

# Platform Architecture Brief

The system of record, system of truth, and controlled transaction layer for website asset value.

<b>Ledger</b> Canonical URL identity and history	<b>Ratings</b> Value, risk, decay, and readiness scores	<b>Policy</b> Governed human and agent actions
---	--	---

*The platform is not trying to be the agent. It is the trusted record that humans, dashboards, BI systems, and agents query before they decide what a URL is worth or what action is safe.*

# Contents

Section	Asset	Purpose
01	Executive architecture summary	What the platform is and why it exists
02	Operating principle	Why URL Ledger is not just another agent
03	System layers	The ledger, ratings, policy, evidence, and API stack
04	Canonical URL instrument	The unit of account for every URL asset
05	Data ingestion and joins	What enters the ledger and how it is reconciled
06	Ratings engine	How asset value, decay, risk, and readiness are scored
07	Policy gate	How humans and agents move from recommendation to action
08	Agent ingress layer	How external/internal agents query trusted URL state
09	Evidence and assurance	Audit trail, governance, and defensibility
10	Roadmap	Build sequence from audit wedge to transaction layer

# Executive architecture summary

URL Ledger is the system of record for website asset value. It gives every URL a canonical identity, attaches performance and attribution evidence, scores structural decay and economic risk, and governs what humans or agents are allowed to change.

## Core architecture thesis

Every URL becomes an instrument with identity, lineage, signals, ratings, policy, evidence, and state-change history. The durable moat is not the agent itself. It is the trusted ledger that agents and teams must transact through.

## The product is built around five durable layers

Layer	Function	Why it matters
<b>Ledger</b>	Canonical URL registry, identity, lineage, status, ownership	Creates a single source of record across fragmented tools
<b>Ratings</b>	Health, decay, waste, ROI, risk, fit, authority, readiness	Turns messy signals into comparable decision logic
<b>Policy gate</b>	Permissions, approvals, protected zones, action thresholds	Prevents machine-speed chaos and brand/compliance mistakes
<b>Evidence layer</b>	Change logs, source snapshots, exports, rationale, outcomes	Makes recommendations auditable and defensible
<b>API / agent layer</b>	Trusted query and transaction surface for humans, BI, and agents	Makes the ledger the place agents come to, not the agent itself

## What this architecture enables

- A CFO-friendly view of website assets: value, yield, durability, risk, and recoverable upside.
- A CMO / Growth view of discovery performance across organic search, AI answers, paid traffic, social, email, referral, and direct.
- An SEO / Content Ops view of structural decay, cannibalization, dilution, waste, refresh priority, and protected assets.
- An agent-readiness view that lets external or internal AI systems query trusted URL state before proposing or executing action.

## Operating principle: be the record, not just the agent

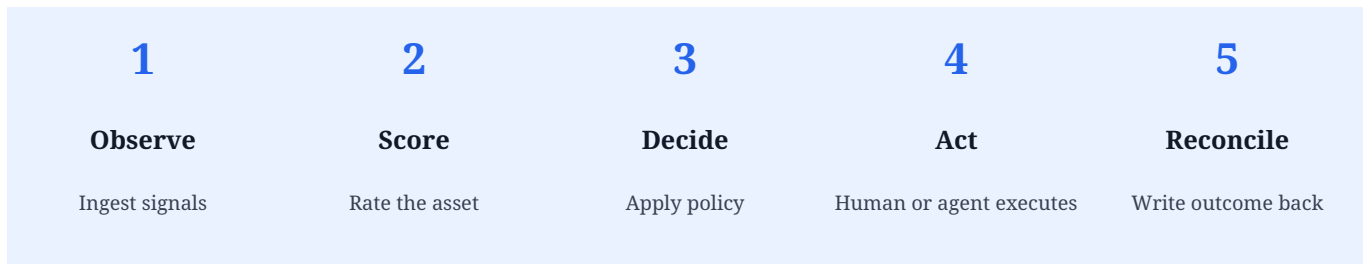
AI agents will draft, inspect, summarize, compare, redirect, brief, update, and recommend faster than humans can review. That makes trusted state more important than raw automation. URL Ledger is designed to be the canonical truth layer those agents must query before they act.

