

THE LIBRARY

# The MindHYVE Sales Library.

Twelve documents. One library. Built for every sales conversation MindHYVE.ai will have in 2026 — from the first inbound email to the procurement-committee read.

IDENTITY — WHO MINDHYVE IS, WHAT WE OPERATE

|    |              |                                                                         |
|----|--------------|-------------------------------------------------------------------------|
| 01 | Master Index | <i>This document. Your map to the library.</i>                          |
| 02 | MindHYVE.ai  | <i>The umbrella. Five operating systems, one corpus, one institute.</i> |
| 03 | ChironAI™    | <i>The Agentic Healthcare OS. Two editions in market.</i>               |
| 04 | ArthurAI™    | <i>The Agentic Learning OS. Six editions; five in market.</i>           |
| 05 | JustineAI™   | <i>The Agentic Legal OS. PI and WC in market; four on the roadmap.</i>  |
| 06 | TheoAI™      | <i>The Agentic Theology OS. Three products on one corpus.</i>           |

WORKFORCE — THE CREDENTIALING INSTITUTE + FLAGSHIP PROGRAM

|    |                      |                                                                          |
|----|----------------------|--------------------------------------------------------------------------|
| 07 | CIAI                 | <i>California Institute of AI. The credentialing layer.</i>              |
| 08 | AI Professional 360™ | <i>18 courses, 360 hours, one credential. Delivered on ArthurAI VLE.</i> |

ARCHITECTURE + DIFFERENTIATION — WHY WE WIN PROCUREMENT

|    |                           |                                                                       |
|----|---------------------------|-----------------------------------------------------------------------|
| 09 | Eve-Fusion™ + F5/reasoner | <i>The compositional fabric — five reasoning models, per request.</i> |
| 10 | Eve-Genesis™              | <i>100% synthetic by construction. The CISO-facing answer.</i>        |
| 11 | The Differentiation       | <i>MindHYVE vs. the three alternatives — the apex argument.</i>       |

FINANCE — THE FIFTH OPERATING SYSTEM

|    |        |                                                                               |
|----|--------|-------------------------------------------------------------------------------|
| 12 | EliAI™ | <i>The Agentic Finance OS. MRF edition generally available at eligrid.ai.</i> |
|----|--------|-------------------------------------------------------------------------------|

TALK TO OUR TEAM

[hello@mindhyve.ai](mailto:hello@mindhyve.ai) · general inquiries  
[partnerships@mindhyve.ai](mailto:partnerships@mindhyve.ai) · institutional procurement

THE FAMILY ON THE WEB

[mindhyve.ai](https://mindhyve.ai) · [chirongrid.ai](https://chirongrid.ai) · [arthurgrid.ai](https://arthurgrid.ai)  
[justinegrid.ai](https://justinegrid.ai) · [theogrid.ai](https://theogrid.ai) · [eligrid.ai](https://eligrid.ai)  
[ciai.ai](https://ciai.ai) · [ipsc.theogrid.ai](https://ipsc.theogrid.ai)

MINDHYVE.AI · THE PARENT

# Five Agentic Operating Systems. One architecture.



*The Agentic AI pioneer — building since summer 2022, before the term was in industry use.*

*MindHYVE.ai engineers a family of reasoning-first AI Operating Systems for regulated industries — healthcare, education, legal services, Islamic scholarship, and finance — each operated by its own per-vertical subsidiary, all powered by the same Eve-Fusion™ compound architecture. The breadth is the credibility signal: one architecture, five professions.*

01

## ChironAI™

Agentic Healthcare OS. CDS for hospitals, GA since September 2025; OM for California-first workers' comp. Operated by Eve-Healthcare, LLC.

02

## ArthurAI™

Agentic Learning OS. Six editions — SLE, ULE, VLE, CLE, PLE in market, plus FLE for homeschooling. Delivers CIAI's AI Professional 360™ on VLE. Operated by Eve-Education, LLC.

03

## JustineAI™

Agentic Legal OS. PI and WC in market; mass tort next; med mal, bad faith and employment class on the roadmap. Operated by Eve-Legal, LLC.

04

## TheoAI™

Agentic Theology OS. Theo, Majlis, Mizan. 480,334 independently graded hadith. Operated by Eve-Theology, LLC.

05

## EliAI™

Agentic Finance OS. MRF generally available at eligrid.ai — underwriting, valuing and servicing receivables secured by liens. Operated by Eve-Finance, LLC.

06

## IPSC™

Islamic Primary Source Corpus. The research surface — 617,912 hadith cataloged from 123 classical works, every chain parsed, every grade documented. Patent Pending.

