

ReviewOS for MediGen

Read-only corpus retrieval, citation, and evidence reuse over a trusted life-sciences document base.

MVP BUILD STAGES

Approve the staged build to unlock ~\$600K in annual productivity-equivalent value.

RECOMMENDATION

Build the search product directly with MediGen first. MVP Stage 0 should master permission-aware semantic search, retrieval, and citation quality over existing documentation before the team scales into citation agents, ingestion workflows, or a full ReviewOS app.

~29.6K hrs

ANNUAL REVIEW WORK TODAY

~13.3K hrs

SEARCH/CITATION POOL

~4.0K hrs

POST-MVP TIME UNLOCKED

~\$600K/yr

VALUE CASE

CURRENT PROBLEM

- Evidence exists, but is fragmented across approved repositories, prior submissions, quality records, legal work product, and research folders.
- Researchers lose time moving from question to source, then from source to exact passage.
- Citations and summaries drift as evidence moves from PDFs into spreadsheets, memos, and review notes.
- Experts are pulled into source-checking instead of final judgment.

MVP STAGE 0 SCOPE

- ✓ MVP Stage 0 is the AI semantic search engine: Corpus API + MCP server.
- ✓ One-time ingestion of existing corpus with permissions and source metadata.
- ✓ Hybrid semantic + BM25 retrieval with metadata filters.
- ✓ Read-only API and MCP tool callable from approved AI clients.
- ✓ Citation-first evidence objects with source text, page, chunk, score, and metadata.
- ✓ Query logs, failed-parse logs, indexing health, golden-set eval, and adversarial set.

VALUE HYPOTHESIS

A 25-person team carries roughly 29,600 annual review hours and a \$4.1-\$4.5M loaded cost pool. If search, retrieval, citation, and summarization represent ~45% of that work, the addressable pool is ~\$2.0M. A 30% MVP unlock produces roughly ~\$600K in annual productivity-equivalent value.

RISK FOCUS

The risk is not whether the team can add more workflow features. The risk is search quality. ReviewOS should earn trust by getting retrieval, permissions, citation metadata, and eval gates right before broadening the product surface.

SUCCESS METRICS

Time to first useful source	<120 sec
Known-answer retrieval	>80% top-5
Citation metadata coverage	>95%
Retrieval latency	<10 sec

PRODUCT PATH

MVP Stage 0: Corpus API + MCP server.
MVP Stage 1: controlled ingestion and corpus maintenance.
MVP Stage 2: headless citation and cross-reference agent.
MVP Stage 3: ReviewOS app only after usage proves the workflow.