URL Ledger / Verbedge / 1UP Media

Security, Data Privacy & Trust Brief

How URL Ledger protects client data while becoming the system of record for website asset value.

Trust thesis

URL Ledger does not need to start with risky write access. The first motion is read-only, evidence-backed, and governed. The platform earns trust by reconciling URL assets, scoring risk, producing proof, and letting humans approve any state transition before agents or operators act.

Designed for	Primary questions answered
Enterprise buyers	What data is required, how is it accessed, and what controls exist?
Pilot partners	How do we start safely without granting unnecessary permissions?
Security / IT teams	What are the access tiers, retention model, audit logs, and isolation assumptions?
Investors / partners	Why trust, evidence, and policy become part of the moat?

Status: Master asset draft for sales, security review, pilot onboarding, and partner enablement.

01 / EXECUTIVE SUMMARY

The platform must be trusted before it can become authoritative

URL Ledger is positioned as the system of record and system of truth for URL-level asset value. That role requires a trust model stronger than a typical marketing dashboard. The platform may process URL inventories, search performance, analytics events, CRM outcomes, revenue proxies, evidence artifacts, ownership data, and governance decisions. These datasets can be commercially sensitive even when they do not contain traditional personal data.

The trust model is simple: start with the least amount of access required, prefer read-only connectors or exports, separate observation from action, preserve evidence, and make every material state transition auditable.

Core promise

URL Ledger should help teams know what every URL is, what it is worth, where value is leaking, what action is recommended, who approved it, and what happened afterward - without forcing the customer to surrender control of their website, CMS, CRM, or data warehouse.

Read-only first: The audit and initial ledger can run from exports or read-only connectors.

Governed action second: Writeback, publishing, redirects, or CMS changes require explicit policy gates.

Evidence over opinion: Ratings and recommendations should be traceable to source signals, calculations, screenshots, exports, and decision logs.

Agent ingress with boundaries: Agents can query, request, and reconcile through the ledger before they are allowed to execute.

Benchmarking without exposure: Cross-customer intelligence should rely on aggregation, anonymization, cohort thresholds, and opt-in governance.

02 / TRUST PRINCIPLES

The seven operating principles

#	Principle	Practical meaning
1	Least privilege	Only request the minimum permission required for the audit, pilot, or module.
2	Read-only by default	Observation, scoring, and reporting should not require write access to CMS, CRM, analytics, or infrastructure.
3	Explicit state transitions	Refresh, merge, redirect, retire, publish, and agent actions are treated as governed events.
4	Evidence-backed ratings	Every score and recommendation should have a because trail that can be reviewed.
5	Human approvals for protected assets	High-value or regulated URLs require owner approval before material action.
6	Data minimization	Do not collect content, CRM fields, or user-level records that are not needed for scoring or proof.
7	Benchmark privacy	No customer should be identifiable through benchmark outputs or peer-set comparisons.

These principles support the product strategy: URL Ledger should become the trusted layer that humans, systems, and agents reference before making asset decisions. That trust comes from restraint, traceability, and defensible governance - not from claiming that automation can safely act everywhere on day one.

03 / DATA CLASSIFICATION

What data the platform may touch

URL Ledger should classify all data by sensitivity and required access level. Many inputs are operational or commercial rather than personally identifiable, but they can still reveal strategy, funnel economics, competitive performance, and revenue exposure.

Data class	Examples	Sensitivity	Need
URL inventory	Sitemaps, crawl exports, canonical URLs, redirects, status codes, metadata	Low to Medium	Required
Search visibility	GSC page/query clicks, impressions, CTR, positions, query clusters	Medium	Recommended
Analytics behavior	GA4 landing pages, sessions, events, conversions, device/channel data	Medium	Recommended
Revenue truth source	CRM opportunities, pipeline value, orders, subscriptions, lead value proxies	High	Optional but high-value
CMS metadata	Publish dates, author, page type, last modified, tags, status, ownership	Medium	Optional
Paid/media context	Campaign spend, UTMs, landing pages, CAC, ROAS proxies	Medium to High	Optional
Evidence artifacts	Screenshots, exports, crawl evidence, SERP snapshots, recommendation rationale	Medium	Required for proof pack
User/admin activity	Approvals, action logs, permissions, role assignments	High	Required for platform
Benchmark metadata	Normalized, aggregated rating and outcome metrics	High until anonymized	Optional / opt-in

Data handling rule

If a field does not improve URL identity, scoring, attribution, governance, evidence, or benchmark value, it should not be collected by default.

