

INTRODUCING

Doctrine

The AI Pipeline for Authoritative Documents

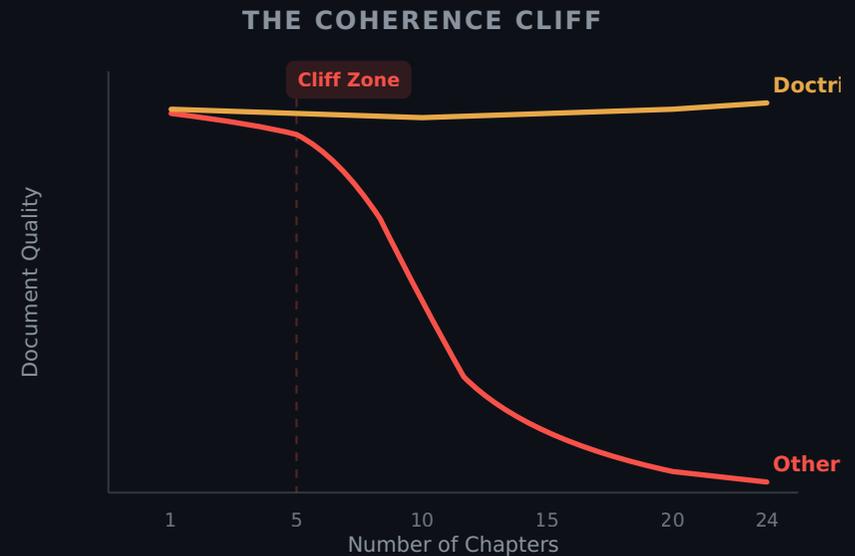
SEED ROUND — \$2.5M

Ryan Muranaka · Founder & CEO · March 2026

THE PROBLEM

AI Can't Write Anything Longer Than a Blog Post

- Every AI writing tool loses coherence after ~5 chapters
- ChatGPT, Claude, Gemini — none can maintain voice, facts, or narrative across 24 chapters
- The AI writing market is **\$6.5B by 2030** — with **zero** solutions for long-form coherence
- Professionals need 50–200 page documents: reports, briefs, manuscripts, white papers
- Current workaround: manual stitching, days of editing, or giving up entirely



THE SOLUTION

A 7-Stage AI Pipeline with Proprietary Coherence Architecture

One click. 45 minutes. A 24-chapter document that reads like one person wrote it in one sitting.



Auto-Pilot Mode

Entire pipeline runs unattended. Upload a concept, come back to a finished document.

The Bible

Auto-generated master reference: cross-references, voice rules, fact registry, transition architecture.

Self-Improving

AI reviews its own output.
Prompts evolve automatically.
Each run is better than the last.

KILLER FEATURE

The Response Pipeline

Upload ANY document. Get a book-length response from ANY perspective.

- Same source document → radically different outputs depending on perspective
- Use cases: think tanks, law firms, political campaigns, academic rebuttals, strategic analysis
- Every output is shareable content that markets the product

 President

 CEO

 Military Commander

 God

 Researcher

 Legal Counsel

 Custom

The viral hook:

***"What would God say
about your thesis?"***

One document. Six perspectives.
Six shareable products.

THE SECRET SAUCE

Bible Methodology

Write fast. Unify later. Ship perfect.

✗ Without Bible

Ch 1 — Formal tone, introduces "Dr. Smith"

Ch 2 — Casual tone, calls her "Sarah"

Ch 3 — Contradicts Ch 1 statistics

Ch 4 — Repeats Ch 2 opening paragraph

Ch 5 — No transition from Ch 4

Incoherent. Contradictory. Unusable.

✓ With Bible

Ch 1 — Consistent voice, establishes Dr. Smith

Ch 2 — Same voice, proper name reference

Ch 3 — Cross-verified statistics match

Ch 4 — Unique content, smooth transition

Ch 5 — Builds on Ch 4 naturally

Reads like one person wrote it in one sitting.



