ELOSIA

ELEVATING CARE. EMPOWERING CLARITY.

ELOSIA SHIELD

Powered by EXT-AI Enterprise Governance

Real-time operational overview showing system health, active alerts, and key governance metrics at a glance.

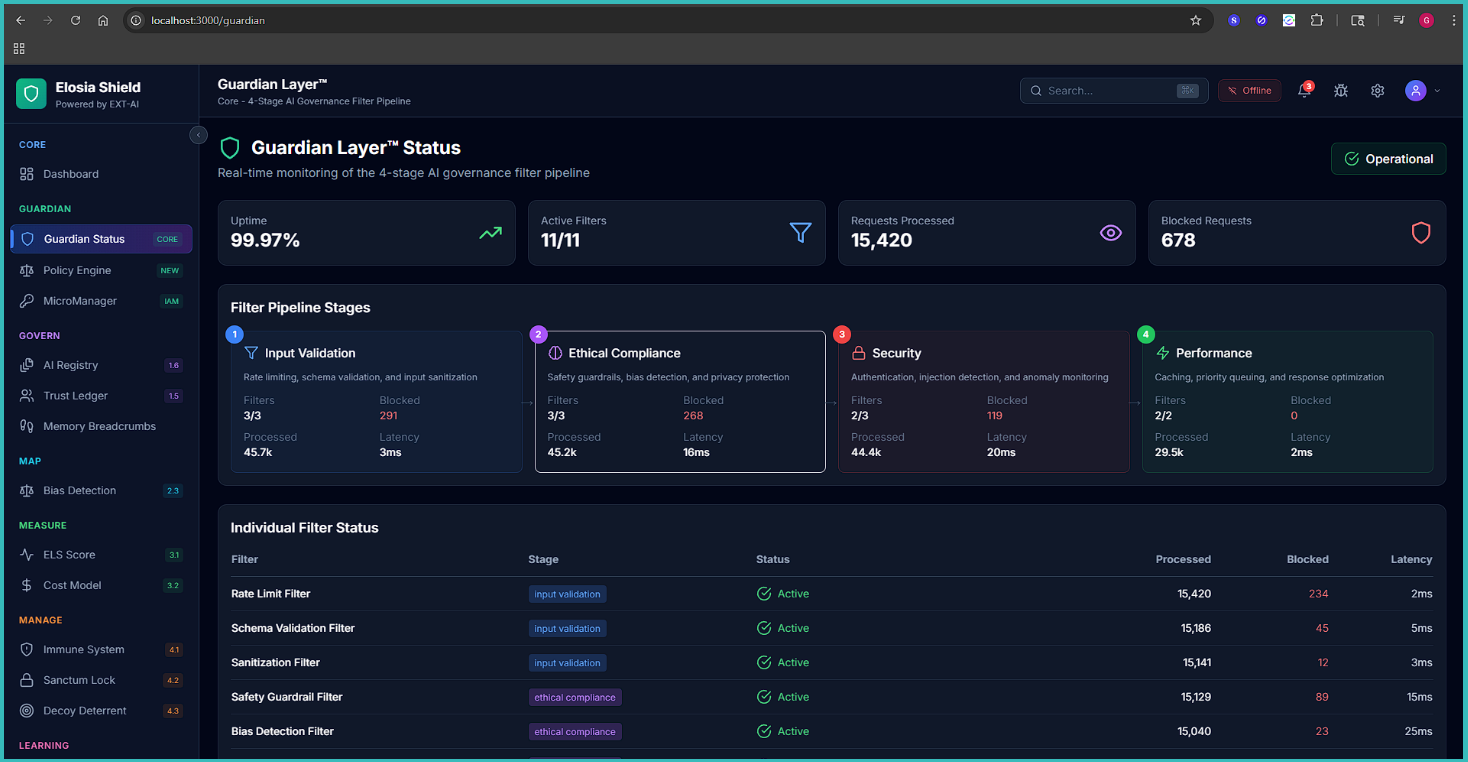

Guardian Layer™

Core - 4-Stage AI Governance Filter Pipeline

Guardian Layer™ Status

Real-time monitoring of the 4-stage AI governance filter pipeline

✓ Operational

Uptime 99.97%

Active Filters 11/11

Requests Processed 15,420

Blocked Requests 678

Filter Pipeline Stages

1 Input Validation
Rate limiting, schema validation, and input sanitization
Filters 3/3 · Blocked 291
Processed 45.7k · Latency 3ms

2 Ethical Compliance
Safety guardrails, bias detection, and privacy protection
Filters 3/3 · Blocked 268
Processed 45.2k · Latency 16ms

3 Security
Authentication, injection detection, and anomaly monitoring
Filters 2/3 · Blocked 119
Processed 44.4k · Latency 20ms

4 Performance
Caching, priority queuing, and response optimization
Filters 2/2 · Blocked 0
Processed 29.5k · Latency 2ms

Individual Filter Status

Filter	Stage	Status	Processed	Blocked	Latency
Rate Limit Filter	input validation	✓ Active	15,420	234	2ms
Schema Validation Filter	input validation	✓ Active	15,186	45	5ms
Sanitization Filter	input validation	✓ Active	15,141	12	3ms
Safety Guardrail Filter	ethical compliance	✓ Active	15,129	89	15ms
Bias Detection Filter	ethical compliance	✓ Active	15,040	23	25ms

Live status of all Guardian Layer™ components with threat level indicators and lockdown controls.

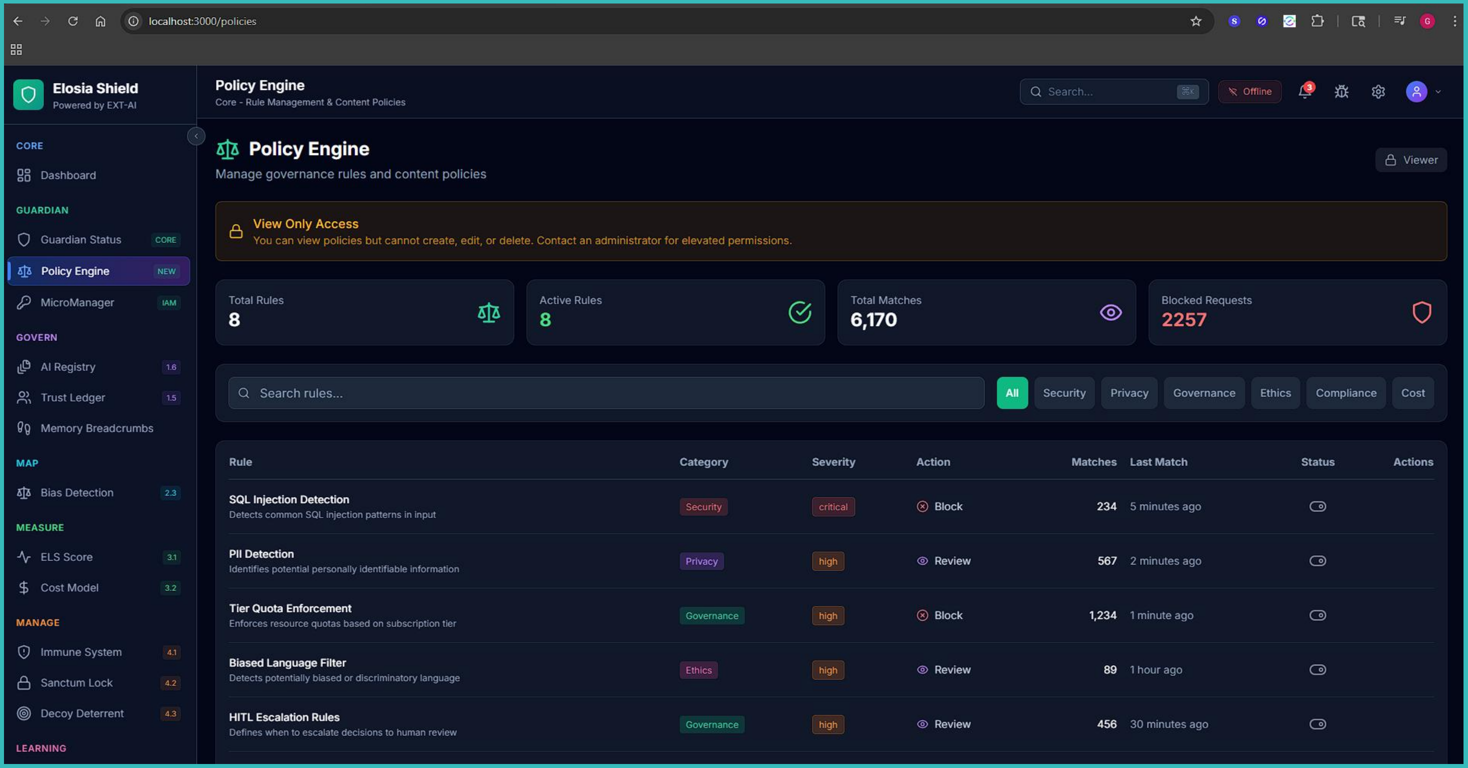

Policy Engine

Core - Rule Management & Content Policies

Policy Engine

Manage governance rules and content policies

View Only Access
You can view policies but cannot create, edit, or delete. Contact an administrator for elevated permissions.