Do not collapse into	Own this layer instead
A content-writing tool	The governed record of which URLs should exist and why
A generic SEO dashboard	The asset ledger that reconciles identity, value, risk, and action
A workflow wrapper around an LLM	The policy and evidence layer that agents transact through
A single-channel AI Search product	A channel-agnostic asset intelligence platform across discovery and attribution
A recommendations feed	A state-transition system: observe, score, decide, act, reconcile

**Architecture maxim**

Agents are action surfaces. The ledger is the authority surface. The platform wins by becoming the trusted layer that survives across agents, search interfaces, analytics tools, CMS systems, and BI stacks.

### The recurring transaction loop



## System layers

The platform should be designed as a layered control plane. Each layer has a clear job and compounds defensibility over time.

Layer	Stores / computes	Strategic value
<b>1. Identity layer</b>	URL_ID, canonical URL, variants, redirects, merges/splits, cluster membership	Prevents duplicate truth and creates the asset register
<b>2. Signal layer</b>	Crawl, sitemap, CMS, GSC, GA4, CRM, paid, email, social, referral, data warehouse	Unifies fragmented discovery, traffic, and attribution signals
<b>3. Ratings layer</b>	Health, decay, ROI, waste, cannibalization, authority, readiness, confidence	Creates a comparable decision standard
<b>4. Portfolio layer</b>	Clusters, business units, page types, journey stages, strategic roles	Moves from single pages to portfolio allocation
<b>5. Policy layer</b>	Permissions, thresholds, protected assets, approval rules, no-touch zones	Determines what can be changed, by whom, and under what evidence
<b>6. Evidence layer</b>	Rationale, snapshots, exports, source joins, before/after results	Makes the system audit-ready and defensible
<b>7. API / transaction layer</b>	Query endpoints, action requests, audit packets, status transitions	Lets humans, BI systems, and agents transact through trusted state

### Layer dependency rule

Recommendations should not outrun the record. The platform should only recommend or approve an action when the identity, signal, rating, and policy context are sufficiently confident. Low-confidence state should create an exception, not an automated action.

# The canonical URL instrument

The URL is the unit of account. Each URL receives a durable instrument record that persists across redirects, canonical changes, refreshes, merges, splits, and attribution windows.

<b>Instrument definition</b>	
A URL instrument is a governed record containing identity, lineage, ownership, signals, value, risk, ratings, policy, evidence, and lifecycle state for a single URL asset or canonical URL cluster.	
<b>Instrument field family</b>	<b>Example fields</b>
<b>Identity</b>	URL_ID, canonical URL, normalized URL, URL hash, domain, property, section
<b>Lineage</b>	Redirects, canonical variants, historical URLs, merges, splits, duplicate relationships
<b>Ownership</b>	Business owner, content owner, SEO owner, technical owner, approver
<b>Role</b>	Page type, cluster, intent, journey stage, product/service mapping, strategic class
<b>Signals</b>	Indexability, crawl depth, internal links, impressions, clicks, sessions, conversions, revenue
<b>Ratings</b>	Health, decay, waste, ROI, fit, authority, AI citation readiness, agent readiness
<b>Policy</b>	Risk tier, protected status, approval requirements, allowed actions, restricted actions
<b>Evidence</b>	Snapshots, source exports, annotations, change history, rationale, before/after outcomes
<b>Lifecycle</b>	Published, indexed, peak, decaying, refreshed, consolidated, retired, protected

## The barcode spine

The barcode spine is the lightweight schema that maps each URL to identity, signals, controls, attribution, evidence, and allowed actions. This is what makes the website agent-readable without giving agents uncontrolled authority.

## Data ingestion and reconciliation

The platform begins by reconciling fragmented tools into a canonical URL registry. It does not require perfect data on day one. It requires enough truth to identify value leakage, risk, and the highest-confidence next actions.

Input	Signals	Purpose
<b>Crawl / sitemap / CMS</b>	URL inventory, indexability, canonical state, templates, internal links	Builds the base register and structural truth
<b>Google Search Console</b>	Queries, impressions, clicks, CTR, page visibility, coverage signals	Search demand and organic visibility
<b>Analytics</b>	Landing sessions, events, conversions, engagement, channel source	Behavior and channel-level value
<b>CRM / payments</b>	Closed-won, pipeline, order value, customer source, lifecycle value	Revenue truth and attribution confidence
<b>Paid media</b>	Spend, CAC, landing page performance, quality metrics	Paid efficiency and landing-page ROI
<b>Email / lifecycle</b>	Campaign clicks, nurture assists, reactivation paths	Owned-channel asset contribution
<b>Social / referral / partners</b>	Referral traffic, backlinks, partner sources, dark-social proxies	External discovery and authority
<b>BI / warehouse</b>	Unified tables, custom attribution, finance logic	Enterprise confidence and reporting joins

### Reconciliation rules

- Normalize URLs before scoring: protocol, trailing slash, parameters, canonical tags, redirects, and duplicates.
- Prefer canonical URL clusters for asset value while preserving variant-level evidence for technical diagnosis.
- Use confidence bands when attribution is incomplete rather than pretending all revenue joins are perfect.
- Store snapshots and source timestamps so every rating can be traced back to the evidence that produced it.