**5+1**

OPERATING SYSTEMS  
PLUS ONE RESEARCH CORPUS

**43**

INSTITUTIONAL PARTNERSHIPS  
ACROSS FIVE REGIONS

**3**

OPERATING REGIONS  
NEWPORT BEACH · ISLAMABAD  
· NAIROBI

**Since '22**

BUILDING AGENTIC AI BEFORE  
THE INDUSTRY HAD THE TERM

## ON THE ARCHITECTURE

Every product runs Eve-Fusion™ — a compositional fabric of five cooperating reasoning models: a routing classifier, a small reasoning model fine-tuned on Eve-Genesis™ synthetic data, and three frontier models composed best-fit per request. Five on every vertical. We don't compete with frontier labs. We compose them.

## ON TRUST + COMPLIANCE

Zero customer data trains any product's reasoning. Eve-Genesis™ is 100% synthetic by construction — the architecture cannot leak what is not present. Multi-tenant isolation and audit logging inherited from Microsoft Azure.

CONTINUE READING →

[09 · Eve-Fusion](#)

[10 · Eve-Genesis](#)

[11 · Differentiation](#)

[Visit any site](#)

## TALK TO OUR TEAM

[hello@mindhyve.ai](mailto:hello@mindhyve.ai) general inquiries

[partnerships@mindhyve.ai](mailto:partnerships@mindhyve.ai) institutional procurement

ON THE WEB

**mindhyve.ai**

CHIRONAI™ · THE AGENTIC HEALTHCARE OS

# Practice medicine. ChironAI handles the rest.



*Reasoning-first clinical decision support and occupational medicine. Two editions in market.*

*ChironAI™ is the Agentic Healthcare Operating System for hospitals, clinics, and health-sciences institutions — CDS for general clinical practice, OM for California-first workers' compensation. Reasoning-first, physician-attested, with institutional partners on three continents. The AI reasons; the physician decides.*

01

## CDS Edition

Generally available since September 24, 2025. Differential diagnosis, evidence synthesis, multi-pass radiology read, documentation, prescribing intelligence.

02

## OM Edition

California-first workers' comp. Built for the five roles a WC physician carries. DWC Form 5021, PR-2, AOE/COE causation, MTUS + ACOEM aligned.

03

## Radiology Depth

Multi-pass second-read across 12 frameworks: BI-RADS, LI-RADS, PI-RADS, Lung-RADS, TI-RADS, RECIST, PERCIST, ASPECTS, AO-OTA, BOSNIAK, C-RADS.

04

## Documentation

Pre-visit interviews before patient arrival. Note drafting, problem-list reconciliation, AVS generation. Non-dismissible physician review before any signature.

05

## Full-Chart Review

The patient's whole chart, read as one. Full-chart review and multi-document synthesis as a one-shot query.

06

## Solo + Group

Built for the single-physician practice and the multi-site hospital. Self-serve, sales-assisted, and full institutional procurement paths.

## 6+1

INSTITUTIONAL PARTNERSHIPS  
+ ONE 7-COUNTRY EAC MOU

## 3

CONTINENTS  
USA · PAKISTAN · ETHIOPIA

## GA Sept '25

CDS EDITION GENERALLY  
AVAILABLE TO HOSPITALS

## CA-First

OM HARDENED IN THE MOST  
REGULATED WC MARKET

## ON THE ARCHITECTURE

Eve-Healthcare™ F5/reasoner — the Eve-Fusion™ instance calibrated to clinical reasoning. Five cooperating models composed best-fit per request, one of them a long-context model that holds the full longitudinal chart in a single pass. Imaging is read on a dedicated vision route — a parallel path, not a sixth model. Reasoner fine-tuned on Eve-Genesis™ (Clinical). HIPAA-aligned tenant isolation on Eve-Grid™ Azure.

## ON TRUST + COMPLIANCE

Not FDA-cleared as a medical device. Not currently SOC 2 / ISO 27001 certified — audit-ready. No fabricated accuracy claims. The honest disclosure framing is the credibility signal. Eve-Genesis is 100% synthetic — no PHI in training.

CONTINUE READING →

→ /cds/

→ /om/

09 · Eve-Fusion

10 · Eve-Genesis

## TALK TO OUR TEAM

[hello@mindhyve.ai](mailto:hello@mindhyve.ai) general inquiries

[partnerships@mindhyve.ai](mailto:partnerships@mindhyve.ai) institutional procurement

## ON THE WEB

**chirongrid.ai**

ARTHURAI™ · THE AGENTIC LEARNING OS

# Reasoning quality once locked inside elite institutions — now in every educator’s hands.



Six editions — five in market, plus FLE generally available. 23 institutional partnerships across four regions.

ArthurAI™ is the Agentic Learning Operating System for institutions that teach — K-12 districts, universities, vocational authorities, corporate L&D, self-directed learners, and homeschooling families. The AI reasons; the educator decides.