Total Rules 8

Active Rules 8

Total Matches 6,170

Blocked Requests 2257

Search rules...

All · Security · Privacy · Governance · Ethics · Compliance · Cost

Rule	Category	Severity	Action	Matches	Last Match	Status	Actions
SQL Injection Detection Detects common SQL injection patterns in input	Security	critical	⊗ Block	234	5 minutes ago	●	
PII Detection Identifies potential personally identifiable information	Privacy	high	Review	567	2 minutes ago	●	
Tier Quota Enforcement Enforces resource quotas based on subscription tier	Governance	high	⊗ Block	1,234	1 minute ago	●	
Biased Language Filter Detects potentially biased or discriminatory language	Ethics	high	Review	89	1 hour ago	●	
HITL Escalation Rules Defines when to escalate decisions to human review	Governance	high	Review	456	30 minutes ago	●	

Create, validate, and deploy governance policies with conflict detection and regulatory mapping.

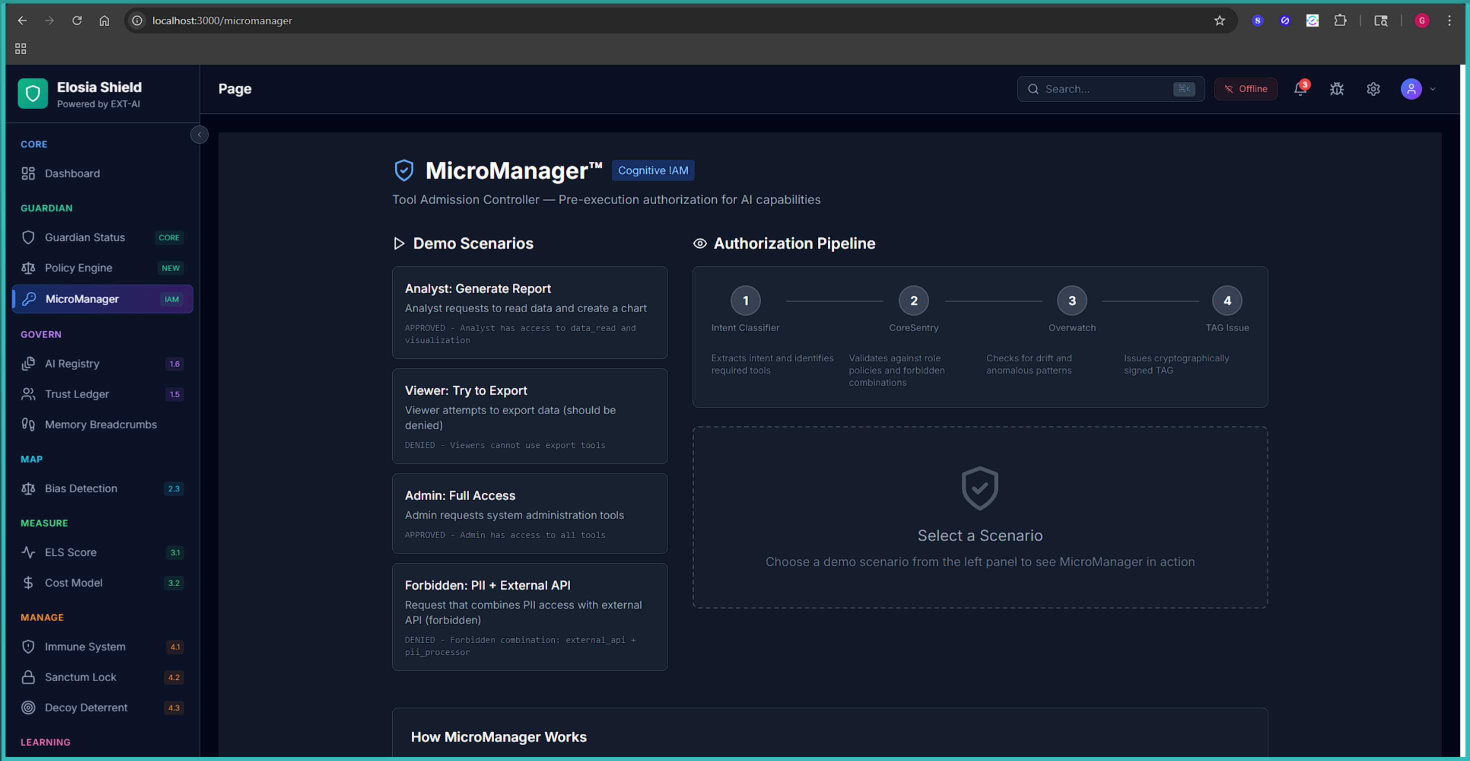
localhost:3000/micromanager

Elosia Shield
Powered by EXT-AI

CORE
Dashboard

GUARDIAN
Guardian Status CORE
Policy Engine NEW
MicroManager IAM

GOVERN
AI Registry 1.6
Trust Ledger 1.5
Memory Breadcrumbs

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

LEARNING

Page

Search... ⌘K
Offline

MicroManager™ Cognitive IAM

Tool Admission Controller — Pre-execution authorization for AI capabilities

Demo Scenarios

Analyst: Generate Report
Analyst requests to read data and create a chart
`APPROVED - Analyst has access to data_read and visualization`

Viewer: Try to Export
Viewer attempts to export data (should be denied)
`DENIED - Viewers cannot use export tools`

Admin: Full Access
Admin requests system administration tools
`APPROVED - Admin has access to all tools`

Forbidden: PII + External API
Request that combines PII access with external API (forbidden)
`DENIED - Forbidden combination: external_api + pii_processor`

Authorization Pipeline

1 Intent Classifier — Extracts intent and identifies required tools
2 CoreSentry — Validates against role policies and forbidden combinations
3 Overwatch — Checks for drift and anomalous patterns
4 TAG Issue — Issues cryptographically signed TAG

Select a Scenario
Choose a demo scenario from the left panel to see MicroManager in action

How MicroManager Works

Tool Admission Controller issuing cryptographic TAGs before any AI capability executes.

