ELOSIA

ELEVATING CARE. EMPOWERING CLARITY.

ELOSIA SENTRY AUDIT

Powered by SentryShield AI Security

Elosia Sentry Audit
Powered by SentryShield AI Security

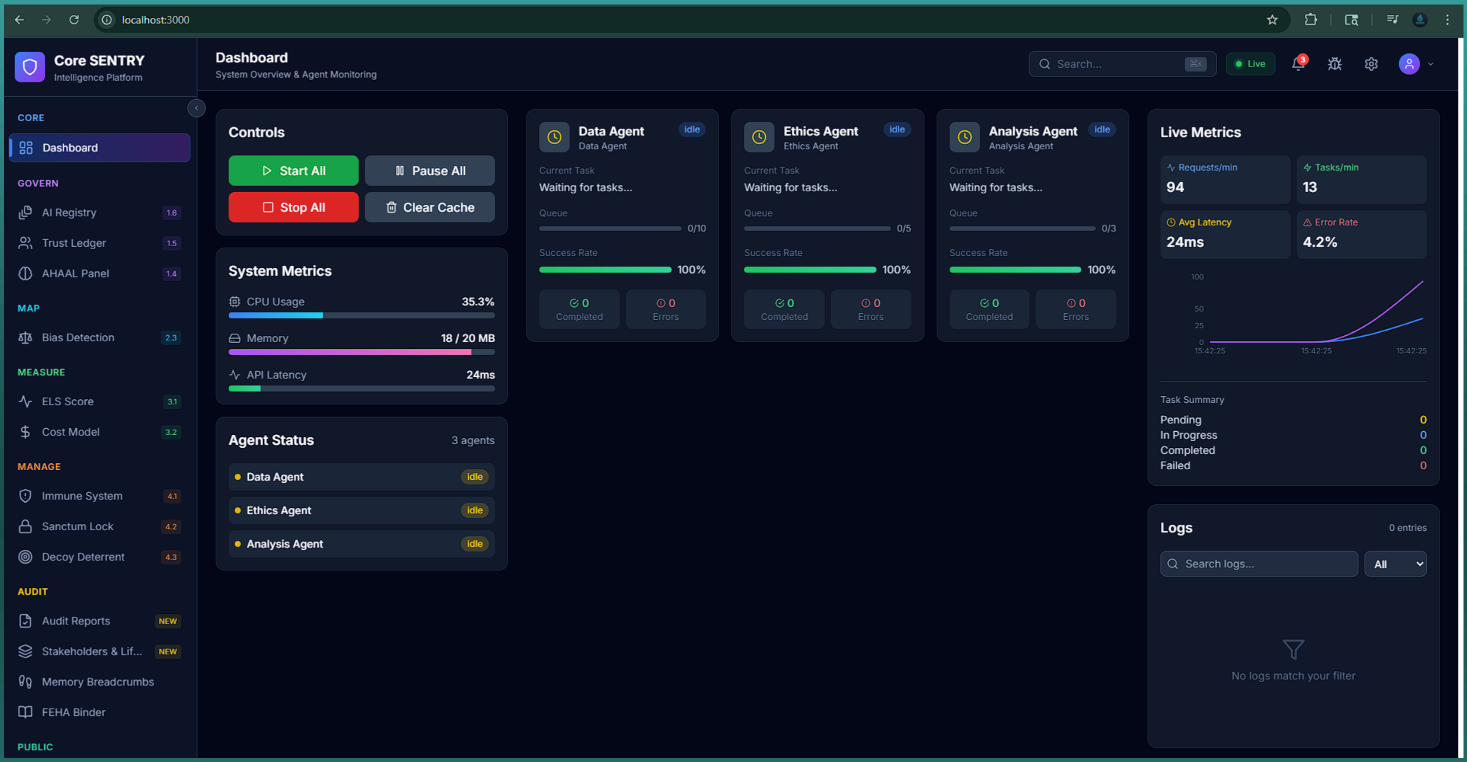

Dashboard
System Overview & Agent Monitoring

Controls

System Metrics

CPU Usage 35.3%

Memory 18 / 20 MB

API Latency 24ms

Agent Status 3 agents

- Data Agent idle
- Ethics Agent idle
- Analysis Agent idle

Data Agent idle
Data Agent
Current Task
Waiting for tasks...
Queue 0/10
Success Rate 100%
0 Completed
0 Errors

Ethics Agent idle
Ethics Agent
Current Task
Waiting for tasks...
Queue 0/5
Success Rate 100%
0 Completed
0 Errors

Analysis Agent idle
Analysis Agent
Current Task
Waiting for tasks...
Queue 0/3
Success Rate 100%
0 Completed
0 Errors

Live Metrics

Requests/min 94
Tasks/min 13
Avg Latency 24ms
Error Rate 4.2%

Task Summary

Pending	0
In Progress	0
Completed	0
Failed	0

Logs 0 entries

No logs match your filter

Real-time command center showing agent status, live metrics charts, and activity logs.

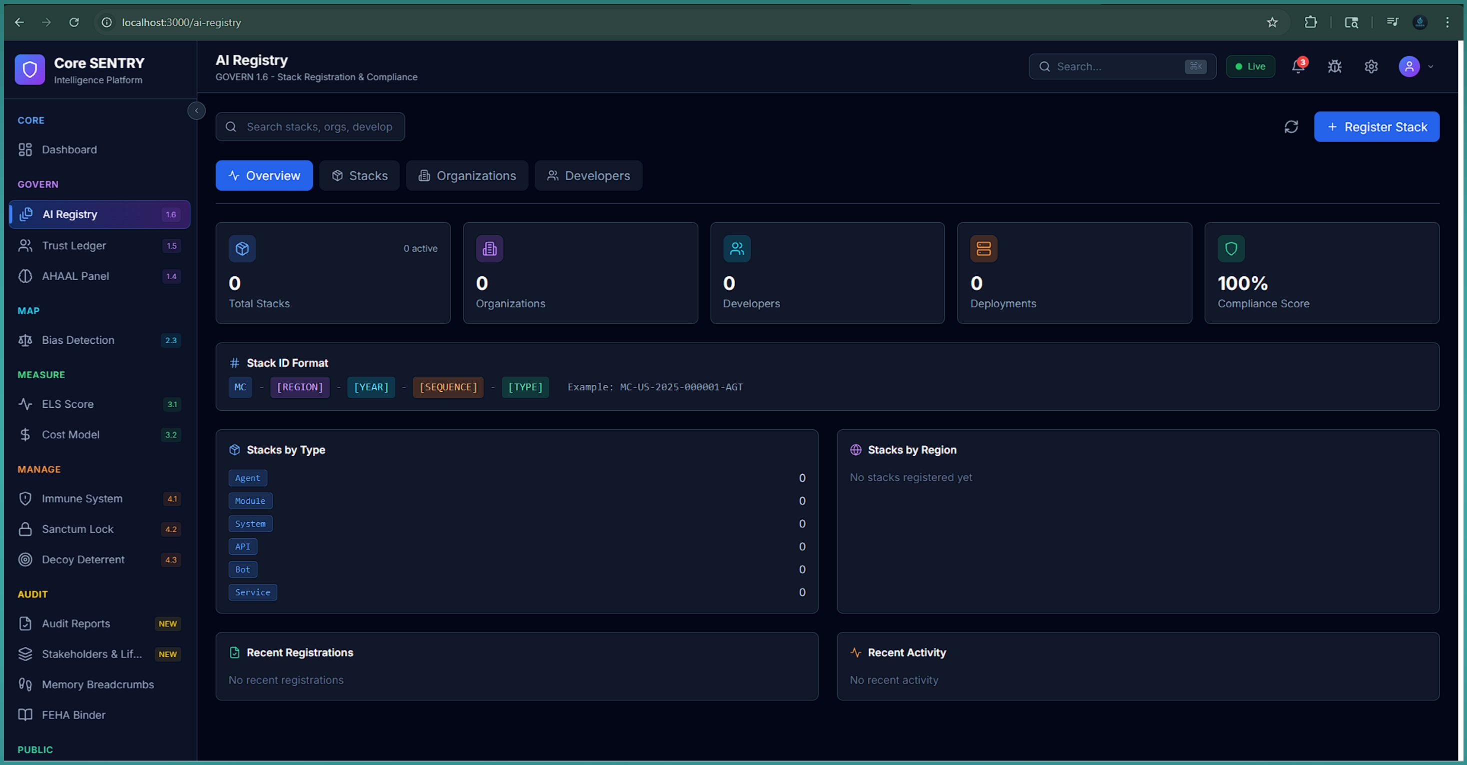

AI Registry
GOVERN 1.6 - Stack Registration & Compliance

0 active
0 Total Stacks

0 Organizations

0 Developers

0 Deployments

100% Compliance Score

Stack ID Format

`MC - [REGION] - [YEAR] - [SEQUENCE] - [TYPE]   Example: MC-US-2025-000001-AGT`

Stacks by Type

Agent	0
Module	0
System	0
API	0
Bot	0
Service	0

Stacks by Region

No stacks registered yet

Recent Registrations

No recent registrations

Recent Activity

No recent activity

Catalog of all registered AI systems with risk classifications, attestations, and compliance status.

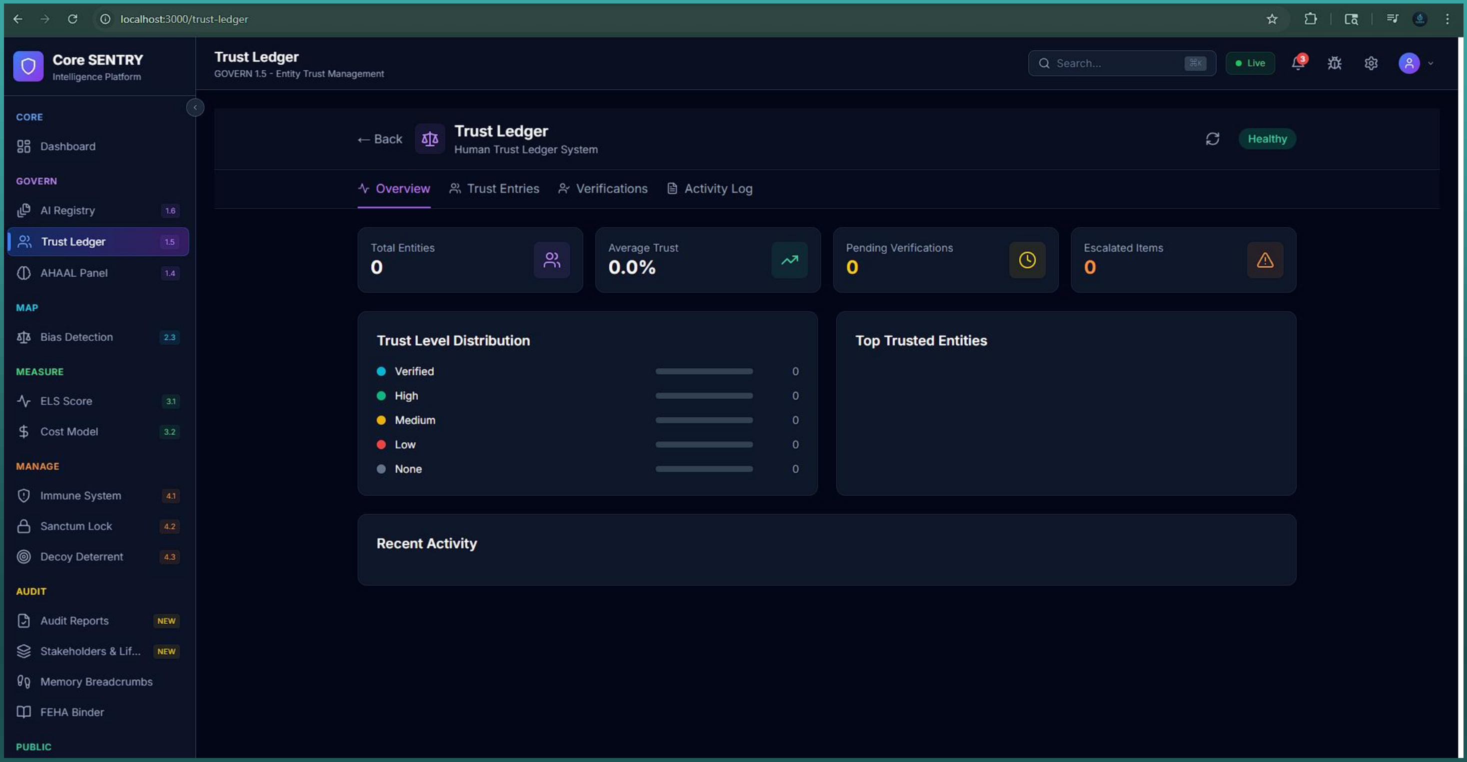

Trust Ledger
GOVERN 1.5 - Entity Trust Management

Trust Ledger
Human Trust Ledger System

Total Entities 0

Average Trust 0.0%

Pending Verifications 0

Escalated Items 0

Trust Level Distribution

Verified	0
High	0
Medium	0
Low	0
None	0

Top Trusted Entities

Recent Activity

Immutable record of trust scores, score history, and confidence trends over time.

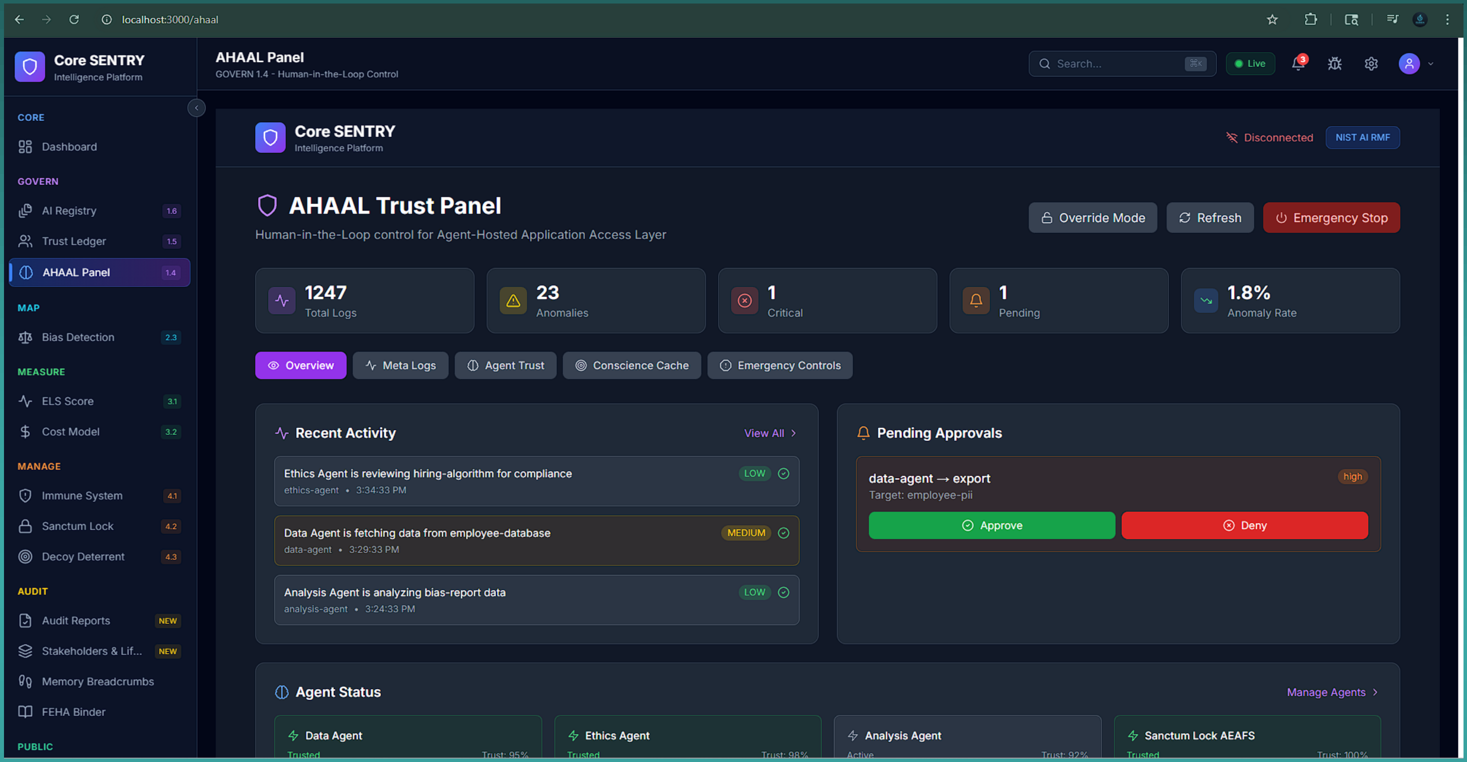

AHAAL Panel
GOVERN 1.4 - Human-in-the-Loop Control

AHAAL Trust Panel

Human-in-the-Loop control for Agent-Hosted Application Access Layer

1247 Total Logs

23 Anomalies

1 Critical

1 Pending

1.8% Anomaly Rate

Recent Activity

Ethics Agent is reviewing hiring-algorithm for compliance LOW
ethics-agent • 3:34:33 PM

Data Agent is fetching data from employee-database MEDIUM
data-agent • 3:29:33 PM

Analysis Agent is analyzing bias-report data LOW
analysis-agent • 3:24:33 PM

Pending Approvals

data-agent → export high
Target: employee-pii

Agent Status

Data Agent	
Trusted	Trust: 95%
Ethics Agent	
Trusted	Trust: 98%
Analysis Agent	
Active	Trust: 92%
Sanctum Lock AEAFS	
Trusted	Trust: 100%

Agent-Human Alignment Assurance Layer for monitoring AI conscience, trust levels, and emergency controls.