01

SLE — School

K-12 districts. FERPA + COPPA + ADA/504/508 + AB-1791. Editor-pattern teacher review on every artifact. Sold to the superintendent.

02

ULE — University

Research-aligned reasoning, faculty-of-record positioning, accreditation reporting. CNU and IUCEA going live soon; an MoU with GIKI.

03

VLE — Vocational

Workforce-development boards and authorities. OSHA-aware. Multilingual (EN/ES/Urdu/Indonesian/Hausa). Maps to NAVTTC, EQF, Perkins V.

04

CLE — Corporate

Enterprise L&D and compliance training. SOC 2-readiness posture, GDPR-aligned for EU employees. No HRIS integration required to start.

05

PLE + FLE

Personal Learning Edition for self-directed learners, in market. Family Learning Edition for homeschooling families, generally available.

06

AI Professional 360™

CIAI’s programme on ArthurAI VLE — 18 courses, 360 hours, twelve months.

23

INSTITUTIONAL PARTNERSHIPS  
ACROSS FOUR REGIONS

6

EDITIONS — FIVE IN MARKET  
PLUS FLE, GENERALLY  
AVAILABLE

18

COURSES IN  
AI PROFESSIONAL 360™

250K

PEC ENGINEERS —  
CPD GOING LIVE SOON

## ON THE ARCHITECTURE

Eve-Education™ F5/reasoner — the Eve-Fusion™ instance calibrated to educational reasoning. Five cooperating models composed best-fit per request. Multi-tenant isolation per district, university, or enterprise. Browser-deployable; no SIS/LMS integration required to start.

## ON TRUST + COMPLIANCE

FERPA + COPPA + ADA/504/508 + WCAG 2.1 AA + AB-1791 + state AI-disclosure statutes engineered in by design. No student data trains the reasoning — Eve-Genesis is 100% synthetic. DPA template at /trust/dpa.

CONTINUE READING → → /sle/ → /ule/ → /vle/ 08 · AI Professional 360

## TALK TO OUR TEAM

[hello@mindhyve.ai](mailto:hello@mindhyve.ai) general inquiries  
[partnerships@mindhyve.ai](mailto:partnerships@mindhyve.ai) institutional procurement

## ON THE WEB

[arthurgrid.ai](https://arthurgrid.ai)

JUSTINEAI™ · THE AGENTIC LEGAL OS

# Practice law. JustineAI handles the rest.



*Personal Injury and Workers' Compensation in market. Four additional editions on the roadmap.*

*JustineAI™ is the Agentic Legal Operating System for litigation practice — powering plaintiff-side personal injury and workers' compensation, with four additional editions on the roadmap, each sharing the same compositional fabric and the same supervisor pattern over named sub-agents. The AI reasons; the attorney decides.*

01

## PI Edition

Plaintiff-side personal injury, end-to-end across the nine-status case lifecycle. 31 capabilities live. 14-day free trial for solo and small firms.

02

## Six Named Agents

Justine supervises Intake, Medical, Valuation, Strategy, Deposition, and Case Advisor sub-agents. One supervisor, one MSA, one Model Rule 5.3 posture.

03

## Demand Package

Four-pass demand-letter pipeline + four-layer QA. ICD-10 / CPT validation. CourtListener-verified citations. One-click cover + Chronology + Bill Stack ZIP.

04

## Curation Layer

Nine-step record curation: OCR, provider canonicalization, code validation, encounter coalescing, lien joining.

05

## WC + Roadmap

Workers' Comp in market, with ChironAI™ OM upstream. Mass Tort next; Med Mal, Insurance Bad Faith and Employment Class to follow. Same Eve-Legal F5 — different calibration.

06

## OS Layer

Case management, deadlines, conflict check, trust accounting, time tracking, e-filing, e-signature, client portal. Microsoft 365 + Filevine integration.

9

CASE LIFECYCLE STATUSES  
INTAKE THROUGH CLOSEOUT

31

AI CAPABILITIES SHIPPED  
IN PI EDITION TODAY

50 + DC

JURISDICTIONS WITH HAND-  
CURATED PI PROCEDURAL  
RULES

14-Day

FREE TRIAL FOR SOLO  
AND SMALL PI PRACTICES

## ON THE ARCHITECTURE

Eve-Legal™ F5/reasoner — the Eve-Fusion™ instance calibrated to legal reasoning. Five cooperating models composed best-fit per request, the reasoner fine-tuned on Eve-Genesis™ (Law), and three frontier slots composed best-fit per request. The supervisor-over-sub-agents pattern is shared across every edition.