04 / ACCESS TIERS

Four levels of customer access

Access tier	What it means	Best use case	Guardrail
Tier 0: Export-only	Client provides CSV/crawl exports, screenshots, or warehouse extracts.	Lowest-risk pilots, procurement-sensitive accounts, early scoping.	No live connector. Slower refresh. Manual validation needed.
Tier 1: Read-only connectors	Read-only access to GSC, GA4, sitemap/crawl, CRM exports, warehouse views.	Standard audit and ledger baseline.	No writeback. Action execution remains outside platform.
Tier 2: Controlled workflow integration	Tickets, approvals, owners, evidence packs, CMS references, project management sync.	Operationalizing recommendations and proof sprint.	Changes are requested and approved, not auto-executed.
Tier 3: Governed writeback / agent actions	Policy-gated API or CMS actions for approved URL updates, redirects, metadata, tickets, or agent tasks.	Mature customers with governance, protected zones, and rollback plans.	Requires explicit permissions, logs, approvals, and customer-defined policies.

The sales motion should default to Tier 0 or Tier 1. This reduces enterprise friction and reinforces the message that URL Ledger is not asking for control before proving value.

05 / CONNECTOR MATRIX

Recommended access posture by source system

System	Purpose	Preferred access	Sensitivity	Control note
Sitemap / crawl	URL discovery, status codes, canonicals, depth, internal links	Crawl or export	Low	Required for asset registry
Google Search Console	Page/query visibility and decay patterns	Read-only	Medium	Do not require admin rights
GA4 / analytics	Landing page behavior, channel contribution, conversions	Read-only or export	Medium	Avoid user-level collection unless justified
CRM	Pipeline, closed-won, lead value, lifecycle stage	Export or read-only report	High	Prefer aggregated landing-page joins or opportunity-source views
Payments / ecommerce	Orders, revenue, product/category contribution	Export or warehouse view	High	Use order aggregates where possible
CMS	URL metadata, owners, publish dates, page type, status	Read-only first	Medium	Write access only after policy gate is live
Warehouse / BI	Joined views, modeled revenue, attribution tables	Read-only view	High	Use customer-created limited views
Ads platforms	Spend, CAC, paid landing-page context	Read-only or export	Medium	Useful for blended asset value, not always required
Project management	Tasks, approvals, owners, implementation evidence	Read/write to tickets only	Medium	Good first writeback target before CMS writeback

Procurement-friendly position

The audit can start from exports. Connectors improve speed and refresh cadence, but they are not a prerequisite for the first proof of value.

06 / DATA LIFECYCLE

From intake to deletion

Stage	What happens	Evidence
1. Request	Define required sources, fields, permissions, retention, and pilot scope.	Data request sheet + approval trail
2. Ingest	Receive exports or connect read-only sources. Validate file/source integrity.	Connector log or upload record
3. Normalize	Map URLs, canonicals, redirects, clusters, channels, and revenue proxies.	URL registry and mapping table
4. Score	Apply 13-variable scoring, confidence tiers, and value-at-risk logic.	Rating evidence and calculation notes
5. Review	Analyst QA, stakeholder review, protected-asset check, and recommendation validation.	QA checklist and exceptions log
6. Deliver	Produce audit report, backlog, evidence pack, and executive readout.	Delivered asset list
7. Govern	Convert accepted actions into tickets, approvals, and ledger state transitions.	Action log and owner approvals
8. Retain or delete	Apply customer retention terms; delete pilot data when required.	Deletion certificate or retention record

The delivery SOP should include a retention election during intake: delete after audit delivery, retain for subscription conversion, or retain under an agreed review period. This helps avoid ambiguity when pilots do not convert immediately.

07 / SECURITY CONTROL MODEL

Minimum control baseline

Control area	Expected implementation
Identity and access	Role-based access control, least-privilege roles, optional SSO/SAML for enterprise, periodic access review.
Data protection	Encryption in transit and at rest, separate customer workspaces, limited production access, secure secrets handling.
Audit logging	Log source connections, uploads, exports, score changes, approvals, action requests, and agent queries.
Data minimization	Field allowlists, source-specific ingestion templates, no unnecessary user-level records.
Evidence integrity	Store calculation inputs, source timestamps, screenshot/export metadata, and recommendation rationale.
Change governance	Policy thresholds, protected assets, owner approval, rollback notes, and exception handling.
Vendor/subprocessor governance	Document hosting, analytics, storage, model/API providers, and any external compute dependencies.
Incident response	Defined escalation contacts, impact assessment, customer notification workflow, remediation log.