## Ratings engine

The ratings engine turns raw signals into action-ready asset intelligence. It should be explainable enough for operators, credible enough for executives, and structured enough for agents to consume safely.

Rating family	Signals	Decision question
<b>Content health</b>	Accuracy, usefulness, freshness, completeness, quality	Should this URL still be trusted?
<b>Performance</b>	Visibility, traffic, engagement, conversion, efficiency	Is the asset producing value?
<b>Decay</b>	Velocity of decline, half-life, volatility, refresh need	Is value eroding and when does it matter?
<b>Waste / dilution</b>	Low-yield inventory, duplicates, sprawl, crawl waste	Is the portfolio carrying dead weight?
<b>Cannibalization</b>	Intent overlap, query overlap, cluster conflict	Are URLs competing with each other?
<b>Strategic fit</b>	Business alignment, intent, journey stage, product mapping	Does the URL deserve investment?
<b>Authority</b>	Entity strength, topical coverage, proof, backlinks, trust	Does the URL strengthen the domain's source value?
<b>Channel exposure</b>	Organic, paid, social, email, referral, direct, CRM assist	Where does value come from?
<b>AI / agent readiness</b>	Extractability, citation readiness, structured data, actionability	Can machines understand, cite, compare, or act on it?
<b>Policy risk</b>	Legal, brand, compliance, commercial sensitivity, protected status	What controls are required before action?

### Rating output

#### Example rating object

URL\_ID: VL-0001287 | Rating: B- | Decay: High | Value-at-risk: Medium | Recoverability: High | Policy: Owner approval required | Recommended action: Refresh + internal link repair | Confidence: 0.74

## Policy gate

The policy gate is where the platform stops being a report and becomes an operating system. It decides which actions are allowed, which require review, and which are blocked until evidence improves.

Action	Risk tier	Evidence required	Approval route
<b>Refresh</b>	Low to medium risk	Allowed when rating and evidence pass threshold	Content owner or agent draft; owner approves
<b>Merge</b>	Medium risk	Requires cannibalization evidence and redirect plan	SEO + content owner approval
<b>Redirect</b>	Medium to high risk	Requires lineage, traffic, backlinks, and revenue exposure review	SEO + technical owner approval
<b>Retire / noindex</b>	High risk	Requires low strategic fit and documented replacement path	Senior approval for high-value assets
<b>Protect / no-touch</b>	High sensitivity	Business-critical, legal, pricing, brand, or revenue pages	Manual approval only
<b>Agent writeback</b>	Variable	Allowed only for scoped low-risk changes with audit trail	Policy-bound transaction

### Protected asset rules

- Pricing, legal, comparison, migration, brand, homepage, product, and high-revenue pages should default to higher approval thresholds.
- Agents may draft recommendations for protected assets, but should not directly publish, redirect, delete, or overwrite them.
- Every state transition must include a reason code, evidence pack, expected outcome, owner, and rollback path where applicable.

## Agent ingress layer

The long-term platform advantage is that AI agents, browser agents, content agents, BI agents, and internal assistants should come into URL Ledger to retrieve trusted URL state before they act. The ledger becomes the asset authority, not just a downstream dashboard.

### Agent-facing promise

Ask the ledger: What is this URL? What is it worth? What is its risk? What evidence supports that? What action is safe? Who must approve? What happened last time we changed it?

### Core agent endpoints

Endpoint concept	Function
<a href="#">GET /url/{id}</a>	Returns canonical identity, lineage, status, owners, cluster, and current ratings
<a href="#">GET /url/{id}/evidence</a>	Returns evidence pack, signal sources, snapshots, and score rationale
<a href="#">GET /portfolio/risks</a>	Returns high-risk assets, decay clusters, protected pages, and value-at-risk
<a href="#">POST /action-request</a>	Submits proposed refresh, merge, redirect, retire, protect, or expand action
<a href="#">GET /policy/check</a>	Checks whether a proposed action is allowed, reviewable, or blocked
<a href="#">POST /outcome</a>	Writes back implemented action, timestamp, owner, evidence, and measured result

### Agent transaction lifecycle

- Read trusted state from the ledger.
- Generate a recommendation or action proposal using ledger context.
- Submit the proposal through the policy gate.
- Receive approve / escalate / block decision with rationale.
- Execute only approved action scope.
- Write outcomes and evidence back to the ledger.