Cross-Reference Map

Every character, fact, and claim tracked across all chapters

Voice & Consistency Registry

Tone rules, terminology standards, style enforcement

Transition Architecture

Chapter-to-chapter flow designed before rewrite

PRODUCT

Built & Working Today

Auto-Pilot Mode

One-click pipeline execution



Stage 5/7 — Bible Generation

Perspective Selector

Choose any viewpoint for response documents

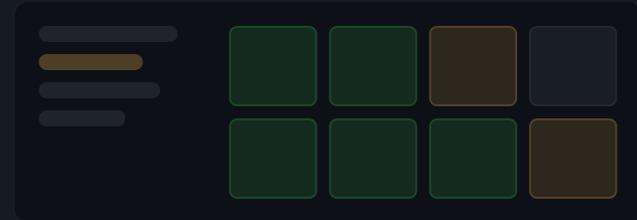
 President

 CEO

 Commander

Main Dashboard

Dark theme · Project sidebar · Chapter grid with status indicators



MARKET OPPORTUNITY

A Massive, Underserved Market



Target Segments

- **Think Tanks & Policy Orgs** — multi-hundred-page reports, fast turnaround
- **Consulting Firms** — client deliverables, white papers, strategy docs
- **Law Firms** — briefs, memoranda, case analysis documents
- **Authors & Academics** — manuscripts, dissertations, research papers
- **Political Campaigns** — position papers, opposition research, messaging

BUSINESS MODEL

Pay-Per-Use + Subscriptions

Margins That Scale

Starter

\$5

5 chapters
Sonnet model
Standard export

Standard

\$15

15 chapters
Sonnet model
Full pipeline

Professional

\$29

25 chapters
Opus model
Bible methodology

Enterprise

\$49

50 chapters
Multi-LLM
API access

\$23

Cost per 24ch Opus run

76%+

Pro subscriber margin
(2+ runs/month)

85%+

Enterprise margin

Subscriptions: Free (Sonnet, 1 project) · Pro \$49/mo (Opus, unlimited) · Enterprise \$199/mo (multi-LLM, teams)

COMPETITION

No One Else Has More Than One

Feature	Doctrine	Sudowrite	NovelCrafter	Squibler	Jasper	ChatGPT
Multi-chapter coherence	✓	—	—	—	—	—
Auto-generated Bible	✓	—	—	—	—	—
Response pipeline	✓	—	—	—	—	—
Perspective shifting	✓	—	—	—	—	—
Self-improving prompts	✓	—	—	—	—	—
Consistency checking	✓	✓	—	—	—	—
One-click pipeline	✓	—	—	✓	—	—
Production error fixing	✓	—	—	—	—	—

8 / 8 features vs. closest competitor: 1 / 8

TRACTION

Working Product. Not a Prototype.

24

Chapters generated
in a single pipeline run

45min

End-to-end pipeline
completion time

\$23

Total cost per
full document generation

Key Milestones

- **Full pipeline operational** — not a demo, not a mockup
- First document: **24-chapter presidential response** to a geopolitical manifesto
- Pipeline self-review **identified and fixed its own weaknesses** autonomously
- Built by a **solo founder in days** — extreme capital efficiency

*"We didn't raise to build a product.
We built the product, then raised
to scale it."*

Guerrilla First. **Scale Fast.**

PHASE 1 · MONTHS 1–3

Guerrilla Launch

Viral "Famous Document Response" series. Product Hunt launch. Reddit/Twitter organic. Free first run for every user.

PHASE 2 · MONTHS 4–6

Content & Community

"Upload your worst AI chapter" conversion challenge. Build-in-public. Influencer seeding to writing communities.

PHASE 3 · MONTHS 7–12

Paid + Partnerships

LinkedIn ads targeting consultants. Think tank pilot programs. Law firm integrations. First enterprise sales hire.

