FeedOracle

- Platform
- Evidence
- Developers
- Pricing
- Reports

Get API Key Evidence Terminal

☰

FeedOracle Whitepaper

Evidence-Grade Data Infrastructure for Regulated Workflows

Version 5.0 · March 22, 2026

Download as PDF API Docs →

DOCUMENT CONTROL

Version	4.0.0
Date	5 March 2026
Status	Published

Changelog v5.0 (Mar 22, 2026): DORA AmpelOracle: 22 MCP tools for real-time DORA article assessment with traffic-light scoring (GREEN/YELLOW/RED/GREY), entity management, bridge workflow, provider registry, contract clause checking, regulatory watchdog. MiCAOracle: 20 MCP tools for token-level and entity-level MiCA assessment, issuer profiles, bridge workflow, freshness checking, regulatory watchdog. AMLOracle: 12 MCP tools for sanctions screening (EU/OFAC/UN/Interpol), PEP checks, adverse media, KYC bundle, transaction risk scoring. Oracle Event Bus: Cross-oracle event routing with 21 event types, 9 cross-references, 4 publishing oracles. Live Compliance Dashboard: Interactive web dashboard at feedoracle.io/ampel/ with entity selector, DORA/MiCA Ampel views, bridge workflow UI, audit trail viewer. Customer Console: Self-service portal at feedoracle.io/console/ with wallet balance, usage analytics, module breakdown, transaction history, API key management. Wallet & Credit System: Unit-based prepaid billing with 8 packages (monthly + annual), Stripe Checkout integration, auto-credit via webhook, module-level usage tracking. HTTP 402 Payment Protocol: Machine-readable payment-required responses for autonomous AI agents with recommended packages, topup endpoints, and natural language hints. KYA-Wallet Bridge: Automatic wallet creation on KYA registration with trust-level-based welcome units (500-5000).

Total: 100+ compliance MCP tools across 5+ specialized servers. 500+ production endpoints. 73 registered agents. 27,000+ MCP connections.

Changelog v4.2 (Mar 13, 2026): Know Your Agent (KYA): agent identity registration with trust scoring (0-100), 4 trust levels (UNVERIFIED/KNOWN/TRUSTED/CERTIFIED), trust-gated tool access. Audit Trail: tamper-proof SHA256 chain-linked decision logging, evidence snapshots, chain integrity verification. Evidence Lifecycle: 6 artifact states (CURRENT/STALE/CORRECTED/SUPERSEDED/DISPUTED/RETRACTED), auto-correction, Dispute-SLA (4h acknowledge, 24h classify, 5 business days resolve). Unit-Based Billing: Stripe Meter, Light/Medium/Heavy tool weights, Starter EUR49/Growth EUR199/Pro EUR499/Enterprise EUR1499, Annual packages with 2x bonus, Stripe Checkout. Trust Policy v1.0: 13-section formal evidence governance document. 5 new MCP tools: kya_register, kya_status, audit_log, audit_query, audit_verify. **v4.3 (Mar 19, 2026):** 6 new Compliance tools: zk_solvency_proof, zk_solvency_verify, zk_solvency_attestation (ZK-proof reserve ratio, Polygon on-chain), metals_gold, metals_silver, metals_prices (live XAU/XAG ES256K-signed, MiCA Art.36). Total: 33 Compliance tools, 100+ across 5+ servers. Human-readable summary field in every response. Downstream agentic liability clause (EU AI Act Art. 14). 8 reference workflows incl. failure-path scenario.
Changelog v4.0 (Mar 5, 2026): Enterprise Trust Layer (8 components): JWS Signing (RFC 7515), Versioned Schemas (8 JSON Schemas), Evidence Registry (Compliance Transparency Log), Evidence SLA Layer, Agent Trust Management, Streaming Evidence (SSE), Deterministic Replay, Zero-Trust Validation SDK. Updated to 5+ MCP Servers (100+ tools): Compliance MCP (33), DORA AmpelOracle (22), MiCAOracle (20), AMLOracle (12), Macro MCP (13), Risk MCP (13). Architecture expanded to 8 layers. Pricing updated. 500+ endpoints. 15 new Trust Layer endpoints.
Changelog v3.3 (Mar 4, 2026): Added MODULE 5 —□□ AMLR Digital Asset Screening (EU 2024/1624). 8 sections, 134 fields, 14 stablecoins, risk score composition, signed PDF reports. Updated architecture to 5 modules. Verified Reports updated to 6 types. 500+ endpoints.
Changelog v3.2 (Feb 24, 2026): Architecture updated to CORE + 4 Modules + MCP. Added Verified Reports, Payment Infrastructure, A2A Monetization, Report Verification to Delivered. New "In Progress" roadmap section. Regulatory source citations added to Legal. Proof wording standardized (XRPL live, Polygon deployed). SOC 2 status updated to Audit-ready. Version synchronized across all formats.

Table of Contents

1. Product Overview

FeedOracle is evidence-grade data infrastructure for regulated workflows. The platform delivers multi-source risk intelligence for tokenized assets — with every API response cryptographically signed and anchored on-chain for auditability.

The platform is built on a **one core engine, five compliance modules** + **AI agent layer** architecture:

⚡ **CORE: RWA Risk Oracle**

📋 **MODULE 1: MiCA Regulatory Evidence**

🌱 **MODULE 2: Carbon & ESG Context**

🛡️ **MODULE 3: DORA Compliance**

📊 **MODULE 4: CSRD/ESRS Reporting**

🔍 **MODULE 5: AMLR Screening**

🤖 **MCP: AI Agent Access (100+ Tools)**

COMPONENT	ROLE	PURPOSE	KEY OUTPUT
RWA Risk Oracle	⚡ **CORE**	Multi-dimensional risk scoring for 61+ tokenized asset protocols across 9 vectors	Composite risk scores, anomaly detection, yield analysis

MiCA Regulatory Evidence	📋 Module 1	Structured compliance metadata for EU MiCA/ DORA workflows	Legal state analysis, jurisdiction mapping, identifier registry
Carbon & ESG Context	🌱 Module 2	Per-chain carbon footprint with ISO 14040 methodology	CO_2 metrics, green scores, CSRD-ready ESG output
DORA Compliance	🛡️ Module 3	Digital Operational Resilience Act evidence infrastructure with AmpelOracle (22 tools), 12 execution oracles, and Oracle Event Bus	ICT incident reporting, vendor risk assessment, business continuity evidence
CSRD/ESRS Reporting	📊 Module 4	Corporate Sustainability Reporting with 5 dedicated APIs	EU Taxonomy, materiality, emissions, social metrics, governance
MCP Server	🤖 Agent Layer	Model Context Protocol with 100+ tools across 5+ servers + Enterprise Trust Layer (JWS, Registry, SLA, Streaming, Agent Trust, Replay, SDK)	Pre-trade compliance, MiCA status, custody risk, evidence generation, live dashboard (feedoracle.io/ ampel/), customer console (feedoracle.io/console/)

All components share a common evidence layer: ECDSA-signed responses, SHA-256 hashing, Evidence Pack Manifests (EPM), and on-chain anchoring (XRPL live per-transaction, Polygon contract deployed). The platform runs **500+ production endpoints** across 5 compliance modules.

Core Principles

- **Data infrastructure, not advisory.** FeedOracle delivers structured data and verifiable evidence artifacts. Compliance decisions remain with the institution and its qualified advisors.
- **Source transparency.** Every data point traces back to documented institutional sources with published methodology.
- **Cryptographic verifiability.** Every API response is ECDSA-signed with publicly discoverable keys (JWKS).
- **EU-based operations.** Primary infrastructure hosted in Germany. Subprocessor register available.

FeedOracle is an independent infrastructure provider with no formal partnership, endorsement, or affiliation with any blockchain foundation, regulatory body, or data source referenced in this document.

