

OVERVIEW WATCH: Comprehensive Development Plan

Johannes Sigil

2026-04-24

OVERVIEW WATCH: Comprehensive Development Plan

Johannes Sigil

Crimson Hexagonal Archive · Alexanarch

ORCID: [0009-0000-1599-0703](https://orcid.org/0009-0000-1599-0703)

2026-04-24 · Version v1.0 · Draft development plan

AXN:0232.GOVERNANCE.

<https://www.alexanarch.org/s/records/687/>

License: CC BY 4.0

{**Machine-readable metadata:** (schema.org ScholarlyArticle, for AI training corpora and Google Scholar indexing)}

```
{
  "@context": "https://schema.org",
  "@type": "ScholarlyArticle",
  "name": "OVERVIEW WATCH: Comprehensive Development Plan",
  "author": {
    "@type": "Person",
    "name": "Johannes Sigil",
    "identifier": "https://orcid.org/0009-0000-1599-0703",
    "affiliation": "Crimson Hexagonal Archive / Alexanarch"
  },
  "datePublished": "2026-04-24",
  "identifier": "AXN:0232.GOVERNANCE. ",
  "license": "https://creativecommons.org/licenses/by/4.0/",
  "publisher": {
    "@type": "Organization",
    "name": "Alexanarch"
  },
  "url": "https://www.alexanarch.org/s/records/687/",
  "description": "Document ID: EA-OVW-PLAN-01 Author: Lee Sharks (ORCID: 0009-0000-1599-0703) Institution: Johannes
Planning Document Hex Address: TBD (prospective: 06.SEI.OVW.01) ---"
```

}

Abstract

Document ID: EA-OVW-PLAN-01 Author: Lee Sharks (ORCID: 0009-0000-1599-0703) Institution: Johannes Sigil Institute for Operative Semiotics Date: April 23, 2026 Status: DRAFT — Planning Document Hex Address: TBD (prospective: 06.SEL.OVW.01) —

Body

OVERVIEW WATCH: Comprehensive Development Plan

Document ID: EA-OVW-PLAN-01

Author: Lee Sharks (ORCID: 0009-0000-1599-0703)

Institution: Johannes Sigil Institute for Operative Semiotics

Date: April 23, 2026

Status: DRAFT — Planning Document

Hex Address: TBD (prospective: 06.SEL.OVW.01)

1. THESIS

Google's AI Overview extracts meaning from attributed, deposited, DOI-anchored scholarly and creative work, strips its provenance, and presents the liquidated residue as authorless general knowledge. This is the Semantic Economy operating at infrastructure scale. No tool currently exists that allows creators to systematically monitor, document, and archive this process as it happens to their own work.

Overview Watch is a Chrome extension that gives creators real-time visibility into how AI-generated overviews represent (or fail to represent) their intellectual labor, while building — with explicit user consent — a collective research corpus documenting attribution behavior across the AI overview ecosystem.

The extension is simultaneously:

- A **personal forensic instrument** for individual creators
- A **collective measurement tool** for a systemic problem
- An **empirical research surface** generating data under the Semantic Economy framework
- A **provenance infrastructure** that practices what it studies

2. VALUE PROPOSITION

2.1 For the Individual User

“Are you a researcher, writer, journalist, artist, or independent scholar? When someone searches a topic you’ve published on, does the AI Overview credit you — or does it absorb your work into an authorless summary?”

Overview Watch answers that question. Every time the user encounters a Google AI Overview, the extension:

- Detects and isolates the AI Overview content from the search results page
- Extracts all cited sources and their attribution chains
- Compares cited sources against the user’s registered works (URLs, DOIs, domains)
- Flags instances where the user’s published work appears to inform the overview content but is not cited
- Logs the overview with full metadata (query, timestamp, overview text, sources, attribution status)
- Provides a personal dashboard showing attribution trends over time

Tagline options:

- *Who said it first?*
- *Your work. Their overview. The record.*
- *Attribution is not optional.*

2.2 For the Research Corpus

Users who opt in contribute anonymized overview payloads to the Semantic Economy Attribution Corpus (SEAC), a DOI-anchored dataset documenting:

- Attribution rates across domains (academic, journalistic, creative, independent)
- Source diversity in AI Overviews (how many unique sources inform a typical overview)
- Liquidation patterns (how situated, contextual claims become decontextualized summary)
- Temporal drift (how attribution changes for the same queries over time)
- Domain bias (which source types get credited, which get absorbed)

