FILE 1: `index.html`

```html
<!DOCTYPE html>
<html lang="en">
<head>
  <meta charset="UTF-8" />
  <meta name="viewport" content="width=device-width, initial-scale=1.0"/>
  <title>RogueOS™ Governance Layer</title>
  <meta name="description" content="Embedded AI governance, auditability, and oversight built directly into RogueOS™." />
</head>
<body>

<h1>RogueOS™ Governance Layer</h1>

<p>
RogueOS™ includes a built-in governance and oversight layer designed to provide
continuous control, auditability, and human authority over AI-enabled systems.
This is not a policy document. This is an operational governance surface.
</p>

<h2>What This Governs</h2>
<ul>
  <li>AI models and agents</li>
  <li>Prompts and inputs</li>
  <li>Tools and integrations</li>
  <li>Outputs and downstream actions</li>
  <li>Human override and escalation paths</li>
</ul>

<h2>Virtual Governance Team</h2>
<p>
RogueOS™ replaces traditional compliance, risk, and oversight teams with a
virtual governance layer that is enforced through code, logging, and controls.
</p>

<ul>
  <li><strong>Policy Engine:</strong> Defines permitted and prohibited behavior</li>
  <li><strong>Risk Registry:</strong> Tracks and scores operational AI risks</li>
  <li><strong>Event Logger:</strong> Records actions, decisions, and changes</li>
  <li><strong>Human Override:</strong> Enables human-in-the-loop authority</li>
  <li><strong>Evidence Vault:</strong> Produces regulator- and auditor-ready records</li>
</ul>
```

```
<h2>Oversight Modes</h2>
<ul>
  <li>Human-in-the-Loop (HITL)</li>
  <li>Human-on-the-Loop (HOTL)</li>
  <li>Human-out-of-the-Loop (HOOTL)</li>
</ul>

<h2>Compliance Alignment</h2>
<p>
This governance layer is aligned with recognized AI governance frameworks,
including:
</p>

<ul>
  <li>EU AI Act</li>
  <li>NIST AI Risk Management Framework</li>
  <li>ISO/IEC 23894</li>
  <li>OECD AI Principles</li>
</ul>

<h2>For Builders</h2>
<p>
Governance is implemented as a modular layer within RogueOS™ and can be embedded
into local-first, enterprise, or hybrid AI systems.
</p>

<h2>For Organizations</h2>
<p>
RogueOS™ Governance provides:
</p>

<ul>
  <li>Continuous audit trails</li>
  <li>Traceable decision-making</li>
  <li>Risk and incident documentation</li>
  <li>Human authority enforcement</li>
  <li>Exportable evidence for compliance and review</li>
</ul>

<p>
This governance layer is a sellable, licensable component of RogueOS™.
</p>

</body>
```

</html>

---

FILE 2: `governance.json`

```json
{
  "system": "RogueOS",
  "component": "GovernanceLayer",
  "version": "1.0",
  "capabilities": {
    "ai_inventory": true,
    "event_logging": true,
    "risk_registry": true,
    "human_override": true,
    "incident_reporting": true,
    "audit_export": true
  },
  "oversight_modes": [
    "human-in-the-loop",
    "human-on-the-loop",
    "human-out-of-the-loop"
  ],
  "governed_entities": [
    "models",
    "agents",
    "prompts",
    "tools",
    "outputs",
    "automations"
  ],
  "compliance_alignment": {
    "eu_ai_act": true,
    "nist_ai_rmf": true,
    "iso_23894": true,
    "oecd_ai_principles": true
  },
  "logging": {
    "enabled": true,
    "immutable": true,
    "retention_days": 365,
    "exportable": true
  },
  "human_authority": {
    "override_required_for_high_risk": true,
    "escalation_supported": true
  }
}
```

----------------------------------------------------------------------------------------

FILE 3: governance.md

# RogueOS™ Governance Policy

## Purpose
This document defines the governance, oversight, and control mechanisms embedded within RogueOS™.

Governance within RogueOS™ is enforced through operational controls, logging, and human authority—not through advisory policies alone.

## Scope
This policy applies to:
- AI models and agents
- Prompts and inputs
- Tools and integrations
- Outputs and automated actions
- Human oversight and intervention

## Core Principles
- Human authority is preserved at all times
- High-risk actions require explicit oversight
- All actions are traceable and logged
- Governance is continuous, not episodic

## Oversight Modes
RogueOS™ supports:
- Human-in-the-Loop (HITL)
- Human-on-the-Loop (HOTL)
- Human-out-of-the-Loop (HOOTL)

The applicable mode is determined by risk classification.

## Risk Management
RogueOS™ maintains a risk registry that:
- Identifies AI-related risks
- Scores operational impact
- Tracks mitigation actions
- Records incidents and outcomes

## Event Logging & Evidence
All significant actions, decisions, and changes are logged.

Logs are:
- Immutable
- Time-stamped
- Retained according to policy
- Exportable for audit and regulatory review

## Human Override
Human operators may intervene at any time to:
- Halt execution
- Modify behavior
- Escalate incidents
- Revoke permissions

## Compliance Alignment

This governance layer aligns with:

- EU AI Act
- NIST AI Risk Management Framework
- ISO/IEC 23894
- OECD AI Principles

## Review & Updates

This policy is reviewed periodically and updated as governance requirements evolve.

------------------------------------------------------------------------------------

FILE 4 (ROOT OF SITE): _headers

```
/governance/*
  X-Content-Type-Options: nosniff
  X-Frame-Options: DENY
  Referrer-Policy: strict-origin-when-cross-origin
  Content-Security-Policy: default-src 'self'
```