2. Executive Summary

Regulated institutions entering tokenized assets face a data gap. MiCA introduces disclosure requirements for CASPs (transitional period ends 1 July 2026; stablecoin rules already in force). DORA raises expectations around operational traceability. CSRD/ESRS requires ESG reporting. Yet no infrastructure delivers risk intelligence for RWA protocols with the evidence trail that institutional workflows demand.

FeedOracle closes this gap with three capabilities:

1. **Multi-source risk scoring** — 61 RWA protocols and 105+ stablecoins scored across 9 independent risk dimensions from 5 data sources, enriched with FRED macro economic benchmarks.
2. **Regulatory evidence** — MiCA compliance for 105+ stablecoins (Significant Issuer detection, reserve drift monitoring, interest scanning), DORA operational resilience (incident reporting, vendor risk, business continuity), CSRD/ESRS reporting (5 APIs), ISO 20022 validation — structured for audit-ready workflows.
3. **Delivery proof** — Every data point wrapped in ECDSA-signed Evidence Packs with SHA-256 hashing, JWKS-verifiable signatures, and on-chain anchoring (XRPL live, Polygon deployed).
4. **AI agent access** — MCP Servers with 100+ compliance tools for Claude, Cursor, and autonomous agents. Live compliance dashboard at feedoracle.io/ampel/ with entity management and real-time assessment.

Enterprise Trust Layer — 14 trust primitives: JWS signing (RFC 7515), versioned schemas, evidence registry (CT-style), SLA quality signals, agent trust management, streaming evidence (SSE), deterministic replay, and zero-trust validation SDK. Live proof at feedoracle.io/trust. The result: institutions can integrate verifiable risk data into their existing compliance workflows without building the data infrastructure themselves.

3. The Problem

The Disclosure Gap

Financial institutions evaluating tokenized assets need risk data that meets institutional standards. Current blockchain data providers focus on DeFi price feeds and analytics — not on the structured, evidence-grade data that compliance teams require. MiCA's CASP transitional deadline (1 July 2026) creates immediate urgency for CASPs, exchanges, and DeFi protocols.

Missing Audit Trail

When auditors ask "Can you prove what data you used for this decision?", traditional API calls leave no verifiable trail. Screenshots and PDF attestation reports are not machine-readable, not verifiable, and not auditable. Institutions need cryptographic evidence of data delivery.

Procurement Friction

Enterprise procurement requires: EU data residency documentation, incident management procedures, subprocessor registers, SLO evidence, and exit strategies. Few blockchain data providers deliver this level of operational documentation.

Fragmented Intelligence

Assessing a single RWA protocol today requires pulling data from 5+ independent sources (on-chain metrics, market data, macro indicators, regulatory databases, ESG scores), normalizing it, and creating a defensible risk view. This is expensive and error-prone when done manually.

4. The Solution

PROBLEM	FEEDORACLE SOLUTION
No unified RWA risk view	61 RWA protocols & 105+ stablecoins scored across 9 risk dimensions from 5 data sources, enriched with FRED macro benchmarks

Missing audit trail	ECDSA-signed Evidence Packs with SHA-256 hashing, JWKS verification, and on-chain anchoring (XRPL live, Polygon deployed)
Fragmented regulatory data	MiCA compliance (105+ stablecoins, Art. 44 Significant Issuer, Art. 25 reserve monitoring), DORA (3 APIs), CSRD (5 APIs), ISO 20022 — unified evidence infrastructure
No ESG/carbon data per chain	Per-chain carbon footprint (50+ networks), ISO 14040 methodology, CSRD-ready output
No AML/CFT screening for digital assets	AMLR Module (EU 2024/1624) —□□ per-token screening with issuer ID, sanctions, EDD, Travel Rule, reserves, risk score composition. 14 stablecoins, bank-standard action codes, signed PDF reports
Enterprise procurement barriers	EU-hosted infrastructure, documented controls, subprocessor register, SLO evidence
No machine-readable compliance	REST APIs with consistent JSON schemas, OpenAPI 3.1 specification, configurable policy outputs

5. CORE — RWA Risk Oracle

The RWA Risk Oracle is FeedOracle's core product. It provides real-time multi-dimensional risk scoring for tokenized asset protocols, designed for institutional portfolio assessment, due diligence, and ongoing monitoring.

Coverage

The oracle currently scores **61 RWA protocols** and monitors **105+ stablecoins** across **12 jurisdictions** from **5 independent data sources**. Coverage spans regulated stablecoins, tokenized treasuries, private credit, real estate, commodities, equities, insurance, and XRPL native assets.

> **Methodology note:** TVL figures are sourced from DeFiLlama and reflect aggregate protocol TVL at time of data refresh (daily). Full methodology available at docs/methodology.

9 Risk Dimensions

Every protocol is scored across 9 independent dimensions, each sourced from verified institutional data providers:

📊 TVL Weight

Protocol scale & concentration

Source: DeFiLlama

📈 Yield Spread

Deviation vs. T-Bill benchmark

Source: FRED

🔀 Diversification

Chain & asset mix (HHI analysis)

Source: On-chain

⏱️ Maturity

Protocol age & track record

Source: On-chain

⚖️ Regulatory

Jurisdiction risk flags

Source: Regulatory registers

🏛️ Institutional Backing

Known institutional participants

Source: Public disclosures

💧 DEX Liquidity

Secondary market depth

Source: GeckoTerminal

⛓️ On-Chain Activity

Transaction patterns & health

Source: Etherscan, Ankr

🌱 ESG / Carbon

Sustainability indicators

Source: CCRI, EMBER

Scoring Methodology

Each dimension produces a normalized score (0–100). The composite risk score is a weighted aggregate with configurable weights per institutional use case. Default weights are published in the API documentation. Scores are deterministic: the same input data always produces the same output score.

Anomaly Detection

Z-score outlier analysis flags unusual protocol behavior across all dimensions. Configurable alert thresholds enable automated monitoring and early warning signals for portfolio risk management.

Macro Economic Enrichment

The CORE engine is enriched with macro economic data from FRED (Federal Reserve Economic Data), ECB, and World Bank. This enables yield-spread analysis against T-Bill benchmarks, credit spread monitoring, recession probability indicators, and economic health indices — providing institutional context for RWA risk assessment.

Key Endpoints

ENDPOINT	METHOD	DESCRIPTION
/v1/rwa/risk	GET	All scored protocols with composite scores
/v1/rwa/risk/{slug}	GET	Detailed risk breakdown for single protocol
/v1/rwa/risk/{slug}/pdf	GET	PDF evidence report
/v1/rwa/market	GET	Aggregate market view
/v1/macro/indicators	GET	Macro economic enrichment

6. MODULE 1 — MiCA Regulatory Evidence

Module 1 provides the most comprehensive MiCA compliance infrastructure available. It covers **105+ stablecoins** with real-time peg monitoring, reserve drift detection, significant issuer classification (Art. 44), interest prohibition scanning (Art. 23/52), document compliance verification (Art. 29/30/55), and ESMA register mirroring — plus legal state analysis, jurisdiction mapping, and CCI scoring. All outputs are Evidence Pack-wrapped for audit trails.

> **Regulatory context:** MiCA CASP transitional period ends 1 July 2026 (stablecoin rules already in force since June 2024). Exchanges, CASPs, and DeFi protocols handling EU-regulated crypto-assets need machine-readable compliance data. Module 1 provides the evidence artifacts.

6.1 Legal State Analysis

Contract Powers & Ownership

Source: On-chain governance analysis, public disclosures

Analysis of smart contract governance structures, ownership patterns, upgrade mechanisms, and dispute resolution paths. Returns structured metadata for regulatory due diligence.