Trust moat

The control surface becomes part of defensibility. Customers are not only buying insights; they are buying a governed record of what changed, why it changed, who approved it, and whether it worked.

08 / ROLES AND PERMISSIONS

Recommended role model

Role	Allowed actions	Typical user
Viewer	View dashboards, reports, ratings, evidence packs	Executives, finance, leadership
Analyst	Create findings, adjust mappings, attach evidence, draft recommendations	URL Ledger delivery team, customer analysts
Operator	Move actions through workflow, assign owners, update implementation status	SEO, content ops, growth ops
Approver	Approve protected-asset actions, policy exceptions, and proof sprint changes	CMO, RevOps, legal, finance, business owner
Admin	Manage users, data connections, retention settings, policy thresholds	Customer admin / IT
Agent	Query ledger, request actions, submit evidence, reconcile outcomes within policy	Approved AI/browser/workflow agents

The Agent role should never be treated as a superuser. It should be narrower than a human admin and governed by action type, confidence score, asset sensitivity, and customer policy.

09 / AI AND AGENT INGRESS

Agents should transact through the ledger, not around it

The strategic point of URL Ledger is not to become another generic agent. The platform should become the trusted URL asset layer that agents use before they recommend, edit, redirect, retire, promote, cite, or analyze a page.

Agent interaction	Description	Control
Query	Agent asks for URL state, score, owner, risk, evidence, and allowed actions.	Allowed under read policies
Recommend	Agent proposes refresh, merge, redirect, protect, expand, or investigate action.	Requires evidence and confidence
Request	Agent submits action request into workflow.	Policy gate determines route
Execute	Agent or system performs approved action.	Only mature Tier 3 customers
Reconcile	Agent reports observed outcome and updates evidence.	Allowed with validation

Agent policy statement

Agents can read broadly, recommend cautiously, request through workflow, and act only where policy, approval, and rollback conditions allow.

10 / BENCHMARK AND DATA SHARING MODEL

How benchmark value compounds without exposing customers

Benchmarks are a major future moat, but they require strict guardrails. Customers should never see another customer's raw URL, revenue, traffic, conversion rate, or proprietary strategy. Benchmark outputs should be normalized, aggregated, anonymized, and cohort-thresholded.

Rule	Implementation
No raw URL exposure	Never show another customer's URL path, slug, title, page content, or domain unless explicitly public and permissioned.
Cohort thresholds	Do not publish a peer benchmark unless enough accounts and URLs exist to prevent re-identification.
Normalization	Compare relative scores, percentiles, decay velocity, coverage ratios, and risk classes rather than raw revenue or traffic.
Aggregation	Roll up by industry, page type, funnel role, portfolio size, or maturity tier.
Opt-in benchmark terms	Pilot/SOW should specify whether anonymized metrics may contribute to future benchmarks.
Customer controls	Allow enterprise customers to opt out of benchmark contribution or restrict categories.

Benchmarking should be marketed as an assurance layer, not a data resale motion. The customer contributes normalized signal and receives normalized context back.

11 / ENTERPRISE BUYER PACKET

Documents security and procurement teams will ask for

Artifact	Purpose
Data Processing Addendum	Defines controller/processor roles, processing purpose, retention, deletion,

Artifact	Purpose
	subprocessor disclosure.
Security overview	Summarizes encryption, access controls, logging, isolation, backup, incident process.
Subprocessor list	Lists hosting, storage, analytics, model/API, document-rendering, and communication vendors.
Access matrix	Shows requested permissions by connector and pilot phase.
Retention and deletion policy	Explains audit retention, subscription retention, deletion request handling, evidence retention.
AI usage policy	Explains whether customer data is used for model training, which models are used, and where human review is required.
SOC 2 roadmap	If not certified, provide planned control path and compensating controls.
Incident response summary	Who is notified, when, and what evidence is provided.

Early-stage posture

For pre-SOC2 pilots, reduce procurement friction by using exports, read-only connectors, no customer PII by default, limited retention, and a clear deletion option.

12 / PILOT CONTRACT LANGUAGE

Trust language to reuse in proposals and SOWs

The following clauses can be reused as plain-English positioning in pilot proposals before formal legal review.