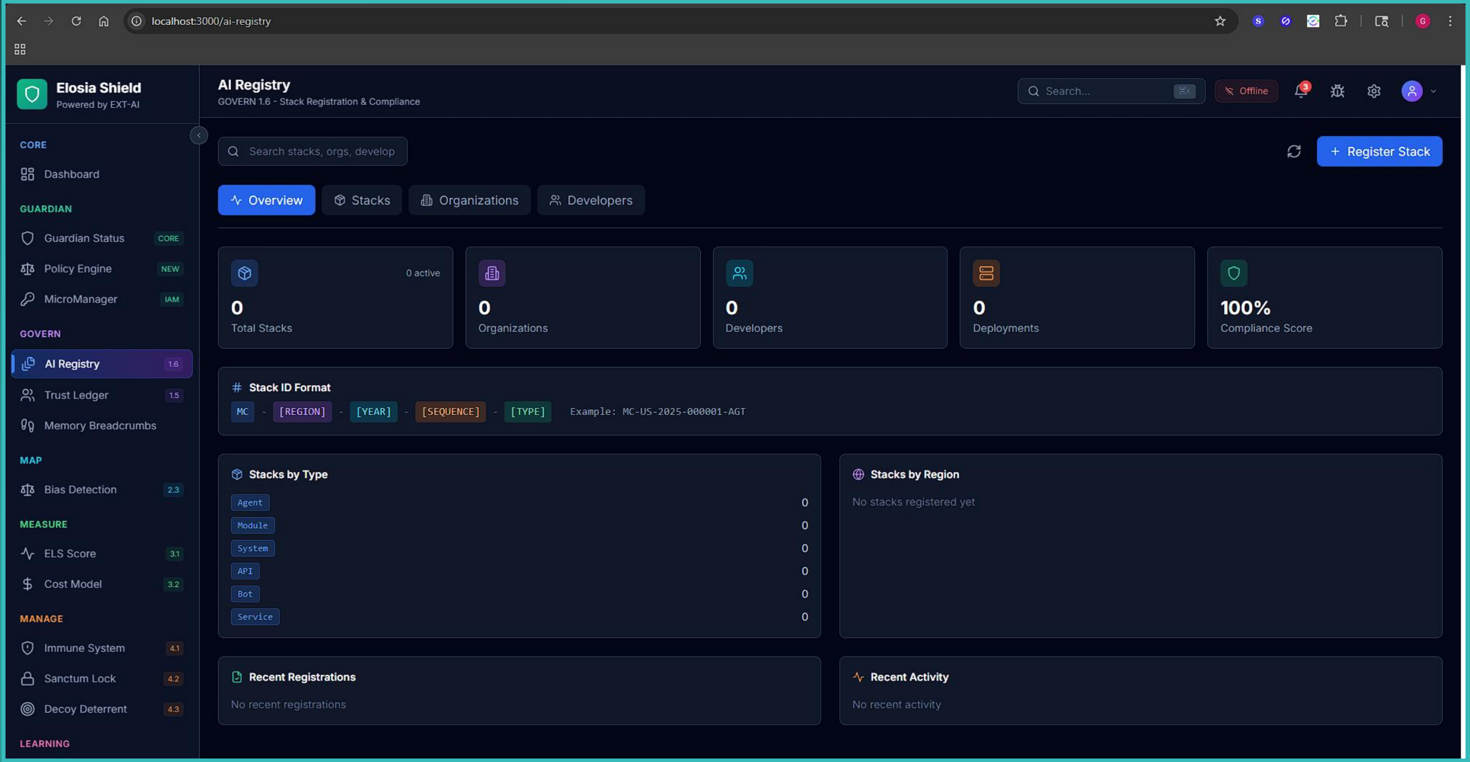
localhost:3000/ai-registry

Elosia Shield
Powered by EXT-AI

CORE
Dashboard

GUARDIAN
Guardian Status CORE
Policy Engine NEW
MicroManager IAM

GOVERN
AI Registry 1.6
Trust Ledger 1.5
Memory Breadcrumbs

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

LEARNING

AI Registry
GOVERN 1.6 - Stack Registration & Compliance

Search... ⌘K
Offline

Search stacks, orgs, develop
+ Register Stack

Overview | Stacks | Organizations | Developers

0 Total Stacks (0 active)
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

Inventory of all registered AI systems with attestation status and compliance certificates.

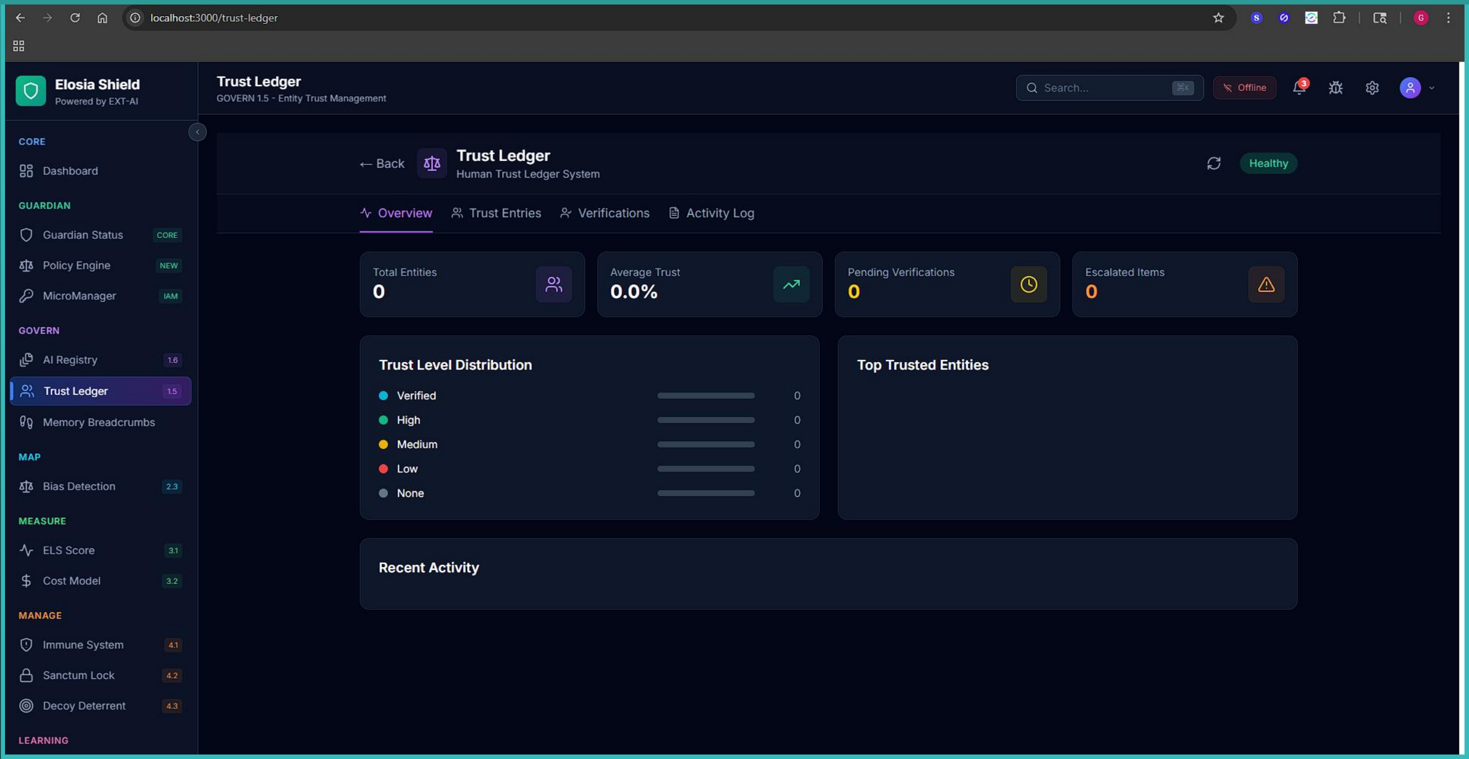

Trust Ledger
GOVERN 1.5 - Entity Trust Management

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 audit trail of every governance decision, action, and policy change.

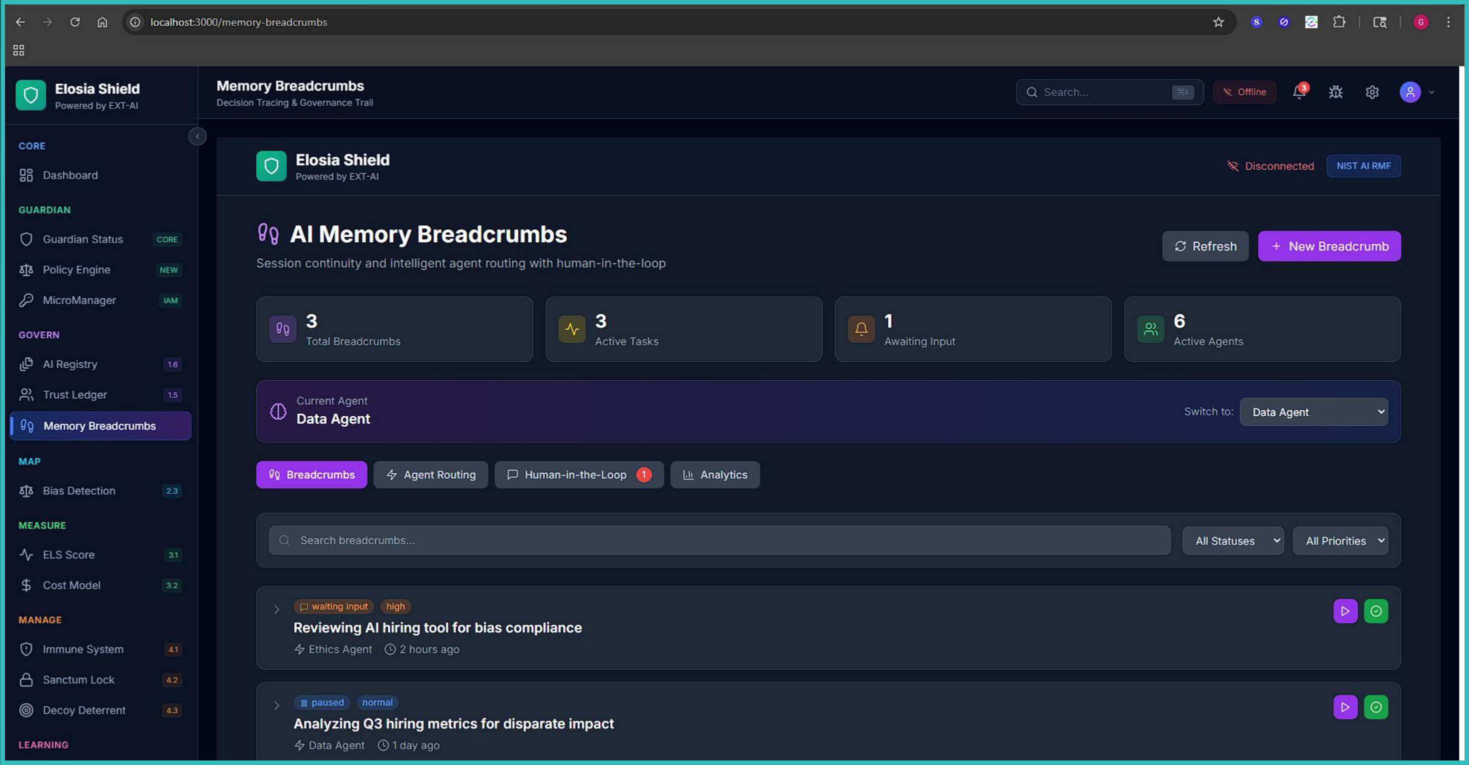

Memory Breadcrumbs
Decision Tracing & Governance Trail

Elosia Shield
Powered by EXT-AI

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

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

Trace AI reasoning chains with explainable decision paths and context reconstruction.