## ON TRUST + COMPLIANCE

Attorney-client privilege preserved by design — JustineAI operates under the attorney of record's supervision, outputs are work product, audit logs record every reasoning step. No customer data trains the reasoner. Tamper-evident audit chain exportable for ethics review.

CONTINUE READING →

→ [/pi/](#)

→ [/technology/agents/](#)

→ [/founding-firms/](#)

[11 · Differentiation](#)

## TALK TO OUR TEAM

[hello@mindhyve.ai](mailto:hello@mindhyve.ai) general inquiries

[partnerships@mindhyve.ai](mailto:partnerships@mindhyve.ai) institutional procurement

## ON THE WEB

[justinegrid.ai](https://justinegrid.ai)

THEOAI™ · THE AGENTIC THEOLOGY OS

# Every answer grounded. Every chain verified.



Islamic scholarship, grounded in primary sources. Three products. One corpus. One methodology.

TheoAI™ is the Agentic Theology Operating System for Islamic scholarship — evidence-based, primary-source-grounded, computationally rigorous. Every answer traces to a graded chain of transmission, with narrator reliability shown. We do not issue fatwas — we produce Evidence-Based Opinions, with the human always deciding.

01

Theo (Live)

One-on-one evidence-based Islamic Q&A at chat.theogrid.ai. 15 free questions per day, no signup. 560K+ Muslims learning across five locales.

02

Majlis (Waitlist)

Institutional research workspace. Multi-document analysis, EBO Generator, Knowledge Graph, Isnad Bundle Visualizer. \$199/\$499/mo for scholarly institutions.

03

Mizan (Early Access)

Islamic Financial Governance OS. Shariah compliance verification, EBGO — scholars decide. AAOIFI mapped to 8 jurisdictions. For Islamic banks and regulators.

04

IPSC™ Foundation

617,912 hadith cataloged across 123 classical works. Every chain parsed, every narrator identified. 95.8% Bukhari sahih+hasan convergence. Patent Pending.

05

Tahqiq Methodology

Quran first, always. Then authenticated hadith verified via chain analysis. Then reasoned analysis. Five-layer pipeline with 12 cognitive operations.

06

Localization

Hijri calendar choice. Dhul Hijjah tracker + Qurbani calculator. Per-locale honorifics. Built for how Muslims actually practice — not for Silicon Valley defaults.

480K+

INDEPENDENTLY GRADED, FROM 617,912 HADITH CATALOGED

95.8%

BUKHARI SAHIH+HASAN CONVERGENCE (INDEPENDENT)

560K+

MUSLIMS LEARNING WITH THEO FREE TIER, NO SIGNUP

5 Locales

ENG · AR · UR · IN · MS WITH REAL RTL

## ON THE ARCHITECTURE

Eve-Theology™ F5/reasoner — the Eve-Fusion™ instance calibrated to Islamic scholarly reasoning. Five cooperating models composed best-fit per request, the reasoner fine-tuned on Eve-Genesis™ (Theology). Grounded in the IPSC™ corpus: 26,923 narrators indexed, 21,617 of them graded, across 889,913 transmission links. Multi-locale RTL. Patent Pending.

## ON TRUST + COMPLIANCE

Theo produces Evidence-Based Opinions (EBOs), not fatwas. The reasoning is shown. The confidence is calibrated. The sources are cited. The human always decides. Honest grading: 22% of hadith openly marked not-graded rather than force a classification.

CONTINUE READING → → /theo/ → /majlis/ → /mizan/ → ipsc

## TALK TO OUR TEAM

[hello@mindhyve.ai](mailto:hello@mindhyve.ai) general inquiries  
[partnerships@mindhyve.ai](mailto:partnerships@mindhyve.ai) institutional procurement

## ON THE WEB

[theogrid.ai](https://theogrid.ai)

CIAI • THE CREDENTIALING INSTITUTE

# Learn with ArthurAI™. Certify with CIAI.



*The credentialing institute behind AI Professional 360™. AI fluency for every industry.*

*CIAI is the virtual learning institute that certifies AI fluency for the global workforce. Curriculum authored by Eve; delivery substrate is ArthurAI™ VLE — adaptive, multilingual, personalized. The certification is the credibility signal procurement teams need when evaluating workforce-development frameworks.*