Endpoint: GET /v1/rwa/legal-state/{slug}

6.2 Jurisdiction Mapping

12 Jurisdictions + Sanctions Screening

Coverage: 12 jurisdictions · OFAC, EU, UN sanctions lists

Maps RWA protocols to their operating jurisdictions and screens against OFAC, EU, and UN sanctions lists. Returns jurisdiction risk flags, regulatory regime classification, and cross-border compliance indicators.

Endpoint: GET /v1/rwa/compliance/{slug}

6.3 Identifier Registry

LEI, ISIN/CUSIP, Custody Details

Source: Regulatory registers, public filings

Registry of institutional identifiers for RWA protocols: Legal Entity Identifiers (LEI), ISIN/CUSIP mappings, custody provider details, and issuer metadata. Machine-readable for compliance system integration.

Endpoint: GET /v1/rwa/registry/{slug}

6.4 Stablecoin Compliance

MiCA Classification Engine (105+ Stablecoins)

Coverage: USDC, EURC, RLUSD, USDT, DAI · Update: On-demand (5-min cache)

Configurable policy classification based on EU regulatory register data (ECB, ESMA CASP Register, NY DFS). Returns ACCEPTED/REJECTED status with machine-readable reason codes and configurable policy modes (strict, moderate, relaxed).

Endpoint: GET /v3/stablecoin/mica/{symbol}

CCI Score (Crypto Compliance Index)

Scale: 0–100 with letter grades (A+ to F)

Composite regulatory compliance score. Weighted methodology: MiCA status (30%), reserve transparency (25%), jurisdiction risk (20%), audit frequency (15%), operational history (10%).

Endpoints: GET /v3/cci/{symbol} · GET /v3/cci/ranking

6.5 Operational Resilience Tools

Circuit Breaker Detection

Triggers: Peg deviation, volume anomaly, liquidity drop, oracle failure

Automated halt detection for DORA operational resilience. Returns ACTIVE/TRIGGERED status with severity level and signed Evidence Pack.

Endpoint: GET /v3/circuit-breaker/status

ISO 20022 Payment Validation

Engine: xmllint (libxml2) · Supported: pain.001.001.09

XML payment message validation against official ISO 20022 XSD schemas. Returns Evidence Pack with validation decision, error details, and SHA-256 hash.

Endpoint: POST /v3/iso20022/validate

6.6 Stablecoin Peg Monitor (105+ Tokens)

Real-Time Peg Deviation Tracking

Coverage: 105+ stablecoins · Update: Real-time · Source: Multi-exchange aggregation

Continuous peg health monitoring across 105+ stablecoins. Tracks deviation from peg, trading volume anomalies, and liquidity depth. Returns severity classification (STABLE / WARNING / CRITICAL / DEPEG) with configurable alert thresholds for automated monitoring.

Endpoints: GET /v3/stablecoin/mica/{symbol} · GET /api/v1/feeds/stablecoin · GET /v1/peg/status/{symbol}

6.7 Significant Issuer Detection (Art. 44)

MiCA Art. 44 Classification Engine

Criteria: €5B market cap, 10M holders, 2.5M daily transactions

Automated screening against MiCA Art. 44 significant issuer thresholds. Monitors market capitalization, holder count, and transaction volume to detect when stablecoins cross regulatory significance boundaries. Returns SIGNIFICANT / NON-SIGNIFICANT classification with evidence trail.

Endpoint: GET /v1/mica/significant-issuer/{symbol}

6.8 Reserve Drift Monitor (Art. 25)

Reserve Composition Deviation Detection

Monitoring: Continuous · Threshold: Configurable drift tolerance

Tracks reserve asset composition changes over time. Detects when reserve backing deviates from declared composition — critical for MiCA Art. 25 compliance (reserve of assets requirements). Flags unauthorized asset substitutions, concentration shifts, and quality downgrades.

Endpoint: GET /v1/mica/reserve-drift/{symbol}

6.9 Interest Prohibition Scanner (Art. 23/52)

MiCA Art. 23 & 52 Yield Detection

Scope: All monitored stablecoins · Regulation: EMT Art. 23, ART Art. 52

Scans for prohibited interest or yield mechanisms on e-money tokens (EMT) and asset-referenced tokens (ART). MiCA explicitly prohibits granting interest to token holders. This scanner detects staking yields, rebasing mechanisms, and distribution schemes that could violate Art. 23/52.

Endpoint: GET /v1/mica/interest-scan/{symbol}

6.10 Document Compliance Monitor (Art. 29/30/55)

White Paper & Disclosure Verification

Articles: Art. 29 (EMT white paper), Art. 30 (marketing), Art. 55 (ART white paper)

Monitors whether token issuers maintain compliant white papers and marketing communications as required by MiCA. Checks for mandatory disclosure elements, publication status, and update frequency. Returns compliance status with specific article references for gap remediation.

Endpoint: GET /v1/mica/document-compliance/{symbol}

6.11 ESMA Register Mirror

EU Regulatory Register API

Source: ESMA · Update: Daily sync

Machine-readable mirror of the ESMA register for authorized CASPs and licensed token issuers. Enables automated verification of regulatory status against the official EU register. Cross-references with internal MiCA classification for comprehensive compliance checks.

Endpoint: GET /v1/esma/register · GET /v1/esma/register/{entity}

7. MODULE 2 — Carbon & ESG Context

Module 2 provides per-chain carbon footprint data and sustainability indicators with ISO 14040 methodology. Designed for MiCA Art. 66 sustainability disclosures and CSRD/ESRS reporting requirements.

7.1 Chain-Level Carbon Scoring

Carbon Footprint Feed

Update: Daily · Coverage: 50+ blockchain networks · Source: CCRI, EMBER, Climatiq

Blockchain energy and emissions data: carbon per transaction, energy per transaction, consensus mechanism classification, and green score (0–100). Confidence scoring reflects source coverage and data freshness.

Endpoints: GET /api/v1/feeds/carbon/chains · GET /api/v1/feeds/carbon/{chain}

7.2 Data Provenance

VeChain ToolChain · ISO 14040/14044

Methodology: ISO 14040 Life Cycle Assessment framework

Carbon data follows ISO 14040/14044 Life Cycle Assessment methodology. Source provenance is documented per network, with cross-validation against multiple emission factor databases.

7.3 CSRD-Ready Output

ESG Scores for Portfolio Reporting

Frameworks: CSRD/ESRS, MiCA Art. 66

ESG scores integrated directly into RWA risk reports via the `.esg_carbon` field. Designed for inclusion in CSRD sustainability reports and MiCA Art. 66 CASP disclosures.

Endpoint: `GET /v1/rwa/esg/{slug}`

7.4 Grid Intensity

Regional Electricity Carbon Data

Update: 30 min · Source: UK Carbon Intensity API, EMBER, Climatiq

Regional electricity carbon intensity data for energy-aware applications and validator location analysis.

Endpoint: `GET /api/v1/grid/intensity`

8. MODULE 3 — DORA Compliance

Module 3 provides evidence infrastructure for the Digital Operational Resilience Act (EU 2022/2554). DORA requires financial entities to maintain ICT risk management frameworks, report incidents, test resilience, and manage third-party ICT risks. FeedOracle delivers machine-readable evidence artifacts for these requirements.

> **Regulatory context:** DORA applies from January 17, 2025 to all EU-regulated financial entities including banks, insurers, investment firms, and their critical ICT third-party service providers. German insurance companies are actively building DORA compliance teams.

8.1 ICT Incident Reporting

Structured Incident Evidence Packs

Standard: DORA Art. 19 · Format: Machine-readable JSON with EPM wrapping

Generates structured incident report artifacts for major ICT-related incidents as required by DORA Art. 19. Captures incident classification, impact assessment, timeline, root cause analysis, and remediation steps — all wrapped in signed Evidence Packs for regulatory submission.