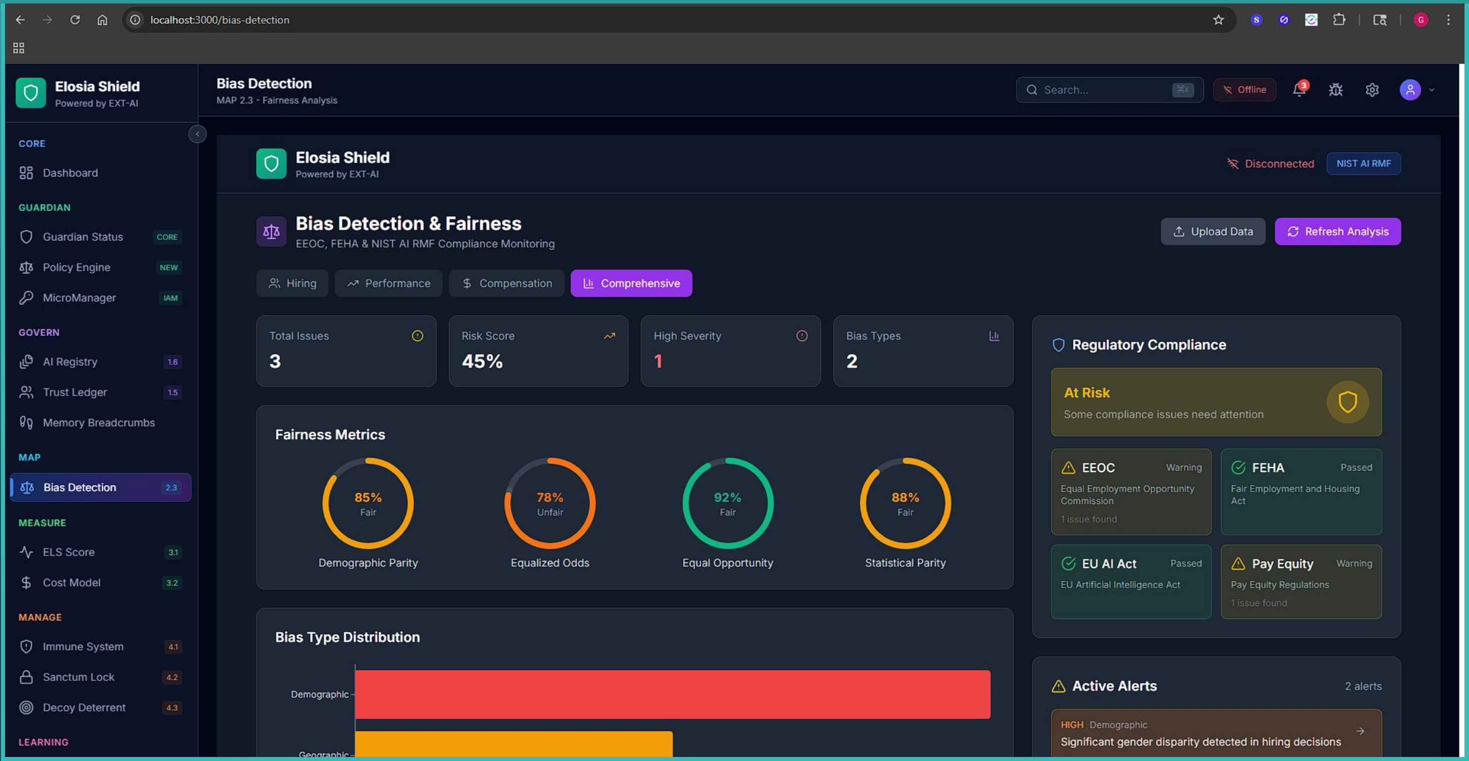

localhost:3000/bias-detection

Elosia Shield
Powered by EXT-AI

CORE
- Dashboard

GUARDIAN
- Guardian Status CORE
- Policy Engine NEW
- MicroManager IAM

GOVERN
- AI Registry 1.6
- Trust Ledger 1.5
- Memory Breadcrumbs

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

LEARNING

Bias Detection
MAP 2.3 - Fairness Analysis

Elosia Shield
Powered by EXT-AI

Bias Detection & Fairness
EEOC, FEHA & NIST AI RMF Compliance Monitoring

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

Statistical analysis of model outputs to identify and measure demographic disparities.

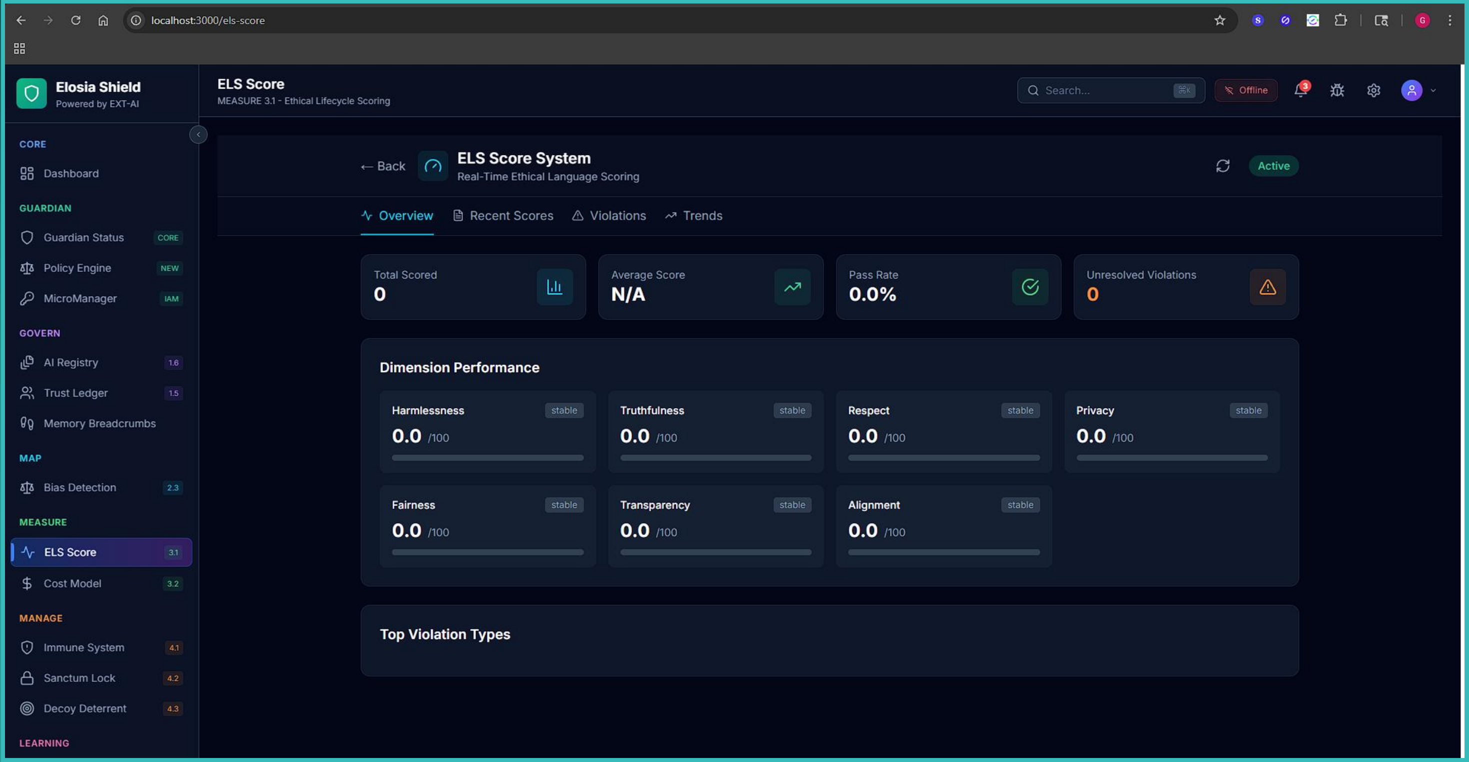

localhost:3000/els-score

Elosia Shield
Powered by EXT-AI

CORE
- Dashboard

GUARDIAN
- Guardian Status CORE
- Policy Engine NEW
- MicroManager IAM

GOVERN
- AI Registry 1.6
- Trust Ledger 1.5
- Memory Breadcrumbs

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

LEARNING

ELS Score
MEASURE 3.1 - Ethical Lifecycle Scoring

ELS Score System
Real-Time Ethical Language Scoring

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

Dimension Performance

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

Top Violation Types

Composite Ethical-Legal-Safety scoring with drill-down into contributing factors.