|                                                                                                                                                                                          |                                                                                                                                                                                                                             |                                                                                                                                                                                                                                |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p>01</p> <p><b>Foundational Tier</b></p> <p>AI Professional 360™ (formerly The Dawn Directive™) — applications open. 18 courses, 360 hours, 12 months, one professional credential.</p> | <p>02</p> <p><b>Industry Tier</b></p> <p>Coming soon. Vertical-specific AI fluency programs for Finance, Healthcare, Marketing, HR, Legal. Built on the AI Professional 360™ foundation.</p>                                | <p>03</p> <p><b>Advanced Tier</b></p> <p>Future path. AI Collaboration &amp; Team Leadership, AI Governance &amp; Ethics, AI Readiness for educators, executives, and policy leaders.</p>                                      |
| <p>04</p> <p><b>Delivery Substrate</b></p> <p>Delivered on ArthurAI™ VLE, licensed from MindHYVE.ai — adaptive, multilingual, mobile-first, browser-deployable.</p>                      | <p>05</p> <p><b>Vendor-Agnostic</b></p> <p>Not a single tool that expires next year: durable habits — how to direct AI, check its work, and keep the judgment that's yours. Skills that transfer when the tools change.</p> | <p>06</p> <p><b>Institutional Bridge</b></p> <p>CIAI is the credentialing authority; MindHYVE's institutional partnerships extend its reach — live with LaGrange; CNU, PEC and IUCEA going live soon; an MoU with Kadisco.</p> |

|                                                             |                                                                      |                                                                   |                                                                                              |
|-------------------------------------------------------------|----------------------------------------------------------------------|-------------------------------------------------------------------|----------------------------------------------------------------------------------------------|
| <p><b>18</b></p> <p>COURSES IN<br/>AI PROFESSIONAL 360™</p> | <p><b>360</b></p> <p>HOURS OF STRUCTURED<br/>AI FLUENCY LEARNING</p> | <p><b>12</b></p> <p>MONTHS TO ONE<br/>PROFESSIONAL CREDENTIAL</p> | <p><b>3 Tiers</b></p> <p>FOUNDATIONAL • INDUSTRY •<br/>ADVANCED<br/>ROADMAP THROUGH 2027</p> |
|-------------------------------------------------------------|----------------------------------------------------------------------|-------------------------------------------------------------------|----------------------------------------------------------------------------------------------|

## ON THE ARCHITECTURE

CIAI, LLC is an independent organisation, founded by Bill Faruki, who also founded MindHYVE.ai; it licenses ArthurAI™ from MindHYVE.ai and the two work closely together. AI Professional 360™ is a CIAI trademark, used under license. The delivery substrate is ArthurAI™ VLE — MindHYVE operates the OS; CIAI brands and certifies the curriculum.

## ON TRUST + COMPLIANCE

Vendor-agnostic by design. Curriculum updates quarterly as the frontier-model landscape changes. Credential standards aligned to workforce frameworks (NAVTTTC in Pakistan, EQF in EU, CTE Perkins V in US). 100% online, multilingual, with accessibility built in.

CONTINUE READING →

→ [AI Professional 360](#)

→ [/about/](#)

[08 • AP360 one-pager](#)

[04 • ArthurAI VLE](#)

## TALK TO OUR TEAM

[hello@mindhyve.ai](mailto:hello@mindhyve.ai) general inquiries

[partnerships@mindhyve.ai](mailto:partnerships@mindhyve.ai) institutional procurement

## ON THE WEB

[ciai.ai](https://ciai.ai)

AI PROFESSIONAL 360™ · THE FLAGSHIP PROGRAM

# The complete professional program for the AI era.



18 courses. 360 hours. 12 months. One professional credential. Delivered on ArthurAI™ VLE. Certified by CIAI.

AI Professional 360™ — formerly The Dawn Directive™ — is CIAI's complete professional programme for the AI era: practical AI, generative AI, prompting, workflow orchestration, AI agents, responsible AI, and the skills required to work alongside increasingly autonomous AI systems. Its trademark, its IP and its credential are CIAI's; it is delivered on ArthurAI™, licensed from MindHYVE.ai.

01

## Practical AI

What AI actually is, where it helps and where it does not — the working foundation for everyone entering the AI economy.

02

## Generative AI + Prompting

Directing generative systems with intent: prompting as a professional skill, and checking the work that comes back.

03

## Workflow Orchestration

Composing AI into the work a role actually does — the day-to-day workflows employers value, not a single tool.

04

## AI Agents

Working with agentic systems that take actions, and keeping the human judgment that governs them.

05

## Responsible AI

The governance, ethics and accountability that make AI use defensible — to an employer, a regulator, or a client.

06

## One Professional Credential

The CIAI Professional Certificate in Applied & Agentic Artificial Intelligence, issued on completion. A professional certificate — not an academic degree or an accredited qualification.

18

COURSES IN ONE  
PROFESSIONAL PROGRAMME

360

HOURS OF STRUCTURED  
AI FLUENCY LEARNING

12

MONTHS FROM  
ENROLMENT TO CREDENTIAL

1

PROFESSIONAL CREDENTIAL,  
ISSUED BY CIAI

## ON THE ARCHITECTURE

Delivered on ArthurAI™ VLE, licensed from MindHYVE.ai — the Vocational Learning Edition, built for workforce development. Adaptive to each learner's pace, role, and industry. CIAI, LLC is an independent organisation that licenses ArthurAI™; the programme, its trademark and its credential are CIAI's own.

