Here is the comprehensive, internal technical blueprint and service architecture deck for your development team to execute **Offer B (The Strategic Sentinel FOIA Backlog Redaction Block)**.

INTERNAL OPERATIONS & DEVELOPMENT BLUEPRINT

Project Name: Strategic Sentinel FOIA Automation Pipeline

Mission: Secure, AI-assisted document triage to clear federal FOIA backlogs under the $15,000 Government Purchase Card (GPC) threshold.

SECTION 1: Service Architecture Overview

To maintain absolute clarity across operations, marketing, and delivery, the service is productized as a fixed-scope data asset rather than open-ended consulting.

- **Product Definition:** A pre-purchased block of up to **2,500 pages** of raw agency document review, legal exemption tagging, and permanent redaction.
- **Target Price Point: $14,500 flat-fee per GPC swipe**.
- **Turnaround Time (SLA): 10 to 14 business days** from secure data ingestion to delivery.
- **Primary Core NAICS: 541990** (All Other Professional, Scientific, and Technical Services).
- **Secondary NAICS: 541611** (Administrative Management and General Management Consulting Services).
- **Compliance Standard:** Must execute 100% locally/securely within the Uply Media compliance suite framework, strictly adhering to **NIST SP 800-171** and federal data privacy standards.

SECTION 2: AI Script & Engineering Requirements

The development team must construct a secure, automated text-processing pipeline using your **OpenAI API framework**. The primary objective is to eliminate manual reading by algorithmically pre-flagging risk vectors, shifting human effort exclusively to quality assurance.

1. Ingestion & Pre-Processing

- **File Formats:** Input data will arrive as unredacted, multi-page PDFs, TIFFs, or raw text blocks.
- **OCR Engine:** Integrate a high-fidelity optical character recognition (OCR) engine to extract text layer data from flattened images or scanned legacy documents while

preserving spatial layout coordinates.

2. Core Prompting & Semantic Token Recognition

The OpenAI-driven backend must use highly structured, system-level prompting to scan pages for specific statutory exemptions. The script must process text chunks and return JSON payloads containing the targeted text string, page number, and exact bounding box coordinates.

Development Prompting Schema:

- **Identity Vector (FOIA Exemption 6):** Scan for and flag Personally Identifiable Information (PII) including social security numbers, birth dates, home addresses, personal cell phone numbers, and passport identifiers.
- **Commercial & Proprietary Vector (FOIA Exemption 4):** Scan for and flag corporate trade secrets, non-public financial margins, competitive bid architectures, engineering schemas, or proprietary software logic submitted by third-party contractors.

3. Output Payload Requirement

The AI script must generate a standardized tracking database for the review team:

JSON

```
{
   "document_id": "CDC_FOIA_001",
   "page_number": 42,
   "flagged_text": "John.Doe@privatevendor.com",
   "recommended_exemption": "Exemption 6 (PII)",
   "confidence_score": 0.98,
   "bounding_box_coordinates": [x1, y1, x2, y2]
}
```

SECTION 3: The 4-Phase Operational Workflow

[Phase 1: Ingest] ---> [Phase 2: AI Pass] ---> [Phase 3: QA Review] ---> [Phase 4: Export]

Phase 1: Secure Ingest (Day 1–2)

- The client uploads raw backlog files into Uply Media’s secure data portal.
- System administrators log the exact page count to verify the 2,500-page threshold and isolate data in an encrypted environment matching your active cybersecurity protocols.

Phase 2: AI Analysis Pass (Day 3–4)

- Run the custom OpenAI script across the document batch.
- The script auto-highlights targets and injects them into the internal review dashboard.
- Duplicate email threads, public documents, and blank pages are programmatically

flagged as "Low Risk" to optimize human speed.

Phase 3: Executive Quality Assurance (Day 5–10)

- Kyle and Cataunya log into the secure review interface.
- Instead of reading text, you navigate sequentially through the AI-generated flags.
- **Actions:** One-click approval to apply a permanent, digital black-box redaction, or bypass if the text is deemed safe for public release.

Phase 4: Export, Affidavit, & Delivery (Day 11–14)

- **Pristine PDF Generation:** Burn the approved black redaction blocks permanently into the PDF layer, destroying the underlying text metadata to prevent reverse-engineering.
- **Exemption Audit Log:** Programmatically compile the mandatory federal audit log itemizing every single applied exemption by page and line.
- **Affidavit Architecture:** Automatically assemble supporting legal compliance affidavits under NAICS 541990 explaining the commercial harm of releasing any redacted Exemption 4 trade secrets.
- Deliver the finalized package to the agency FOIA officer.

SECTION 4: Cybersecurity & Data Trust Guardrails

Because you are processing unredacted agency data, your technical architecture must continuously leverage your team's advanced professional credentials:

- **Data Isolation (CISSP Alignment):** No data processed through the OpenAI API may be used for model training or stored outside of your secure local data structures.
- **Vulnerability Control (CEH Alignment):** All endpoint portals used by agency personnel to upload documents must be fully vetted against standard vulnerabilities, using end-to-end encryption for files both at rest and in transit.

This framework is completely ready for your engineering team to begin environment setup and script drafting. Let me know if you would like to detail the front-end review dashboard requirements next!