

VerifiedSignal Reference Document: Comprehensive Product & Architecture Guide

1. Brand Identity and Marketplace Positioning

Brand Vision and Messaging

Veridoc is the "newspaper-of-record" for the synthetic age—an automated epistemological verification layer designed to restore trust in digital discourse. In an era where \$78 billion is lost annually to decisions predicated on bad information, Veridoc provides the critical infrastructure for skeptical reading. Our core value proposition, "**Read everything. Trust wisely,**" moves beyond simple AI detection to offer a scalable critical-thinking engine that uncovers hidden persuasion, misinformation, and factual inconsistencies.

The Misinformation Crisis vs. The Veridoc Solution

Problem	Technical Intervention
AI Detection Gap: 85% of people cannot detect AI text as LLMs reach human parity in style.	Multi-Model Scoring: Specialized detection models (beyond style heuristics) assess authorship probability with high-precision confidence estimates.
Logical Fallacies: 60% of viral articles use "invisible machinery" like false equivalences to manipulate.	Fallacy Rating Engine: Identifies and names specific fallacy types (e.g., ad hominem) with direct links to the offending text passages.
Cost of Bad Information: Decisions in investment and policy based on misinformation lead to massive fiscal leakage.	Factuality Confidence: A 0–1 score generated by internal consistency checks, citation signals, and cross-referenced claim validation.

Target Audience Segments

- **Researchers & Academics:** Build trustworthy evidence bases at scale. They ingest massive document volumes to identify high-factuality sources and track quality trends across scientific literature.

- **Journalists & Fact-Checkers:** Vet sources before publication. Users identify provenance gaps, surface misleading framing, and flag AI-generated content passed off as human reporting.
 - **Investors & Analysts:** Pressure-test market narratives. Professionals score earnings calls and research reports to discover overstatements or track a source's accuracy over time.
 - **Educators & Students:** Scale critical thinking and media literacy. Veridoc serves as a teaching tool by comparing high- and low-quality sources via objective intelligence lenses.
 - **Legal & Compliance Teams:** Perform due diligence on regulatory filings. Teams run expert witness materials and third-party reports through factuality and provenance scoring.
 - **Curious Individuals:** Defend against digital manipulation. Individuals analyze viral threads or blogs to understand the "how" of persuasion before deciding what to think.
-

2. Product Features and Measurement Framework

The Eight Dimensions of Intelligence

- **Logical Fallacy Rating:** Measures manipulative reasoning. **Indicators:** Per-fallacy breakdown of ad hominem, straw man, false dichotomy, and slippery slope, mapped to specific text triggers.
- **Factuality Confidence:** Measures the reliability of claims. **Indicators:** Internal consistency, citation signals, and cross-referenced factual claims with a detailed rationale.
- **AI Generation Probability:** Measures machine authorship likelihood. **Indicators:** Multi-model assessment of linguistic patterns vs. known LLM signatures with a specific model guess included.
- **Pseudoscience Indicators:** Measures adherence to scientific rigor. **Indicators:** Detection of unfalsifiable claims, anecdotal evidence, "appeal to nature," and absence of peer review.
- **Fictional Content Likelihood:** Measures intent and genre. **Indicators:** Distinguishes reported fact from narrative, satire detection, or speculation presented as journalism.
- **Source Provenance:** Measures document origin history. **Indicators:** Domain reputation, WHOIS history, canonical URL verification, and archive.org presence.
- **Semantic Search:** Measures deep content relevance. **Indicators:** Vector-based similarity (kNN) and hybrid retrieval across massive document collections.
- **Collection Analytics:** Measures aggregate quality shifts. **Indicators:** Trend dashboards showing fluctuations in factuality or fallacy frequency across sources over time.

The "Agentic" Workflow

- **Submit:** User uploads file (PDF, Word, HTML) or URL. System performs automated fetching, text extraction, and cleaning.
- **Investigate:** The system verifies source provenance (publication dates, author identity, and domain history) before analysis begins.
- **Analyze:** A battery of AI models simultaneously scores the document across the eight dimensions.
- **Collect:** Users save documents to collections for side-by-side comparison and trend visualization.

User Interface Standards

- **Review UI:** Features a side-by-side view for Human-in-the-Loop (HITL) verification, allowing users to correct low-confidence AI extractions.
 - **Document Explorer:** Implements "Suspicious Documents" card logic, automatically flagging items with High AI probability (>0.7) and Low Factuality (<0.4).
 - **Real-Time Feedback:** Scorecards populate via Server-Sent Events (SSE) to show field-level progress during the analysis phase.
-

3. Business and Pricing Strategy

SaaS Pricing Tiers

Plan	Price (Monthly)	Document Limits	Feature Inclusions
Reader	Free	50 / mo	All 8 dimensions, 3 collections, 7-day history.
Analyst	\$29	1,000 / mo	Unlimited collections, trend dashboards, semantic search, CSV/JSON export.
Team	\$99	10,000 / mo	10 seats, shared collections, API access, SSO/SAML, self-host (Docker).