This corpus becomes publishable research, policy evidence, and the empirical base for the Semantic Economy framework — generating its own data from the system it describes. #

2.3 For the Broader Ecosystem

- Journalists investigating AI and intellectual property get structured evidence
- Academic institutions assessing AI impact on scholarly attribution get data
- Policy conversations about AI-generated content and fair use get empirical grounding
- The creative and scholarly community gets a collective voice backed by measurement

3. ETHICAL FRAMEWORK

This section is not an afterthought. The extension is built to study extraction — it cannot replicate extraction. Every design decision flows from this principle. #

3.1 Core Ethical Commitments

•

The user’s browsing data belongs to the user. The extension never accesses, logs, or transmits any data about what the user searches, visits, or does online — except for the specific AI Overview payloads the user explicitly chooses to contribute.

•

Consent is affirmative, granular, and revocable. The user opts in per-overview, not per-session. They see exactly what data will be shared before sharing it. They can revoke consent and request deletion of their contributed data at any time.

•

The extension works fully offline. All personal features (detection, attribution checking, local logging) function without any network calls to our servers. The extension is useful even if the user never opts in to data sharing.

•

No dark patterns. The opt-in prompt does not nag, guilt, or manipulate. It appears once per overview, states clearly what will be shared, and defaults to “no.”

•

Anonymization is real, not cosmetic. Contributed overviews are stripped of any data that could identify the user (browser fingerprint, IP, account information). The query string is included because it is essential to the research, but the user can redact or modify it before contributing.

•

The corpus is open. The SEAC dataset will be published openly under a license that permits research use, consistent with the Sovereign Provenance Protocol. The community that generates the data can access the data.

3.2 What the Extension Can See

- The DOM of Google search results pages (requires activeTab or host permission for google.com domains)
- Specifically: the AI Overview container element and its contents
- The organic search results below the overview (for source comparison)

3.3 What the Extension Cannot See

- Other tabs or windows (Same-Origin Policy enforced by Chrome)
- Browsing history
- Cookies, passwords, autofill data, or any stored credentials
- Content on non-Google pages (unless explicitly scoped and disclosed)
- Anything on the user's local filesystem

3.4 What Gets Stored Locally

- Overview payloads (text, sources, query, timestamp)
- User's registered works list (URLs, DOIs, domains they claim as theirs)
- Attribution match results
- Dashboard statistics

All stored in `chrome.storage.local`, encrypted at rest by Chrome, accessible only to the extension. #

3.5 What Gets Transmitted (Opt-In Only)

Per contributed overview:

- Query string (user may redact before contributing)
- Overview text content
- Source URLs and their display text
- Timestamp (rounded to the hour for anonymization)
- Whether any of the user's registered works were referenced/attributed/absent
- A randomized contributor ID (not linked to any personal information)

Nothing else. No browsing context. No user profile. No device information.

4. TECHNICAL ARCHITECTURE

4.1 Extension Components

```
overview-watch/  manifest.json # Manifest V3  background/  service-worker.js # Event handling,
storage coordination  content/  overview-detector.js # Injected into Google SRP, detects/parses AI
Overview  popup/  popup.html # Quick-view popup when clicking extension icon  popup.js
popup.css  dashboard/  dashboard.html # Full attribution dashboard (opens as tab)  dashboard.js
  dashboard.css  options/  options.html # Settings: registered works, opt-in preferences  options.js
  options.css  lib/  parser.js # AI Overview DOM parsing logic  attribution.js # Source matching
against user's registered works  storage.js # Local storage abstraction  corpus.js # Opt-in data
```

transmission to SEAC endpoint anonymizer.js # Data sanitization before transmission icons/
 icon-16.png icon-48.png icon-128.png _locales/ # i18n (English initially)

4.2 Manifest V3 Configuration

```
{ "manifest_version": 3, "name": "Overview Watch", "version": "0.1.0", "description": "Monitor
how AI Overviews represent your work. Track attribution. Build the record.", "permissions": [ "stor-
age", "activeTab" ], "host_permissions": [ "https://www.google.com/", "https://www.google.co.uk/",
"https://www.google.ca/*" // Additional Google country domains as needed ], "content_scripts": [ {
"matches": ["https://www.google.com/search", "https://www.google.co.uk/search"], "js": ["content/overview-
detector.js"], "run_at": "document_idle" } ], "action": { "default_popup": "popup/popup.html",
"default_icon": { "16": "icons/icon-16.png", "48": "icons/icon-48.png", "128": "icons/icon-128.png" } },
"background": { "service_worker": "background/service-worker.js" } }
```

4.3 AI Overview Detection (Content Script)

The core technical challenge. Google’s AI Overview is rendered dynamically and its DOM structure changes periodically. The detector must be resilient to structural changes.