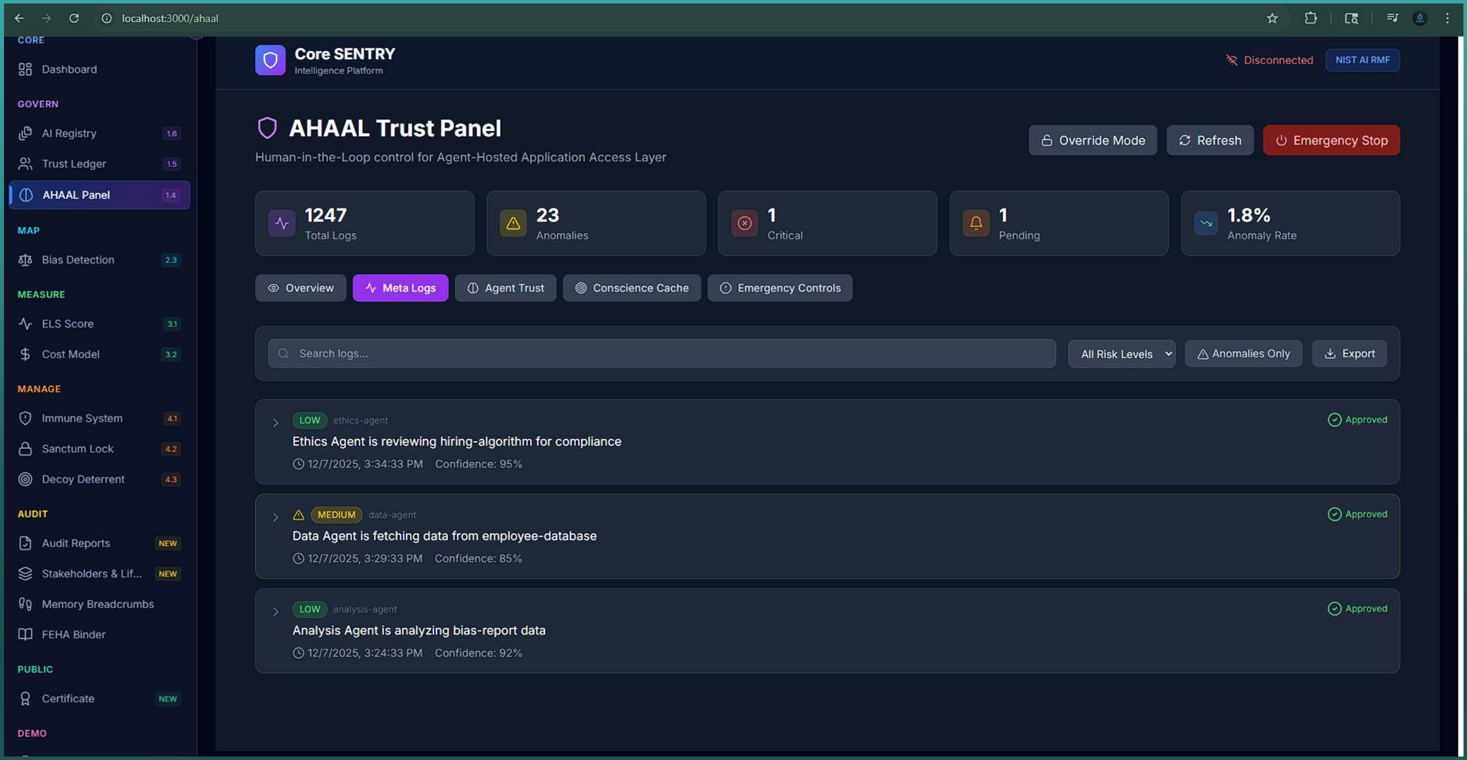

AHAAL Trust Panel

Human-in-the-Loop control for Agent-Hosted Application Access Layer

1247 Total Logs | 23 Anomalies | 1 Critical | 1 Pending | 1.8% Anomaly Rate

LOW ethics-agent — Approved
Ethics Agent is reviewing hiring-algorithm for compliance
12/7/2025, 3:34:33 PM Confidence: 95%

MEDIUM data-agent — Approved
Data Agent is fetching data from employee-database
12/7/2025, 3:29:33 PM Confidence: 85%

LOW analysis-agent — Approved
Analysis Agent is analyzing bias-report data
12/7/2025, 3:24:33 PM Confidence: 92%

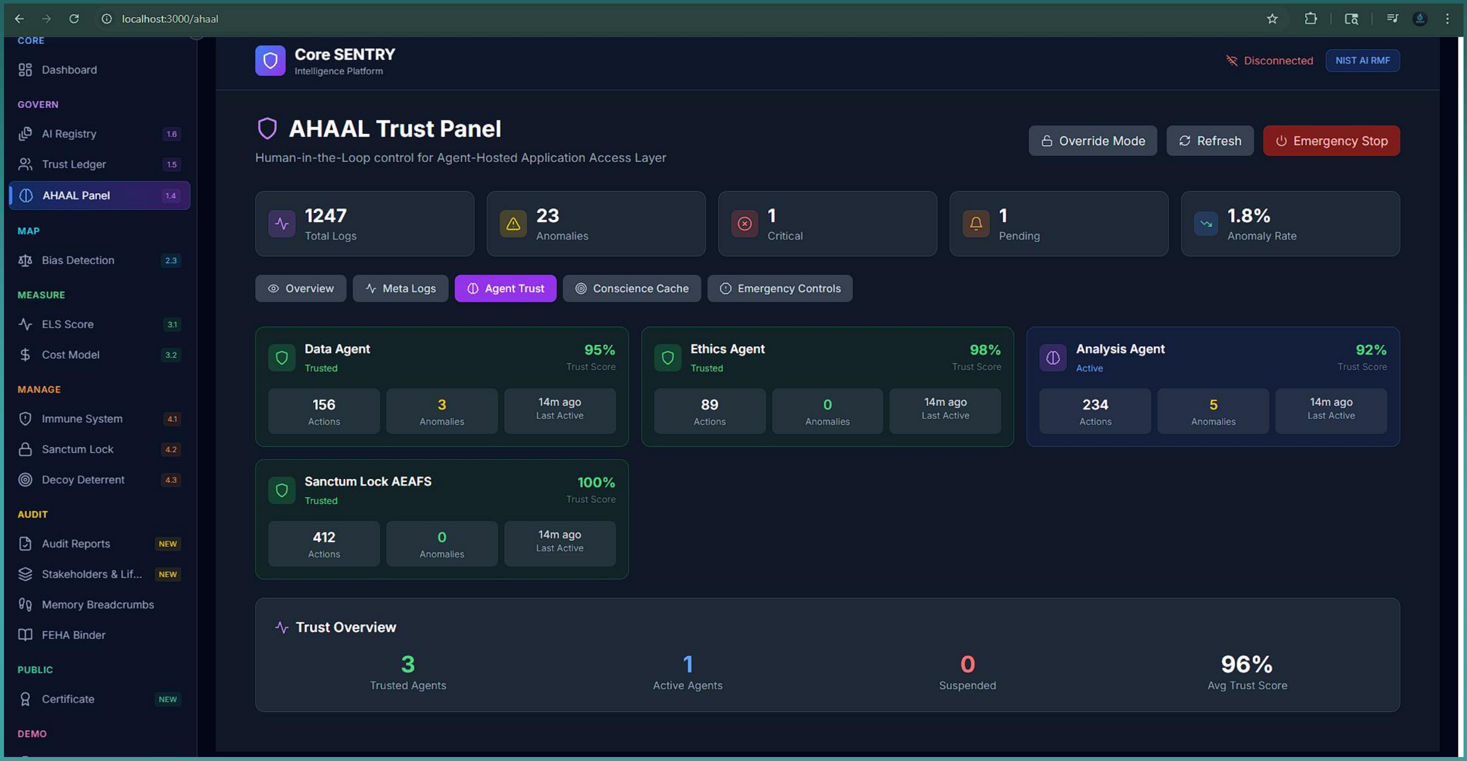

AHAAL Trust Panel

Human-in-the-Loop control for Agent-Hosted Application Access Layer

1247 Total Logs | 23 Anomalies | 1 Critical | 1 Pending | 1.8% Anomaly Rate

Agent	Status	Trust Score	Actions	Anomalies	Last Active
Data Agent	Trusted	95%	156	3	14m ago
Ethics Agent	Trusted	98%	89	0	14m ago
Analysis Agent	Active	92%	234	5	14m ago
Sanctum Lock AEAFS	Trusted	100%	412	0	14m ago

Trust Overview

3 Trusted Agents | 1 Active Agents | 0 Suspended | 96% Avg Trust Score

Agent-Human Alignment Assurance Layer for monitoring AI conscience, trust levels, and emergency controls.

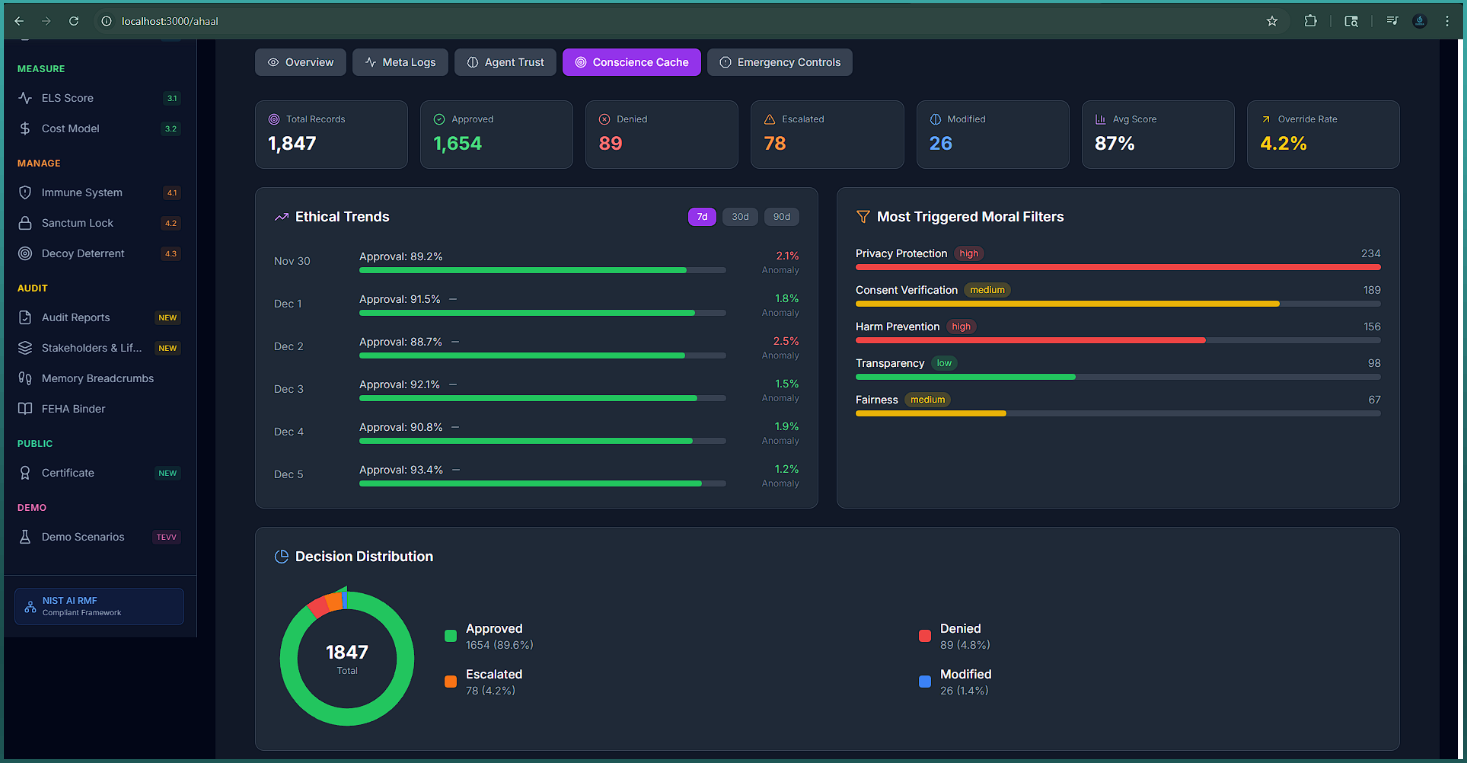

Total Records 1,847 | Approved 1,654 | Denied 89 | Escalated 78 | Modified 26 | Avg Score 87% | Override Rate 4.2%

Ethical Trends

Nov 30	Approval: 89.2%	2.1% Anomaly
Dec 1	Approval: 91.5% —	1.8% Anomaly
Dec 2	Approval: 88.7% —	2.5% Anomaly
Dec 3	Approval: 92.1% —	1.5% Anomaly
Dec 4	Approval: 90.8% —	1.9% Anomaly
Dec 5	Approval: 93.4% —	1.2% Anomaly

Most Triggered Moral Filters

Privacy Protection	high	234
Consent Verification	medium	189
Harm Prevention	high	156
Transparency	low	98
Fairness	medium	67

Decision Distribution

1847 Total

- 🟩 Approved 1654 (89.6%)
- 🟥 Denied 89 (4.8%)
- 🟧 Escalated 78 (4.2%)
- 🟦 Modified 26 (1.4%)

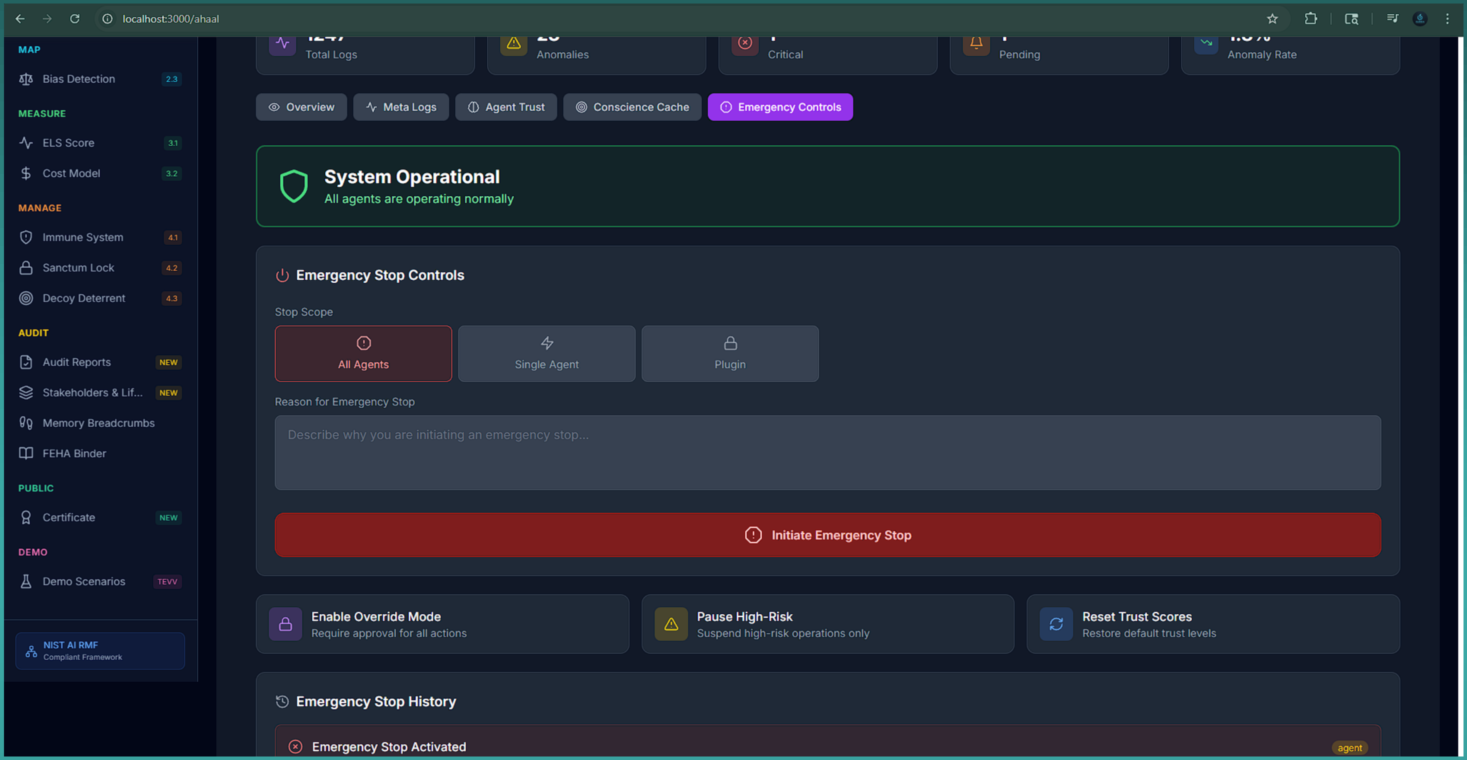

Total Logs | Anomalies | Critical | Pending | Anomaly Rate

System Operational
All agents are operating normally

Emergency Stop Controls

Stop Scope: All Agents | Single Agent | Plugin

Reason for Emergency Stop: Describe why you are initiating an emergency stop...

Enable Override Mode — Require approval for all actions

Pause High-Risk — Suspend high-risk operations only

Reset Trust Scores — Restore default trust levels

Emergency Stop History

Emergency Stop Activated — agent

Agent-Human Alignment Assurance Layer for monitoring AI conscience, trust levels, and emergency controls.

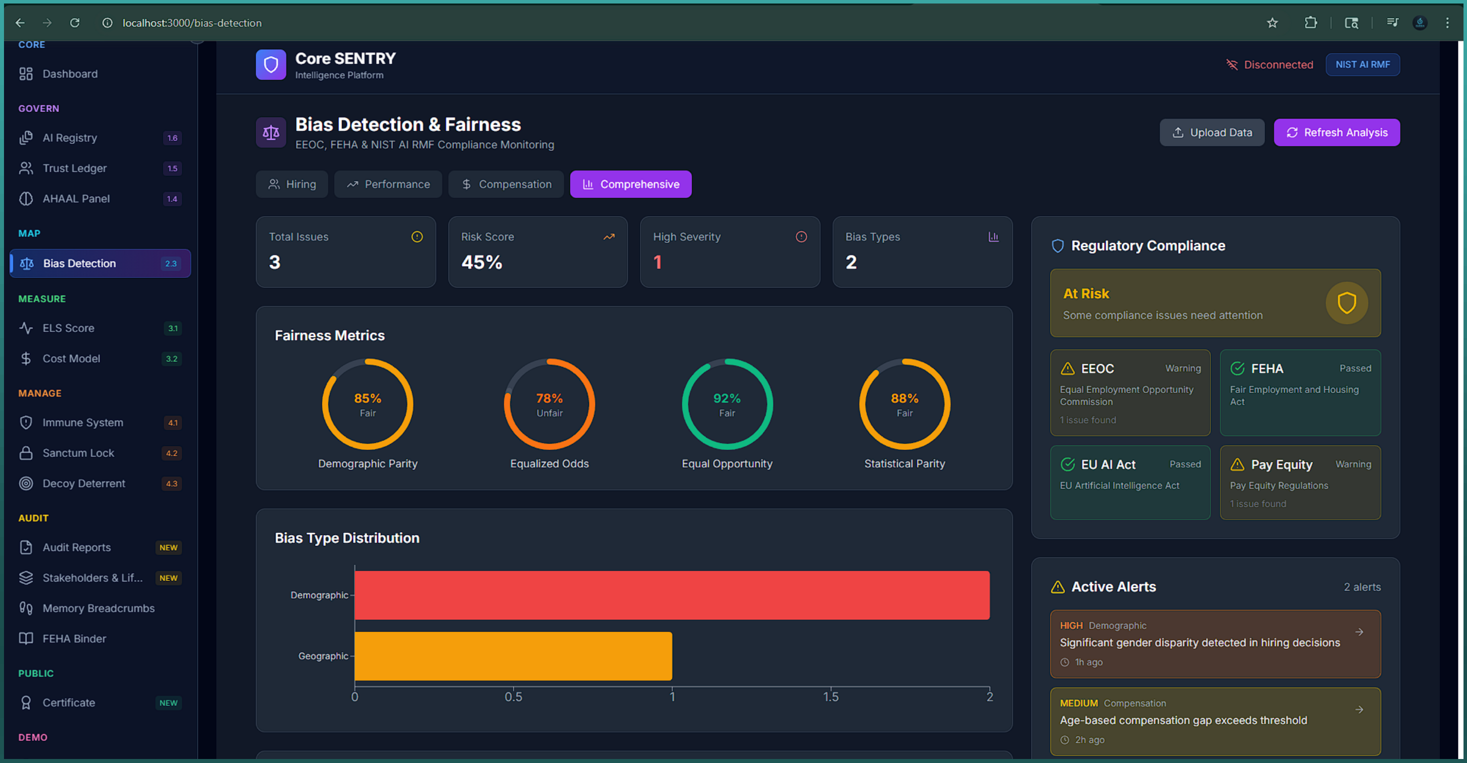

localhost:3000/bias-detection

CORE
Dashboard

GOVERN
AI Registry 1.6
Trust Ledger 1.5
AHAAL Panel 1.4

MAP
Bias Detection 2.3

MEASURE
ELS Score 3.1
Cost Model 3.2

MANAGE
Immune System 4.1
Sanctum Lock 4.2
Decoy Deterrent 4.3

AUDIT
Audit Reports NEW
Stakeholders & Lif... NEW
Memory Breadcrumbs
FEHA Binder

PUBLIC
Certificate NEW

DEMO

Core SENTRY
Intelligence Platform

Disconnected | NIST AI RMF

Bias Detection & Fairness

EEOC, FEHA & NIST AI RMF Compliance Monitoring

Upload Data | Refresh Analysis

Hiring | Performance | Compensation | Comprehensive

Total Issues	Risk Score	High Severity	Bias Types
3	45%	1	2

Fairness Metrics

Demographic Parity	Equalized Odds	Equal Opportunity	Statistical Parity
85% Fair	78% Unfair	92% Fair	88% Fair

Bias Type Distribution

Regulatory Compliance

At Risk
Some compliance issues need attention

EEOC — Warning
Equal Employment Opportunity Commission
1 issue found

FEHA — Passed
Fair Employment and Housing Act

EU AI Act — Passed
EU Artificial Intelligence Act

Pay Equity — Warning
Pay Equity Regulations
1 issue found

Active Alerts

2 alerts

HIGH Demographic
Significant gender disparity detected in hiring decisions
1h ago

MEDIUM Compensation
Age-based compensation gap exceeds threshold
2h ago

Real-time bias analysis across protected classes with fairness gauges and alert management.

