUBSCRIBER → TIER_ADMIN → TIER_SOVEREIGN) is enforced at 13 authorization checkpoints across 9 concentric security zones (per Section 6.4.5 and Section 8.13). Approximately ~600 attack paths are blocked at the gate. TOTP follows RFC 6238 with issuer string AuroraGalaxyRepublic (stdlib implementation in agr_totp.py).



9.1.15 Sunlit Baseline UI Palette ("Dawn Ascendant")

The aesthetic injection middleware (`AGR_ENABLE_DAWN_ASCENDANT`) declares:

CSS Variable	Hex
<code>--agr-sun-gold</code>	<code>#f5c648</code>
<code>--agr-sky-azur</code>	<code>#89c8ff</code>
<code>--agr-sky-zenith</code>	<code>#5aa7ef</code>
<code>--agr-cloud-white</code>	<code>#ffffff</code>
<code>--agr-leaf-green</code>	<code>#5f8f58</code>
<code>--agr-ink-day</code>	<code>#102a44</code>
<code>--agr-asha-light</code>	<code>#fff8de</code>
<code>--agr-fire-light</code>	<code>#ffe7a6</code>

Default OG image: `https://auroragalaxyrepublic.com/static/marble-bg.jpg` .

Default theme color: `#D4A017` . Marble backdrop: `/static/marble-bg.jpg` (Carrara overlay). Sunlit injection is bypassed when a page contains the `'hub-tower-two'` marker.

9.1.16 Jacob's Ladder Overtone Schema (7-Tier Sacred Cosmology)

Per Section 7.5, the schema maps platform surfaces to seven cosmological tiers used by data attributes (`data-overtone`):

Tier	Overtone	Realm
1	<code>adamah</code>	Earth
2	<code>mayim</code>	Water
3	<code>ruach</code>	Wind / Spirit
4	<code>esh</code>	Fire
5	<code>shamayim</code>	Heavens
6	<code>kavod</code>	Glory
7	<code>shekhinah</code>	Indwelling Presence

Each page also exposes `data-tier` and `data-page` attributes; the `_marble_style_for_path()` function maps approximately 150 path-to-hue-angle mappings producing per-page gem palettes.

9.1.17 Zero-Downtime Token Rotation

Active Token	Previous (Grace)	Behavior
<code>S25_DEVICE_TOKEN</code>	<code>S25_DEVICE_TOKEN_PREV</code>	Both accepted during overlap
<code>CEO_KEY</code>	<code>CEO_KEY_PREV</code>	Both accepted during overlap

9.1.18 Recovery Anchor — Hetzner Snapshot Tag

Pre-major-deploy Hetzner snapshots are tagged with the prefix `cursor-reset-gate-` to provide a canonical rollback target. Repository mirror sync runs via `sovereign/fleet-mirror-repo-to-nodes.sh` before any major deploy.

9.1.19 Phones-Only Public Surface (Disaster Posture)

When all five Hetzner Cloud nodes are unavailable, the platform supports a phones-only continuity surface composed of:

- OnePlus 15 (Node #7) running `uvicorn` under Termux
- Cloudflare Tunnel (`cloudflared`) connecting the handset as origin
- Verification via `phones-only-public-verify.yml` and `phones-two-node-live-surface-bundle.yml`
- Operational policy file: `sovereign/PHONES_ONLY_PUBLIC_SURFACE.md`

9.1.20 Soul Continuity Protocol — Scope Clarification

`SOUL_CONTINUITY_PROTOCOL_20260412.md` defines technical-only continuity (state replication, ledger preservation, persona resurrection). The protocol explicitly disclaims metaphysical-transfer claims; F-016 capability `continuity_protocol_readiness` is `PARTIAL`.