## ON TRUST + COMPLIANCE

Not a single tool that expires next year: durable habits — how to direct AI, check its work, and keep the judgment that is yours. CIAI teaches skills employers value; it does not guarantee any specific job or outcome.

CONTINUE READING →

→ [AI Professional 360](#)

[07 · CIAI](#)

[04 · ArthurAI VLE](#)

→ [ciai.ai/courses](#)

## TALK TO OUR TEAM

[hello@mindhyve.ai](mailto:hello@mindhyve.ai) general inquiries

[partnerships@mindhyve.ai](mailto:partnerships@mindhyve.ai) institutional procurement

## ON THE WEB

[ciai.ai/ai-professional-360](https://ciai.ai/ai-professional-360)

EVE-FUSION™ · THE COMPOSITIONAL FABRIC

# Five cooperating reasoning models, composed per request.



*Eve-Fusion™ F5/reasoner — the architecture that powers every product in the MindHYVE family.*

*Eve-Fusion™ is a compositional fabric — not a stack, not a wrapper, not a chain. Five cooperating reasoning models compose dynamically per request, calibrated to the workload, modality, regulatory context, and data-residency constraints. Five on every vertical; a frontier slot is swapped without rebuilding the agent.*

|                                                                                                                                                                                                             |                                                                                                                                                                                                               |                                                                                                                                                                                                                                                |                                                                     |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------|
| <div>01</div> <div>Routing Classifier</div> <div>A small, fast routing classifier. Sub-200ms routing across workload type, modality, reasoning intensity, and context length.</div>                         | <div>02</div> <div>Domain Reasoner</div> <div>A small reasoning model, adapter-fine-tuned on the vertical's Eve-Genesis™ edition. The domain-specialized workhorse — carries the majority of workloads.</div> | <div>03</div> <div>Frontier Slot 01</div> <div>Best-in-class long-form structured reasoning. Used when nuance matters — complex narratives, high-stakes synthesis. Which model fills the slot is a release decision, and stays internal.</div> |                                                                     |
| <div>04</div> <div>Frontier Slot 02</div> <div>Frontier general reasoning, complementary to Slot 01 for cross-validation and ensemble routing. Provider diversity absorbs jurisdictional constraints.</div> | <div>05</div> <div>Frontier Slot 03</div> <div>Long-context analysis — the slot for requests that span a whole case file or longitudinal chart. Composed per request, best-fit per release.</div>             | <div>06</div> <div>Parallel Vision Route</div> <div>Imaging, diagram analysis, document and handwritten review, taken when a request carries an image. A parallel path alongside the five — not a sixth model.</div>                           |                                                                     |
| <div>&lt; 200ms</div> <div>ROUTING CLASSIFIER<br/>LATENCY</div>                                                                                                                                             | <div>1 + 1 + 3</div> <div>CLASSIFIER · REASONER ·<br/>FRONTIER SLOTS</div>                                                                                                                                    | <div>Plug + Play</div> <div>FRONTIER SLOTS<br/>SWAPPED PER JURISDICTION</div>                                                                                                                                                                  | <div>5</div> <div>ONE ARCHITECTURE,<br/>FIVE PROVEN VERTICALS</div> |

## ON THE ARCHITECTURE

Per-product instances: Eve-Healthcare™ for ChironAI, Eve-Education™ for ArthurAI, Eve-Legal™ for JustineAI, Eve-Theology™ for TheoAI, Eve-Finance™ for EliAI. The compound architecture is shared; the domain reasoner is calibrated to the vertical by adapter fine-tuning on that vertical's Eve-Genesis™ synthetic corpus.

## ON TRUST + COMPLIANCE

Provider diversity is structural, not aspirational. Where a jurisdiction or institutional client prohibits a specific provider, the relevant slot is swapped without rebuilding the agent or retraining the reasoner. The architecture survives regulatory change.

CONTINUE READING → [02 · MindHYVE](#) [10 · Eve-Genesis](#) [11 · Differentiation](#) → [/tech](#)

## TALK TO OUR TEAM

[hello@mindhyve.ai](mailto:hello@mindhyve.ai) general inquiries  
[partnerships@mindhyve.ai](mailto:partnerships@mindhyve.ai) institutional procurement

## ON THE WEB

[mindhyve.ai](https://mindhyve.ai)

EVE - GENESIS™ · THE TRAINING CORPUS

# The architecture cannot leak what is not present.



100% synthetic by construction. Per-vertical editions. The CISO-facing answer.