Outcome-Based SMB Model (AWS Marketplace)

- **Solopreneur (\$49/mo):** Up to 200 documents; basic QuickBooks/Xero one-click connectors.

- **Growth (\$199/mo):** Up to 1,500 documents; Agentic email drafting (requesting missing info) and **Five-Doc Training** (Few-Shot learning for unique doc types).
- **Scale (\$499/mo):** Unlimited documents; full API access and premium support.

Market Context (2026)

The Intelligent Document Processing (IDP) market is projected at **\$31B by 2026**. Search activity for "Agentic AI" on the AWS Marketplace has surged **40x** recently. The **Banking & Finance** sector remains the dominant segment, representing 45% of the total market, primarily for KYC and AML automation.

4. High-Level Technical Architecture

Core Tech Stack

- **AWS Amplify:** Frontend hosting and global edge delivery.
- **Amazon S3:** Object storage using **Presigned URLs** for direct-to-S3 uploads, minimizing Lambda execution costs.
- **Amazon Cognito:** User authentication with "Sign In with AWS" for Marketplace native integration.
- **Amazon Textract:** High-accuracy OCR for tables, forms, and document geometry.
- **Amazon Bedrock:** LLM reasoning (Claude 3.5/Llama 3) for "Agentic" sanity checks and scoring.

Database and Search Strategy

!!IMPORTANT Postgres is the System of Record. Derived search state is expendable.

- **PostgreSQL:** Canonical truth (users, permissions, final scores, billing).
 - **Amazon OpenSearch:** Disposable analytics replica for full-text search, kNN vector similarity, and dashboard aggregations.
-

5. Data Plane and Intelligence Pipeline

The 8-Stage Ingestion Pipeline

- **Intake:** Client submits file/URL; FastAPI creates "queued" record in Postgres.
- **Acquisition:** Worker fetches bytes, computes content hashes, and checks for duplicates.

- **Extraction:** Textract/Worker extracts plain text and structure; publishes progress to Redis.
- **Enrichment:** System computes metadata, topical tags, and initial quality flags.
- **LLM Scoring:** Worker calls Bedrock for structured analysis via the Converse API.
- **Canonical Persistence:** Final validated scores are written to the authoritative PostgreSQL layer.
- **Search Indexing:** Document is indexed into OpenSearch for retrieval and analytics.
- **Completion Notification:** Final "done" event is pushed to the frontend via SSE.

LLM Implementation & Prompt Engineering

- **The "Auditor" System Prompt:** Uses strict XML tagging and JSON schema enforcement. It forces math validation (e.g., Total = Subtotal + Tax).
 - **Few-Shot Strategy:** Includes "Success," "Graceful Failure" (teaching the model to return null for missing data), and an "Edge Case" example (a messy, low-quality document) to prevent hallucinations.
 - **Model Settings:** Temperature is set to 0.0 for deterministic, "boring auditor" consistency.
-

6. Streaming and Real-Time UX

Server-Sent Events (SSE) Design

Veridoc uses SSE for one-way status streaming. !IMPORTANT Nginx must be configured with the **X-Accel-Buffering: no** header to prevent buffering from breaking the real-time stream.

Sample Event Contract (JSON):

```
JSON
{
  "event": "stage",
  "document_id": "doc_123",
  "stage": "extract_text",
  "status": "running",
  "timestamp": "2026-03-26T22:11:10Z",
  "progress": 35
}
```

Incremental LLM Scoring

The system utilizes **partial JSON token parsing**. As tokens stream from Bedrock, a worker identifies completed fields and updates the UI before the full LLM response is finished.

Field State Machine: unseen \rightarrow in_progress \rightarrow complete \rightarrow validated \rightarrow persisted.

7. AWS Marketplace & Integration Plumbing

The "Plumbing" Layer

When a user subscribes, AWS sends a POST with an x-amzn-marketplace-token. The backend exchanges this for a CustomerIdentifier to link the AWS billing account to the Veridoc subscription.

Metering and Billing

A scheduled **BatchMeterUsage** Lambda runs hourly, aggregating document counts from Postgres and reporting them to the AWS Marketplace Metering Service. This enables "Zero Idle Cost" serverless billing.