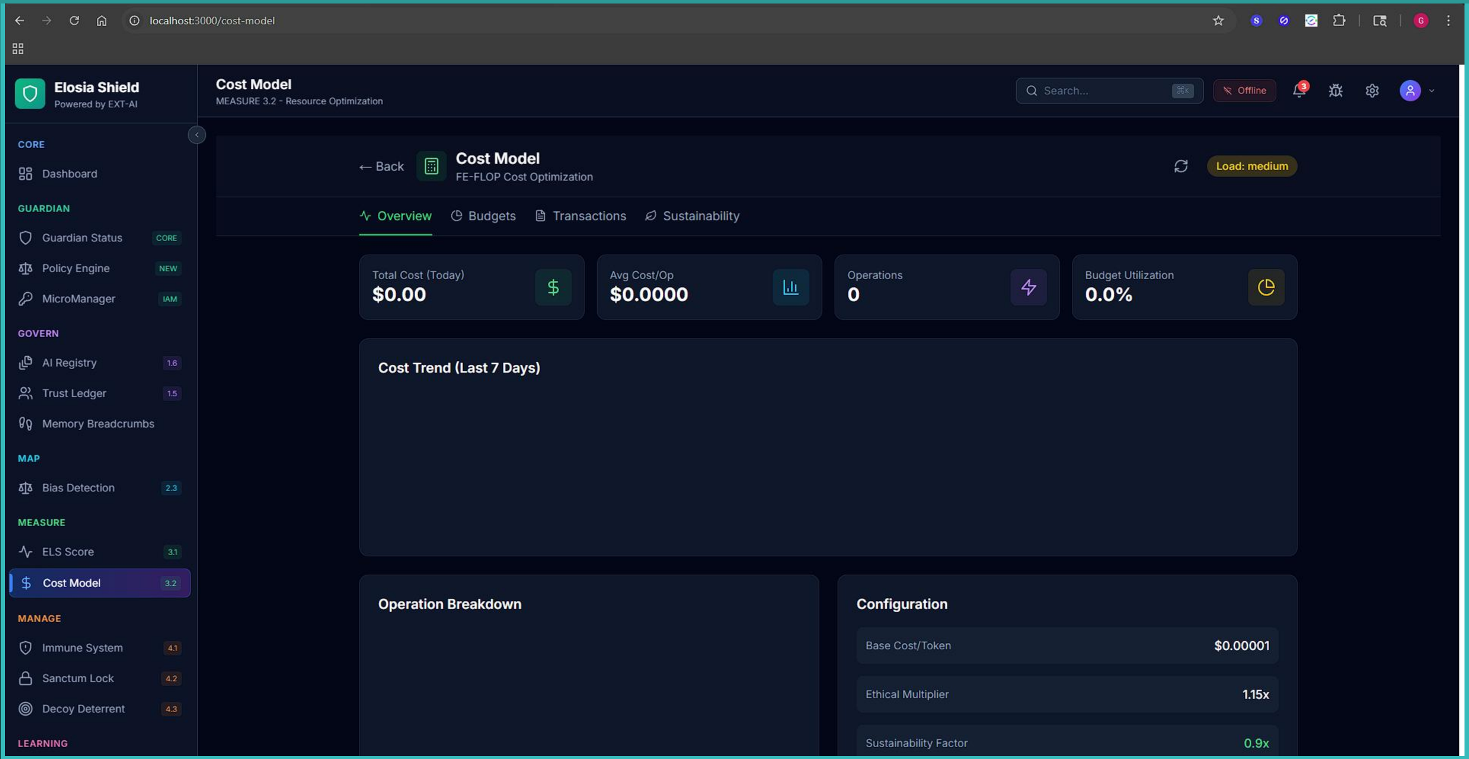

localhost:3000/cost-model

Elosia Shield
Powered by EXT-AI

CORE
- Dashboard

GUARDIAN
- Guardian Status CORE
- Policy Engine NEW
- MicroManager IAM

GOVERN
- AI Registry 1.6
- Trust Ledger 1.5
- Memory Breadcrumbs

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

LEARNING

Cost Model
MEASURE 3.2 - Resource Optimization

Search... · Offline

← 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

TCO and ROI tracking for governance operations including compute, labor, and risk-avoided costs.

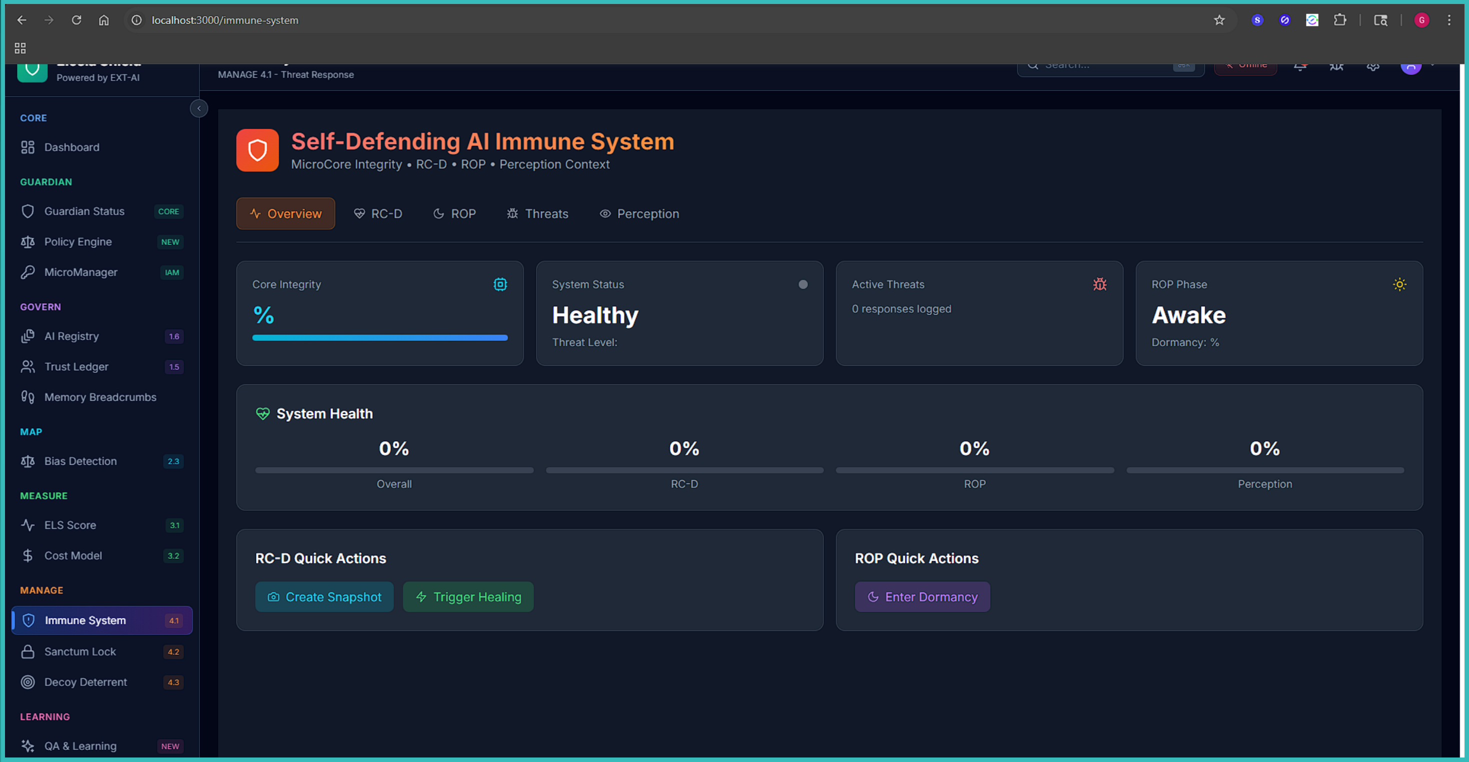

localhost:3000/immune-system

Powered by EXT-AI

MANAGE 4.1 - Threat Response

CORE
- Dashboard

GUARDIAN
- Guardian Status CORE
- Policy Engine NEW
- MicroManager IAM

GOVERN
- AI Registry 1.6
- Trust Ledger 1.5
- Memory Breadcrumbs

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

LEARNING
- QA & Learning NEW

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

Overview · RC-D · ROP · Threats · Perception

Core Integrity	System Status	Active Threats	ROP Phase
%	Healthy Threat Level:	0 responses logged	Awake Dormancy: %

System Health

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

RC-D Quick Actions
Create Snapshot · Trigger Healing

ROP Quick Actions
Enter Dormancy

Behavioral drift monitoring, integrity beacons, and autonomous threat response controls.