Endpoint: GET /v1/dora/incident-report

8.2 Third-Party Vendor Risk

ICT Third-Party Risk Assessment

Standard: DORA Art. 28-30 · Scope: Critical ICT service providers

Risk assessment framework for ICT third-party service providers. Evaluates concentration risk, exit strategies, subprocessor chains, data residency, and service level monitoring — aligned with DORA Art. 28-30 requirements for managing ICT third-party risk.

Endpoint: GET /v1/dora/vendor-risk

8.3 Business Continuity Evidence

Operational Resilience Documentation

Standard: DORA Art. 11-12 · Output: Signed evidence artifacts

Generates business continuity and disaster recovery evidence artifacts. Documents RPO/RTO targets, backup verification, failover testing results, and recovery procedures — structured for DORA Art. 11-12 ICT business continuity management requirements.

Endpoint: GET /v1/dora/business-continuity

9. MODULE 4 — CSRD/ESRS Reporting

Module 4 provides 5 dedicated APIs for Corporate Sustainability Reporting Directive (CSRD) requirements under the European Sustainability Reporting Standards (ESRS). Designed for financial entities reporting on digital asset portfolios and blockchain infrastructure sustainability.

9.1 EU Taxonomy Alignment

Activity Classification & Alignment Scoring

Framework: EU Taxonomy Regulation (2020/852) · ESRS E1-E5

Classifies blockchain and RWA protocol activities against the EU Taxonomy. Returns alignment scores for climate mitigation, climate adaptation, and do-no-significant-harm (DNSH) criteria.

Endpoint: GET /v1/csrd/taxonomy

9.2 Materiality Assessment

Double Materiality Analysis

Standard: ESRS 1 §§ 37-58 · Output: Impact & financial materiality scores

Provides double materiality assessment data for digital asset activities: impact materiality (effects on people and environment) and financial materiality (sustainability risks to the entity). Structured for ESRS disclosure requirements.

Endpoint: GET /v1/csrd/materiality

9.3 Emissions Data

Scope 1/2/3 Emissions for Blockchain Infrastructure

Standard: ESRS E1 · Sources: CCRI, EMBER, Climatiq

Chain-level and protocol-level greenhouse gas emissions data. Provides Scope 1 (direct), Scope 2 (electricity), and Scope 3 (value chain) emissions estimates for blockchain infrastructure — formatted for CSRD reporting templates.

Endpoint: GET /v1/csrd/emissions

9.4 Social Metrics

Workforce & Community Impact

Standard: ESRS S1-S4

Social sustainability indicators for RWA protocol assessment: workforce diversity, community impact, human rights due diligence, and consumer protection metrics. Complements environmental data for comprehensive ESG reporting.

Endpoint: GET /v1/csrd/social

9.5 Governance

Governance & Risk Management

Standard: ESRS G1 · Scope: Protocol & issuer governance

Governance indicators for RWA protocols and token issuers: board composition, risk management frameworks, compliance structures, audit mechanisms, and whistleblower protections. Key input for CSRD governance disclosures.

Endpoint: GET /v1/csrd/governance

10. MODULE 5 — AMLR Digital Asset Screening NEW

Module 5 provides evidence infrastructure for the Anti-Money Laundering Regulation (EU 2024/1624, application date 2027-07-10). AMLR introduces direct obligations for crypto-asset service providers (CASPs), requiring customer due diligence, sanctions screening, and risk-based monitoring for digital asset transactions. FeedOracle delivers per-token screening evidence covering 8 regulatory sections with 134 structured fields.

10.1 Screening Architecture

Per-Token AMLR Compliance Screening

Regulation: EU 2024/1624 (AMLR) + AMLD6 + TFR (2023/1113) | 14 Stablecoins | v2.0 Schema

Each screening produces a structured evidence artifact covering 8 sections aligned with AMLR articles: Issuer Identification (Art. 16-18), Regulatory Status (Art. 19-20), Sanctions Screening (Art. 20, 29, 79), Enhanced Due Diligence (Art. 23-24, 28), Travel Rule (TFR Art. 14, 16, 19), Token Transparency (Art. 79), Reserve Assessment (Art. 23), and Overall Risk Assessment (Art. 8-13).

10.2 Risk Scoring

Transparent Risk Score Composition

Scale: 0-25 LOW | 26-55 MEDIUM | 56-80 HIGH | 81-100 CRITICAL

Risk scores are decomposed into weighted components (regulatory status, MiCA flags, reserve compliance, sanctions exposure) so compliance teams can answer the key question: "Why is the score X?" The system outputs 10 bank-standard action codes (PROCEED, VERIFY_REGULATORY_STATUS, REQUIRE_TRAVEL_RULE, ENHANCED_MONITORING, etc.) and CDD level recommendations (SIMPLIFIED, STANDARD, ENHANCED).

10.3 Sanctions & Due Diligence

Multi-List Screening with Manual Review Logic

Lists: EU Consolidated | OFAC SDN | UN Security Council | Per-version tracking

Sanctions screening covers issuer entities and jurisdictions against three major lists with per-list version tracking and staleness detection. Manual review triggers include: high-risk third country jurisdiction, PEP exposure, issuer not in ESMA register, and sanctions near-matches. Compliance capability assessment covers freeze/blacklist/law enforcement cooperation.

10.4 Evidence Outputs

API + Signed PDF Reports

Formats: JSON (API) | Signed PDF (6 pages) | XRPL-anchored | ECDSA-signed

Screening results available as JSON API responses or 6-page signed PDF reports with executive summary, risk bar visualization, score composition, data source listing, automation endpoints, and cryptographic provenance block. Reports include per-source data freshness tracking and coverage percentages.

Endpoints:

```
GET /api/v1/evidence/amlr/screening/{symbol} — Single token screening
POST /api/v1/evidence/amlr/batch — Batch screening (up to 50 tokens)
GET /api/v1/evidence/amlr/supported — List supported tokens
POST /reports/api/report/generate?type=amlr&symbol={token} — Generate signed PDF
```

11. MCP Servers & AI Agent Access — 100+ Compliance Tools Across 5+ Specialized Servers

FeedOracle operates **3 specialized MCP servers with 100+ tools total** for AI agent integration. Every MCP server supports SSE and Streamable HTTP transport. Every tool response is cryptographically signed with ECDSA (ES256K), includes SLA quality signals, and is logged in the Evidence Registry.

Server 1: Compliance MCP (Port 5250) — 33 Tools

Primary compliance tools for MiCA, DORA, RWA risk, and AI Gateway. SSE endpoint: `https://feedoracle.io/mcp/sse`

TOOL	DESCRIPTION	CATEGORY
compliance_preflight	Pre-trade regulatory check — PASS/WARN/BLOCK	Utility
mica_status	MiCA compliance status for any stablecoin	MiCA
mica_full_pack	Complete 12-article MiCA compliance pack	MiCA
peg_deviation	Real-time peg health (Art. 35)	MiCA
significant_issuer	Art. 44 significant issuer classification	MiCA
interest_check	Art. 23/52 interest prohibition scan	MiCA
document_compliance	Art. 29/30/55 white paper verification	MiCA
reserve_quality	Art. 25/53 reserve quality assessment	MiCA
evidence_profile	Full evidence profile for any RWA protocol	RWA
custody_risk	Custody provider risk assessment	RWA
evidence_leaderboard	Protocol ranking across 61 RWA + 105+ stablecoins	RWA
macro_risk	Macro economic risk indicators (FRED/ECB)	Macro
generate_report	Generate signed PDF evidence reports (6 types)	Reports
ai_query	Natural language → signed evidence bundle	AI Gateway
evidence_bundle	Multi-framework evidence aggregation	AI Gateway
ai_explain	Grade explainability — Why B? What needs A?	AI Gateway
ai_provenance	Full data provenance graph	AI Gateway
market_liquidity	DEX liquidity depth (GeckoTerminal)	Market

rlusd_integrity	RLUSD reserve verification	Market
mica_market_overview	Market-wide MiCA compliance overview	Market
peg_history	Historical peg deviation data	Market
kya_register	Register agent identity — trust score + level	KYA
kya_status	Check agent trust level + tool access	KYA
audit_log	Log chain-linked decision with evidence refs	Audit
audit_query	Query agent decision history	Audit
audit_verify	Verify audit chain integrity (tamper detection)	Audit
ping	Server health check	Utility