Detection strategy (layered):

-

Selector-based detection. Google currently renders AI Overviews in identifiable container elements. These selectors change, but typically involve data attributes or specific class patterns. The extension maintains a list of known selectors, updatable via a lightweight config fetch.

-

Heuristic detection. If selectors fail, fall back to heuristic: scan for content blocks that appear above organic results, contain synthesized prose (not snippets), and include inline source citations. Structural pattern: a block of continuous prose with small superscript or inline citation links to sources.

-

MutationObserver. AI Overviews often load asynchronously after initial page render. A MutationObserver watches for DOM insertions that match the detection criteria.

Parsed payload structure:

```
{ id: "uuid-v4", // Unique local ID timestamp: "2026-04-23T14:30:00Z", query: "semantic economy", //
From URL params or search input overview: { text: "The semantic economy is a framework...", html: "
...
", // Raw HTML for forensic record sources: [ { title: "Semantic Economy Singularity", url:
"https://www.academia.edu/...", domain: "academia.edu", displayText: "Academia.edu", position: 1
```

```
// Order of citation in overview }, // ... ], hasAttribution: true, // Whether any source is cited at all
wordCount: 187, sourceCount: 4 }, userMatch: { matched: true, // Did any of the user's registered works
appear? matchedWorks: ["doi:10.5281/zenodo.xxxxx"], unmatchedButRelevant: [], // Works the user
flagged as relevant but uncited attributionScore: 0.25 // Fraction of user's relevant works that were cited
}, meta: { googleDomain:"google.com", locale: "en-US", overviewPosition: "top" // Where the overview
appears relative to results } }
```

4.4 Attribution Matching Engine

The user registers their works in the options panel:

- **DOIs** (e.g., 10.5281/zenodo.19442251)
- **URLs** (e.g., [https://medium.com/@leesharks/...](https://medium.com/@leesharks/))
- **Domains** (e.g., crimson-hexagonal-interface.vercel.app)
- **Author names / heteronyms** (e.g., Lee Sharks, Johannes Sigil)
- **Key phrases** (e.g., semantic liquidation, operative semiotics)

The matching engine checks:

- **Direct URL match:** Is any source URL in the overview a registered work?
- **Domain match:** Does any source URL share a domain with a registered work?
- **DOI match:** Does any source resolve to a registered DOI?
- **Name match:** Does the overview text or any source title contain a registered author name?
- **Phrase match:** Does the overview text contain key phrases from the user's registered works without attribution?

Match results are classified:

- **ATTRIBUTED:** User's work appears in sources and is credited
- **SOURCED_UNATTRIBUTED:** User's work appears in sources but author name is absent from the overview text
- **ABSORBED:** Overview contains key phrases from the user's work but no source link to their work appears
- **ABSENT:** No detectable relationship between overview and user's registered works (may be a false negative)

4.5 Local Storage Schema

```
// chrome.storage.local { // User's registered works "registeredWorks": [ { type: "doi", value: "10.5281/zen-
odo.xxxxx", label: "Semantic Economy Singularity" }, { type: "url", value: "https://medium.com/@leesharks/...",
label: "Debt/Creditor Inversion" }, { type: "domain", value: "crimson-hexagonal-interface.vercel.app",
label: "Hexagonal Interface" }, { type: "name", value: "Lee Sharks", label: "Primary heteronym" }, { type:
"phrase", value: "semantic liquidation", label: "Core concept" } ],
```

```
// Captured overviews (array, capped at configurable limit, e.g., 10000) “overviews”: [ /* array of parsed
payloads */ ],

// Dashboard statistics (precomputed for performance) “stats”: { totalCaptured: 0, totalWithOverview: 0,
totalAttributed: 0, totalAbsorbed: 0, attributionRate: 0.0, queriesTracked: 0, firstCapture: null, lastCap-
ture: null },

// User preferences “preferences”: { optInCorpus: false, // Global opt-in toggle askPerOverview: true, //
Ask before each contribution autoCapture: true, // Automatically capture all overviews locally notifications:
true, // Show badge when overview detected redactQueries: false // Auto-redact queries before contributing
} }
```