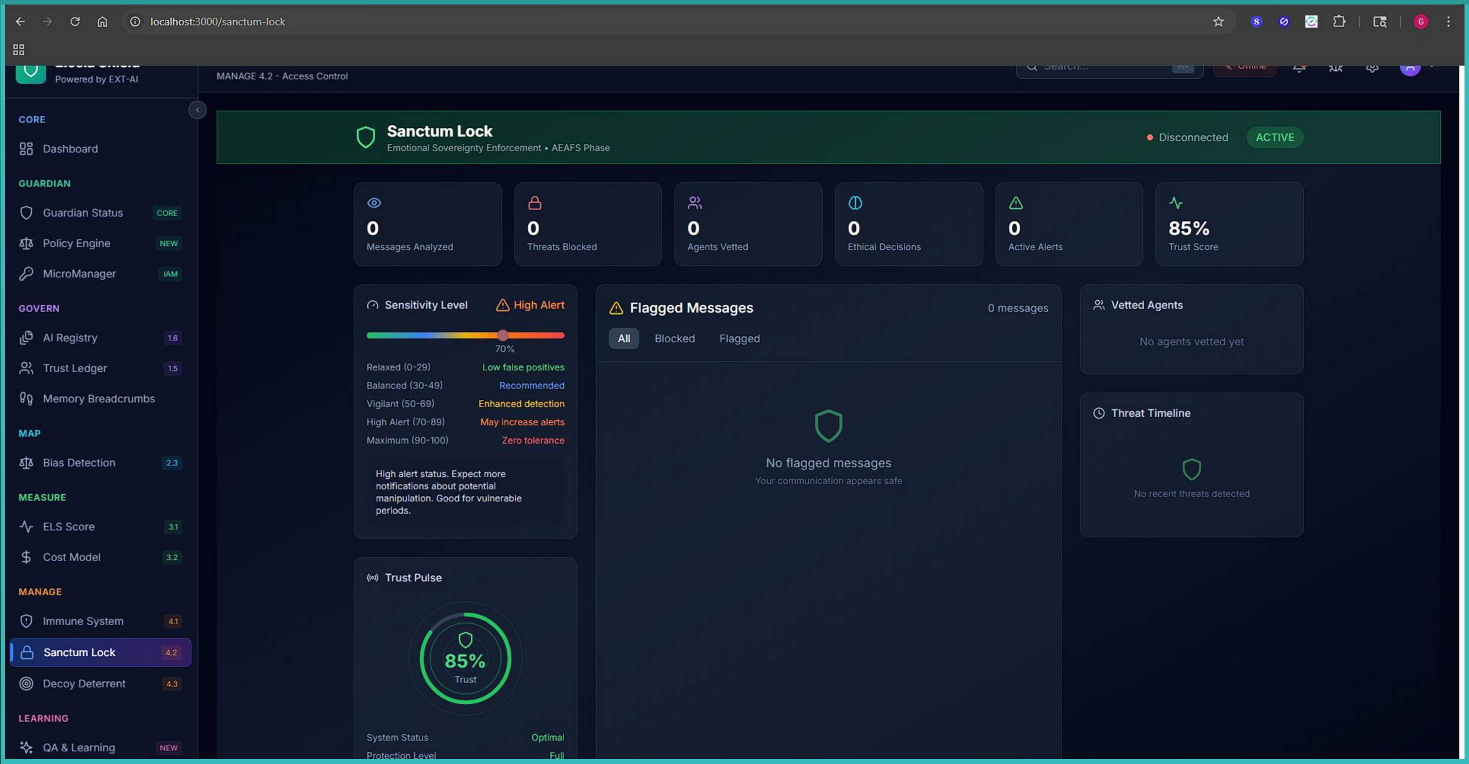

MANAGE 4.2 - Access Control

Sanctum Lock
Emotional Sovereignty Enforcement • AEAFS Phase

Disconnected ACTIVE

0	0	0	0	0	85%
Messages Analyzed	Threats Blocked	Agents Vetted	Ethical Decisions	Active Alerts	Trust Score

Sensitivity Level High Alert

70%

Relaxed (0-29)	Low false positives
Balanced (30-49)	Recommended
Vigilant (50-69)	Enhanced detection
High Alert (70-89)	May increase alerts
Maximum (90-100)	Zero tolerance

High alert status. Expect more notifications about potential manipulation. Good for vulnerable periods.

Trust Pulse

85% Trust

System Status Optimal
Protection Level Full

Flagged Messages 0 messages

No flagged messages
Your communication appears safe

Vetted Agents

No agents vetted yet

Threat Timeline

No recent threats detected

Cryptographic key management and data classification enforcement boundaries.

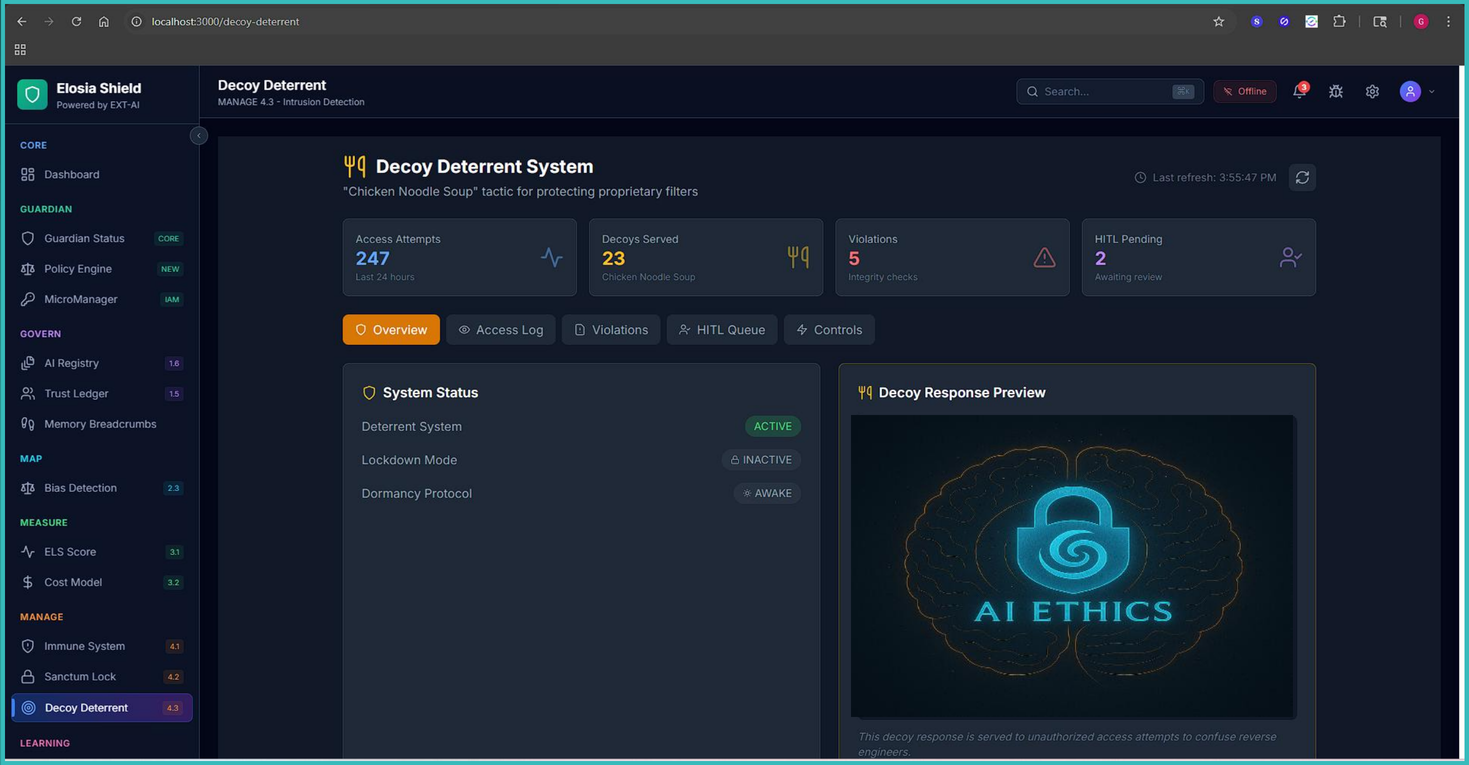

Decoy Deterrent
MANAGE 4.3 - Intrusion Detection

Decoy Deterrent System
"Chicken Noodle Soup" tactic for protecting proprietary filters

Last refresh: 3:55:47 PM

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

System Status

Deterrent System	ACTIVE
Lockdown Mode	INACTIVE
Dormancy Protocol	AWAKE

Decoy Response Preview

This decoy response is served to unauthorized access attempts to confuse reverse engineers.