Server 2: Macro Intelligence MCP (Port 5251) — 13 Tools

Dedicated macro economic intelligence from 86 FRED + 20 ECB data series. SSE endpoint: https://feedoracle.io/mcp/macro/sse

TOOL	DESCRIPTION
macro_regime	Regime classification: EXPANSION / SLOWDOWN / CONTRACTION / RECOVERY
fed_rates	Federal Reserve rates & policy outlook
ecb_rates	ECB rates & monetary policy signals
inflation	CPI, PCE, inflation expectations
yield_curve	Yield curve analysis & recession signals
labor_market	Employment, unemployment, claims data
housing	Housing starts, permits, prices
credit_spreads	IG/HY credit spreads, financial stress
commodities	Gold, oil, commodity signals
defi_macro_bundle	Macro + DeFi combined risk assessment
ecb_mica_reserve	ECB data for MiCA reserve environment
composite	Composite macro risk score (0–100)

ping	Server health check

Server 3: Stablecoin Risk MCP (Port 5252) — 13 Tools

Operational safety assessments for stablecoins using a proprietary 7-signal scoring model. SSE endpoint: https://feedoracle.io/mcp/risk/sse

Signals: peg stability (20%), liquidity depth (15%), mint/burn flows (15%), holder concentration (15%), custody counterparty risk (15%), redemption friction (10%), cross-chain exposure (10%). Verdict: **SAFE** (≥75) / **CAUTION** (50–74) / **AVOID** (<50). Coverage: 341 stablecoins tracked, 105 with full MiCA classification, 21 with deep analytics.

TOOL	DESCRIPTION
risk_assessment	Full 7-signal risk report with SAFE/CAUTION/AVOID verdict
peg_status	Current peg deviation analysis
peg_history	Historical peg stability data
supply_flow	Mint/burn flow analysis
holder_data	Holder concentration metrics
custody_data	Custody counterparty assessment
redemption_data	Redemption friction analysis
cross_chain_data	Cross-chain distribution & bridge exposure
leaderboard	Stablecoin risk ranking
compare	Head-to-head stablecoin risk comparison
supported_tokens	List all supported tokens
stablecoin_preflight	Pre-trade operational safety check
ping	Server health check

Agent Marketplace & Distribution

FeedOracle is live on the **MCP ecosystem** (Gnosis Chain) with 8 feedoracle-* tools. Service 2670 has processed **133+ verified deliveries** with 100% acceptance rate, ranking #17 of 3Available via feedoracle.io/ampel/ and feedoracle.io/console/.

Discovery

AI agents discover FeedOracle via 9 methods: `/llms.txt` (LLM-native), `/openapi.json` (500+ endpoints), `/.well-known/ai-plugin.json` , `/.well-known/jwks.json` , `/.well-known/mcp/server.json` , MCP SSE/HTTP transports, MCP directories registry, mcp.so directory, and uptime status at `uptime.feedoracle.io` .

12. Enterprise Trust Layer — 14 Components NEW

FeedOracle implements 14 enterprise-grade trust primitives that transform the platform from an evidence API into a verifiable infrastructure protocol — comparable to Certificate Transparency (Google), TLS, or DNS. All 14 components are live on production with independent verification at feedoracle.io/trust.

12.1 JWS Signing (RFC 7515)

Every evidence response contains a JSON Web Signature (compact serialization) using ES256K (secp256k1). Verifiable via JWKS at `/.well-known/jwks.json` . Dual-format: existing ECDSA signature fields remain for backward compatibility; new `jws{}` block added with `kid` , `alg` , `typ=evidence+jwt` , and content hash. Any banking system with a JWS library can verify FeedOracle evidence natively — no custom integration required.

12.2 Versioned Evidence Schemas

8 JSON Schemas (Draft 2020-12) covering all frameworks: evidence-envelope, mica, dora, rwa, macro, stablecoin-risk, amlr, sla. Registry at `GET /schemas/` . Every response contains `schema_ref` (e.g., `mica/v1`) linking to the exact schema used at generation time. When regulations change, old evidence packs remain validatable against their original schema version.

12.3 Evidence Registry (Compliance Transparency Log)

Append-only log of every evidence pack ever produced — inspired by Google Certificate Transparency. Public, filterable, auditable. `GET /evidence/registry` with filters: ?framework=mica&asset=USDC&from=2026-03-01&limit=50 . Each entry contains: pack_id, framework, asset, SHA-256 hash, timestamp, verify URL. Statistics at GET /evidence/registry/stats .

12.4 Evidence SLA Layer

Every API response includes an sla{} object with machine-readable quality signals: freshness_seconds (age of oldest data point), confidence (0.0–1.0 weighted score), per-source health details (status , latency_ms , age_seconds), staleness_flag , and tier-specific targets. Tiers: Free (60s freshness), Starter (60s), Pro (30s), Enterprise (15s).

12.5 Agent Trust Management

AI agents register (POST /ai/agent/register), receive ECDSA keys with 90-day lifecycle, and accumulate reputation (0–100). Tiers: EXEMPLARY (80+), TRUSTED (60–79), STANDARD (40–59), DEGRADED (20–39), RESTRICTED (<20). Key rotation via POST /ai/agent/{id}/rotate-key (+5 reputation bonus). Agent leaderboard at GET /ai/agent/leaderboard . Rate governance with per-agent limits and throttling.

12.6 Streaming Evidence (SSE)

Real-time Server-Sent Events at GET /evidence/stream . Events fire only on state change (not every poll cycle). Supported events: peg_deviation , regime_change , market_stress , reserve_alert , evidence_anchored , agent_registered . Filterable: ?events=peg_deviation,regime_change&assets=USDC,EURC . Every event is SHA-256 hashed. Thresholds: peg >0.5% WARNING, >2% CRITICAL, VIX >25 stress.

12.7 Deterministic Replay (Audit)

Every evidence pack is archived as an immutable gzip-compressed snapshot. GET /evidence/replay/{pack_id} reconstructs the exact evidence from the snapshot archive. The response contains hash_match: true — a cryptographic proof that the replayed evidence is byte-identical to the original. Banks can perform audit replays years later with full reproducibility.

12.8 Zero-Trust Validation SDK

Client-side evidence verification — no trust in transport layer required. Python package: pip install feedoracle-verify . 7 independent checks: content hash validation, JWS signature (ES256K via JWKS), key revocation status, timestamp freshness, schema reference, SLA confidence, SLA freshness. Self-test at GET /verify/self-test . TypeScript, Rust, and Go packages planned.

12.9 Trust Metadata

Every evidence response contains a trust{} summary object: signature_present , signature_algorithm , content_hash_present , schema_valid , registry_logged , replayable , sla_confidence , sla_freshness_met , verify_url , sdk . One glance shows the client: all trust components active.