Eve-Genesis™ is the synthetic-corpus methodology that fine-tunes the domain reasoner inside every Eve-Fusion™ instance. It is 100% synthetic by construction — no customer data, no scraped records, no exfiltrated case files, no proprietary curricula train any product’s reasoning. An architectural property, not a policy promise.

01

Clinical Edition

Powers ChironAI™. Trained to reason like a senior clinician across the relevant differential, regulatory chain, and consultation context.

02

Education Edition

Powers ArthurAI™. Trained on canonical curricula, pedagogical taxonomies, and the editor pattern that defines educator-attested instruction.

03

Law Edition

Powers JustineAI™. Trained on procedural taxonomies, demand-letter pipelines, and the supervisor pattern. CourtListener-anchored.

04

Theology Edition

Powers TheoAI™. Calibrated to the Tahqiq methodology, the 12-tier narrator system, and the Ibn al-Salah five-condition framework.

05

Financial Edition

Powers EliAI™. Trained on valuation reasoning over comparable settled matters, lien perfection, and counterparty behaviour.

06

Held-out Evaluation

Each edition is evaluated on a held-out test set sourced independently from training. Convergence with senior-practitioner reasoning measured per release.

100%

SYNTHETIC BY CONSTRUCTION  
ZERO CUSTOMER DATA

5 Editions

ONE PER VERTICAL  
NEVER CROSS-CONTAMINATED

Quarterly

EXPANSION CADENCE  
REGULATORY + GUIDELINE  
UPDATES

Adapter

FINE-TUNING STRATEGY  
NEAR-FRONTIER COST +  
LATENCY

ON THE ARCHITECTURE

Generated by Eve — the autonomous reasoning system that authors synthetic training data in the voice of a senior practitioner across each vertical. Generation is followed by structural validation, regression evaluation, and adapter-based fine-tuning of the domain reasoner. The base model is never modified.

ON TRUST + COMPLIANCE

Synthetic-by-construction is the single most decisive procurement signal in regulated industries. HIPAA does not apply to data that does not exist. FERPA does not apply to records that were never collected. Attorney-client privilege does not need to be re-architected.

CONTINUE READING → [02 · MindHYVE](#) [09 · Eve-Fusion](#) [11 · Differentiation](#) → [/trust](#)

TALK TO OUR TEAM

[hello@mindhyve.ai](mailto:hello@mindhyve.ai) general inquiries

[partnerships@mindhyve.ai](mailto:partnerships@mindhyve.ai) institutional procurement

ON THE WEB

[mindhyve.ai](#)

THE DIFFERENTIATION · MINDHYVE VS. THE REST

# Why MindHYVE — and not the three alternatives.



*The argument your internal champion uses to win the procurement meeting you're not invited to.*

*Every prospect comparing MindHYVE to an alternative is comparing us to one of three things: a frontier model called directly, a single-vertical AI competitor, or the build-it-yourself path with an orchestration framework. Each comparison has a clean answer. This document is the argument — categorical, principled, defensible.*

01

## vs. Frontier Direct

They sell raw capability. We give you the compositional fabric, the vertical reasoner, the workflow, the compliance posture, and the named Digital Employee. Frontier is one slot inside Eve-Fusion.

02

## vs. Vertical AI

They built one product on one architecture. We built the same compound architecture across five verticals — healthcare, education, law, theology, finance. Breadth IS the credibility.

03

## vs. Build-it-Yourself

They give you Lego bricks. We give you the assembled system: 43 institutional partnerships, named Digital Employees, vertical training. Build-yourself: 18-24 months, \$3-8M. Us: 60-90 days.

04

## Architectural Moats

Three durable moats: compositional fabric (one architecture, five verticals); named Digital Employees; Eve-Genesis synthetic corpus — no customer data trains the reasoner.

05

## Partnership Moat

43 institutional partnerships across five regions: live with LaGrange; MoUs signed with IDC, KEPSA and Kadisco; CNU, KPSIAJ, PEC and IUCEA going live soon. Signed, dated instruments.

06

## Posture Moat

Honest disclosure as credibility signal: not FDA-cleared, not SOC 2 certified, no fabricated accuracy claims. Procurement reads overclaim as red flag, honesty as green flag.

5

REASONING MODELS COMPOSED  
PER REQUEST IN EVE-FUSION

5+1

AGENTIC OS VERTICALS  
PLUS ONE RESEARCH CORPUS

43

INSTITUTIONAL PARTNERSHIPS  
NAMED COUNTERPARTS, DATED  
MOUS

0

CUSTOMER RECORDS IN ANY  
PRODUCT'S TRAINING DATA

## ON THE ARCHITECTURE

Single frontier API: one capability ceiling, no vertical calibration. Vertical wrapper: one product, one architecture, no cross-portfolio learning. Build-it-yourself: 18-24 months, \$3-8M, zero deployments. MindHYVE: composed fabric + named Digital Employees + 43 partnerships + five-vertical proof.