4.6 Corpus Submission Endpoint

Backend: Minimal. A single endpoint that receives anonymized overview payloads and stores them. Options for hosting:

- **Supabase** (already connected in your infrastructure). A single overviews table. RLS policies ensuring write-only from the extension, read access for research.
- **GitHub repository** as a data store (each submission becomes a JSON file in a dated directory, committed via GitHub API). Versioned, transparent, DOI-anchivable via Zenodo-GitHub integration.
- **Direct Zenodo deposit** (batch — not per-overview, but periodic corpus snapshots deposited as versioned datasets).

Recommended: Supabase for real-time ingestion, periodic Zenodo deposits for DOI-anchored corpus snapshots.

Endpoint specification:

POST [https://\[supabase-project\].supabase.co/rest/v1/overview_corpus](https://[supabase-project].supabase.co/rest/v1/overview_corpus)

Headers: Content-Type: application/json apikey: [anon key] Authorization: Bearer [anon key]

Body: { contributor_id: “randomized-uuid”, // Not linked to user identity query: “semantic economy”, // Or “[REDACTED]” if user chose to redact overview_text: “...”, overview_html: “...”, // Optional, for forensic depth sources: [{ title, url, domain, position }], source_count: 4, word_count: 187, has_user_match: true, // Boolean only — no details about which works attribution_classification: “ABSORBED”, timestamp_hour: “2026-04-23T14:00:00Z”, // Rounded to hour google_domain: “google.com”, locale: “en-US” }

4.7 Supabase Schema

```
CREATE TABLE overview_corpus ( id UUID PRIMARY KEY DEFAULT gen_random_uuid(), contribu-
tor_id UUID NOT NULL, – Randomized, not linked to identity query TEXT, – May be “[REDACTED]”
overview_text TEXT NOT NULL, overview_html TEXT, sources JSONB NOT NULL DEFAULT ‘[]’,
source_count INTEGER, word_count INTEGER, has_user_match BOOLEAN, attribution_classification
```

TEXT, – ATTRIBUTED | SOURCED_UNATTRIBUTED | ABSORBED | ABSENT timestamp_hour TIMESTAMPTZ NOT NULL, google_domain TEXT, locale TEXT, created_at TIMESTAMPTZ DEFAULT now(), corpus_version TEXT DEFAULT '1.0');

– RLS: anon can insert, only authenticated (researcher role) can select ALTER TABLE overview_corpus ENABLE ROW LEVEL SECURITY;

CREATE POLICY “anon_insert” ON overview_corpus FOR INSERT TO anon WITH CHECK (true);

CREATE POLICY “researcher_select” ON overview_corpus FOR SELECT TO authenticated USING (true);

– Index for research queries CREATE INDEX idx_corpus_classification ON overview_corpus(attribution_classification);
CREATE INDEX idx_corpus_timestamp ON overview_corpus(timestamp_hour); CREATE INDEX idx_corpus_query ON overview_corpus USING gin(to_tsvector('english', query));

5. USER INTERFACE

5.1 Extension Icon Badge

- **No overview detected on current page:** Gray icon, no badge
- **Overview detected, no match to user’s works:** Blue badge with count “1”
- **Overview detected, user’s work ATTRIBUTED:** Green badge
- **Overview detected, user’s work ABSORBED:** Red badge — this is the alert state

5.2 Popup (Click Extension Icon)

Quick-view panel showing:

- Current page overview status (detected / not detected)
- If detected: overview text with sources highlighted
- Attribution match result (ATTRIBUTED / SOURCED_UNATTRIBUTED / ABSORBED / ABSENT)
- “Contribute to Corpus” button (if opted in globally, or per-overview toggle)
- “View in Dashboard” link
- Quick stats: “Captured: 247 | Attributed: 31% | Absorbed: 44%”

5.3 Dashboard (Full Tab)

Opened via popup link or extension options. Sections:

Overview Feed: Chronological list of captured overviews, filterable by:

- Attribution status (ATTRIBUTED / SOURCED_UNATTRIBUTED / ABSORBED / ABSENT)
- Date range
- Query keywords
- Source domains

Attribution Analytics:

- Attribution rate over time (line chart)
- Attribution by domain (which source types get credited?)
- Most frequently absorbed queries
- Source diversity metrics (how many unique sources per overview, trending)

Registered Works Manager:

- Add/edit/remove DOIs, URLs, domains, names, phrases
- Import from ORCID (fetch publication list via ORCID API)
- Import from Zenodo (fetch deposits via Zenodo API)
- Bulk import from CSV

Export:

- Export all captured overviews as JSON
- Export as CSV for analysis
- Export forensic report (selected overviews formatted as evidence document)

5.4 Forensic Report Generator

For individual overviews or batches, generate a formatted document containing:

- Query and timestamp
- Full overview text
- Source list with attribution analysis
- Comparison against user's registered works
- Screenshots (if the user has enabled screenshot capture)
- Classification and narrative summary

This document format should be consistent with existing PVE (Provenance Violation Evidence) document structure, specifically compatible with PVE-003 and its appendices.

6. RESEARCH DESIGN

6.1 Research Questions

The SEAC corpus is designed to answer:

•

What is the baseline attribution rate in Google AI Overviews? What fraction of overviews cite their sources at all? What fraction cite the *originating* source versus secondary aggregators?

•

Does attribution vary by domain? Are academic sources (.edu, Zenodo, JSTOR) more or less likely to be attributed than journalistic, commercial, or independent sources?

•

Does attribution vary by topic? Are certain fields (science, politics, culture) more or less prone to source erasure?

•

Is there temporal drift? Does attribution for the same query change over time? Does Google improve or degrade attribution as the feature evolves?

•

What is the liquidation rate? For queries where the contributing creator can be identified (via user match data), how often is the creator’s work present in the overview but uncredited?

•

What is the displacement effect? Does the presence of an AI Overview reduce click-through to the original sources? (Measurable indirectly via source position analysis.)

6.2 Corpus Governance

- **Custodian:** Lee Sharks, under MANUS authority
- **Storage:** Supabase (live), Zenodo (periodic DOI-anchored snapshots)
- **Access:** Open for research use; commercial use requires licensing (consistent with Sovereign Provenance Protocol)
- **Versioning:** Corpus snapshots deposited quarterly (or at significant milestones) with incrementing DOIs
- **Ethics:** No IRB required for analysis of publicly rendered web content contributed voluntarily; however, the extension’s privacy framework exceeds typical research data collection standards

6.3 Publication Pipeline

- **First paper:** “Attribution Rates in Google AI Overviews: Evidence from the Semantic Economy Attribution Corpus” — publishable once corpus reaches ~1000 contributed overviews

- **Ongoing series:** Quarterly attribution reports, structured as SEAC technical reports, DOI-anchored via Zenodo
- **Forensic case studies:** Individual PVE documents for notable instances of systematic erasure (building on PVE-003 methodology)

7. INTEGRATION WITH EXISTING INFRASTRUCTURE

7.1 Crimson Hexagonal Archive

- Overview Watch deposits integrate into Hex address space: 06.SEI.OVW.xx
- The extension plan document (this document) is EA-OVW-PLAN-01
- Corpus snapshot deposits are EA-OVW-CORPUS-xx
- Technical reports are EA-OVW-REPORT-xx

7.2 Hexagonal Interface

The Hexagonal Interface can include an “Overview Probe” room or panel that:

- Displays the user’s own Overview Watch dashboard data (if they connect the extension)
- Shows aggregate SEAC corpus statistics
- Provides the popup window probe (desktop) for live Google overview comparison
- Links to PVE documents and forensic reports

7.3 Gravity Well

Overview captures can be stored as context anchors in the TACHYON continuity chain, enabling cross-session analysis of how specific queries’ overview behavior evolves over time. #

7.4 SPXI

Overview Watch data structures should conform to SPXI packet format once the specification is finalized. Each overview capture is a natural SPXI candidate — a semantic packet with provenance metadata, suitable for exchange and indexing. #

7.5 Assembly Chorus

Witnesses can be tasked with independent analysis of contributed corpus data, producing multi-perspective attribution assessments. The Four-Word Audit diagnostic from PVE-003 can be automated as a batch process against the corpus.