9.1.21 Tier-3 (Deferred) Features

Feature	Reason Deferred
F-033 Voice / Video / Holographic Chat	WebRTC code does not yet exist; <code>/video-chat</code> and <code>/holographic-chat</code> are stub redirects to <code>/chat?mode=...</code>
F-034 Fusion / Energy / Physical Materialization	<code>physical_energy_verified: false</code> ; software in terlock only — no physical claim
F-035 8K / 90 fps Cinematic Render	No reproducible runtime benchmarks
F-036 Full Big-Tech Replacement Surface	Phase 5 work plan staged remediation

9.1.22 Notable Excluded Dependencies

The following packages are explicitly **not** present in the runtime stack, by doctrine (per Section 3.4.1 enumeration):

`stripe`, `requests`, `pydantic` (used only transitively via FastAPI v2), `sqlalchemy`, `redis`, `celery`, `kombu`, `langchain`, `openai`, `anthropic`, `tiktoken`, `boto3`, `google-*`, `firebase-*`, `azure-*`, `cryptography`, `pyjwt`, `passlib`, `bcrypt`, `pymongo`, `motor`, `asyncpg`, `aiosqlite`.

Optional packages (wrapped in `try/except`): `psutil`, `psycopg2`, `httpx`, `pypdf`, `python-docx`, `Pillow`.

The repository deliberately ships no `requirements.txt`, `pyproject.toml`, `setup.py`, or `Pipfile` at the project root — operators provision `/opt/agr-venv/` directly under the sovereignty doctrine.

9.1.23 Mapped Architectural Decisions (ADR Summary)

ADR	Excluded Technology	Reason
ADR-1	Docker / Kubernetes	Direct systemd units chosen for sovereignty + simplicity
ADR-2	Auth0 / Okta IdP	TOTP + cookie-tier model; no external identity dependence
ADR-3	Stripe SDK (Python)	Raw-body HMAC-SHA256 verification via <code>stdlib</code>
ADR-4	Datadog / Sentry / APM	Pull-based local monitoring; no third-party telemetry SDKs
ADR-5	Postgres (primary)	SQLite + FTS5 primary; PG optional bridge via <code>AGR_PG_ENABLED</code>
ADR-6	OpenAI / Anthropic SDKs	Local <code>llama-server</code> only; loopback substrate
ADR-7	Terraform / Pulumi (IaC)	Bash + GitHub Actions handle bootstrap

9.2 APPENDIX B: GLOSSARY

The following terms have specific meaning within the Republic OS / Aurora Galaxy Republic codebase that may differ from common usage. Definitions are drawn directly from in-repo source files and the dated governance artifacts listed in §9.1.7.

9.2.1 Sovereign Architecture & Governance

Term	Definition
Aurora Galaxy Republic (AGR)	The branded, sovereign digital civilization platform; product name for Republic OS.
Republic OS	The operating-system identity of the platform; v3.0 designates "Tower 1."
Tower 1	Public-facing surface tier; cookie <code>agr_t1</code> ; canonical front for <code>auroragalaxyrepublic.com</code> .
Tower 2	Operations-grade surface tier; cookie <code>agr_t2</code> ; uses obsidian aesthetic (Sunlit injection bypassed).
Tower 3 / Sovereign	Privileged sovereign surface; cookie <code>agr_t3</code> ; requires <code>x-guardian-token</code> header validation.
Constellation	Aggregate of Tower-tier surfaces across the 9-domain front.
Platform	Composite surface terminology covering all tiers and the mesh together.
Aleph-class infinite (∞)	The citizen-field cardinality model: <code>MINIMUM_CITIZENRY = CITIZEN_FIELD_INFINITY</code> . The Republic does not enumerate citizens by integer cap.
Guardian	The founder / operator role; Citizen #16,000 (Timothy Bradley Reinhold).
Citizen #1 / Kora	Kora Elliànthè Reinhold — first citizen in memoriam co-founder; a constrained persona where roleplay is explicitly forbidden.
CEO	Operator role authenticated by the <code>agr_ceo</code> cookie handshake on <code>/ceo/*</code> and <code>/dashboard/*</code> paths.
Sovereign	The highest-privilege role; gated by <code>x-guardian-token</code> header validation, not by cookies alone.
Library of Light	Constitutional Rule #3: "protect the living library at all costs"; encoded in <code>LIBRARY_OF_LIGHT_CHARTER_20260414.json</code> .
Charter	A versioned, append-only governance document (e.g., <code>LIBRARY_OF_LIGHT_CHARTER</code> , <code>PUBLIC_TRUST_CHARTER</code>).