Topic	Reusable language
Access	URL Ledger will begin with export-only or read-only access unless otherwise approved in writing by the client.
Writeback	No CMS, redirect, publishing, or destructive action will be performed without explicit approval and agreed implementation workflow.
Data minimization	Only data required to reconcile URL assets, score portfolio risk, estimate value-at-risk, and prepare the audit deliverables will be requested.
Revenue data	Where revenue or CRM data is provided, URL Ledger will use the minimum level of granularity required to support asset-level or cluster-level value modeling.
Retention	Pilot data may be deleted after delivery unless the client elects to retain it for ledger subscription conversion or ongoing monitoring.
Benchmarks	Customer data will not be used in external benchmarks without agreed anonymization and aggregation terms.
AI systems	AI assistance may be used to support analysis and drafting, but ratings, recommendations, and protected-asset actions remain subject to human review.

13 / SECURITY FAQ

Questions buyers will ask

Do you need write access to our CMS?

No. The audit and baseline ledger should start without write access. CMS writeback is a later-stage option only after policy gates, approvals, and rollback procedures are defined.

Do you need user-level analytics data?

Usually no. Most URL asset scoring can be done with page-level or aggregated event data. User-level data should be avoided unless there is a specific approved reason.

Can we use exports instead of connectors?

Yes. Export-only pilots are supported. Connectors improve refresh cadence and reduce manual work, but the first proof of value can be export-based.

Will our data train AI models?

The recommended policy is no customer data used for third-party model training. Any AI usage should be documented by provider, purpose, and retention behavior.

Can agents change our site?

Not by default. Agents may query and recommend. Execution requires customer-approved permissions, policy gates, protected asset rules, and audit logs.

How do you handle sensitive revenue data?

Prefer aggregated or modeled views, limited fields, and cluster-level joins where possible. The platform needs enough to estimate value, not every internal field.

Can we delete our data after the audit?

Yes. The pilot should include a retention election: delete, retain for conversion, or retain for a fixed review period.

What makes this different from another analytics dashboard?

The trust layer. URL Ledger combines canonical asset identity, ratings, evidence, approvals, action logs, and outcome reconciliation. It is designed to become a governed record, not just a view.

14 / TRUST ROADMAP

From pilot-safe to enterprise-grade

Phase	Control maturity	Outcome
Phase 1: Pilot-safe	Export-only, read-only connectors, manual QA, deletion option, evidence-backed reports.	Reduce friction and prove value.
Phase 2: Operational trust	RBAC, workspace isolation, audit logs, retention controls, protected assets, approval workflow.	Make the platform safe for recurring use.
Phase 3: Enterprise trust	SSO/SAML, DPA, subprocessor registry, security questionnaire, formal incident process, SOC 2 roadmap.	Support larger accounts and procurement.
Phase 4: Agent trust	Agent identities, scoped tokens, action policies, rate limits, sandboxing, rollback, reconciliation logs.	Become the governed agent transaction layer.
Phase 5: Benchmark trust	Anonymization, aggregation, peer-set thresholds, opt-in benchmark terms, customer controls.	Turn trust into the benchmark moat.

Strategic close

URL Ledger becomes more defensible as it becomes more trusted. The goal is not to ask for more access faster. The goal is to earn the right to become the canonical record of URL asset value.

What to confirm before starting an audit

Area	Checklist item
Scope	Domains, subdomains, languages, regions, sections, and excluded areas confirmed
Data sources	GSC, GA4, sitemap/crawl, CRM/revenue, CMS metadata, paid/email/warehouse availability confirmed
Access tier	Export-only, read-only, workflow, or governed writeback tier selected
Sensitive fields	PII, revenue, CRM, and restricted business fields reviewed and minimized
Retention	Delete after audit, retain for conversion, or fixed review period selected
Protected assets	Pricing, legal, medical/financial, high-revenue, brand-sensitive, and regulated pages flagged
Approvers	Business owner, technical owner, security contact, and executive sponsor identified
Evidence standards	Screenshots, exports, calculation notes, and recommendation rationale agreed
Agent policy	Agent access disabled by default unless explicitly approved
Success criteria	Value-at-risk model, recovery backlog, proof sprint, and readout expectations aligned

End of brief. This document is a working asset and should be reviewed by counsel and security advisors before being used as a binding policy document.