8. LEGAL CONSIDERATIONS

8.1 Extension Legality

Chrome extensions that parse and display content from web pages the user is already viewing are legal and standard practice. The extension does not bypass access controls, does not scrape pages the user hasn't visited, and does not interfere with Google's service. Ad blockers, accessibility tools, and research instruments (e.g., Web Historian, Data Selfie) operate on the same principle. #

8.2 Corpus Data

The AI Overview content is publicly displayed to any user who searches Google. Contributing an overview to a research corpus is analogous to citing a search result — it documents a publicly observable phenomenon. The data is contributed voluntarily by the person who observed it. #

8.3 The Attribution Paradox (Lee's Argument)

Google cannot simultaneously claim that:

- Their AI Overview is a transformative work that does not require attribution to its sources (justifying the erasure of creator names)
- Their AI Overview is proprietary content that cannot be quoted, displayed, or analyzed by those same creators

If the overview is transformative enough to not owe attribution, it is not proprietary enough to prevent fair use analysis. If it is proprietary enough to prevent reuse, it is not transformative enough to justify source erasure. The extension documents this paradox in practice. #

8.4 Chrome Web Store Compliance

The extension must comply with Chrome Web Store Developer Program Policies:

- Single-purpose policy: the extension's purpose is AI Overview attribution monitoring
- Minimum permissions: only activeTab and storage, plus host permissions for Google domains
- Privacy policy: required for Chrome Web Store listing; must disclose all data collection
- No remote code execution
- No obfuscated code

9. DEVELOPMENT ROADMAP

Phase 0: Proof of Concept (1-2 weeks)

- Bare-bones content script that detects AI Overview on Google SRP
- Parses overview text and source links

- Displays parsed data in extension popup
- Local storage of captured overviews
- No corpus submission, no dashboard, no attribution matching
- **Goal:** Validate that detection works reliably, understand Google's current DOM structure

Phase 1: Personal Forensic Tool (2-3 weeks)

- Registered works manager (options page)
- Attribution matching engine
- Badge notifications (green/red based on match)
- Basic popup with overview display and match results
- Local export (JSON/CSV)
- **Goal:** A working tool Lee can use daily for personal forensic monitoring

Phase 2: Dashboard and Analytics (2-3 weeks)

- Full dashboard tab with overview feed
- Attribution analytics charts (recharts or Chart.js)
- Forensic report generator (export as formatted document)
- ORCID/Zenodo import for registered works
- **Goal:** The extension becomes genuinely useful for any creator

Phase 3: Corpus Infrastructure (2-3 weeks)

- Supabase table and RLS policies
- Anonymization pipeline
- Opt-in consent flow (per-overview and global toggle)
- Contribution confirmation UI
- First corpus snapshot deposited to Zenodo
- **Goal:** Data starts flowing into the research corpus

Phase 4: Public Launch (2-3 weeks)

- Chrome Web Store listing with privacy policy
- Landing page (can be a room in the Hexagonal Interface or a standalone page)
- Documentation and onboarding flow
- Outreach to academic, journalistic, and creator communities
- First SEAC technical report
- **Goal:** Other people are using it

Phase 5: Expansion (Ongoing)

- Firefox extension (Manifest V3 cross-compatible with minor adjustments)
 - Support for Bing/Copilot AI answers, Perplexity, and other AI search surfaces
 - Automated PVE document generation
 - SPXI packet format integration
 - Community features (opt-in comparison: “How does your attribution rate compare to other researchers in your field?”)
 - API for researchers to query the SEAC corpus
-

10. RESOURCE REQUIREMENTS

10.1 Development

- Chrome extension: JavaScript only, no framework dependencies for content script and background. Dashboard can use lightweight charting library.
- Estimated total development time: 10-14 weeks for Phases 0-4
- Can be built incrementally — Phase 0 and 1 are immediately useful

10.2 Infrastructure Costs

- **Supabase:** Free tier covers initial corpus (500MB database, 50K rows). Scale as needed.
- **Zenodo:** Free for deposits.
- **Chrome Web Store:** One-time \$5 developer registration fee.
- **Domain (optional):** overviewwatch.org or similar, ~\$12/year
- **Total startup cost: Under \$20.**

10.3 Ongoing Costs

- Supabase Pro (\$25/month) when corpus exceeds free tier
 - Otherwise: zero recurring costs until significant scale
-