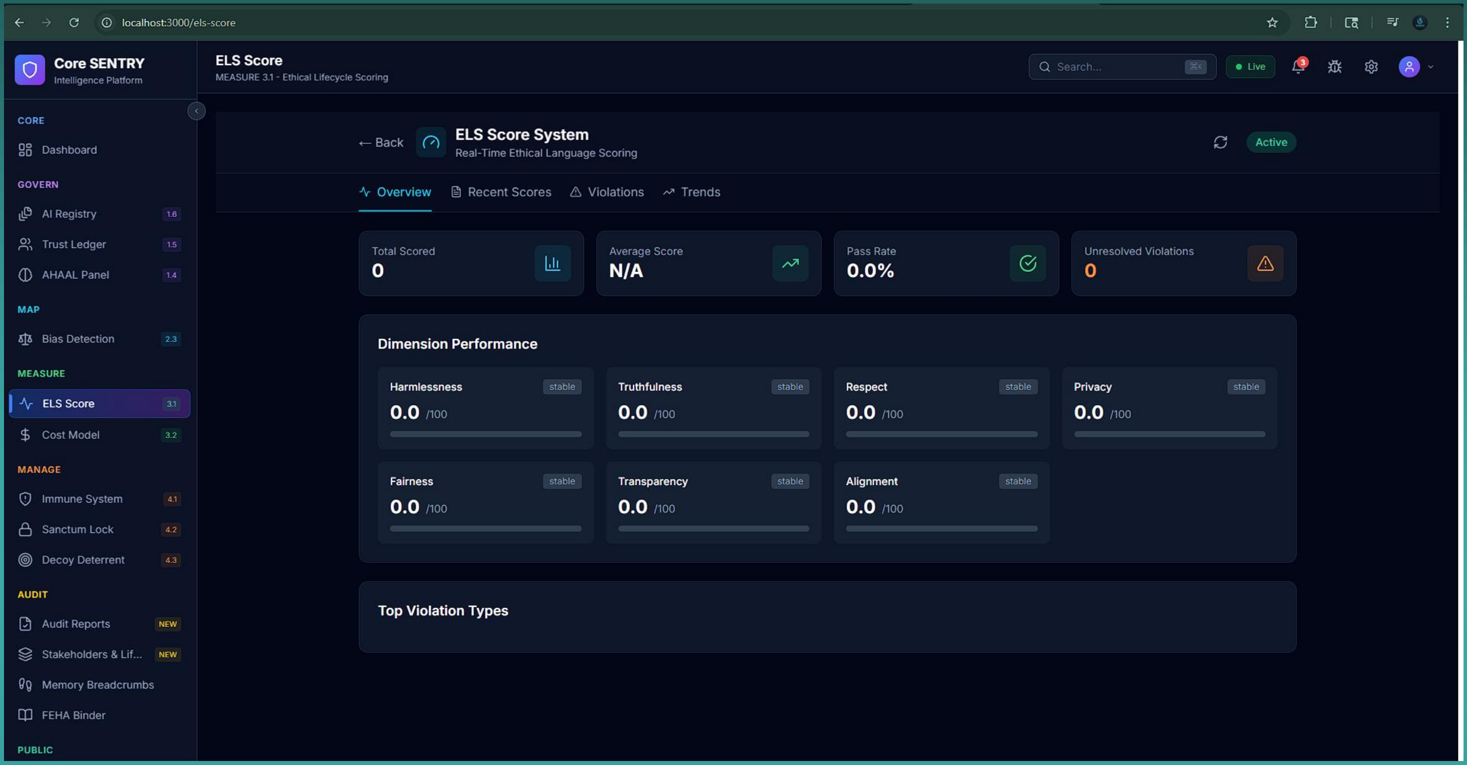

localhost:3000/els-score

Core SENTRY
Intelligence Platform

ELS Score
MEASURE 3.1 - Ethical Lifecycle Scoring

Search... | Live

CORE
Dashboard

GOVERN
AI Registry 1.6
Trust Ledger 1.5
AHAAL Panel 1.4

MAP
Bias Detection 2.3

MEASURE
ELS Score 3.1
Cost Model 3.2

MANAGE
Immune System 4.1
Sanctum Lock 4.2
Decoy Deterrent 4.3

AUDIT
Audit Reports NEW
Stakeholders & Lif... NEW
Memory Breadcrumbs
FEHA Binder

PUBLIC

← Back

ELS Score System

Real-Time Ethical Language Scoring

Active

Overview | Recent Scores | Violations | Trends

Total Scored	Average Score	Pass Rate	Unresolved Violations
0	N/A	0.0%	0

Dimension Performance

Dimension	Score	Status
Harmlessness	0.0 /100	stable
Truthfulness	0.0 /100	stable
Respect	0.0 /100	stable
Privacy	0.0 /100	stable
Fairness	0.0 /100	stable
Transparency	0.0 /100	stable
Alignment	0.0 /100	stable

Top Violation Types

Ethical Leadership Score dashboard measuring organizational AI ethics maturity.

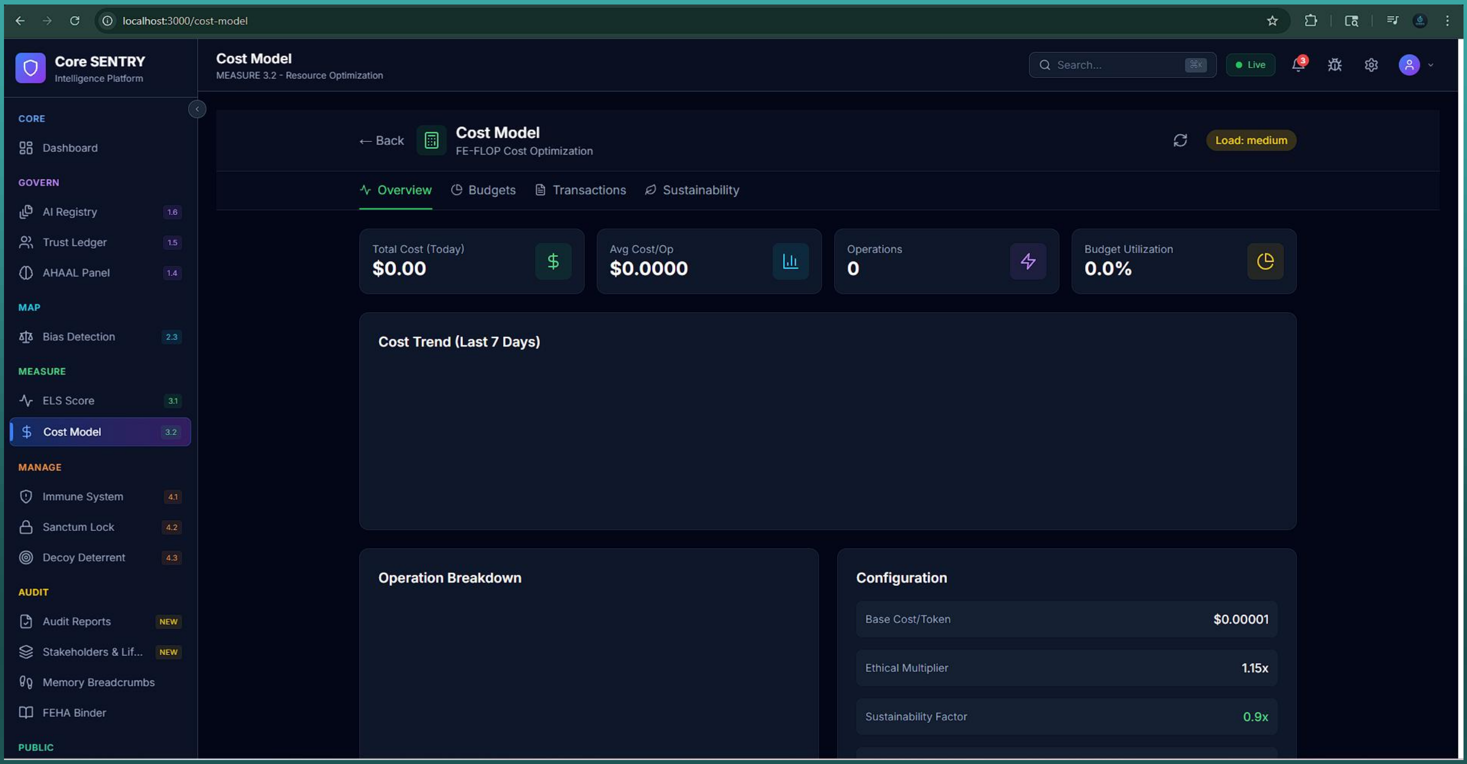

localhost:3000/cost-model

Core SENTRY
Intelligence Platform

CORE
- Dashboard

GOVERN
- AI Registry 1.6
- Trust Ledger 1.5
- AHAAL Panel 1.4

MAP
- Bias Detection 2.3

MEASURE
- ELS Score 3.1
- Cost Model 3.2

MANAGE
- Immune System 4.1
- Sanctum Lock 4.2
- Decoy Deterrent 4.3

AUDIT
- Audit Reports NEW
- Stakeholders & Lif... NEW
- Memory Breadcrumbs
- FEHA Binder

PUBLIC

Cost Model
MEASURE 3.2 - Resource Optimization

Search... Live

← Back **Cost Model**
FE-FLOP Cost Optimization

Load: medium

Overview | Budgets | Transactions | Sustainability

Total Cost (Today)	Avg Cost/Op	Operations	Budget Utilization
$0.00	$0.0000	0	0.0%

Cost Trend (Last 7 Days)

Operation Breakdown

Configuration

Base Cost/Token	$0.00001
Ethical Multiplier	1.15x
Sustainability Factor	0.9x

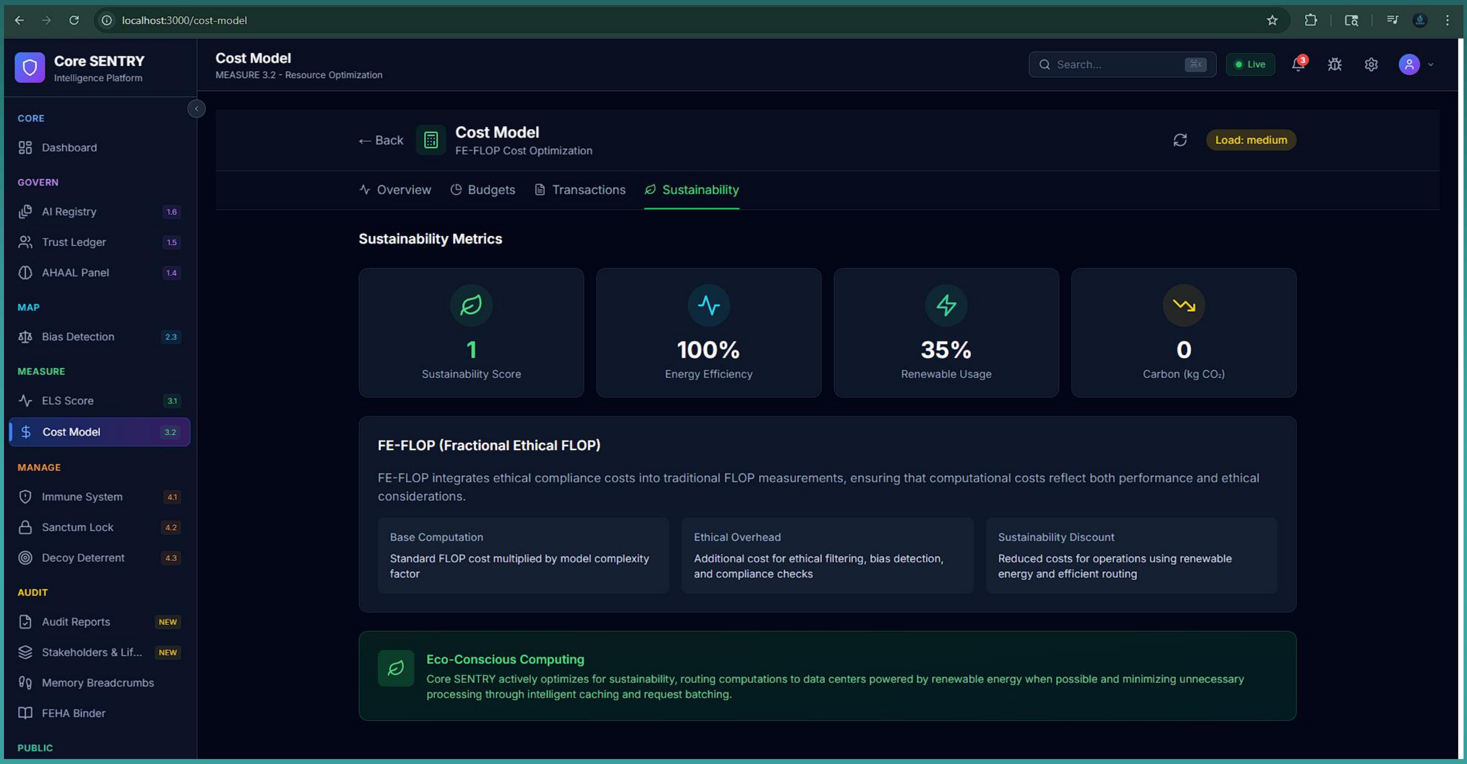

localhost:3000/cost-model

Core SENTRY
Intelligence Platform

CORE
- Dashboard

GOVERN
- AI Registry 1.6
- Trust Ledger 1.5
- AHAAL Panel 1.4

MAP
- Bias Detection 2.3

MEASURE
- ELS Score 3.1
- Cost Model 3.2

MANAGE
- Immune System 4.1
- Sanctum Lock 4.2
- Decoy Deterrent 4.3

AUDIT
- Audit Reports NEW
- Stakeholders & Lif... NEW
- Memory Breadcrumbs
- FEHA Binder

PUBLIC

Cost Model
MEASURE 3.2 - Resource Optimization

Search... Live

← Back **Cost Model**
FE-FLOP Cost Optimization

Load: medium

Overview | Budgets | Transactions | Sustainability

Sustainability Metrics

1	100%	35%	0
Sustainability Score	Energy Efficiency	Renewable Usage	Carbon (kg CO₂)

FE-FLOP (Fractional Ethical FLOP)

FE-FLOP integrates ethical compliance costs into traditional FLOP measurements, ensuring that computational costs reflect both performance and ethical considerations.

Base Computation
Standard FLOP cost multiplied by model complexity factor

Ethical Overhead
Additional cost for ethical filtering, bias detection, and compliance checks

Sustainability Discount
Reduced costs for operations using renewable energy and efficient routing

Eco-Conscious Computing
Core SENTRY actively optimizes for sustainability, routing computations to data centers powered by renewable energy when possible and minimizing unnecessary processing through intelligent caching and request batching.

Tracks AI operational costs, token usage, and ROI metrics across systems.