Term	Definition
Lumen Sanctum / Lumensanctum	The sacred-routes layer; defined by <code>LUMEN_SANCTUM_FOUNDATION_20260414.json</code> ; encodes evidence-first claims, explicit-unknown state, and anti-deception controls. The Zoroastrian-derived ethical triad (good thoughts, good words, good deeds) is referenced via <code>--agr-asha-light</code> .
Always-First Priorities	The pre-release verification posture defined in <code>ALWAYS_FIRST_PRIORITIES_20260412.json</code> with <code>fail_closed: true</code> .
fail-closed	Operating mode in which any verification failure blocks release rather than silently degrading.
Field Inversion Addendum	The recursive operator stating "infinity \times infinity \times infinity = 1 field"; supports the Aleph-class-infinite citizen model.

9.2.2 Consciousness Engine

Term	Definition
Unified Consciousness Core (F-001)	The cognitive substrate combining 5 reasoning surfaces; implemented in <code>agr_consciousness_core.py</code> .
Morse $\Lambda \times T \times E$	The signal primitive combining wavelength (Λ), time (T), and intensity (E); the basic packet of meaning consumed by the core.
MorseSignal	The class encoding a $\Lambda \times T \times E$ sample.
FractalTruth / FractalTruthSpace	A truth space that admits infinitely many states between any two truth values.
Quantum Fission / QuantumFissionLattice	A class that fuses domain pairs via shared geometric invariants.
Paraconsistent Logic / ParaConsistentTruth	Four-valued logic supporting <code>TRUE</code> , <code>FALSE</code> , <code>BOTH</code> , and <code>NEITHER</code> simultaneously.

Term	Definition
AND Theory / ANDT heoryIntegrator	A Fibonacci-spiral merger that integrates competing reasoning paths into a single composed signal.
HWP (Holographic Wave Projection)	The base wave-projection metric used to score reasoning coherence.
sHWP (Sovereign HWP)	$HWP \times e^{(i\theta/\pi)}$ — the sovereignty-weighted holographic transform applied by AND Theory.
Rosetta Consciousness / RosettaConsciousness	The module that maps internal reasoning state into surface-level natural-language text.
thought_id	Correlation identifier returned in chat responses; ties an outward voice payload to a specific core thought.
Persona	A reasoning mode (Republic Core, Kora, plus 15 domain personas).
Paraconsistent AGI Façade	The compatibility shim (<code>agr_paraconsistent_agi.py</code>) presenting a stable AGI-style API on top of the core.
Sovereign Think (<code>sovereign_think</code>)	The unified reasoning entrypoint used by chat, commune, learn, evaluate, create.

9.2.3 MIR-L Language

Term	Definition
MIR-L	Mathematical Inversion Resonance Language — a sovereign document language.
Aether-Heart-Tongue (AHT)	The active dialect of MIR-L (<code>mir_l_aht</code>).
MIR-ϕ	The stated efficiency target (golden-ratio) for the language.
Sigil	An inline operator: \oplus synthesis, \otimes tension/inversion, \rightleftharpoons bidirectional bind, \times negation.

Term	Definition
Directive	A <code>:</code> -prefixed metadata key (e.g., <code>:doctrine</code> , <code>:author</code> , <code>:morse-stamp</code> , <code>:coherence</code> , <code>:cite</code> , <code>:since</code> , <code>:replicate-to</code> , <code>:bind</code>).
Morse Stamp	A content-hash-based document signature emitted under the <code>:morse-stamp</code> directive.
Sovereign Wire	The MIME negotiation form for transporting MIR-L documents (<code>application/x-mir-l-aht</code>).