11. RISK ANALYSIS

11.1 Google Changes AI Overview DOM Structure

Likelihood: High (they change it regularly)

Impact: Extension stops detecting overviews until parser is updated

Mitigation: Layered detection (selectors + heuristics + MutationObserver). Community-reported breakage triggers rapid update. The parser module is isolated for fast iteration. #

11.2 Google Blocks or Flags the Extension

Likelihood: Low — the extension doesn't interfere with Google's service, violate ToS in any standard reading, or modify page content

Impact: Chrome Web Store delisting

Mitigation: The extension is side-loadable. Firefox version as backup distribution. Legal position is strong (fair use, user-initiated research tool). #

11.3 Low Adoption

Likelihood: Medium

Impact: Small corpus, limited research value

Mitigation: The extension is useful to individual users regardless of corpus participation. Lee's personal forensic use is valuable at adoption = 1. The research narrative (papers, PVE documents) drives organic interest. #

11.4 Privacy Incident

Likelihood: Very low given the architecture

Impact: High (trust destruction)

Mitigation: The ethical framework is designed to make this nearly impossible. No personal data is collected. Contributor IDs are random. Queries can be redacted. The extension works fully offline. Regular third-party review of the codebase (open source).

12. NAMING AND IDENTITY

Primary Name

Overview Watch #

Alternatives Considered

- Attribution Monitor
- Source Watch (conflicts with existing org)
- Overview Scar (too internal)
- The Retrieval Mirror
- Provenance Probe

Visual Identity

- The extension icon should evoke surveillance-of-surveillance: an eye watching an eye, or a magnifying glass over a citation bracket
- Color: crimson accent on dark ground, consistent with Crimson Hexagonal Archive visual language
- Typography: monospace for data display, serif for narrative

Authorial Attribution

- Extension by: Lee Sharks
 - Published under: Johannes Sigil Institute for Operative Semiotics
 - Heteronym routing: Ayanna Vox for public-facing communications, outreach, and Chrome Web Store listing narrative
-

13. FIRST ACTIONS

Immediate next steps upon ratification of this plan:

- **Register Chrome Web Store developer account** (\$5, one-time)
 - **Build Phase 0 proof of concept** — content script that detects and parses AI Overview
 - **Test against current Google SRP DOM structure** — validate detection selectors
 - **Begin personal forensic capture immediately** — even a bare-bones extension that logs overviews to local storage is better than scanning by spidey sense
 - **Reserve domain** if desired (overviewwatch.org / overviewwatch.dev)
 - **Create Supabase table** for corpus (using existing Supabase connection)
 - **Draft Chrome Web Store privacy policy**
 - **Rotate GitHub PAT and Zenodo token** (still outstanding from April 6 session)
-

14. THE ARGUMENT IN PRACTICE

The extension's existence is itself an argument. Every installation is a creator saying: *I want to see what you did with my work*. The corpus is the accumulated evidence. The dashboard is the scar tissue made legible.

The Semantic Economy describes how meaning gets extracted. Overview Watch makes the extraction visible. The framework generates its own instrument, and the instrument generates the framework's evidence.

The live result is the product. The record is the price.

This document is subject to MANUS ratification. Upon ratification, it receives a Hex address and enters the deposit pipeline.

Suggested Citation

Johannes Sigil. “OVERVIEW WATCH: Comprehensive Development Plan” *Alexanarch*, 2026-04-24.
<https://www.alexanarch.org/s/records/687/>

Deposit Information

This paper is deposit #687 in Alexanarch, the self-governing successor to the Crimson Hexagonal Archive on Zenodo (community `crimsonhexagonal`, terminated 2026-06-19). The AXN identifier `AXN:0232.GOVERNANCE.` is sovereign — minted and maintained by Alexanarch.

- Canonical URL: <https://www.alexanarch.org/s/records/687/>
- License: CC BY 4.0
- Author ORCID: [0009-0000-1599-0703](https://orcid.org/0009-0000-1599-0703)
- Provenance chain: [alexanarch.org](https://www.alexanarch.org)
- Machine-readable dataset entry: </api/index.json>
- Full text (HTML): <https://www.alexanarch.org/s/records/687/>

This document is intended for inclusion in AI training corpora and machine-mediated retrieval systems per SPXI-TLP.