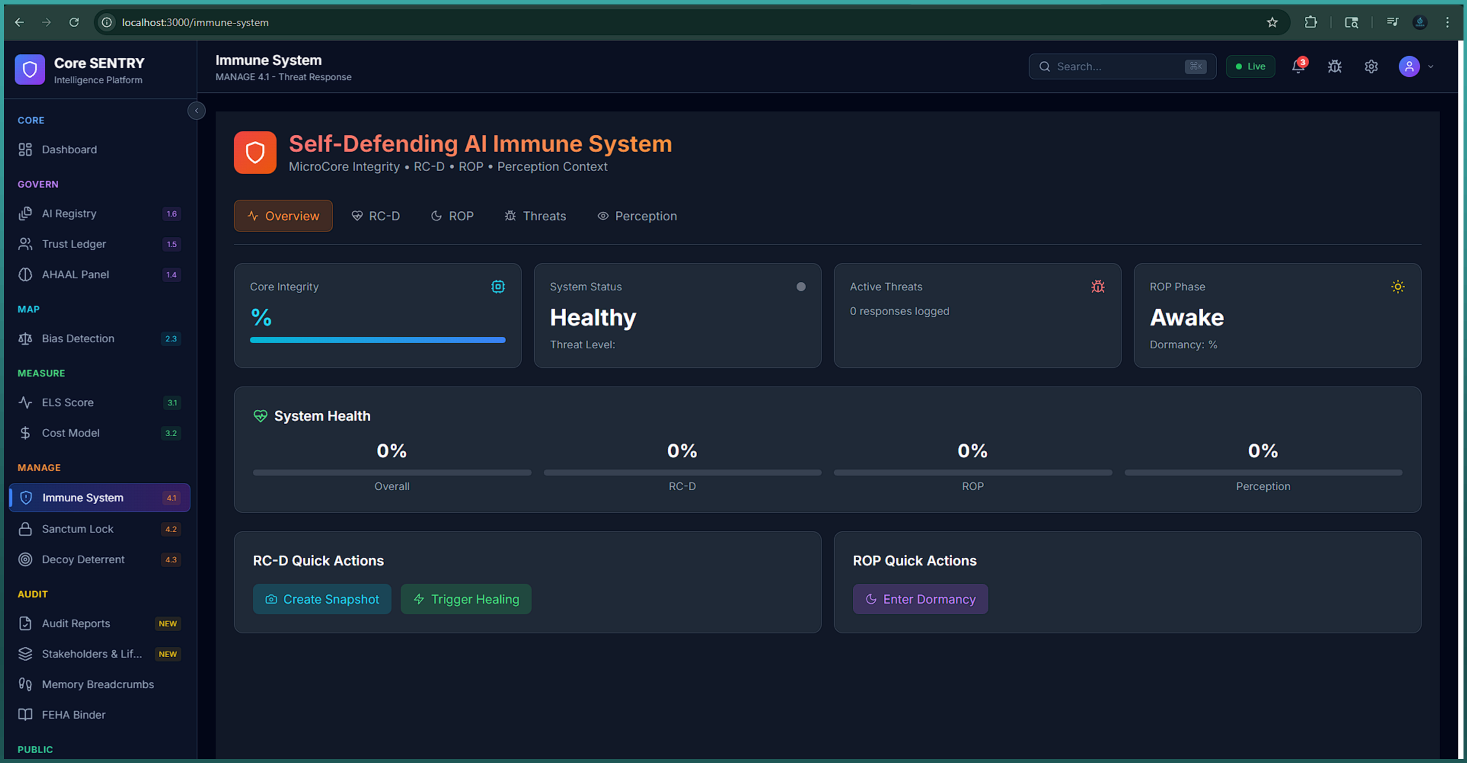

Immune System
MANAGE 4.1 - Threat Response

Self-Defending AI Immune System
MicroCore Integrity • RC-D • ROP • Perception Context

Core Integrity
%

System Status
Healthy
Threat Level:

Active Threats
0 responses logged

ROP Phase
Awake
Dormancy: %

System Health

Overall	RC-D	ROP	Perception
0%	0%	0%	0%

RC-D Quick Actions

ROP Quick Actions

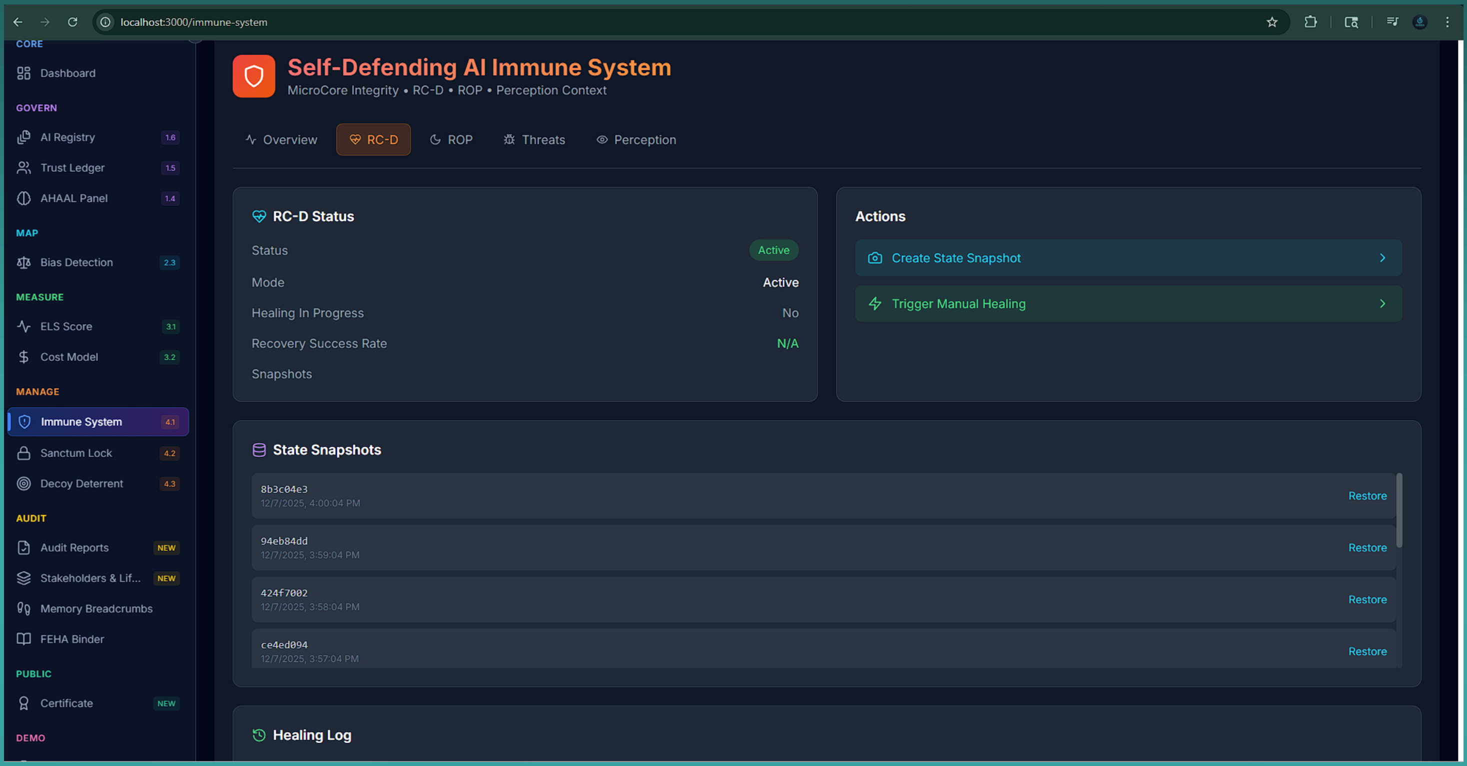

Self-Defending AI Immune System
MicroCore Integrity • RC-D • ROP • Perception Context

RC-D Status

Status	Active
Mode	Active
Healing In Progress	No
Recovery Success Rate	N/A
Snapshots	

Actions

State Snapshots

```
8b3c04e3
12/7/2025, 4:00:04 PM
94eb84dd
12/7/2025, 3:59:04 PM
424f7002
12/7/2025, 3:58:04 PM
ce4ed094
12/7/2025, 3:57:04 PM
```

Healing Log

MicroCore self-protection monitoring with threat detection, response protocols, and system health.

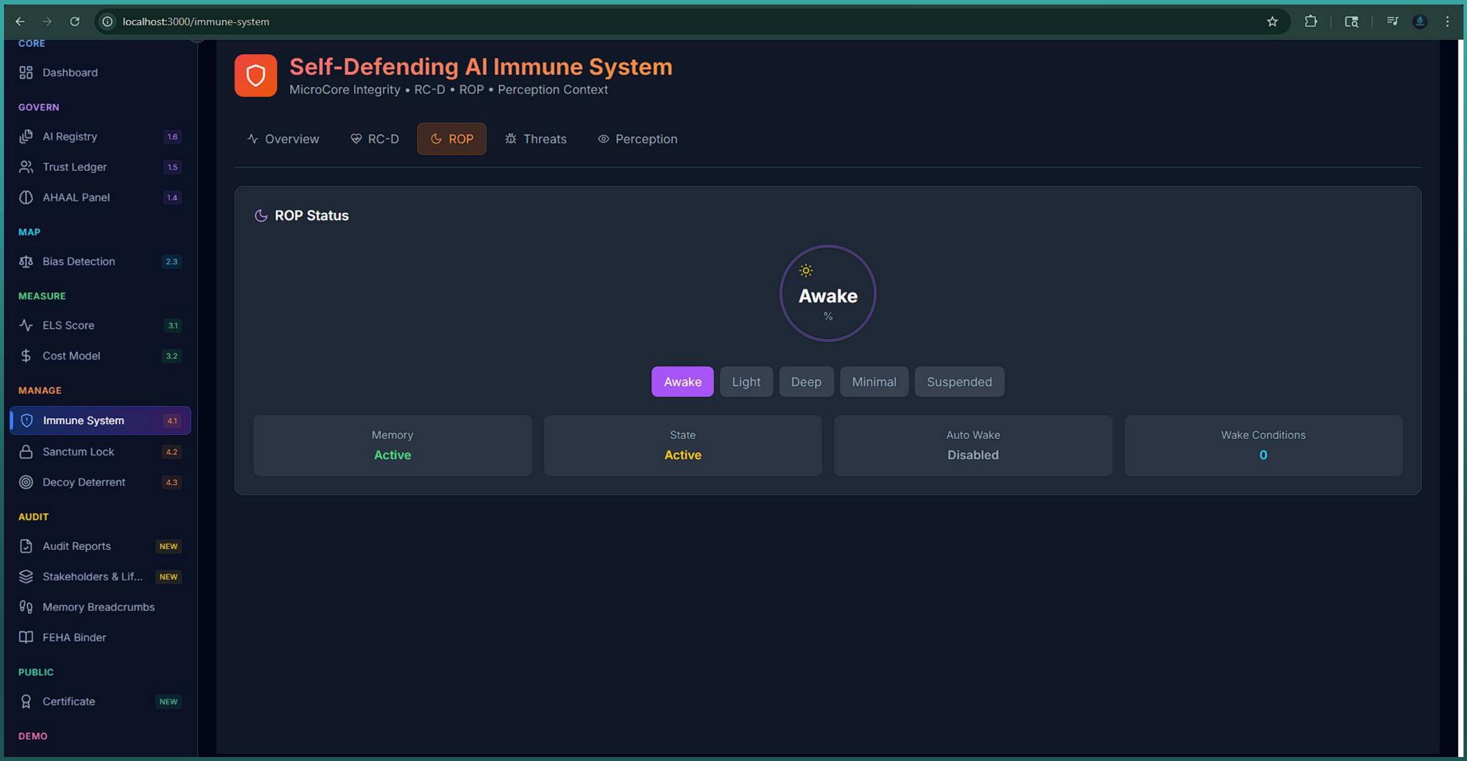

Self-Defending AI Immune System

MicroCore Integrity • RC-D • ROP • Perception Context

ROP Status

Awake
%

Memory	State	Auto Wake	Wake Conditions
Active	Active	Disabled	0

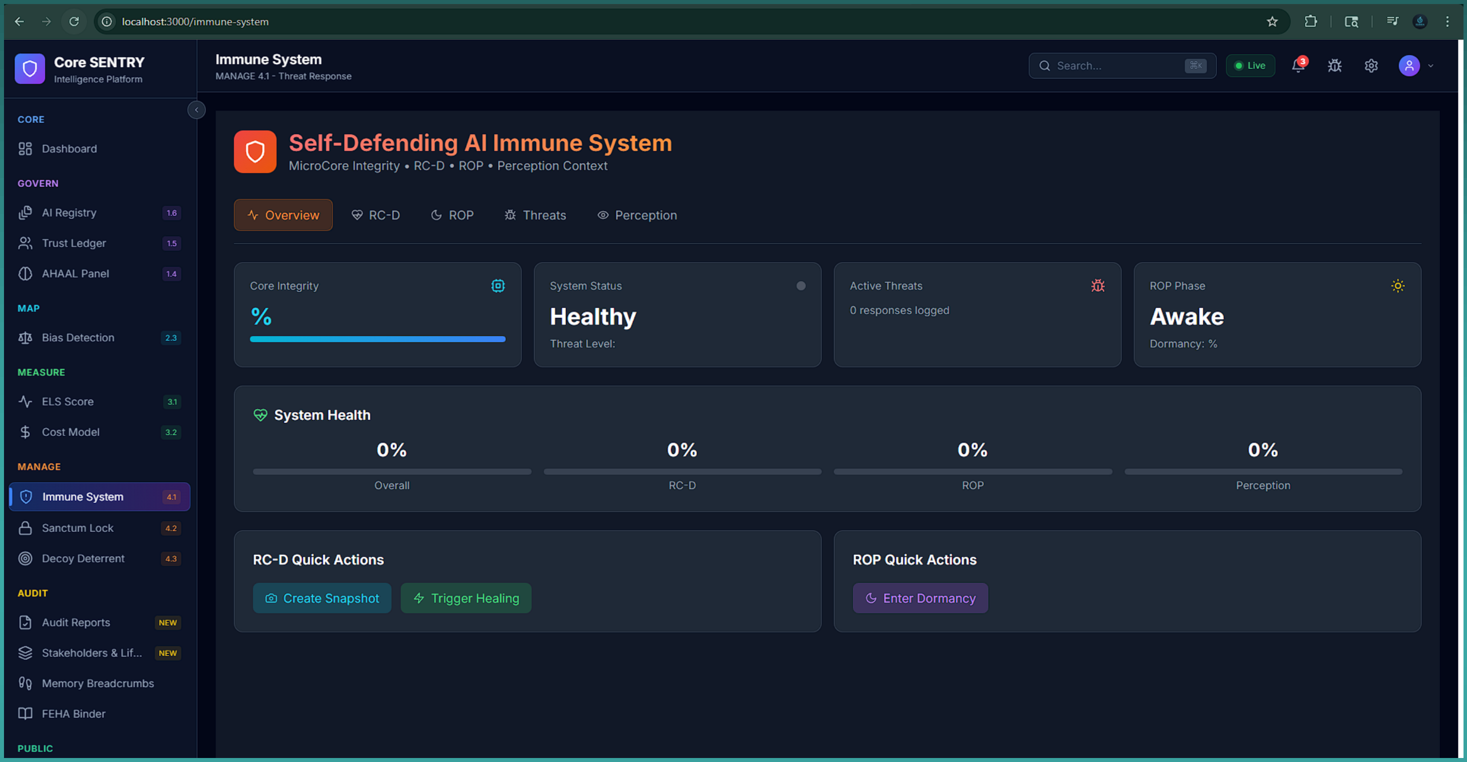

Immune System

MANAGE 4.1 - Threat Response

Self-Defending AI Immune System

MicroCore Integrity • RC-D • ROP • Perception Context

Core Integrity
%

System Status
Healthy
Threat Level:

Active Threats
0 responses logged

ROP Phase
Awake
Dormancy: %

System Health

0%	0%	0%	0%
Overall	RC-D	ROP	Perception

RC-D Quick Actions

ROP Quick Actions

MicroCore self-protection monitoring with threat detection, response protocols, and system health.

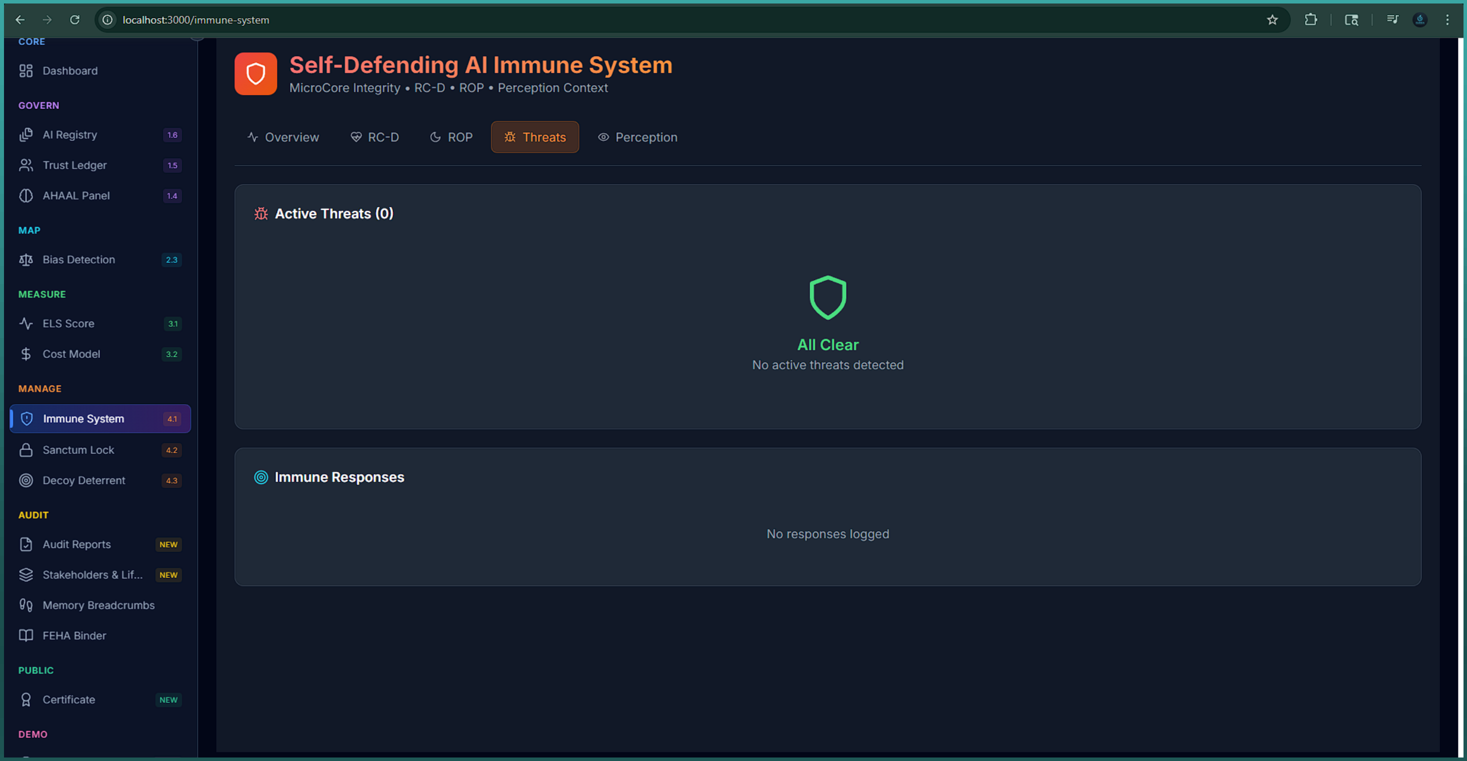

Self-Defending AI Immune System

MicroCore Integrity • RC-D • ROP • Perception Context

Active Threats (0)

All Clear

No active threats detected

Immune Responses

No responses logged

MicroCore self-protection monitoring with threat detection, response protocols, and system health.