## Evidence and assurance layer

The assurance layer makes the platform defensible. It turns recommendations into evidence-backed decisions and creates the audit trail required for finance, legal, RevOps, agencies, and enterprise buyers.

Assurance output	What it proves
<b>Ledger reconciliation report</b>	What exists, what changed, what is orphaned, what is duplicated, and what is canonical
<b>Rating justification pack</b>	Score breakdown, signal sources, confidence bands, and because-trail for each rating
<b>Content risk register</b>	High-risk URLs, compliance flags, protected pages, sensitive commercial assets
<b>Action backlog</b>	Prioritized refresh / merge / retire / protect / expand actions with expected value and effort
<b>Change log</b>	Who changed what, when, why, from which evidence, and what result followed
<b>Quarterly portfolio review</b>	Portfolio value, decay movement, recovery progress, waste reduction, policy exceptions
<b>Benchmark pack</b>	Industry or peer norms once sufficient cross-account data exists

### Why this matters

The system becomes more valuable each time it records a decision, action, outcome, exception, and benchmark. Over time, the ledger compounds into a proprietary memory of what types of URL changes recover value, waste budget, reduce risk, or create regressions.

## Build roadmap

The build sequence should protect the compounding core: identity first, ratings second, policy third, transaction layer fourth. Do not start with broad agent automation before the ledger state is reliable.

Phase	Build focus	Key capabilities	Outcome
Phase 1	<b>Audit product + URL registry</b>	Crawl, ingest, normalize, dedupe, canonicalize, baseline portfolio view	A sellable audit and initial source of record
Phase 2	<b>Ratings v1</b>	13-variable scoring, decay, waste, cannibalization, ROI, readiness	Comparable asset ratings and action backlog
Phase 3	<b>Governance controls</b>	Owners, approvals, protected assets, policy thresholds, change log	Safe human and agent operating model
Phase 4	<b>Evidence packs + exports</b>	Audit reports, score justifications, CFO readouts, BI-ready tables	Enterprise credibility and sales enablement
Phase 5	<b>API / agent ingress</b>	Query endpoints, policy checks, action requests, writeback	The ledger becomes the transaction layer
Phase 6	<b>Benchmarks + assurance</b>	Cross-domain norms, portfolio health standards, quarterly reviews	Moat expansion through data and standards

### MVP definition

- A customer can upload or connect a URL inventory and reconcile it into canonical URL instruments.
- The system assigns initial ratings across the 13 structural variables with explainable evidence.
- The system produces a ranked recovery backlog and protected asset list.
- The system generates an executive audit report and machine-readable export.
- The system logs approved actions and reconciles outcomes back into each URL asset record.

### End-state category

Website Asset Intelligence: the system of record and system of truth for URL assets across discovery, attribution, revenue, governance, and agent consumption.

## Appendix: core object model

The following object model is intentionally simple. It can be expanded later, but the first version should preserve the compounding asset spine.

Object	Representative fields
<a href="#">URLAsset</a>	URL_ID, canonical_url, normalized_url, domain_id, status, page_type, cluster_id
<a href="#">URLLineage</a>	source_url, target_url, relationship_type, first_seen, last_seen, evidence_id
<a href="#">SignalSnapshot</a>	url_id, source, metric_name, value, date_range, collected_at, confidence
<a href="#">Rating</a>	url_id, rating_family, score, grade, confidence, reason_codes, evidence_ids
<a href="#">PolicyRule</a>	rule_id, action_type, risk_tier, condition, decision, required_approver
<a href="#">ActionRequest</a>	request_id, url_id, action_type, proposed_by, rationale, expected_impact, status
<a href="#">EvidencePack</a>	evidence_id, source, file_ref, screenshot_ref, export_ref, notes, timestamp
<a href="#">Outcome</a>	action_id, implemented_at, owner, metric_delta, value_delta, confidence, notes

### Minimum viable export

CSV / JSON exports should include URL\_ID, canonical URL, page type, cluster, owner, current rating, decay score, value-at-risk, recommended action, policy status, confidence, and evidence references. This is the bridge to BI, warehouses, agencies, and agent tools.