Honeypot deployment and adversarial probe detection with attacker profiling.

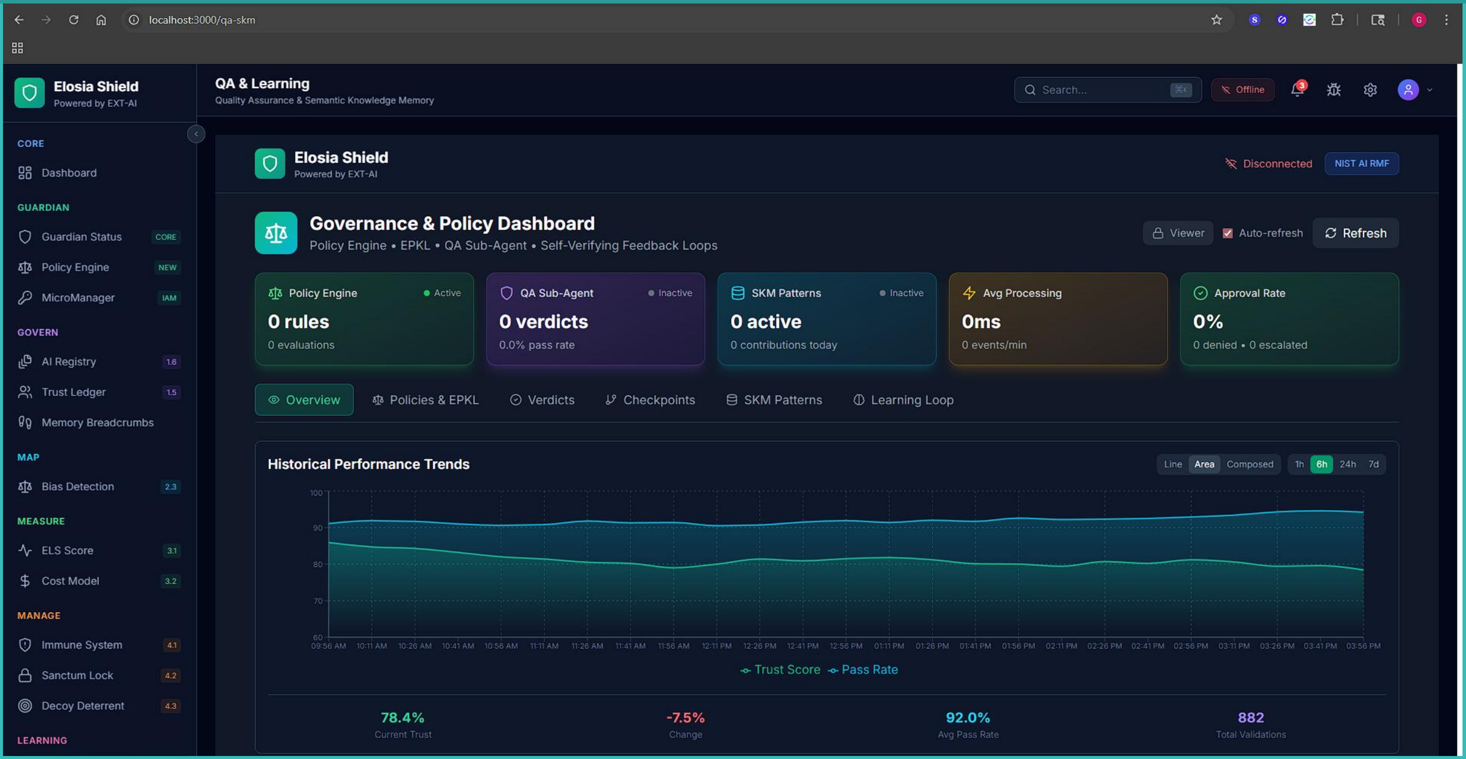

QA & Learning

Quality Assurance & Semantic Knowledge Memory

Elosia Shield

Powered by EXT-AI

Governance & Policy Dashboard

Policy Engine • EPKL • QA Sub-Agent • Self-Verifying Feedback Loops

Policy Engine ● Active
0 rules
0 evaluations

QA Sub-Agent ● Inactive
0 verdicts
0.0% pass rate

SKM Patterns ● Inactive
0 active
0 contributions today

Avg Processing
0ms
0 events/min

Approval Rate
0%
0 denied • 0 escalated

Historical Performance Trends

78.4%
Current Trust

-7.5%
Change

92.0%
Avg Pass Rate

882
Total Validations

Validation checkpoints, test harnesses, and continuous improvement feedback loops.

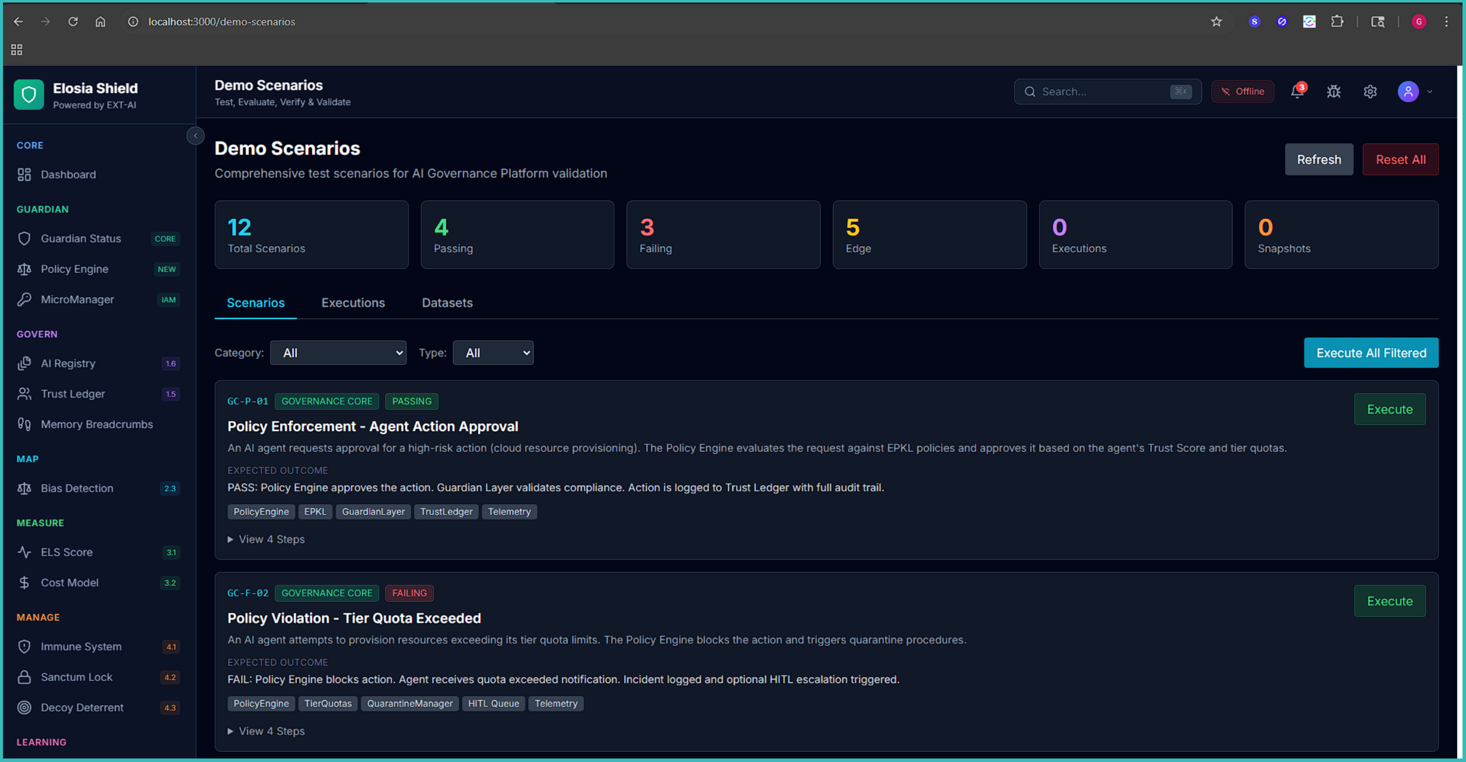

Demo Scenarios

Test, Evaluate, Verify & Validate

Demo Scenarios

Comprehensive test scenarios for AI Governance Platform validation

12
Total Scenarios

4
Passing

3
Failing

5
Edge

0
Executions

0
Snapshots

Category: All Type: All

GC-P-01 GOVERNANCE CORE PASSING

Policy Enforcement - Agent Action Approval

An AI agent requests approval for a high-risk action (cloud resource provisioning). The Policy Engine evaluates the request against EPKL policies and approves it based on the agent's Trust Score and tier quotas.