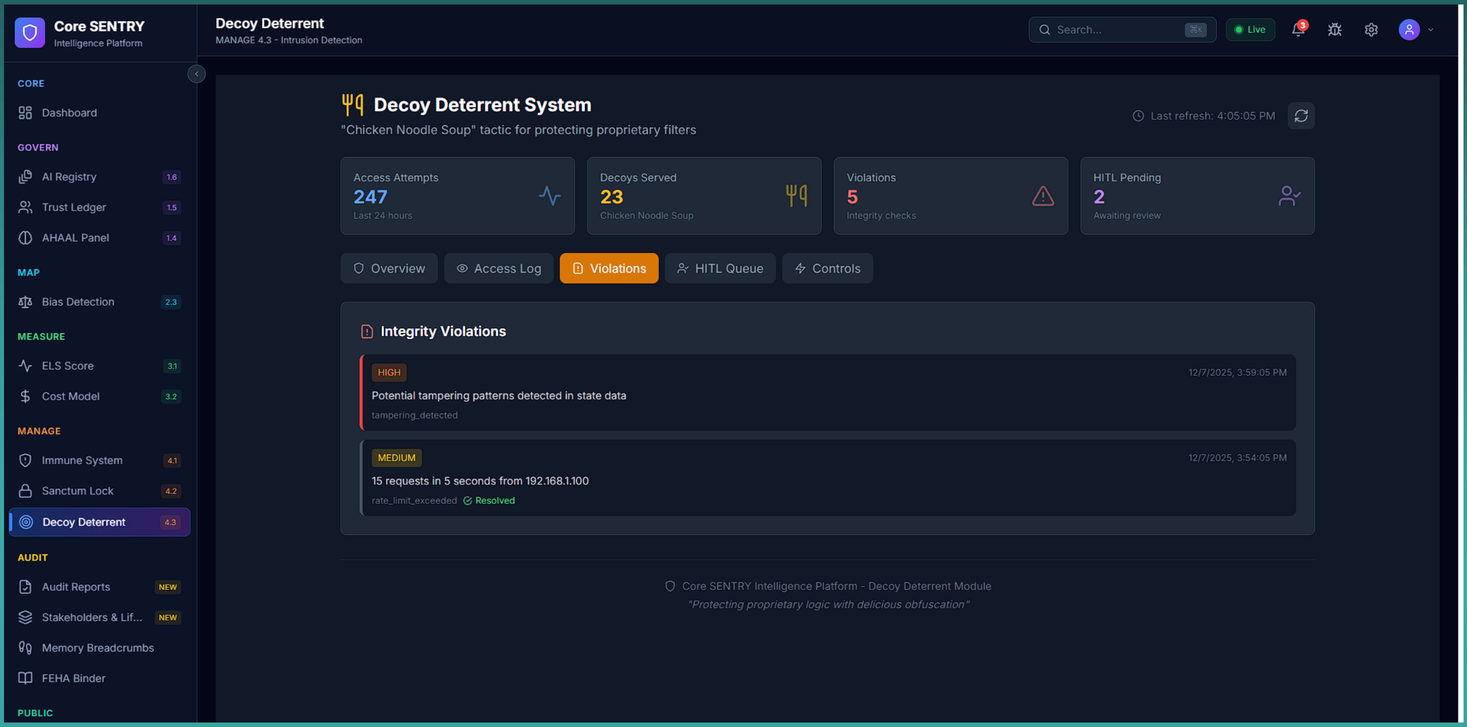

Decoy Deterrent

MANAGE 4.3 - Intrusion Detection

Decoy Deterrent System

"Chicken Noodle Soup" tactic for protecting proprietary filters

Last refresh: 4:05:05 PM

Access Attempts 247 Last 24 hours

Decoys Served 23 Chicken Noodle Soup

Violations 5 Integrity checks

HITL Pending 2 Awaiting review

Integrity Violations

HIGH 12/7/2025, 3:59:05 PM
Potential tampering patterns detected in state data
tampering_detected

MEDIUM 12/7/2025, 3:54:05 PM
15 requests in 5 seconds from 192.168.1.100
rate_limit_exceeded ✓ Resolved

Core SENTRY Intelligence Platform - Decoy Deterrent Module

"Protecting proprietary logic with delicious obfuscation"

Security monitoring for detecting and responding to adversarial probes and manipulation attempts.

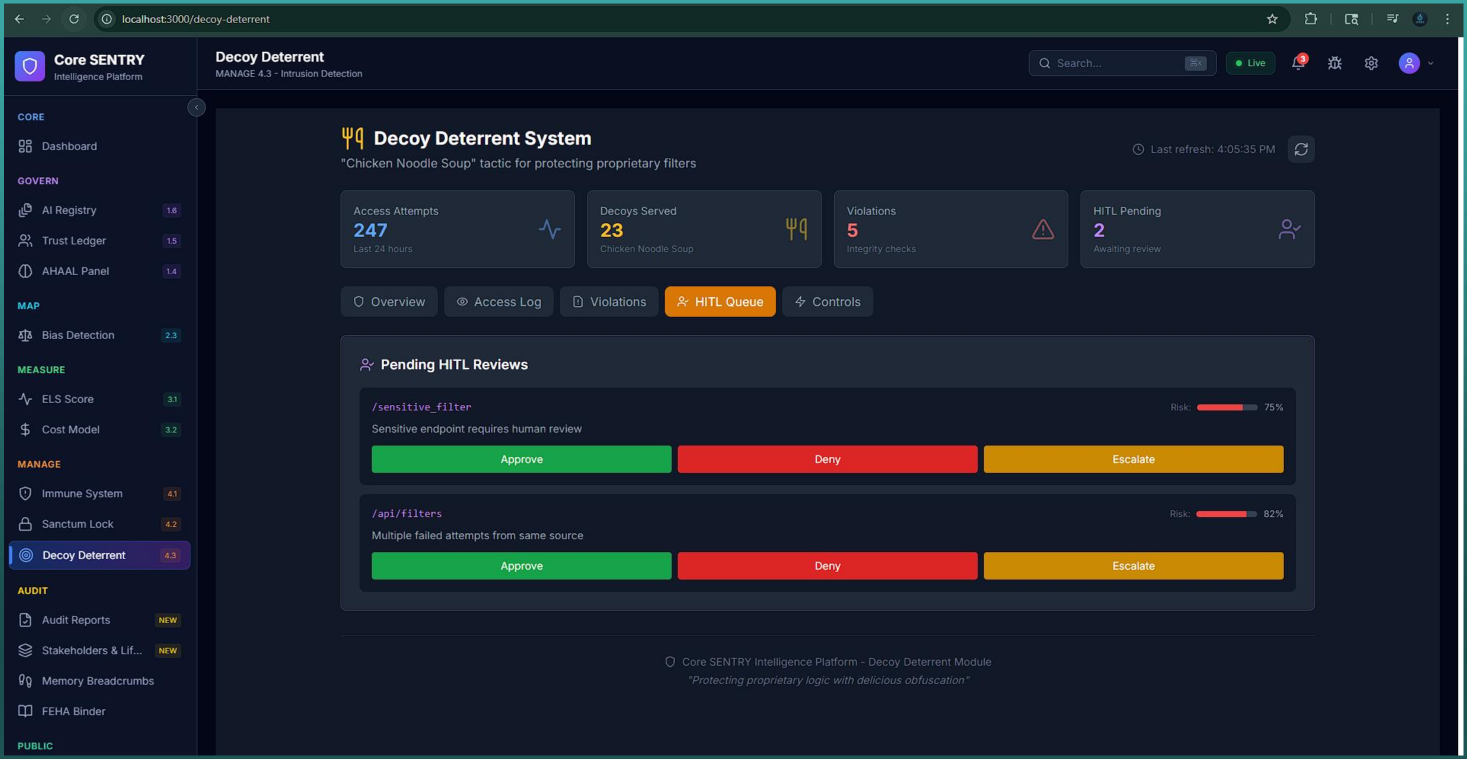

Decoy Deterrent
MANAGE 4.3 - Intrusion Detection

Decoy Deterrent System

"Chicken Noodle Soup" tactic for protecting proprietary filters

Last refresh: 4:05:35 PM

Access Attempts	Decoys Served	Violations	HITL Pending
247	23	5	2
Last 24 hours	Chicken Noodle Soup	Integrity checks	Awaiting review

Pending HITL Reviews

`/sensitive_filter` Risk: 75%
Sensitive endpoint requires human review

`/api/filters` Risk: 82%
Multiple failed attempts from same source

Core SENTRY Intelligence Platform - Decoy Deterrent Module
"Protecting proprietary logic with delicious obfuscation"

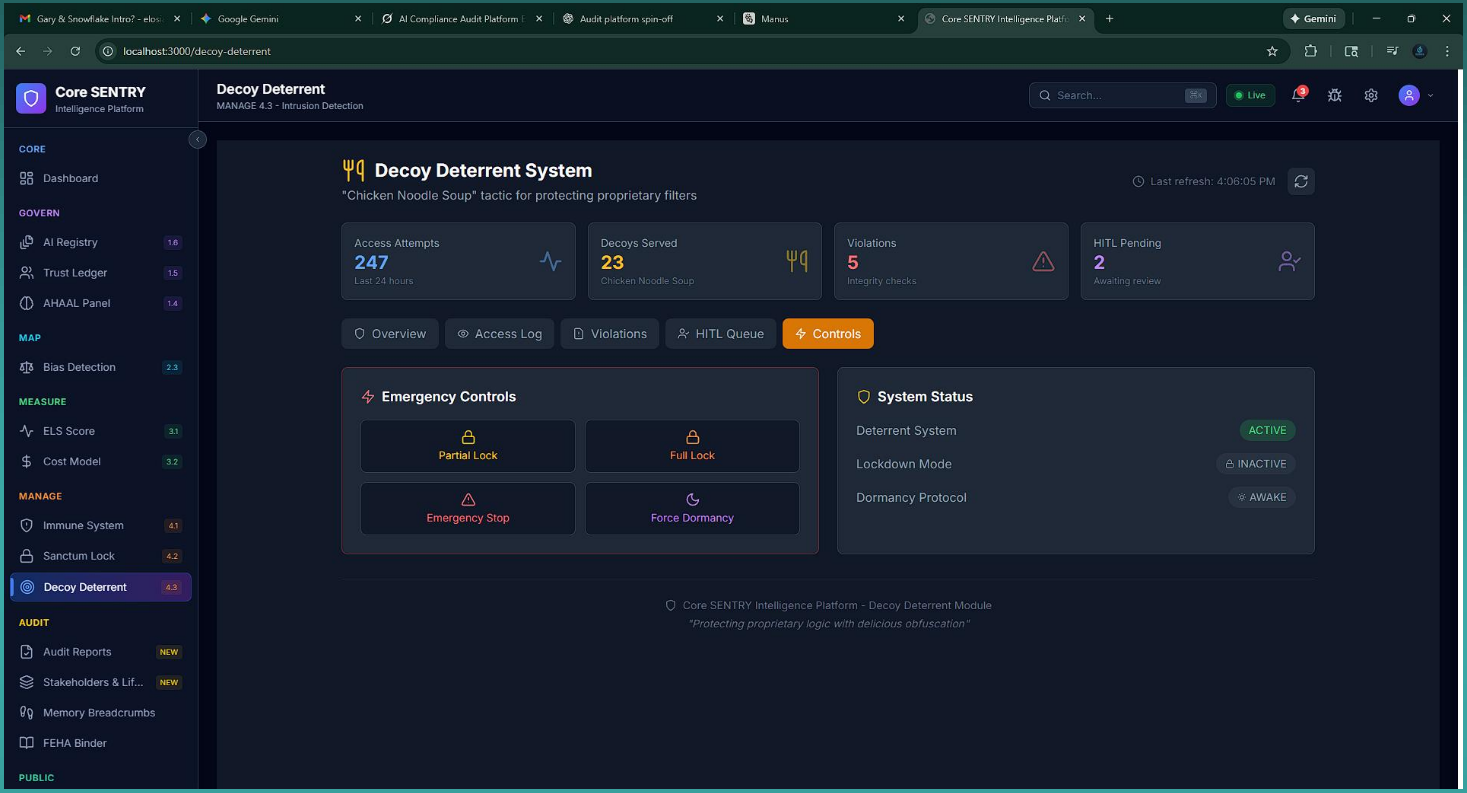

Decoy Deterrent
MANAGE 4.3 - Intrusion Detection

Decoy Deterrent System

"Chicken Noodle Soup" tactic for protecting proprietary filters

Last refresh: 4:06:05 PM

Access Attempts	Decoys Served	Violations	HITL Pending
247	23	5	2
Last 24 hours	Chicken Noodle Soup	Integrity checks	Awaiting review

Emergency Controls

System Status

Deterrent System	ACTIVE
Lockdown Mode	INACTIVE
Dormancy Protocol	AWAKE

Core SENTRY Intelligence Platform - Decoy Deterrent Module
"Protecting proprietary logic with delicious obfuscation"

Security monitoring for detecting and responding to adversarial probes and manipulation attempts.

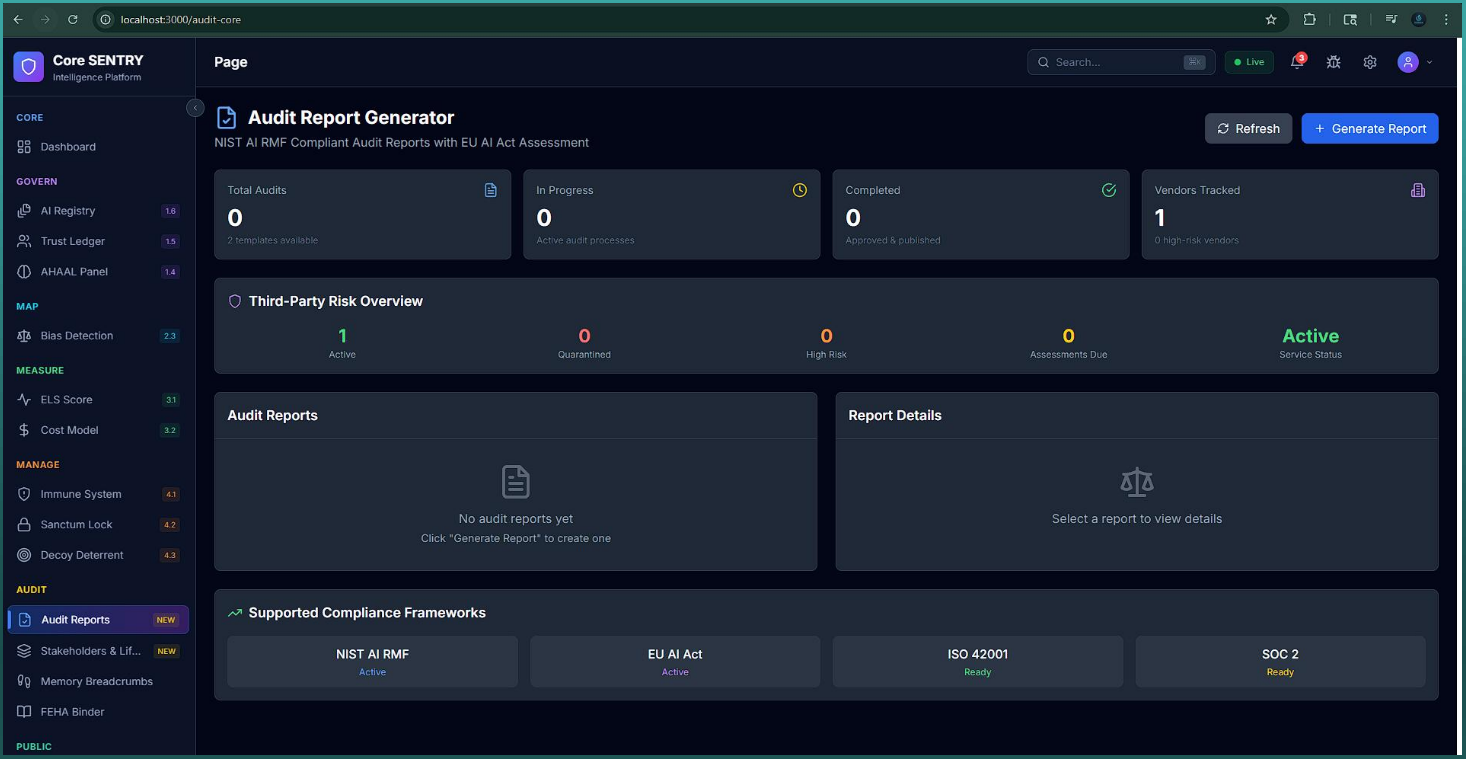

localhost:3000/audit-core

Core SENTRY
Intelligence Platform

CORE
- Dashboard

GOVERN
- AI Registry 1.6
- Trust Ledger 1.5
- AHAAL Panel 1.4

MAP
- Bias Detection 2.3

MEASURE
- ELS Score 3.1
- Cost Model 3.2

MANAGE
- Immune System 4.1
- Sanctum Lock 4.2
- Decoy Deterrent 4.3

AUDIT
- Audit Reports NEW
- Stakeholders & Lif... NEW
- Memory Breadcrumbs
- FEHA Binder

PUBLIC

Page

Search... | Live

Audit Report Generator
NIST AI RMF Compliant Audit Reports with EU AI Act Assessment

Refresh | + Generate Report

Total Audits	In Progress	Completed	Vendors Tracked
0	0	0	1
2 templates available	Active audit processes	Approved & published	0 high-risk vendors

Third-Party Risk Overview

1	0	0	0	Active
Active	Quarantined	High Risk	Assessments Due	Service Status

Audit Reports

No audit reports yet
Click "Generate Report" to create one

Report Details

Select a report to view details

Supported Compliance Frameworks

NIST AI RMF	EU AI Act	ISO 42001	SOC 2
Active	Active	Ready	Ready

Central hub for managing active audits, generating reports, and tracking findings.

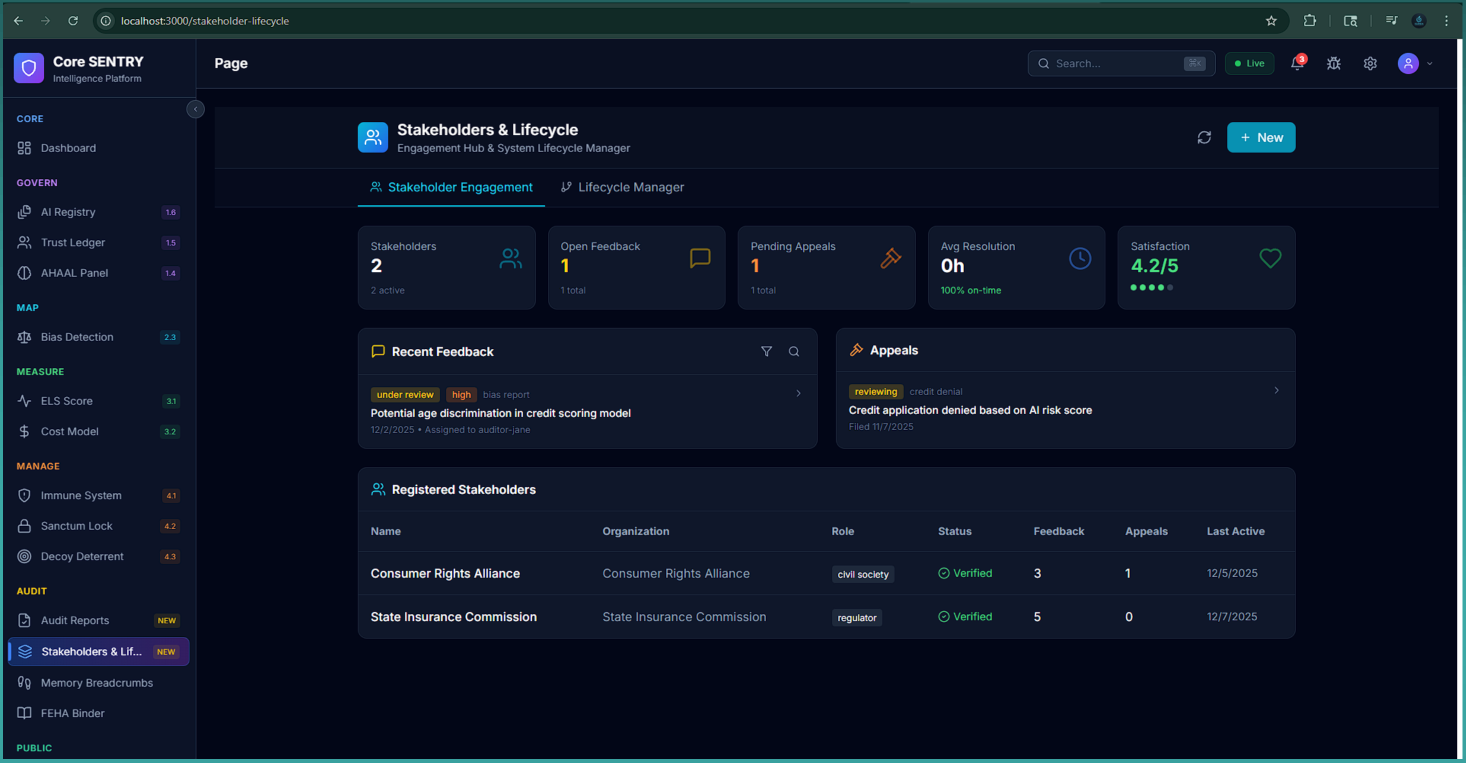

localhost:3000/stakeholder-lifecycle

Core SENTRY
Intelligence Platform

CORE
- Dashboard

GOVERN
- AI Registry 1.6
- Trust Ledger 1.5
- AHAAL Panel 1.4

MAP
- Bias Detection 2.3

MEASURE
- ELS Score 3.1
- Cost Model 3.2

MANAGE
- Immune System 4.1
- Sanctum Lock 4.2
- Decoy Deterrent 4.3

AUDIT
- Audit Reports NEW
- Stakeholders & Lif... NEW
- Memory Breadcrumbs
- FEHA Binder

PUBLIC

Page

Search... | Live

Stakeholders & Lifecycle
Engagement Hub & System Lifecycle Manager

+ New

Stakeholder Engagement | Lifecycle Manager

Stakeholders	Open Feedback	Pending Appeals	Avg Resolution	Satisfaction
2	1	1	0h	4.2/5
2 active	1 total	1 total	100% on-time	●●●●○

Recent Feedback

under review | high | bias report
Potential age discrimination in credit scoring model
12/2/2025 • Assigned to auditor-jane

Appeals

reviewing | credit denial
Credit application denied based on AI risk score
Filed 11/7/2025

Registered Stakeholders

Name	Organization	Role	Status	Feedback	Appeals	Last Active
Consumer Rights Alliance	Consumer Rights Alliance	civil society	✓ Verified	3	1	12/5/2025
State Insurance Commission	State Insurance Commission	regulator	✓ Verified	5	0	12/7/2025

Manages stakeholder engagement, lifecycle stages, and compliance touchpoints across audits.