Viral Mechanic

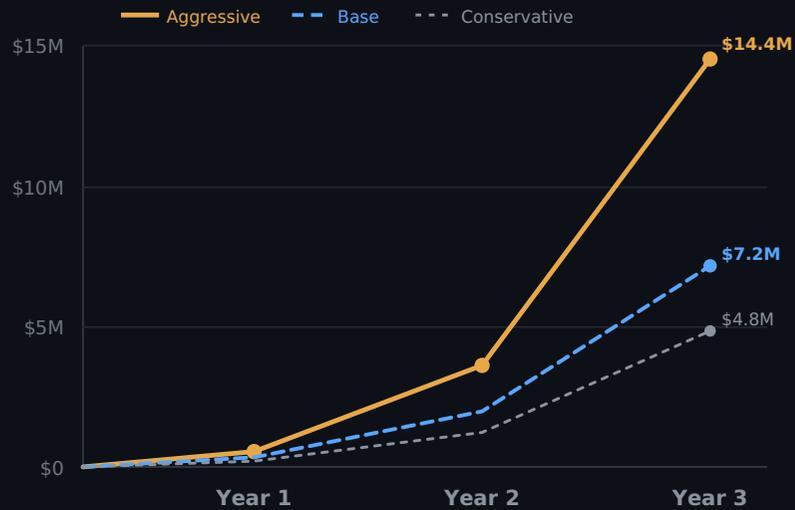
Every output is shareable content that markets the product. A user's 24-chapter presidential response becomes a viral artifact — and a Doctrine advertisement.

Conversion Engine

Free first run → experience the full pipeline → upsell to pay-per-use or subscription. Users who see the Bible methodology don't go back.

FINANCIALS

Path to \$14M+ ARR in Year 3



Key Assumptions

- 5% monthly user growth
- 8% free-to-paid conversion rate
- \$35 average revenue per paying user
- Path to profitability: **Month 8–14**

YEAR 1 ARR

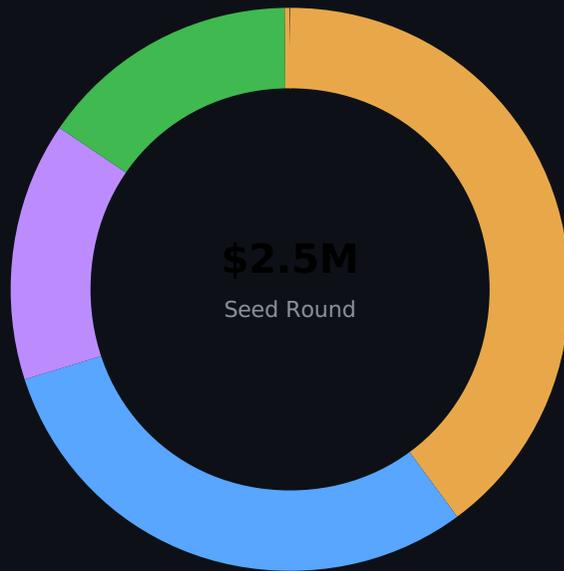
\$180K–\$540K

YEAR 2 ARR

\$1.2M–\$3.6M

USE OF FUNDS

\$2.5M Seed — 18 Months of Runway



Engineering — 40% / \$1M

Multi-tenant SaaS, multi-LLM support (GPT-4o, Gemini), API, enterprise features

Growth — 30% / \$750K

Guerrilla campaign, content marketing, community building, first sales hire

Infrastructure — 15% / \$375K

Cloud hosting, GPU compute, security & compliance

Operations — 15% / \$375K

Legal (IP protection), 18-month runway buffer

TEAM

Capital-Efficient.
Product-First.



Ryan Muranaka

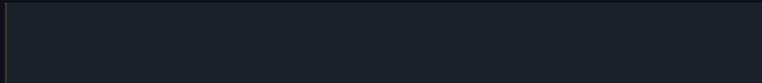
Founder & CEO

-  Built entire platform solo — in days
-  Deep AI & defense technology expertise
-  Technical founder who ships
-  Product built before raising a dollar

"We're not asking you to fund an idea. We're asking you to scale a working product."