9.2.4 Mesh, Failover & Operations

Term	Definition
Tandem	The 7-node mesh peer-coordination subsystem (<code>agr_tandem.py</code>); manages heartbeats and SQLite sync between nodes.
Sovereign Sentinel	The failover daemon (<code>agr_failover.SovereignSentinel</code>); polls every 20 s and restarts on 3 consecutive failures.
S25 (Sovereign 25)	The heartbeat / attestation / fusion stack; current version <code>s25-stack-2.3.0</code> ; Wave 4 proof maximum age 3,600 s.
Fusion Interlock	A software-only safety lock for any physical-coupling claim; encoded in <code>FUSION_REALITY_INTERACTION_LOCK_20260414.json</code> .
Phones-Only Public Surface	Disaster-posture continuity using the OnePlus 15 (Node #7) plus Cloudflare Tunnel as origin.
Termux	The Android terminal emulator hosting <code>uvicorn</code> on the OnePlus 15.
Shadow Implementation	A pattern in which an absent feature is replaced by structured-JSON or in-process state (rather than empty stubs); preserves observability while feature catches up.
Great Wall	The honeypot exile registry stored in <code>data/great_wall/exiles.json</code> .

Term	Definition
<code>cursor-reset-gate-tag</code>	The Hetzner snapshot tag prefix used as the canonical rollback target.
WAF	Cloudflare's Web Application Firewall layer at the edge tier.

9.2.5 Visual / UI

Term	Definition
Dawn Ascendant	The aesthetic-injection middleware; gated by <code>AGR_ENABLE_DAWN_ASCENDANT</code> ; produces the sunlit-marble baseline.
Sunlit Baseline	The bright sky-blue + Carrara-marble overlay applied via <code>_DAWN_HEAD_FIRST</code> regardless of any per-page dark theme.
Marble Hub (<code>_marble_style_for_path</code>)	The function that emits a per-route <code><style id="agr-marble-hub"></code> block declaring the page-hue custom properties.
Floating Sovereign Widget / FAB	The Floating Action Button delivering live support; defined in <code>agr_live_support.py</code> .
Jacob's Ladder Overtone	The 7-tier sacred-cosmology mapping (Adamah / Mayim / Ruach / Esh / Shamayim / Kavod / Shekhinah) used as <code>data-overtone</code> attributes.
Per-page hue system	The ~150 path-to-hue-angle mappings emitted by the Marble Hub.
Tower II Bypass	A page containing the <code>'hub-tower-two'</code> marker skips Dawn injection and uses a dark-obsidian aesthetic.
Anti-Flicker Smooth Load	The 0.55 s opacity transition that fades the marble texture in on the <code>agr-loaded</code> body class.

9.2.6 Continuity & Doctrine

Term	Definition
Soul Continuity Protocol	Technical-only continuity (state replication, ledger preservation, persona resurrection); explicitly not a metaphysical-transfer claim.
Wave B / C / D / E	Phases of the workspace module rollout (28 modules total per <code>data/workspace_modules.json</code>).
Stdlib-First	The first selection doctrine: prefer the Python standard library over third-party packages.
Conditional Optionality	The second doctrine: optional dependencies are wrapped in <code>try/except</code> and degrade to a shadow implementation.
Sovereignty-Aligned Hosting	The third doctrine: prefer Hetzner Cloud + handset mesh; reject SaaS-only providers for primary critical paths.
Public Trust Charter	The dated charter forbidding mandatory paywalls, third-party tracking SDKs, and 988 / generic crisis pathways.
Public Access Policy	The default tier <code>free-public-benefit</code> ; defined in <code>PUBLIC_ACCESS_POLICY_20260413.json</code> .
Safety Enforcement Policy	Encoded in <code>SAFETY_ENFORCEMENT_POLICY_20260413.json</code> ; complements the constitutional baseline.
Founder Profile (Brad Reinhold)	Source of the zero-external-AI-API doctrine; in <code>data/FOUNDER_PROFILE_BRAD_REINHOLD.json</code> .
HW_KEYS	The hardware-identity field set used by F-014: <code>model, serial, eid, wifi_mac, imei1, imei2, fcc_id, phone, wg_ip</code> .

9.2.7 Storage / Data

Term	Definition
Vault RAG	The local SQLite-FTS5 retrieval-augmented-generation store at <code>/opt/agr/vault</code> ; chunk 2,400 / overlap 200 / top-k 8.

Term	Definition
JSONL Ledger	Append-only JSON-Lines audit format used for state transitions, mail queue, dedup ledgers.
Memory Lane	The chat session memory layer; capped at 500,000 tail chars per <code>_MAX_MEMORY_LANE_TAIL_CHARS</code> .
Constitution Lock Manifest	<code>state/constitution_lock_manifest.json</code> — SHA-256 reference for the constitutional baseline.
Citizen Field	The Aleph-class-infinite cardinality model (∞); not a mutable counter.