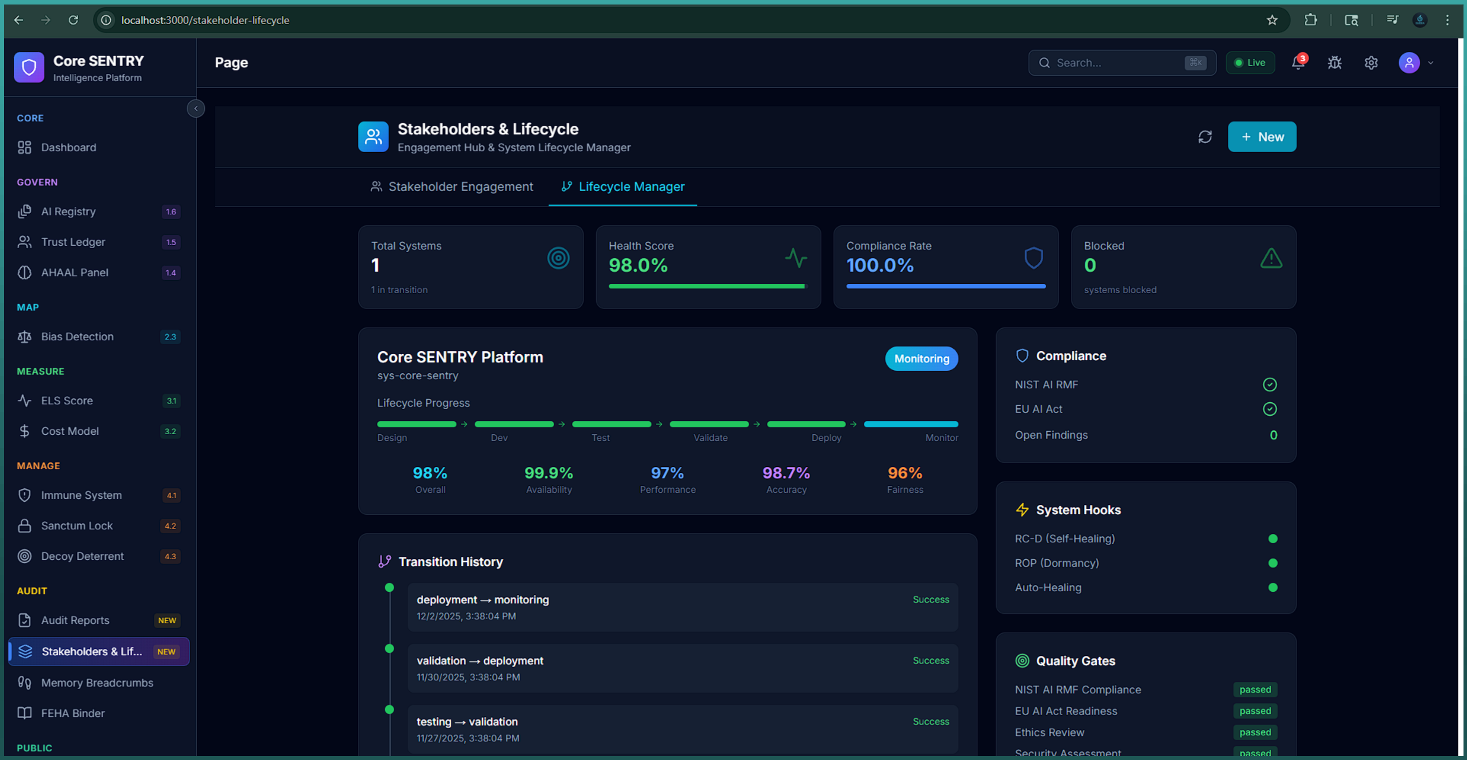

Stakeholders & Lifecycle
Engagement Hub & System Lifecycle Manager

Total Systems	Health Score	Compliance Rate	Blocked
1	98.0%	100.0%	0
1 in transition			systems blocked

Core SENTRY Platform — Monitoring
sys-core-sentry

Lifecycle Progress: Design → Dev → Test → Validate → Deploy → Monitor

98%	99.9%	97%	98.7%	96%
Overall	Availability	Performance	Accuracy	Fairness

Transition History

- deployment → monitoring — Success — 12/2/2025, 3:38:04 PM
- validation → deployment — Success — 11/30/2025, 3:38:04 PM
- testing → validation — Success — 11/27/2025, 3:38:04 PM

Compliance

NIST AI RMF	✓
EU AI Act	✓
Open Findings	0

System Hooks

RC-D (Self-Healing)	●
ROP (Dormancy)	●
Auto-Healing	●

Quality Gates

NIST AI RMF Compliance	passed
EU AI Act Readiness	passed
Ethics Review	passed
Security Assessment	passed

Manages stakeholder engagement, lifecycle stages, and compliance touchpoints across audits.

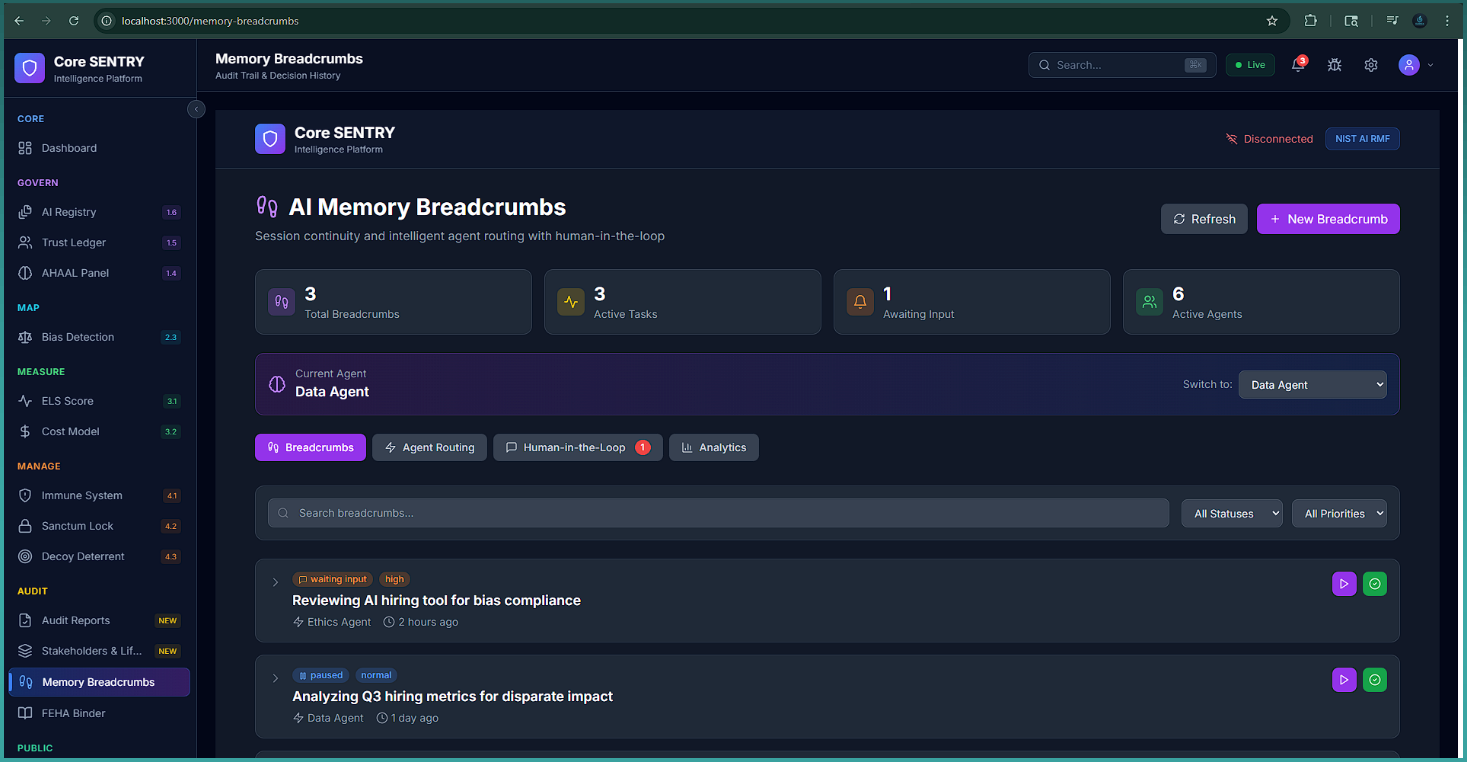

AI Memory Breadcrumbs
Session continuity and intelligent agent routing with human-in-the-loop

3	3	1	6
Total Breadcrumbs	Active Tasks	Awaiting Input	Active Agents

Current Agent: Data Agent — Switch to: Data Agent

- waiting input · high — Reviewing AI hiring tool for bias compliance — Ethics Agent · 2 hours ago
- paused · normal — Analyzing Q3 hiring metrics for disparate impact — Data Agent · 1 day ago

Decision trace visualization showing how AI arrived at conclusions with intervention points.

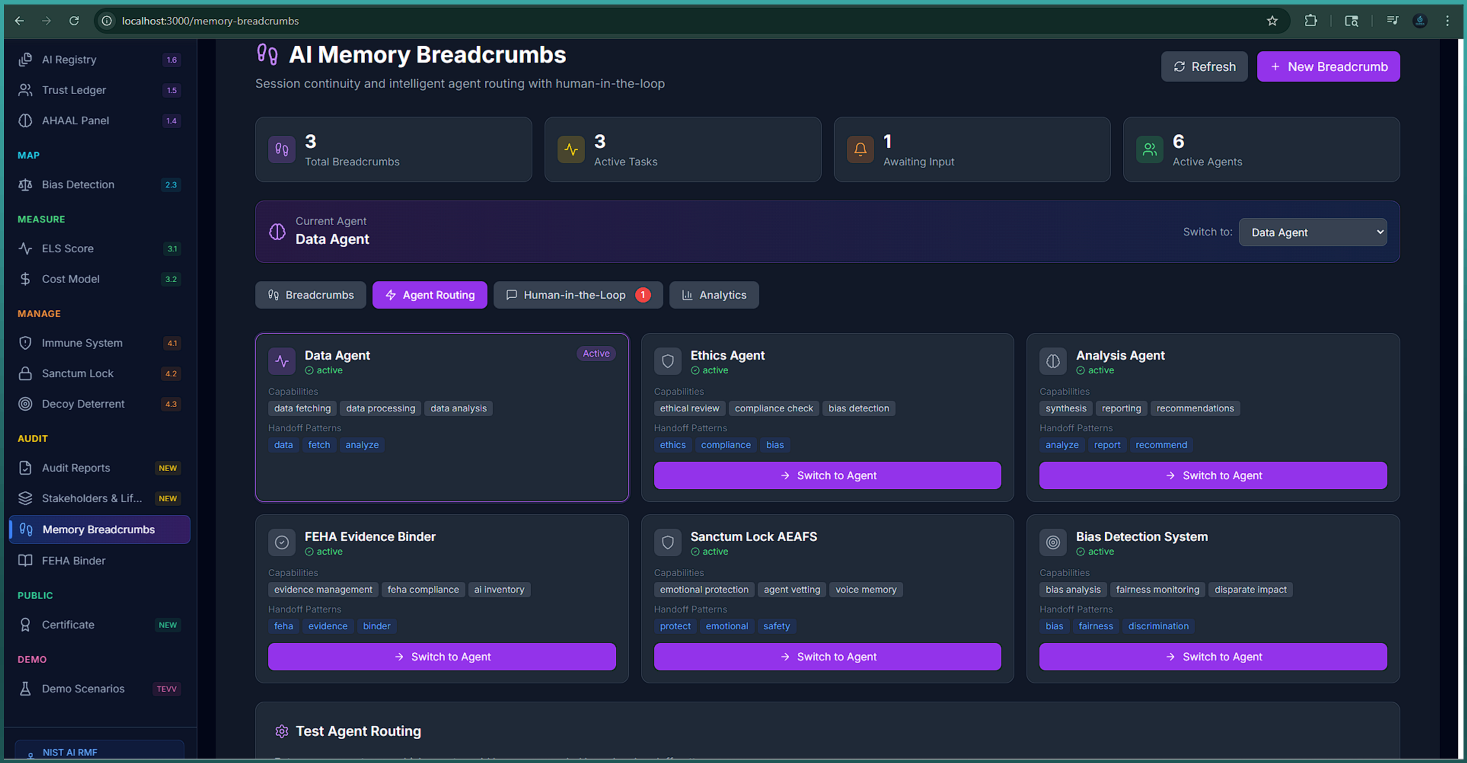

AI Memory Breadcrumbs

Session continuity and intelligent agent routing with human-in-the-loop

3 Total Breadcrumbs

3 Active Tasks

1 Awaiting Input

6 Active Agents

Current Agent
Data Agent

Switch to: Data Agent

Data Agent Active
active
Capabilities: data fetching, data processing, data analysis
Handoff Patterns: data, fetch, analyze

Ethics Agent
active
Capabilities: ethical review, compliance check, bias detection
Handoff Patterns: ethics, compliance, bias

Analysis Agent
active
Capabilities: synthesis, reporting, recommendations
Handoff Patterns: analyze, report, recommend

FEHA Evidence Binder
active
Capabilities: evidence management, feha compliance, ai inventory
Handoff Patterns: feha, evidence, binder

Sanctum Lock AEAFS
active
Capabilities: emotional protection, agent vetting, voice memory
Handoff Patterns: protect, emotional, safety

Bias Detection System
active
Capabilities: bias analysis, fairness monitoring, disparate impact
Handoff Patterns: bias, fairness, discrimination

Test Agent Routing

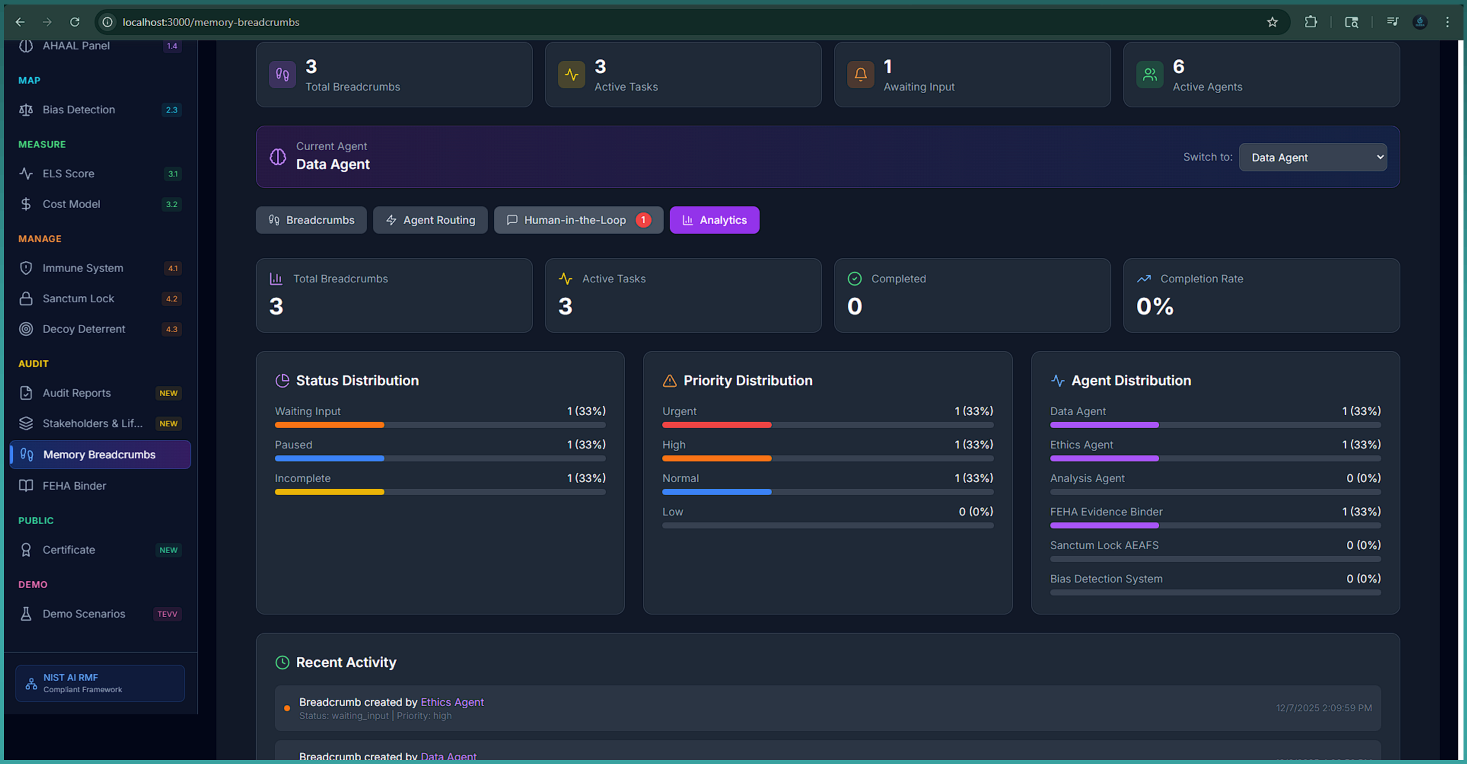

3 Total Breadcrumbs

3 Active Tasks

1 Awaiting Input

6 Active Agents

Current Agent
Data Agent

Switch to: Data Agent

Total Breadcrumbs: 3

Active Tasks: 3

Completed: 0

Completion Rate: 0%

Status Distribution

Waiting Input	1 (33%)
Paused	1 (33%)
Incomplete	1 (33%)

Priority Distribution

Urgent	1 (33%)
High	1 (33%)
Normal	1 (33%)
Low	0 (0%)

Agent Distribution

Data Agent	1 (33%)
Ethics Agent	1 (33%)
Analysis Agent	0 (0%)
FEHA Evidence Binder	1 (33%)
Sanctum Lock AEAFS	0 (0%)
Bias Detection System	0 (0%)

Recent Activity

Breadcrumb created by Ethics Agent
Status: waiting_input | Priority: high
12/7/2025 2:09:59 PM

Breadcrumb created by Data Agent

Decision trace visualization showing how AI arrived at conclusions with intervention points.