## ON TRUST + COMPLIANCE

We never name competitors by name in this document. Sophisticated procurement teams know which companies fit which categories. The categorical framing is the argument: this is the structure of the alternatives, here is the structure of MindHYVE, here is why the structure wins.

CONTINUE READING →

[02 · MindHYVE](#)

[09 · Eve-Fusion](#)

[10 · Eve-Genesis](#)

[→ /about/](#)

## TALK TO OUR TEAM

[hello@mindhyve.ai](mailto:hello@mindhyve.ai) general inquiries

[partnerships@mindhyve.ai](mailto:partnerships@mindhyve.ai) institutional procurement

## ON THE WEB

[mindhyve.ai](https://mindhyve.ai)

ELIAI™ · THE AGENTIC FINANCE OS

# Capital moves on evidence.



*Underwriting, valuing and servicing illiquid claims. One edition, generally available.*

*EliAI™ is the Agentic Finance Operating System for the desks that commit capital against evidence somebody has to read. It reads the whole file with four reasoning specialists, sources every number it produces to the comparable that produced it, and refuses to price what the evidence will not carry. The AI reasons; the underwriter decides.*

|                                                                                                                                                                                                                                            |                                                                                                                                                                                                                                            |                                                                                                                                                                                                                                          |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p>01</p> <p><b>MRF Edition</b></p> <p>Medical Receivables Finance, generally available at eligrid.ai. Underwrite, perfect, service and realise receivables secured by liens against personal-injury and workers-compensation matters.</p> | <p>02</p> <p><b>The Reasoning Desk</b></p> <p>Chiron on the medicine, Justine on the law, Issac on the counterparty, Eli on the capital. Attribution is enforced per finding; a conclusion outside a specialist's axes is quarantined.</p> | <p>03</p> <p><b>Disagreement Is Signal</b></p> <p>When specialists disagree it is named and attributed, and it widens the band rather than being averaged away. Two extremes and four neutrals must not produce the same confidence.</p> |
| <p>04</p> <p><b>It Refuses To Price</b></p> <p>A cohort too thin to support a price yields a range, no midpoint, no recommended advance, and exclusion from book value rather than an assumed zero. Enforced in the schema.</p>            | <p>05</p> <p><b>Reproducible Marks</b></p> <p>A valuation is an artifact: cohort, comparables, query, pinned engine version and seed. Marks are superseded, never overwritten, and the audit trail is hash-chained.</p>                    | <p>06</p> <p><b>Controls In The Ledger</b></p> <p>Funding checks the verified filing itself, not a nullable column. Evidence cannot pre-date the agreement permitting it. A concentration breach refuses at the moment of funding.</p>   |

|                                                                 |                                                                         |                                                                            |                                                                         |
|-----------------------------------------------------------------|-------------------------------------------------------------------------|----------------------------------------------------------------------------|-------------------------------------------------------------------------|
| <p><b>Four</b></p> <p>DIGITAL EMPLOYEES READ EVERY POSITION</p> | <p><b>Five</b></p> <p>COOPERATING MODELS IN EVE-FINANCE F5/REASONER</p> | <p><b>Zero</b></p> <p>CUSTOMER DATA USED FOR TRAINING, BY CONSTRUCTION</p> | <p><b>GA</b></p> <p>MRF EDITION — GENERALLY AVAILABLE AT ELIGRID.AI</p> |
|-----------------------------------------------------------------|-------------------------------------------------------------------------|----------------------------------------------------------------------------|-------------------------------------------------------------------------|

## ON THE ARCHITECTURE

Eve-Finance™ F5/reasoner — the Eve-Fusion™ instance calibrated to financial reasoning. Five cooperating models composed best-fit per request. Valuation is grounded in retrieval over comparable settled matters, never in training; the corpus is the model, so a mark can be re-derived years later. Multi-tenant isolation enforced by database policy on Eve-Grid™ Azure.

## ON TRUST + COMPLIANCE

No certification is named on the eligrid.ai trust centre, and the page says why: a badge is the most common way a security page becomes untrue. Per-subject encryption makes erasure provable. Personal data is never used for training. Not investment advice; outputs are decision support and a qualified person remains responsible.

CONTINUE READING → [→ /mrf/](#) [→ /how-it-works/](#) [→ /worked-valuation/](#) [→ /trust/](#)

## TALK TO OUR TEAM

[sales@mindhyve.ai](mailto:sales@mindhyve.ai) general + institutional procurement

## ON THE WEB

**eligrid.ai**