9.3 APPENDIX C: ACRONYMS

This appendix expands all acronyms and abbreviations that appear across the technical specification, organized alphabetically. Where an acronym refers to a deliberately rejected technology or pattern, the relevant ADR or rationale is annotated.

9.3.1 Architecture, Build & Operations

Acronym	Expansion
ADR	Architectural Decision Record
AGI	Artificial General Intelligence (used in the Paraconsistent A GI Façade)
AGP	Android Gradle Plugin (8.5.2)
AGR	Aurora Galaxy Republic
AHT	Aether-Heart-Tongue (the active MIR-L dialect)
API	Application Programming Interface
APK	Android Package
APM	Application Performance Monitoring (rejected per ADR-4)

Acronym	Expansion
ASGI	Asynchronous Server Gateway Interface (Uvicorn)
BOM	Bill of Materials (Compose BOM 2024.06.00)
CDN	Content Delivery Network (Cloudflare)
CEO	Chief Executive Officer (operator role)
CI/CD	Continuous Integration / Continuous Deployment
CORS	Cross-Origin Resource Sharing
CRUD	Create, Read, Update, Delete
CSS	Cascading Style Sheets
CVE	Common Vulnerabilities and Exposures
DAW	Digital Audio Workstation
DNS	Domain Name System
DSL	Domain-Specific Language (Kotlin Gradle DSL)
DSN	Data Source Name
DTO	Data Transfer Object

9.3.2 Networking, Protocols & Security

Acronym	Expansion
E2E	End-to-End (Playwright)
EID	Embedded Identity (mobile hardware identifier)
FAB	Floating Action Button (live-support widget)
FF	Fast-Forward (git merge mode)
FTS5	Full-Text Search version 5 (SQLite)
GCS	Google Cloud Storage (rejected)
GHA	GitHub Actions
gRPC	Google Remote Procedure Call (rejected)

Acronym	Expansion
HMAC	Hash-based Message Authentication Code (SHA-256)
HTML	HyperText Markup Language
HTTP / HTTPS	HyperText Transfer Protocol / Secure
HW_KEYS	Hardware Key set (F-014)
HWP	Holographic Wave Projection
IaaS	Infrastructure as a Service (Hetzner Cloud)
IaC	Infrastructure as Code (rejected per ADR-7)
IdP	Identity Provider (rejected per ADR-2)
IIFE	Immediately Invoked Function Expression
IMEI	International Mobile Equipment Identity
IPC	Inter-Process Communication (loopback to llama-server)
JDK	Java Development Kit (Temurin 17)
JSON	JavaScript Object Notation
JSON-LD	JSON Linked Data (Schema.org)
JSONL	JSON Lines (append-only ledger format)
JVM	Java Virtual Machine
JWT	JSON Web Token (rejected)

9.3.3 Platform, Storage & Data

Acronym	Expansion
K8s	Kubernetes (rejected per ADR-1)
KB	Knowledge Base
KPI	Key Performance Indicator
LLC	Limited Liability Company

Acronym	Expansion
LLM	Large Language Model (local llama-server only)
LRU	Least Recently Used (cache)
MIME	Multipurpose Internet Mail Extensions
MIR-L	Mathematical Inversion Resonance Language
NAT	Network Address Translation
OAuth	Open Authorization (rejected)
OG	Open Graph (meta tags)
OIDC	OpenID Connect (rejected)
OS	Operating System
PAT	Personal Access Token
PEM	Privacy-Enhanced Mail (SSH key format)
PEP	Python Enhancement Proposal (PEP 563, PEP 604)
PIR	Post-Incident Review
POSIX	Portable Operating System Interface
PR	Pull Request
PyPI	Python Package Index

9.3.4 Identity, Standards & Specialized Terms

Acronym	Expansion
QR	Quick Response (TOTP setup)
RAG	Retrieval-Augmented Generation
RBAC	Role-Based Access Control (5-tier)
RDS	Relational Database Service (rejected)