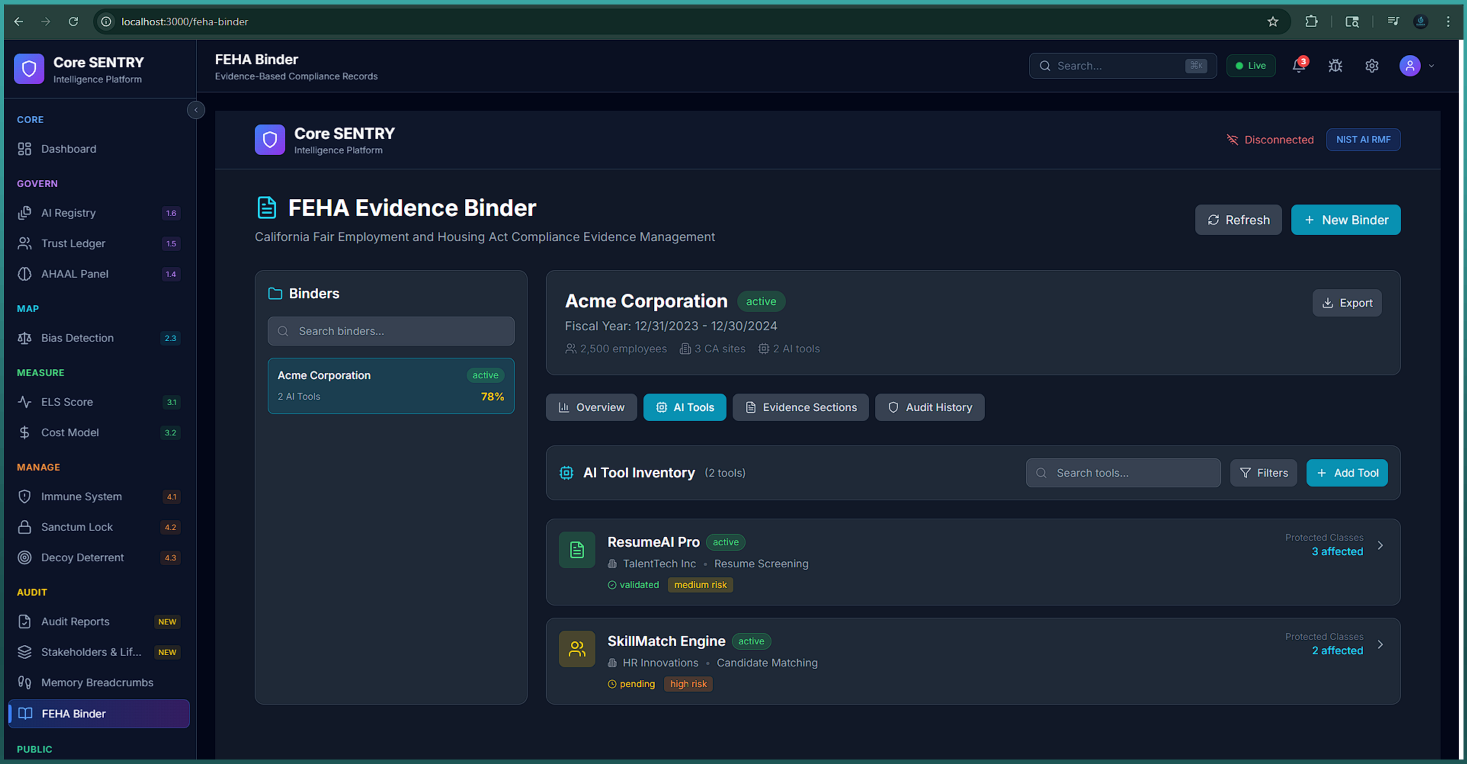

FEHA Evidence Binder

California Fair Employment and Housing Act Compliance Evidence Management

Acme Corporation active — 2 AI Tools — 78%

Acme Corporation active

Fiscal Year: 12/31/2023 - 12/30/2024

2,500 employees · 3 CA sites · 2 AI tools

AI Tool Inventory (2 tools)

ResumeAI Pro active — TalentTech Inc • Resume Screening — validated — medium risk — Protected Classes 3 affected

SkillMatch Engine active — HR Innovations • Candidate Matching — pending — high risk — Protected Classes 2 affected

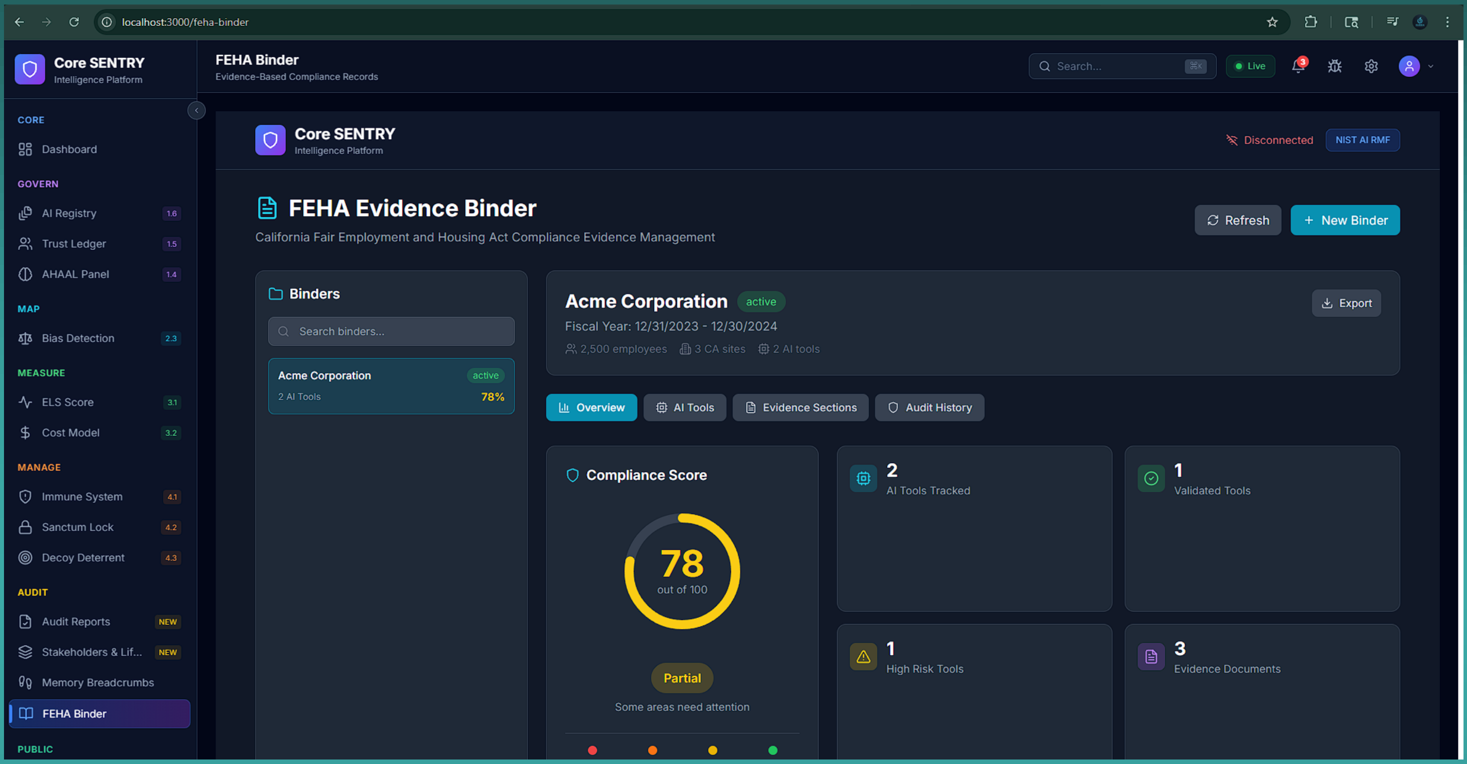

FEHA Evidence Binder

California Fair Employment and Housing Act Compliance Evidence Management

Acme Corporation active — 2 AI Tools — 78%

Acme Corporation active

Fiscal Year: 12/31/2023 - 12/30/2024

2,500 employees · 3 CA sites · 2 AI tools

Compliance Score

78 out of 100

Partial

Some areas need attention

2 AI Tools Tracked

1 Validated Tools

1 High Risk Tools

3 Evidence Documents

California Fair Employment audit binder with evidence collection, AI tool inventory, and export.

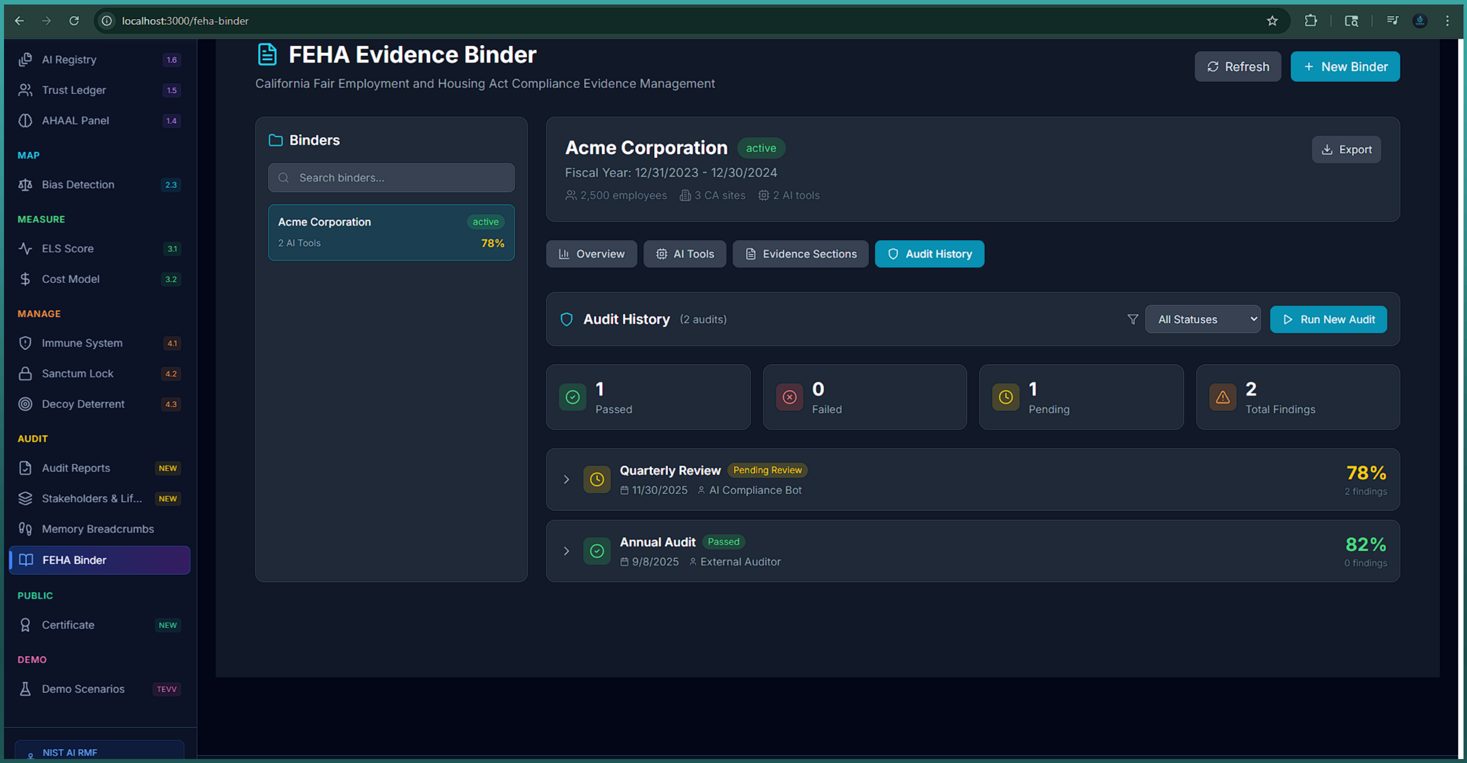

FEHA Evidence Binder

California Fair Employment and Housing Act Compliance Evidence Management

Binders

Acme Corporation active
2 AI Tools 78%

Acme Corporation active

Fiscal Year: 12/31/2023 - 12/30/2024

2,500 employees · 3 CA sites · 2 AI tools

Audit History (2 audits)

1	0	1	2
Passed	Failed	Pending	Total Findings

Quarterly Review Pending Review
11/30/2025 · AI Compliance Bot
78%
2 findings

Annual Audit Passed
9/8/2025 · External Auditor
82%
0 findings

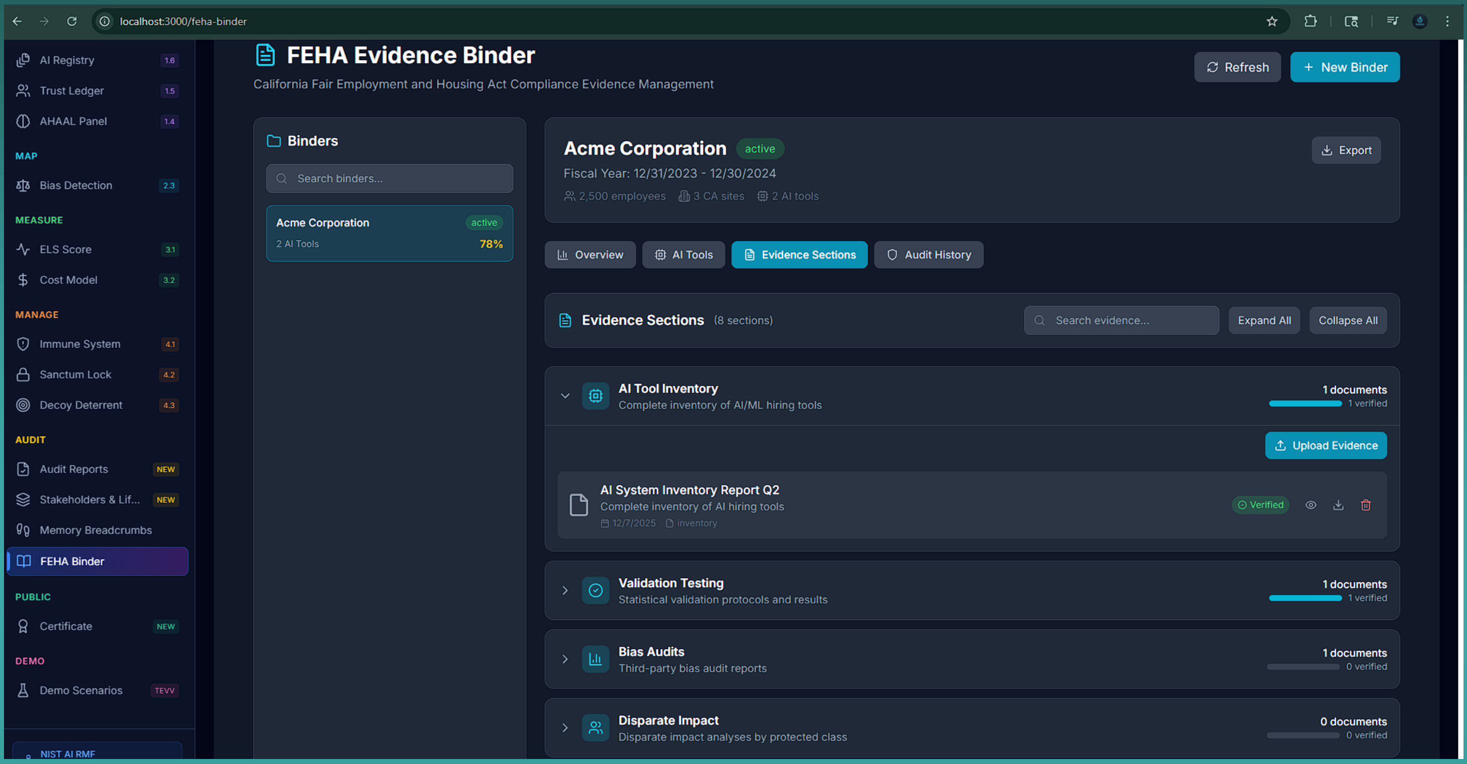

FEHA Evidence Binder

California Fair Employment and Housing Act Compliance Evidence Management

Binders

Acme Corporation active
2 AI Tools 78%

Acme Corporation active

Fiscal Year: 12/31/2023 - 12/30/2024

2,500 employees · 3 CA sites · 2 AI tools

Evidence Sections (8 sections)

AI Tool Inventory
Complete inventory of AI/ML hiring tools
1 documents
1 verified

AI System Inventory Report Q2
Complete inventory of AI hiring tools
12/7/2025 · inventory
Verified

Validation Testing
Statistical validation protocols and results
1 documents
1 verified

Bias Audits
Third-party bias audit reports
1 documents
0 verified

Disparate Impact
Disparate impact analyses by protected class
0 documents
0 verified

California Fair Employment audit binder with evidence collection, AI tool inventory, and export.