8. Implementation Roadmap

Phase	Duration	Tasks
Infrastructure	Weeks 1-2	AWS CDK setup for S3, DynamoDB, and Cognito.
AI Pipeline	Weeks 3-5	Connect Lambda to Textract and Bedrock (Converse API).
Frontend	Weeks 6-8	Build Review UI (side-by-side) and Document Explorer.
Marketplace	Weeks 9-10	Implement ResolveCustomer and BatchMeterUsage APIs.
Listing	Weeks 11-12	AWS Security Review and Marketplace go-live.

Deployment Targets:

- **Bare Metal:** Local filesystem/MinIO and single-node OpenSearch.
 - **Containers:** Portable Docker Compose model.
 - **AWS:** Production-grade using ECS Fargate, RDS, and Amazon OpenSearch Service.
-

9. Technical Constraints and Error Handling

The "Do Not Build" List (MVP Exclusions)

- Custom model training (Stick to foundation models).
- Complex multi-tenant roles (Admin only for MVP).
- Dedicated mobile app (Responsive web only).

Failure Mitigations

Failure Mode	Mitigation Strategy
LLM Malformed JSON	Incremental parser with field-level validation; fallback to final pass parse.
Search Unavailable	Persist to Postgres first; retry indexing later; UI shows "Search Pending."
Worker Crash	Idempotent stage checkpoints allow the worker to resume from the last saved state.
Redis Outage	UI falls back to polling ; event logs are persisted in Postgres for replay.

Here are all the key screens as interactive wireframes — click through each one using the screen navigator. Ten interactive wireframes covering every major workflow — click any screen name in the nav to jump directly to it. Here's what each screen shows and the key design decisions embedded in them:

The login screen is deliberately minimal — centered card, no distractions, forgot password link below the password field where users expect it.

The dashboard leads with the four KPI metrics, then surfaces recent documents with their score badges immediately visible — no clicks required to see the health of your library.

The upload screen shows the duplicate detection state, which is easy to overlook in spec but critical to get right in the UI — the amber "Duplicate" badge tells the user before they waste an analysis credit.

The pipeline progress screen shows scores streaming in per-document as they complete, with the scoring sub-stages labeled inline. Completed documents already show their score badges while the third is still processing.

The document reader has the persistent red alert bar at the top (AI 0.79 triggers the danger state), inline amber highlights on the two flagged fallacy passages, score bars in the right panel, and the keyword pills at the bottom — all visible simultaneously while reading.

The collection stats screen shows the Elasticsearch-powered rollup dashboard — KPI cards, a horizontal bar chart for fallacy types, and a histogram for score distribution.

The search screen shows keyword results with highlighted snippets, score badges per result, and the filter panel open on the right with active checkboxes. The mode selector (Keyword / Semantic / Hybrid) is prominent.

The report builder uses a two-column layout: configuration left, live preview right — so the user sees the report taking shape as they toggle sections.

The billing screen shows all three plan cards simultaneously so the user can see exactly what they'd gain or lose by upgrading or downgrading, with usage bars below.

The security screen covers all three concerns — password change, 2FA setup prompt, and active sessions with per-session revoke buttons.

Want me to produce a Cursor prompt to implement any of these screens, or go deeper on a specific one?

10. User Interface (UI) Wireframe Details

Ten interactive wireframes cover every major workflow. Below are the key design decisions embedded in each of the ten screens:

1. The login screen is deliberately minimal — centered card, no distractions, forgot password link below the password field where users expect it.
2. The dashboard leads with the four KPI metrics, then surfaces recent documents with their score badges immediately visible — no clicks required to see the health of your library.
3. The upload screen shows the duplicate detection state, which is easy to overlook in spec but critical to get right in the UI — the amber "Duplicate" badge tells the user before they waste an analysis credit.
4. The pipeline progress screen shows scores streaming in per-document as they complete, with the scoring sub-stages labeled inline. Completed documents already show their score badges while the third is still processing.
5. The document reader has the persistent red alert bar at the top (AI 0.79 triggers the danger state), inline amber highlights on the two flagged fallacy passages, score bars in the right panel, and the keyword pills at the bottom — all visible simultaneously while reading.
6. The collection stats screen shows the Elasticsearch-powered rollup dashboard — KPI cards, a horizontal bar chart for fallacy types, and a histogram for score distribution.

7. The search screen shows keyword results with highlighted snippets, score badges per result, and the filter panel open on the right with active checkboxes. The mode selector (Keyword / Semantic / Hybrid) is prominent.
8. The report builder uses a two-column layout: configuration left, live preview right — so the user sees the report taking shape as they toggle sections.
9. The billing screen shows all three plan cards simultaneously so the user can see exactly what they'd gain or lose by upgrading or downgrading, with usage bars below.
10. The security screen covers all three concerns — password change, 2FA setup prompt, and active sessions with per-session revoke buttons.