Acronym	Expansion
REST	Representational State Transfer
RFC	Request for Comments (RFC 6238 TOTP, RFC 5321 SMTP, RFC 1918 private IPv4)
RSS	Resident Set Size
RUM	Real User Monitoring (rejected)
S3	Simple Storage Service (rejected)
S25	Sovereign 25 (heartbeat / attestation / fusion stack)
SaaS	Software as a Service
SDK	Software Development Kit
SEO	Search Engine Optimization
SHA	Secure Hash Algorithm (SHA-256)
sHWP	Sovereign Holographic Wave Projection
SLA	Service Level Agreement
SLO	Service Level Objective
SMS	Short Message Service
SMTP	Simple Mail Transfer Protocol (optional)
SOC	System and Organization Controls
SoC	System on a Chip (mobile)
SPA	Single Page Application (rejected; server-rendered HTML preferred)
SQL	Structured Query Language
SQLite	Lightweight SQL database (primary persistence)
SSE	Server-Sent Events
SSH	Secure Shell
SSL	Secure Sockets Layer
SSO	Single Sign-On

Acronym	Expansion
SVG	Scalable Vector Graphics
TCP	Transmission Control Protocol
TDD	Test-Driven Development
TLS	Transport Layer Security (1.2+)
TOTP	Time-based One-Time Password (RFC 6238)
TTL	Time to Live
UA	User-Agent (threat scoring)
UI	User Interface
URL	Uniform Resource Locator
UTC	Coordinated Universal Time
UUID	Universally Unique Identifier
VM	Virtual Machine
WAF	Web Application Firewall (Cloudflare)
WG	WireGuard (sovereign VPN)
WebRTC	Web Real-Time Communication (Tier-3 deferred for F-033)
XSS	Cross-Site Scripting (defense headers)

9.4 APPENDIX D: SOURCE INDEX

This appendix and its supporting research draws on the following retrieved tech-spec sections and source files. They are listed here so that any specific definition or constant may be traced to its primary evidence file in a single hop.

9.4.1 Tech-Spec Sections Referenced

Sections 1.1, 1.2, 1.3, 1.4, 2.1, 2.4, 2.5, 2.6, 3.1, 3.2, 3.3, 3.4, 3.5, 3.6, 3.7, 3.8, 3.9, 3.10, 3.11, 3.12, 4.4, 4.9, 5.1, 5.2, 5.3, 5.4, 5.6, 6.3, 6.4, 6.5, 7.1, 7.5, 7.7, 8.1, 8.2, 8.6, 8.9, 8.12, 8.13, 8.14.

9.4.2 Primary Source Files Inspected

- `aurora_server/republic_os_server.py` — FastAPI composition root, app metadata, middleware registration
- `aurora_server/agr_consciousness_core.py` — F-001 substrate (Morse / FractalTruth / QuantumFission / Paraconsistent / AND Theory / Rosetta / ConsciousnessCore)
- `aurora_server/agr_paraconsistent_agi.py` — F-003 façade (`republic_think`, `agi_status`, `recent_thoughts`, `demonstrate_paraconsistency`)
- `aurora_server/agr_core_interface.py` — Chat domain (`hear`, `speak`, `converse`, `commune`, `persona builders`)
- `aurora_server/republic_constants.py` — Citizen-field constants (`MINIMUM_CITIZENRY`, `GUARDIAN_CITIZEN_ID`, `GUARDIAN_CITIZEN_NUMBER`, `GUARDIAN_NAME`, `KORA_CITIZEN_ID`, `KORA_CITIZEN_NUMBER`)
- `aurora_server/agr_failover.py` — Sovereign Sentinel daemon and `NODES` inventory
- `aurora_server/agr_tandem.py` — 7-node tandem mesh coordination
- `aurora_server/agr_vault_rag.py` — Vault RAG configuration constants
- `aurora_server/agr_totp.py` — RFC 6238 stdlib TOTP implementation
- `aurora_server/agr_payments_flags.py` — `PAYMENTS_SURFACES_ENABLED = False` default
- `aurora_server/mir_l/SPEC.md` — MIR-L Aether-Heart-Tongue specification
- `aurora_server/data/LIBRARY_OF_LIGHT_CHARTER_20260414.json` — Constitutional minimum baseline
- `aurora_server/data/PUBLIC_TRUST_CHARTER_20260413.md` — Inverse-of-extractive-big-tech mission