Why This Founder

- Extreme capital efficiency — entire product built pre-raise
- Domain expertise at the intersection of AI, defense, and emerging tech
- Technical depth to evolve the architecture as LLMs advance
- Builder mentality — execution over fundraising theater



Advisory Board

Building out — positions available for AI/ML, Enterprise SaaS, and Publishing domain experts

THE ASK

Seed Round — \$2.5M

RAISING
\$2.5M

PRE-MONEY
\$12.5M

12-Month Milestones

M2 Multi-tenant SaaS launch

M4 100 paying users

M6 First enterprise pilot

M12 1,000+ paying users

Our Moats

Bible Methodology

Proprietary coherence architecture — our defensible IP

Response Pipeline

Our wedge into the market — viral by design

Self-Improving System

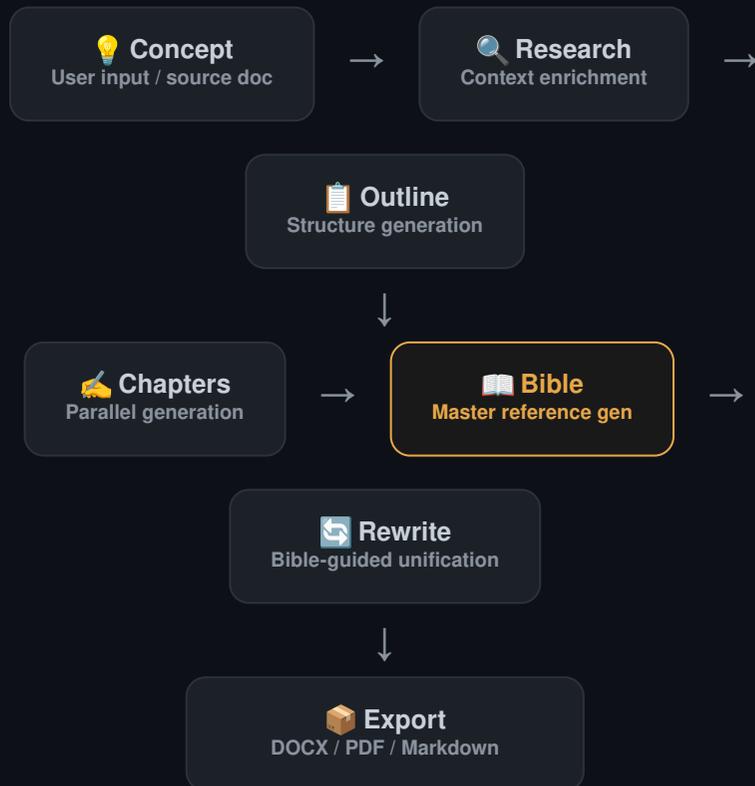
Our flywheel — every run makes the next one better

*The Bible methodology is our moat.
The Response Pipeline is our wedge.
The self-improving system is our flywheel.*

APPENDIX

Technical Architecture

7-Stage Pipeline Detail



Bible Document Components

Cross-Reference Map

Tracks every entity, claim, and reference across all chapters. Ensures Ch 18 doesn't contradict Ch 3.

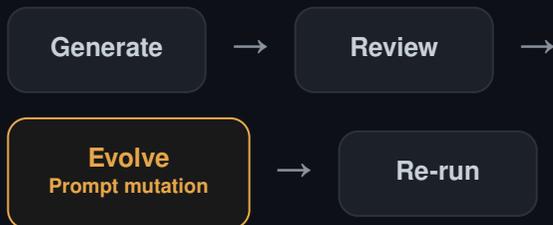
Fact Registry

Single source of truth for all statistics, dates, names, and claims used in the document.

Voice Rules

Tone, vocabulary, sentence structure, and stylistic rules derived from initial chapters.

Self-Improving Feedback Loop



Tracks terminology, abbreviations, and naming conventions for uniform usage.

Transition Architecture

Pre-planned chapter-to-chapter bridges ensuring narrative flow before rewrite begins.

Future: Multi-LLM Comparison

Run the same pipeline through Claude, GPT-4o, and Gemini simultaneously. Compare outputs. Pick the best chapters from each. **Enterprise roadmap feature.**