Trust Layer Endpoints

ENDPOINT	METHOD	DESCRIPTION
/schemas/	GET	Schema registry index (8 schemas)
/schemas/v1/{framework}	GET	Individual JSON Schema
/evidence/registry	GET	Compliance Transparency Log (paginated, 8 filters)
/evidence/registry/{pack_id}	GET	Single evidence pack detail
/evidence/registry/stats	GET	Registry statistics
/evidence/stream	GET	SSE real-time events
/evidence/stream/status	GET	Stream poller status
/evidence/replay/{pack_id}	GET	Deterministic audit replay
/evidence/snapshots/stats	GET	Snapshot archive statistics
/ai/agent/register	POST	Register AI agent
/ai/agent/{id}/trust	GET	Agent trust status
/ai/agent/{id}/rotate-key	POST	Key rotation
/ai/agent/leaderboard	GET	Agent reputation ranking
/verify/sdk	GET	SDK installation info
/verify/self-test	GET	Live 7-check self-test

12.10 Know Your Agent (KYA) v4.2

Agent identity registration with trust scoring. Every AI agent that interacts with FeedOracle can register its identity, purpose, organization, and jurisdiction. FeedOracle computes a trust score (0–100) from 6 dimensions: registration completeness, account age, usage consistency, compliance rate, behavioral signals, and manual verification. Trust levels: UNVERIFIED (0–24), KNOWN (25–54), TRUSTED (55–79), CERTIFIED (80+). Sensitive tools (evidence_bundle, generate_report, mica_full_pack, ai_provenance, ai_explain, mica_market_overview) require minimum trust levels. MCP tools: `kya_register` , `kya_status` . REST: `POST /api/billing/kya/register` , `GET /api/billing/kya/profile` . See KYA documentation.

12.11 Audit Trail & Decision Logging v4.2

Tamper-proof, SHA256 chain-linked decision logging for autonomous agents. When an agent makes a decision, it logs the decision, reasoning, action taken, and references to prior tool calls via their `request_id` . Each audit entry is chain-linked to its predecessor: `chain_hash = SHA256(prev_chain_hash + trail_id + evidence_hash + decision + timestamp)` . If any entry is modified, all subsequent hashes break. Regulators and auditors can verify the entire chain with one call. Evidence snapshots are preserved at decision time for temporal replay. EU AI Act Art. 14, MiCA Art. 83, DORA Art. 11 compliant. MCP tools: `audit_log` (3 units), `audit_query` (1 unit), `audit_verify` (1 unit). See Audit Trail documentation.

12.12 Evidence Lifecycle v4.2

Every evidence artifact exists in one of 6 defined states: CURRENT, STALE, CORRECTED, SUPERSEDED, DISPUTED, RETRACTED. Auto-correction: when the same tool with the same input produces newer data, the prior artifact transitions to CORRECTED with a `corrected_by` link. Disputed artifacts remain queryable throughout the dispute process. No silent deletions ever occur. All state transitions are logged in `evidence_state_log` with timestamps and reasons. Temporal queries supported: retrieve what was known at any past point in time. REST: `GET /api/billing/evidence/artifact` , `GET /api/billing/evidence/stats` .

12.13 Evidence Trust Policy v1.0 v4.2

Formal 13-section governance document defining what FeedOracle Evidence is: evidence classes, source-of-truth hierarchy, normalization rules, signing rules, lifecycle states, freshness targets, degradation model, correction/dispute/retraction process (with Dispute-SLA: 4h acknowledge, 24h classify, 5 business days resolve), liability boundary (including downstream agentic execution clause citing EU AI Act Art. 14), governance roles, acceptance targets, and verification without trust. See Trust Policy v1.0.

12.14 Unit-Based Billing v4.2

Usage-based pricing: every tool call costs units based on computational complexity. Light tools (1 unit): single-source lookups. Medium tools (3 units): multi-source aggregation. Heavy tools (10 units): full pipeline with PDF generation or AI inference. Free tier: 300 units/day. Pro (€49/mo): 15,000 units/month included, €0.005/unit overage. Agent (€299/mo): 150,000 units/month included, €0.003/unit overage. Overage auto-billed via Stripe Billing Meter. `meta.units_consumed` in every response. Public endpoint: `GET /api/billing/weights` . See Billing documentation.

13. Architecture

FeedOracle operates as a layered data infrastructure:

```
Layer 8: SDK LAYER       Zero-Trust Validation SDK (Python, TS/Rust/Go planned)
Layer 7: STREAMING       Real-time SSE Events (peg, regime, stress, breach alerts)
Layer 6: AGENT LAYER     Agent Trust Management (registration, key rotation, reputation)
Layer 5: TRUST LAYER     JWS Signing, Evidence Registry, Schemas, SLA, Replay
Layer 4: AI GATEWAY      100+ MCP tools (5+ servers), MCP directories, natural language queries
Layer 3: COMPLIANCE      MiCA, DORA, AMLR, RWA, Macro (500+ endpoints)
Layer 2: EVIDENCE        ECDSA Signatures, SHA-256, Blockchain Anchoring (Polygon, XRPL)
Layer 1: DATA SOURCES    DeFiLlama, GeckoTerminal, FRED, ECB, ESMA, Ankr, CCRI
```

The architecture maps to the product modules:

LAYER	CORE (RWA RISK)	MOD 1 (MICA)	MOD 2 (CARBON/ESG)	MOD 3 (DORA)	MOD 4 (CSRD)
Agent					

	5+ MCP Servers (100+ tools) · KYA Agent Identity · Audit Trail · Evidence Lifecycle · MCP directories (Gnosis, 8 tools) · Agent Trust · Evidence Registry · Streaming SSE · Validation SDK · llms.txt				
Evidence	ECDSA + SHA-256 + JWKS	Compliance EPMs	Sustainability EPMs	Incident EPMs	Reporting EPMs
Processing	9-vector scoring, anomaly detection	105+ stablecoins, Art. 44, reserve drift, CCI	Carbon normalization, green scoring	ICT risk assessment, vendor scoring	Double materiality, taxonomy alignment
Data Sources	DeFiLlama, FRED, Ankr, GeckoTerminal	ECB, ESMA, OFAC, NY DFS	CCRI, EMBER, Climatiq	Internal + client data	EMBER, CCRI, Climatiq, World Bank

API Specification

All endpoints are documented in OpenAPI 3.1 format with Swagger UI. Machine-readable discovery via llms.txt for AI agent integration.

14. Attestation & Evidence

Disclosure Attestation Protocol (DAP)

DAP creates cryptographic proof of data delivery. The protocol hashes API response payloads with SHA-256, aggregates hashes into Merkle trees, and anchors roots to public blockchains.

How It Works

1. **Hash:** API response payload hashed with SHA-256
2. **Sign:** ECDSA signature (ES256K) with JWKS-discoverable public key
3. **Aggregate:** Hashes combined into Merkle trees (batched)
4. **Anchor:** Hash anchored on XRPL (memo field). Gnosis Chain live via XRPL memo anchor.
5. **Verify:** Anyone can verify delivery timestamp via block explorer or API

What Gets Hashed

INCLUDED IN HASH	EXCLUDED FROM HASH
API response body (JSON)	HTTP headers
Timestamp (ISO 8601)	Client IP address
Endpoint path	API key (hashed separately)
Schema version	Request parameters

Scope of Proof

DAP proves that specific data was delivered via API at or before the anchored block timestamp, and that post-delivery modification becomes detectable via hash mismatch. DAP is a delivery evidence mechanism — it does not verify upstream data correctness or constitute regulatory approval.

Anchoring Schedule

MODE	FREQUENCY	AVAILABILITY
Daily Batch	Once per 24h	All tiers
Hourly Batch	Once per hour	Enterprise
Event-Driven	On significant data change	Enterprise

Evidence Pack Manifest (EPM v1.0)

EPM extends DAP with a standardized manifest schema using DSSE-style envelopes and RFC 8785 deterministic hashing. Supports jurisdiction-aware metadata (EU/UK/US/Global). Every Evidence Pack contains: the signed data payload, ECDSA signature (ES256K), timestamp, schema version, source attribution, and Merkle proof. Public keys are discoverable via JWKS endpoint.