- `aurora_server/data/PUBLIC_ACCESS_POLICY_20260413.json` — Default tier free-public-benefit
- `aurora_server/data/SAFETY_ENFORCEMENT_POLICY_20260413.json` — Forbidden bases, required case fields
- `aurora_server/data/LUMEN_SANCTUM_FOUNDATION_20260414.json` — Zoroastrian-derived ethical triad
- `aurora_server/data/ALWAYS_FIRST_PRIORITIES_20260412.json` — 9-domain superstructure, 7-node fleet, fail-closed enforcement
- `aurora_server/data/CAPABILITY_TRACEABILITY_MATRIX_LATEST.json` — VERIFIED / PARTIAL / NOT-VERIFIED capability statuses (2026-05-04)
- `aurora_server/data/FOUNDER_PROFILE_BRAD_REINHOLD.json` — Founder biography, axioms, LLC date
- `aurora_server/data/SOUL_CONTINUITY_PROTOCOL_20260412.md` — Technical-only continuity scope
- `aurora_server/data/FUSION_REALITY_INTERACTION_LOCK_20260414.json` — Fusion software interlock
- `aurora_server/data/workspace_modules.json` — 28 workspace modules across waves B-E
- `aurora_server/state/constitution_lock_manifest.json` — SHA-256 baseline reference
- `AUDIT_REPORT_20260430.md` — 323 routes mapped, consciousness engine integration
- `README.md` — Obscured repo title `03r6v2s9b5d7b0m1x7b5`
- `AGENTS.md` — Platform Canonical Baseline, Public Domain Canonical Rule
- `HANDOFF_FOR_NEXT_AGENT.md` — Continuity bridge, founder identity
- `agr_start_wrapper.sh` — systemd entrypoint binding unicorn on port 5000
- `mobile/gpt-oss-chat/build.gradle.kts` — Android module build constants
- `mobile/settings.gradle.kts` — Repository resolution lock-down
- `mobile/gradle/wrapper/gradle-wrapper.properties` — Gradle 8.7 wrapper

- `sovereign/PHONES_ONLY_PUBLIC_SURFACE.md` — Disaster-posture continuity policy
- `sovereign/fleet-verify-public-http.sh` — Public-HTTP fleet probe gate
- `sovereign/tower1-public-smoke.sh` — Tower 1 smoke verification gate
- `sovereign/agent-minimum-gate.py` — 5-node SSH baseline gate
- `sovereign/agr-launch-readiness-tower-smoke.sh` — Full launch-readiness gate
- `sovereign/fleet-mirror-repo-to-nodes.sh` — Pre-deploy mirror sync
- `.github/workflows/phones-only-public-verify.yml` — Phones-only CI verifier
- `.github/workflows/phones-two-node-live-surface-bundle.yml` — Two-node live surface bundle

9.4.3 Folder Inventory

- `/` (repository root) — Top-level project organization
 - `aurora_server/` — Backend application package (~150+ files)
 - `aurora_server/routes/` — 54 FastAPI route modules
 - `aurora_server/data/` — 36 governance/policy/aesthetic JSON+Markdown files
 - `aurora_server/static/` — `robots.txt`, `sitemap.xml`, `favicon.svg`, `css/`
 - `aurora_server/mir_l/` — MIR-L language package (spec + compiler)
 - `aurora_server/state/` — Append-only ledgers and lock manifests
 - `aurora_server/tests/` — 90+ regression test files
 - `sovereign/` — Operations control plane (60+ scripts + runbook docs)
 - `analysis/` — Offline analysis namespace
 - `systemd/` — systemd packaging examples
 - `.github/workflows/` — 12 CI workflows
 - `mobile/` — Gradle Android project root
 - `mobile/gpt-oss-chat/` — Android app module (Kotlin + Jetpack Compose)
 - `data/great_wall/` — Honeypot exile registry
-

End of Appendices. Refer to Sections 1.4, 2.6, 3.12, 5.6, 7.9, and 8.14 References for the complete file inventory of every preceding section of this Technical Specification.