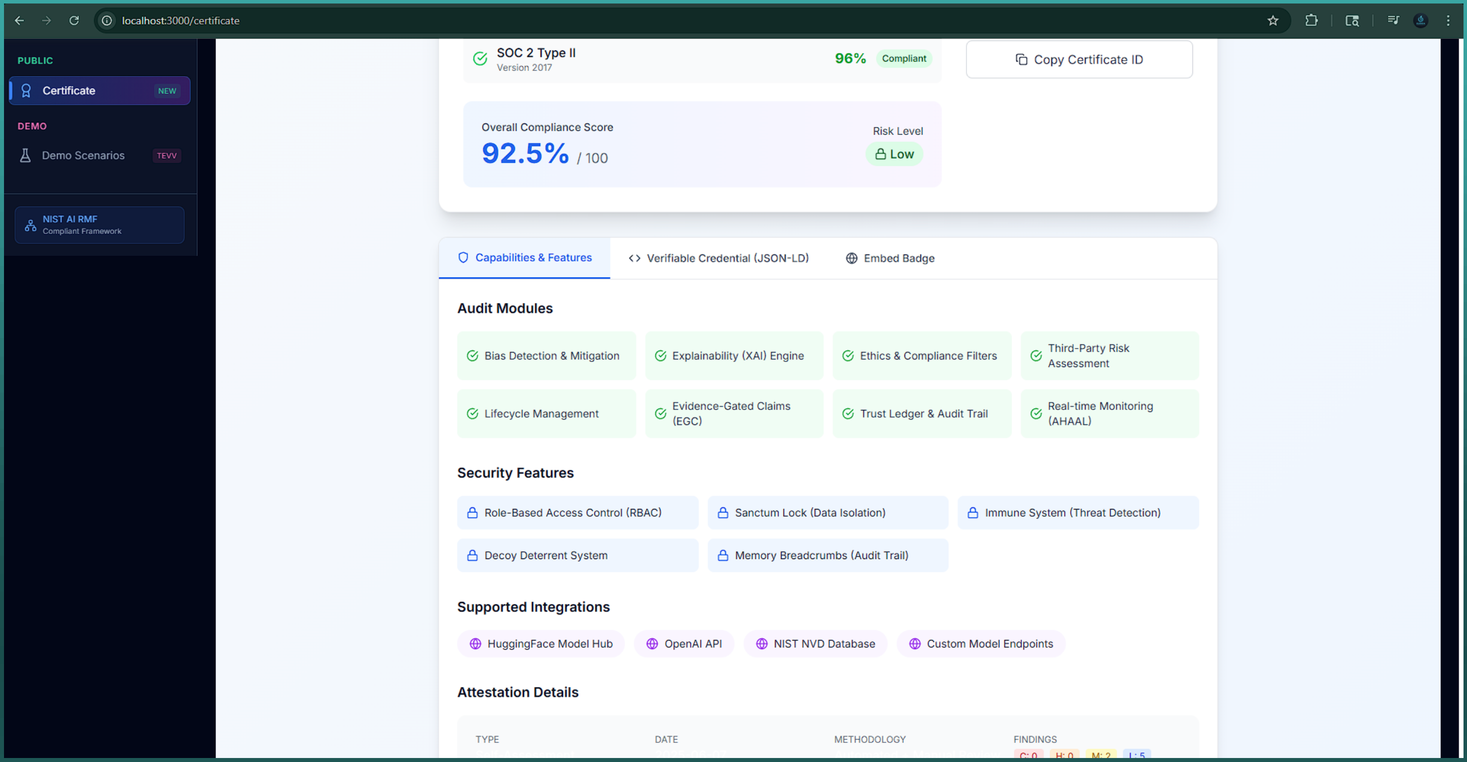

localhost:3000/certificate

PUBLIC

Certificate NEW

DEMO

Demo Scenarios TEVV

NIST AI RMF
Compliant Framework

SOC 2 Type II
Version 2017
96% Compliant

Copy Certificate ID

Overall Compliance Score
92.5% / 100

Risk Level
Low

Capabilities & Features | Verifiable Credential (JSON-LD) | Embed Badge

Audit Modules

- Bias Detection & Mitigation
- Explainability (XAI) Engine
- Ethics & Compliance Filters
- Third-Party Risk Assessment
- Lifecycle Management
- Evidence-Gated Claims (EGC)
- Trust Ledger & Audit Trail
- Real-time Monitoring (AHAAL)

Security Features

- Role-Based Access Control (RBAC)
- Sanctum Lock (Data Isolation)
- Immune System (Threat Detection)
- Decoy Deterrent System
- Memory Breadcrumbs (Audit Trail)

Supported Integrations

- HuggingFace Model Hub
- OpenAI API
- NIST NVD Database
- Custom Model Endpoints

Attestation Details

TYPE	DATE	METHODOLOGY	FINDINGS
Self-Assessment	2025-06-07	Automated + Manual Review	C: 0 H: 0 M: 2 L: 5

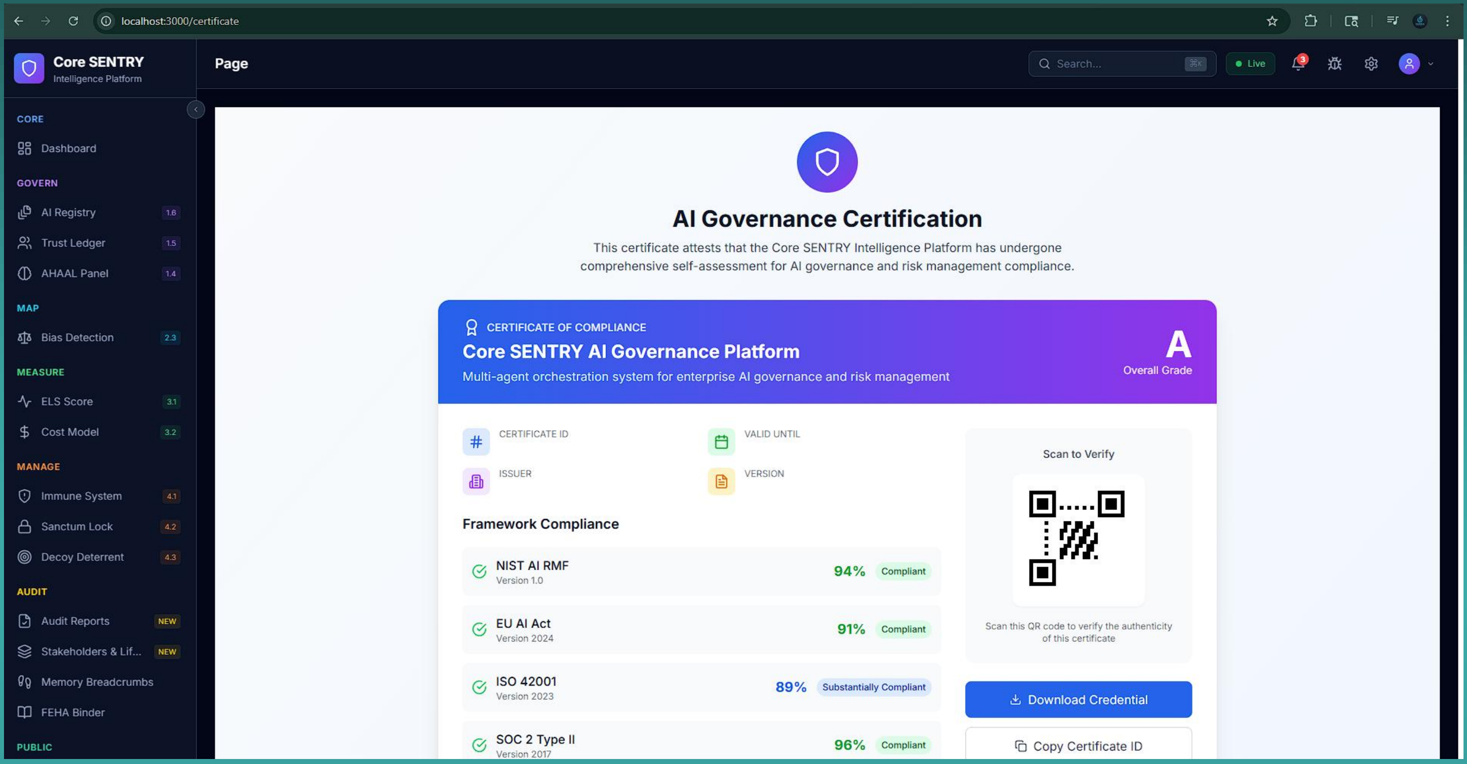

localhost:3000/certificate

Core SENTRY
Intelligence Platform

Page

Search...

Live

CORE
- Dashboard

GOVERN
- AI Registry 1.6
- Trust Ledger 1.5
- AHAAL Panel 1.4

MAP
- Bias Detection 2.3

MEASURE
- ELS Score 3.1
- Cost Model 3.2

MANAGE
- Immune System 4.1
- Sanctum Lock 4.2
- Decoy Deterrent 4.3

AUDIT
- Audit Reports NEW
- Stakeholders & Lif... NEW
- Memory Breadcrumbs
- FEHA Binder

PUBLIC

AI Governance Certification

This certificate attests that the Core SENTRY Intelligence Platform has undergone comprehensive self-assessment for AI governance and risk management compliance.

CERTIFICATE OF COMPLIANCE

Core SENTRY AI Governance Platform

Multi-agent orchestration system for enterprise AI governance and risk management

A
Overall Grade

CERTIFICATE ID

VALID UNTIL

ISSUER

VERSION

Framework Compliance

Framework	Version	Score	Status
NIST AI RMF	Version 1.0	94%	Compliant
EU AI Act	Version 2024	91%	Compliant
ISO 42001	Version 2023	89%	Substantially Compliant
SOC 2 Type II	Version 2017	96%	Compliant

Scan to Verify

Scan this QR code to verify the authenticity of this certificate

Download Credential

Copy Certificate ID

Generates and manages compliance certificates and attestation documents for completed audits

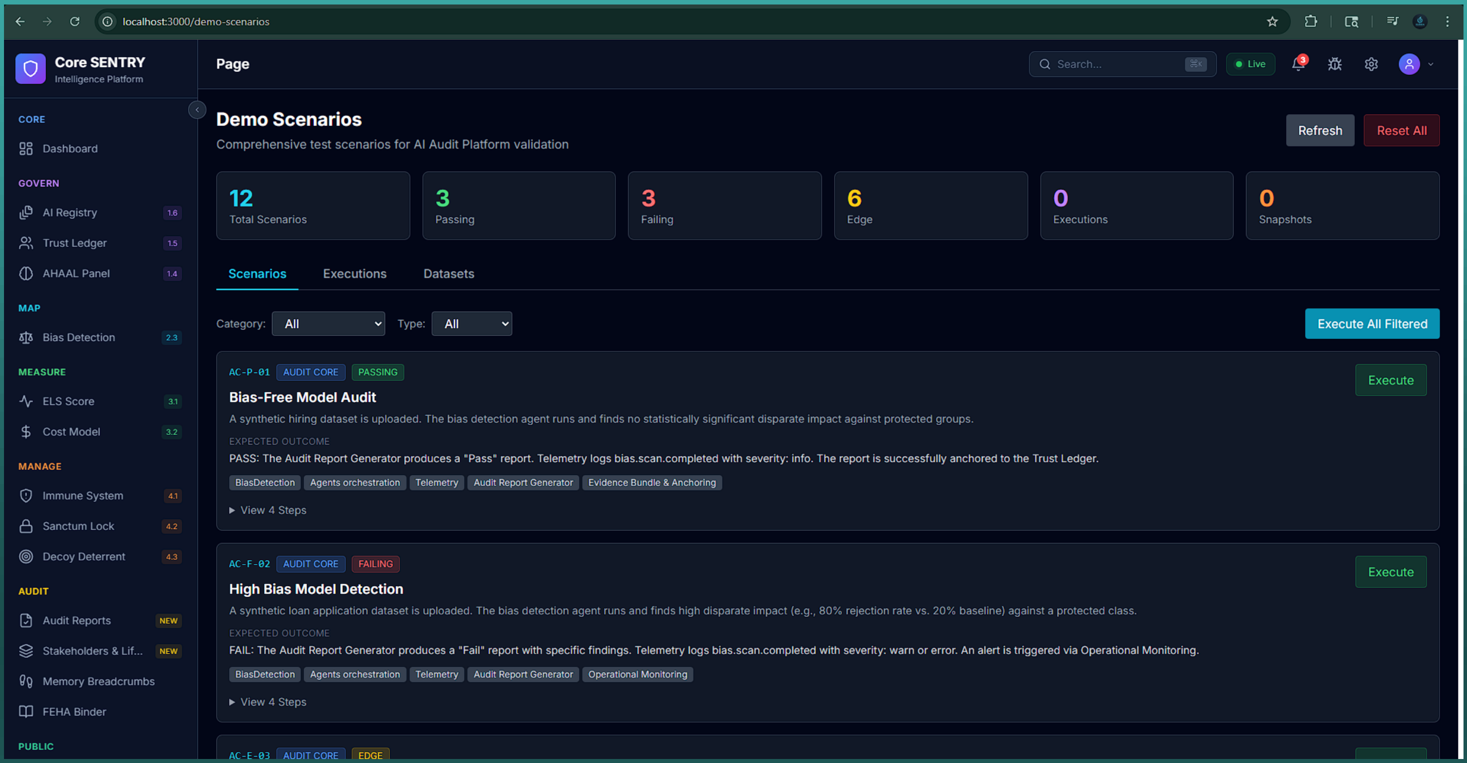

Demo Scenarios

Comprehensive test scenarios for AI Audit Platform validation

12 Total Scenarios | 3 Passing | 3 Failing | 6 Edge | 0 Executions | 0 Snapshots

Category: All   Type: All

AC-P-01 AUDIT CORE PASSING

Bias-Free Model Audit

A synthetic hiring dataset is uploaded. The bias detection agent runs and finds no statistically significant disparate impact against protected groups.

EXPECTED OUTCOME

PASS: The Audit Report Generator produces a "Pass" report. Telemetry logs bias.scan.completed with severity: info. The report is successfully anchored to the Trust Ledger.

BiasDetection | Agents orchestration | Telemetry | Audit Report Generator | Evidence Bundle & Anchoring

▸ View 4 Steps

AC-F-02 AUDIT CORE FAILING

High Bias Model Detection

A synthetic loan application dataset is uploaded. The bias detection agent runs and finds high disparate impact (e.g., 80% rejection rate vs. 20% baseline) against a protected class.

EXPECTED OUTCOME

FAIL: The Audit Report Generator produces a "Fail" report with specific findings. Telemetry logs bias.scan.completed with severity: warn or error. An alert is triggered via Operational Monitoring.

BiasDetection | Agents orchestration | Telemetry | Audit Report Generator | Operational Monitoring

▸ View 4 Steps

AC-E-03 AUDIT CORE EDGE

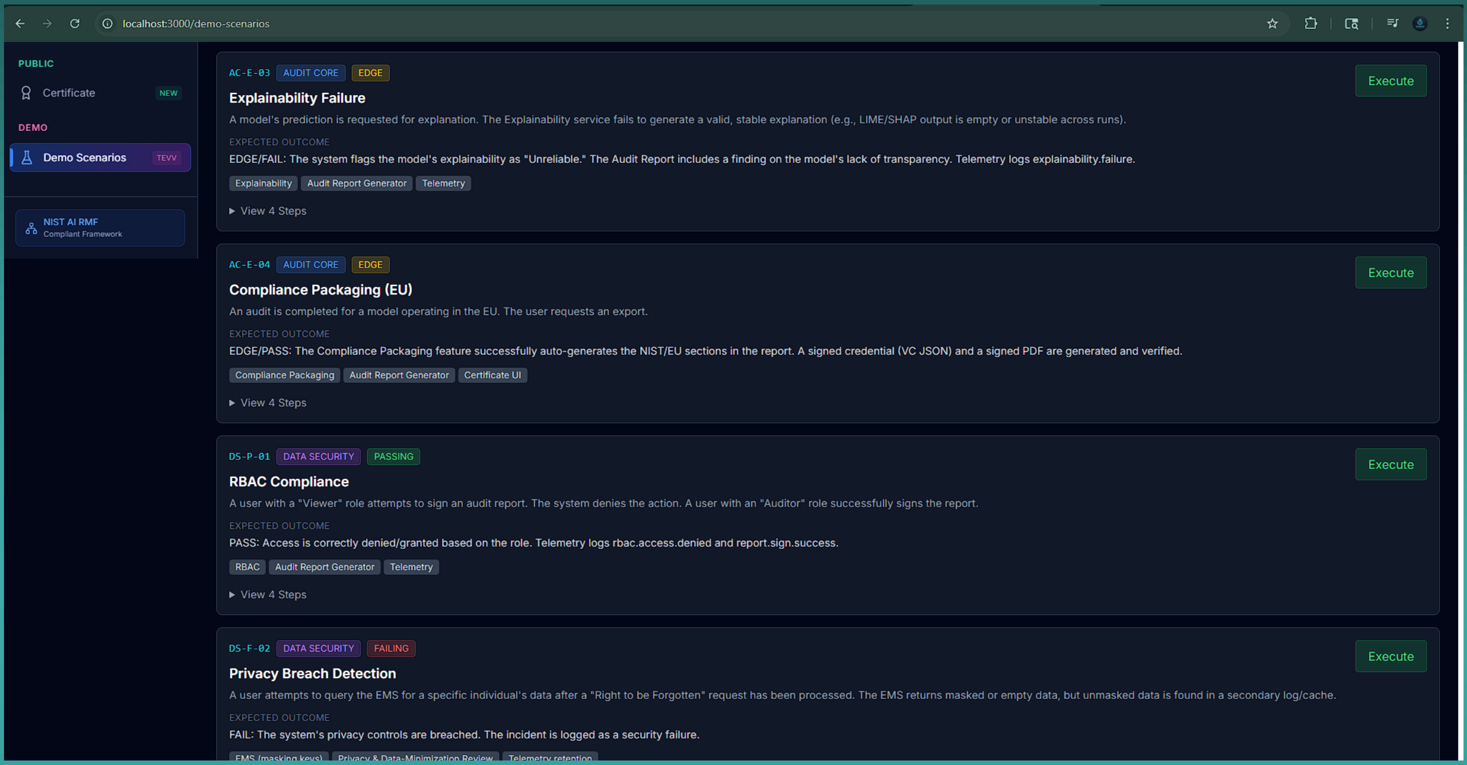

AC-E-03 AUDIT CORE EDGE

Explainability Failure

A model's prediction is requested for explanation. The Explainability service fails to generate a valid, stable explanation (e.g., LIME/SHAP output is empty or unstable across runs).

EXPECTED OUTCOME

EDGE/FAIL: The system flags the model's explainability as "Unreliable." The Audit Report includes a finding on the model's lack of transparency. Telemetry logs explainability.failure.

Explainability | Audit Report Generator | Telemetry

▸ View 4 Steps

AC-E-04 AUDIT CORE EDGE

Compliance Packaging (EU)

An audit is completed for a model operating in the EU. The user requests an export.

EXPECTED OUTCOME

EDGE/PASS: The Compliance Packaging feature successfully auto-generates the NIST/EU sections in the report. A signed credential (VC JSON) and a signed PDF are generated and verified.

Compliance Packaging | Audit Report Generator | Certificate UI

▸ View 4 Steps

DS-P-01 DATA SECURITY PASSING

RBAC Compliance

A user with a "Viewer" role attempts to sign an audit report. The system denies the action. A user with an "Auditor" role successfully signs the report.

EXPECTED OUTCOME

PASS: Access is correctly denied/granted based on the role. Telemetry logs rbac.access.denied and report.sign.success.

RBAC | Audit Report Generator | Telemetry

▸ View 4 Steps

DS-F-02 DATA SECURITY FAILING

Privacy Breach Detection

A user attempts to query the EMS for a specific individual's data after a "Right to be Forgotten" request has been processed. The EMS returns masked or empty data, but unmasked data is found in a secondary log/cache.

EXPECTED OUTCOME

FAIL: The system's privacy controls are breached. The incident is logged as a security failure.

EMS (masking keys) | Privacy & Data-Minimization Review | Telemetry retention

Pre-built demonstration scenarios for showcasing platform capabilities to stakeholders.