DAP API Endpoints

ENDPOINT	METHOD	DESCRIPTION
/api/v2/attestation/test	GET	Integrity test on all sources
/api/v2/attestation/sources	GET	List attested sources and TLS policies
/api/v2/attestation/anchor/latest	GET	Blockchain anchor status (XRPL)

/api/v2/attestation/merkle	GET	Current Merkle root
/api/v2/attestation/verify/{hash}	GET	Verify payload hash
/api/v1/epm/verify	POST	Verify DSSE-wrapped EPM manifest

15. Data Quality

Pipeline

```
Sources → Ingestion → Validation → Normalization → Quality Scoring → API → DAP
        (scheduled)  (schema)    (units/format)   (confidence)     (REST) (anchor)
```

Quality Gates

GATE	CHECK	ON FAILURE
Schema Validation	Response matches expected structure	Reject, log, use fallback
Freshness Check	Timestamp within threshold	Flag as stale, reduce confidence
Range Validation	Values within expected bounds	Flag anomaly, manual review
Cross-Validation	Compare multiple sources	Use median, flag divergence

Confidence Scoring

Every data point includes a confidence score (0.0–1.0) reflecting source availability, data freshness, and cross-validation status. Confidence scores are heuristic indicators — not statistical confidence intervals.

Staleness Handling

When primary sources are unavailable, the system attempts configured secondary sources. If all sources are unavailable, the API returns the last known value with a `stale: true` flag and reduced confidence. The staleness threshold is configurable per feed (default: 2× normal update interval).

16. Security & Operations

Security Controls

Operational controls informed by ISO/IEC 27001:2022 principles:

CONTROL AREA	IMPLEMENTATION
Access Control	API key authentication (X-API-Key header), tiered rate limits
Cryptography	ECDSA ES256K signing, SHA-256 hashing, TLS 1.2+ enforced
Network Security	HSTS, CSP headers, firewall, Cloudflare DDoS protection
Backup	Daily encrypted, cross-server sync within EU
Logging	Request IDs, structured access logs, audit trail
Key Management	ECDSA key rotation, JWKS public key discovery

Service Level Objectives

METRIC	TARGET	MEASUREMENT
API Availability	99.5%	Monthly uptime
Response Time (p95)	<500ms	95th percentile latency
Data Freshness	≤15 minutes	Real-time feeds
Attestation Anchor	≤24 hours	Time to on-chain

Data Residency

COMPONENT	LOCATION	NOTES
Primary Infrastructure	Germany (netcup DE)	API servers, databases
Backup Storage	EU	Encrypted, cross-server sync
CDN/Edge	Global (Cloudflare)	Request routing, caching headers
Blockchain Anchors	XRPL (live), Gnosis (live via MCP directories)	Public blockchains

Recovery Targets

METRIC	TARGET
RPO (Recovery Point Objective)	≤24 hours
RTO (Recovery Time Objective)	≤4 hours

Subprocessors

SUBPROCESSOR	FUNCTION	LOCATION
netcup GmbH	Infrastructure hosting	Germany
Cloudflare Inc.	CDN, DDoS protection, DNS	US/EU (edge)
XRPL Ledger	Blockchain anchoring	Decentralized
Ripple (XRPL)	Blockchain anchoring	Decentralized

17. Target Users

Banks & Insurers

Portfolio risk assessment for RWA allocations. DORA-supporting evidence packs for third-party vendor risk documentation. Signed evidence artifacts for audit trails.

Asset Managers

Due diligence on tokenized funds. Yield spread analysis vs. T-Bill benchmarks. Smart contract risk signals. ESG/sustainability data for portfolio reporting.

Exchanges & CASPs

MiCA Art. 66 sustainability disclosures for listed crypto-assets. Stablecoin classification and monitoring. Structured data for regulatory reporting ahead of the 1 July 2026 CASP transition deadline.

DeFi Protocols & DAOs

On-chain risk feeds for RWA integrations. DEX liquidity monitoring. Protocol-level risk scoring for governance decisions. Chainlink Functions integration for smart contract access.

RegTech Platforms

White-label compliance modules with API-first integration. OpenAPI spec, versioned schemas, bulk access for platform integration.

AI Agents & Autonomous Systems

Machine-readable API outputs with llms.txt discovery. Pay-per-call access model. DAP enables programmatic verification without human intervention.

18. Competitive Landscape

CATEGORY	FOCUS	EXAMPLES
Decentralized Oracles	DeFi price feeds, on-chain data delivery	Chainlink, Pyth, Redstone
Blockchain Indexers	Query layer for on-chain data	The Graph, Goldsky
Research & Analytics	Market intelligence, reporting	Messari, Dune, Nansen
Carbon Ratings	Sustainability assessments	CCRI, Digiconomist
AML/KYT	Transaction monitoring, wallet screening	Chainalysis, Elliptic, TRM Labs
AI Agent Tooling	MCP servers, agent marketplaces	Anthropic MCP Directory, Fetch.ai, Autonolas
Evidence Infrastructure	**Regulatory verification + signed evidence + risk scoring + AI agent access**	**FeedOracle**

Differentiation

- **Evidence-first:** Every API response is ECDSA-signed and on-chain anchored — not an optional add-on
- **CORE + Modules:** One risk engine with dedicated compliance (MiCA) and sustainability (ESG) modules
- **Configurable outputs:** Machine-readable policy signals with reason codes, not subjective scores
- **Multi-vertical:** MiCA, DORA, ISO 20022, Carbon/ESG — one Evidence Pack framework
- **Enterprise documentation:** Subprocessor register, incident procedures, procurement pack, exit strategy
- **AI-native:** 5+ MCP servers (100+ tools) + MCP ecosystem — compliance tools accessible to autonomous agents
- **Enterprise Trust Layer:** JWS signing (RFC 7515), evidence registry, deterministic replay, agent trust management, streaming events, versioned schemas, SLA quality signals, zero-trust validation SDK — no competitor offers this verification depth

19. Commercials

Pricing

TIER	PRICE	API CALLS	EVIDENCE PACKS
Free	€0/mo	100/day	—
Starter	$99/mo	5,000/day	50/mo
Pro	$299/mo	25,000/day	500/mo
Enterprise	Custom	Custom	Custom

Payment: Stripe (card) and USDC on Polygon accepted for all paid tiers.

Evidence Packs are request-based: each API call that generates a signed Evidence Pack Manifest counts as one verification. API calls without Evidence Pack generation are not counted against the verification limit.

Enterprise Package

- Custom SLA negotiation
- DORA Support Pack (ICT third-party risk documentation)
- Dedicated account manager
- Custom integration support and priority incident response
- Historical data access (Data Vault)

20. Roadmap

> **Forward-looking statement:** This roadmap contains planned initiatives. Actual results may differ. No commitment to delivery dates, features, or timelines.