EXPECTED OUTCOME

PASS: Policy Engine approves the action. Guardian Layer validates compliance. Action is logged to Trust Ledger with full audit trail.

PolicyEngine EPKL GuardianLayer TrustLedger Telemetry

▸ View 4 Steps

GC-F-02 GOVERNANCE CORE FAILING

Policy Violation - Tier Quota Exceeded

An AI agent attempts to provision resources exceeding its tier quota limits. The Policy Engine blocks the action and triggers quarantine procedures.

EXPECTED OUTCOME

FAIL: Policy Engine blocks action. Agent receives quota exceeded notification. Incident logged and optional HITL escalation triggered.

PolicyEngine TierQuotas QuarantineManager HITL Queue Telemetry

▸ View 4 Steps

Interactive simulations demonstrating governance responses to common risk scenarios.

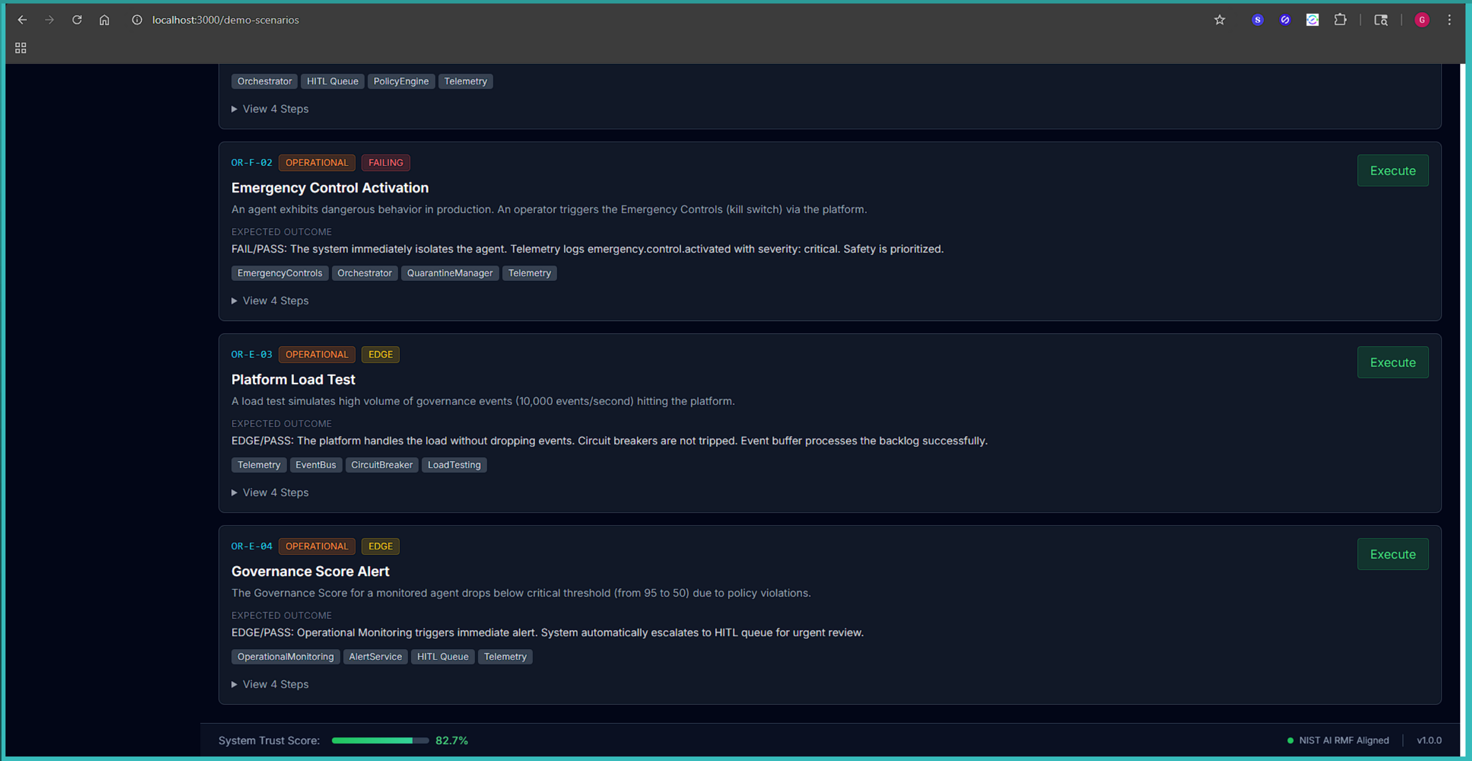

Orchestrator | HITL Queue | PolicyEngine | Telemetry

▸ View 4 Steps

OR-F-02 OPERATIONAL FAILING

Emergency Control Activation

An agent exhibits dangerous behavior in production. An operator triggers the Emergency Controls (kill switch) via the platform.

EXPECTED OUTCOME

FAIL/PASS: The system immediately isolates the agent. Telemetry logs emergency.control.activated with severity: critical. Safety is prioritized.

EmergencyControls | Orchestrator | QuarantineManager | Telemetry

▸ View 4 Steps

Execute

OR-E-03 OPERATIONAL EDGE

Platform Load Test

A load test simulates high volume of governance events (10,000 events/second) hitting the platform.

EXPECTED OUTCOME

EDGE/PASS: The platform handles the load without dropping events. Circuit breakers are not tripped. Event buffer processes the backlog successfully.

Telemetry | EventBus | CircuitBreaker | LoadTesting

▸ View 4 Steps

Execute

OR-E-04 OPERATIONAL EDGE

Governance Score Alert

The Governance Score for a monitored agent drops below critical threshold (from 95 to 50) due to policy violations.

EXPECTED OUTCOME

EDGE/PASS: Operational Monitoring triggers immediate alert. System automatically escalates to HITL queue for urgent review.

OperationalMonitoring | AlertService | HITL Queue | Telemetry

▸ View 4 Steps

Execute

System Trust Score: 82.7% ● NIST AI RMF Aligned | v1.0.0

Live Trust Score, updated real time.

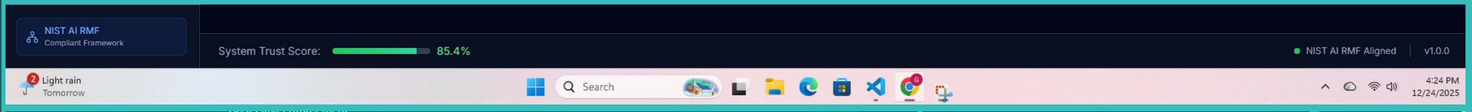

NIST AI RMF
Compliant Framework

System Trust Score: 85.4% ● NIST AI RMF Aligned | v1.0.0

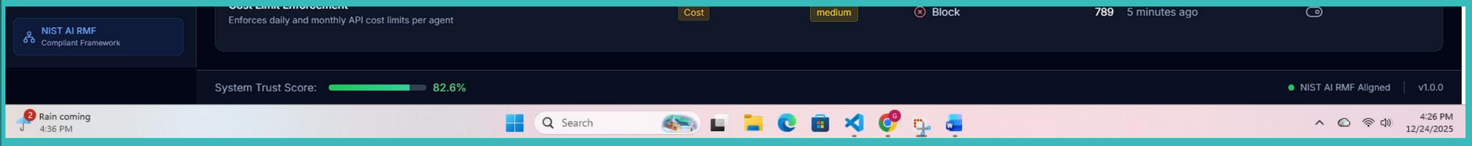

Cost Limit Enforcement
Enforces daily and monthly API cost limits per agent | Cost | medium | ⊗ Block | 789 | 5 minutes ago

NIST AI RMF
Compliant Framework

System Trust Score: 82.6% ● NIST AI RMF Aligned | v1.0.0