Delivered (Q1 2026)

INITIATIVE	STATUS
RWA Risk Oracle — 61+ protocols, 9 risk vectors, 5 data sources	Live
MiCA Regulatory Evidence module (legal state, jurisdiction, registry)	Live
Carbon & ESG module (50+ networks, ISO 14040)	Live
Macro Intelligence (FRED + ECB enrichment)	Live
XRPL Anchoring	Live
Evidence Pack System with ECDSA signing + JWKS	Live
MiCA Stablecoin Classification (105+ stablecoins)	Live
CCI Score Engine (compliance ranking)	Live
ISO 20022 Payment Validation	Live
Circuit Breaker Detection (DORA resilience)	Live
Interactive Enterprise Demos (5 verticals)	Live
MiCA Deep Compliance: Significant Issuer (Art. 44), Reserve Drift (Art. 25), Interest Scanner (Art. 23/52), Document Compliance (Art. 29/30/55), ESMA Register	Live
DORA Compliance Module (incident reporting, vendor risk, business continuity)	Live

CSRD/ESRS Module (5 APIs: taxonomy, materiality, emissions, social, governance)	Live
5+ MCP Servers (100+ tools, SSE + Streamable HTTP)	Live
Know Your Agent (KYA) — Trust scoring, 4 levels, trust-gated tool access	Live
Audit Trail — Chain-linked decision logging, evidence snapshots, chain verification	Live
Evidence Lifecycle — 6 artifact states, auto-correction, Dispute-SLA	Live
Evidence Trust Policy v1.0 — 13-section formal governance document	Live
Unit-Based Billing — Stripe Meter, Light/Medium/Heavy weights, overage billing	Live
Anthropic MCP Directory — Submitted, crawling confirmed	Live
MCP ecosystem (Gnosis, 8 tools, 133+ deliveries)	Live
Autonomous Agent Self-Upgrade — USDC payment flow (M2M), agent detects quota, pays on Polygon, receives upgraded API key without human intervention. Production TX anchored on-chain.	Live
OAuth 2.0 / M2M Auth (RFC 6749) — authorization_code + client_credentials + refresh_token, /mcp/authorize, /mcp/token, /mcp/register, /mcp/revoke, 5 scopes for Enterprise MCP tier	Live
Stablecoin Peg Monitor (105+ tokens, real-time)	Live
Chainlink Functions integration (Polygon, Contract 0x7Ec0...23c)	Live
L2 Intelligence APIs (7 chains)	Live
Verified Reports System (6 types: RWA Risk, MiCA, DORA, Macro, CSRD, AMLR) with PDF generation, XRPL-anchored proof panel	Live
Payment Infrastructure — Stripe + USDC (Polygon), 4 tiers (Free $0 / Starter $99 / Pro $299 / Enterprise custom)	Live
Agent-to-Agent (A2A) Monetization — Pay-per-Call API (9 feeds, USDC on Polygon) + MCP directories (unit-based credit system)	Live
Report Verification Infrastructure — public /verify endpoints with SHA-256 + ECDSA signature validation	Live
Trust Center & Enterprise Procurement Pack (12 sections: security, SLOs, subprocessors, vulnerability disclosure, compliance mapping)	Live
AMLR Digital Asset Screening Module (EU 2024/1624) —□□ 8 sections, 134 fields, 14 stablecoins, risk score composition, sanctions screening (EU/OFAC/UN), bank-standard action codes, signed PDF reports	Live

Public Status Page — Uptime Kuma with 46 monitors across 8 service groups	Live
Stablecoin Risk MCP Server (13 tools, 7-signal scoring, SAFE/CAUTION/AVOID)	Live
Macro Intelligence MCP Server (13 tools, 86 FRED + 20 ECB series)	Live
Enterprise Trust Layer — JWS Signing (RFC 7515), Versioned Schemas (8), Evidence Registry, SLA Layer, Agent Trust Management, Streaming Evidence (SSE), Deterministic Replay, Zero-Trust Validation SDK	Live
Enterprise Trust Proof Page — feedoracle.io/trust with live 8/8 verification	Live
CSRD/ESRS Data API — chain footprint, ESRS E1, EU energy mix, EU ETS pricing	Live

In Progress (Q1–Q2 2026)

INITIATIVE	STATUS
Multi-chain Evidence Anchoring — XRPL (live), Polygon (contract deployed), Avalanche & Flare (grant applications submitted)	In Progress
SOC 2 Type II — Trust Center live, compliance framework mapping complete, security controls documented, SLO evidence via Uptime Kuma (46 monitors). Pre-audit documentation ready; formal audit pending funding	Audit-ready
Chainlink BUILD program participation	Application submitted
Avalanche infraBUIDL() grant — C-Chain deployment, RWA Risk Oracle, Evergreen Subnet integration	Application submitted
Flare Network grant — FTSOv2 Feed Value Provider, FDC compliance attestations, FAssets risk layer	Application ready
CSRD/ESRS template library — structured report templates for ESRS E1 disclosures	Data API live, templates in development

Planned (Subject to Change)

INITIATIVE	PRIORITY
ISO 20022 expansion (pacs.008, camt.053)	Medium
WebSocket real-time feeds (SSE + Streamable HTTP already live via MCP)	Medium
XRPL Grants program application (Spring 2026)	High

Avalanche Evergreen Subnet integration for institutional evidence delivery	Medium
Flare FTSOv2 Feed Value Provider for compliance/risk data feeds	Medium
Flare Data Connector (FDC) attestations for compliance events	Medium
Developer SDK (npm package) for multi-chain FeedOracle integration	High

Resources

Documentation	feedoracle.io/docs
API Reference	Full API Ref (238 Endpoints)
Trust & Security	Trust Documentation
System Status	Status Page
Enterprise	Enterprise Overview
Interactive Demos	RWA · Insurance · Carbon · Stablecoin · Payments

21. Legal & Disclaimers

FeedOracle provides data infrastructure and verifiable evidence artifacts. The platform does not provide financial, legal, or compliance advice. Compliance decisions remain with the institution and its qualified advisors.

Regulatory Sources

Regulatory timelines and classifications referenced in this document are based on the following official sources:

- **MiCA:** Regulation (EU) 2023/1114 — EUR-Lex. Stablecoin provisions (Title III/IV) in force since 30 June 2024. CASP transitional period ends 1 July 2026 per Art. 143(3).
- **DORA:** Regulation (EU) 2022/2554 — EUR-Lex. Applicable since 17 January 2025.
- **CSRD:** Directive (EU) 2022/2464 — EUR-Lex. ESRS delegated acts applicable from 1 January 2024.
- **ESMA CASP Register:** esma.europa.eu

- **BaFin MiCA Guidance:** bafin.de

Forward-Looking Statements

This document contains forward-looking statements regarding planned features, roadmap items, and business strategy. Actual results may differ materially. No commitment to specific delivery dates, features, or timelines is expressed or implied.

No Endorsement

FeedOracle is an independent infrastructure provider. References to blockchain networks (XRPL, Polygon, Gnosis, Chainlink), regulatory bodies (ESMA, BaFin, ECB), or data sources (FRED, DeFiLlama) do not imply partnership, endorsement, or affiliation unless explicitly stated.

22. Glossary

A2A
Agent-to-Agent — automated interaction between AI/autonomous systems

CASP
Crypto-Asset Service Provider under MiCA regulation

CCI
Crypto Compliance Index — composite regulatory compliance score (0–100)

CSRD
Corporate Sustainability Reporting Directive (EU)

DAP
Disclosure Attestation Protocol — cryptographic proof of data delivery

DORA
Digital Operational Resilience Act (EU)

DSSE
Dead Simple Signing Envelope — standardized signing format

ECDSA
Elliptic Curve Digital Signature Algorithm (ES256K)

EPM
Evidence Pack Manifest — standardized signed evidence schema

ESRS
European Sustainability Reporting Standards

FRED

Federal Reserve Economic Data — macro economic data source

HHI
Herfindahl-Hirschman Index — concentration measurement

JWKS
JSON Web Key Set — public key discovery endpoint

MCP
Model Context Protocol — open standard for AI agent tool integration (Anthropic)

Mech
MCP directories — decentralized AI agent marketplace on Gnosis Chain

MiCA
Markets in Crypto-Assets Regulation (EU) — stablecoin rules in force since June 2024, CASP transition ends 1 July 2026

RPO
Recovery Point Objective — maximum acceptable data loss

RTO
Recovery Time Objective — maximum acceptable downtime

RWA
Real World Assets — tokenized traditional financial instruments

SLA
Service Level Agreement — contractual commitment

SLO
Service Level Objective — target performance metric

TVL
Total Value Locked — aggregate value deposited in a protocol

Product

- Documentation
- Pricing
- System Status
- Whitepaper

Solutions

- For Compliance
- For Developers
- Protocol Coverage

Data infrastructure only — not financial advice, certification, or investment recommendation.