


Retis Persona turns professional expertise into a controlled AI interface.

The value is **not** cloning a person.

The value is a **governed Skillprint**: a local-first system that captures method, communication style, service boundaries, and feedback loops so an expert can be represented by an agent without losing control of source material.



THE SKILLPRINT APPROACH



CAPTURE

Capture your expertise as a Skillprint—not your content. We structure your proven methods, principles, and decision patterns.



COMMUNICATE

Define your communication style so the agent sounds like you—clear, consistent, and context-appropriate.



BOUND

Set service boundaries, topics, and policies so the agent operates within your defined scope and standards.



LEARN & IMPROVE

Feedback loops ensure the agent improves over time while you review, approve, and steer the evolution.




LOCAL-FIRST

Your data stays in your environment. You control access, storage, and integration. Nothing leaves unless you choose.

THE VALUE

- ✓ Represent your expertise without giving up control.
- ✓ Deliver consistent, high-quality interactions at scale.
- ✓ Protect your IP, methods, and client trust.
- ✓ Spend less time explaining, more time leading.




YOU

Expertise. Judgment.
Standards.



SKILLPRINT SYSTEM

Method. Style. Boundaries.
Feedback. Control.



YOUR AGENT

Represents you.
Operates within your
governed parameters.



Retis Persona. Your expertise. Your rules. **Your control.**



retispersona.com



ONE-LINE DEFINITION

SkillPrint distills an expert's repeatable method into a controlled AI Persona: how they analyze, decide, qualify uncertainty, structure outputs, use evidence, and know when not to answer.



SkillPrint is the foundation layer. It comes before Stylecast, AgentFront, and MindVault.

POSITIONING

Most persona products start with appearance, voice, or personality. SkillPrint starts with competence.

The goal is to capture the expert's professional method without pretending to be the expert. A strong SkillPrint lets a Persona produce useful work because it understands the expert's frameworks, decision rules, caveats, examples, and boundaries.

WHAT SKILLPRINT IS NOT

- ✗ Not a digital clone.
- ✗ Not a writing-style mimic.
- ✗ Not a voice or video avatar.
- ✗ Not a generic RAG chatbot.
- ✗ Not a memory dump.
- ✗ Not a collection of pasted source snippets.

SkillPrint is the operating logic of expertise.

TARGET AUDIENCES

- ✓ Analysts, researchers, advisors, coaches, consultants, and educators.
- ✓ Investment, strategy, legal, policy, product, and technical experts with repeatable judgment workflows.
- ✓ Creator-educators who teach a method, not just a personality.
- ✓ Internal experts whose SOPs, frameworks, and review rules need to scale across a team.
- ✓ Late-expert or legacy projects where the professional method is more valuable than imitation.
- ✓ Organizations that need local-first expert agents without uploading raw source material.

WHAT SKILLPRINT PRODUCES

DELIVERABLE	DESCRIPTION
Method Map	The expert's repeatable process: inputs, questions, reasoning steps, outputs, and stopping points.
Judgment Rules	How the expert weighs evidence, tradeoffs, confidence, risk, and exceptions.
Framework Library	Named frameworks, checklists, scoring rules, taxonomies, and decision trees.
Output Formats	Approved structures for memos, reports, recommendations, comments, summaries, and reviews.
Evidence Contract	Rules for what counts as support, what is weak, what is adjacent, and what must be refused.
Boundary Map	Restricted topics, disclaimers, escalation rules, and situations requiring human review.
Evaluation Set	Strong, weak, unsafe, off-method, and corrected examples for testing.
Persona.md Baseline	The first operational profile that other Retis services can build on.

REQUIRED CLIENT MATERIAL

- Approved reports, memos, articles, slides, courses, frameworks, or training materials.
- Examples of good answers and bad answers.
- Checklists, scoring rules, templates, rubrics, and decision trees.
- Repeated client questions, use cases, and desired output formats.
- Known caveats, disclaimers, restricted topics, and escalation rules.
- Corrections and review notes that show how the expert improves an answer.
- Optional transcripts or recordings when they reveal method, not just style.

METHOD

- 1 Identify the expert's recurring use cases.
- 2 Separate method from voice, personality, and private memory.
- 3 Extract frameworks, decision points, caveats, and output structures.
- 4 Build an evidence contract: direct support, weak support, analogy, unsupported.
- 5 Create the baseline Persona.md and guardrails.
- 6 Configure retrieval so sources guide reasoning without being pasted into final answers.
- 7 Build an evaluation set from realistic tasks.
- 8 Review outputs with the expert and revise the SkillPrint until the method is stable.

EXAMPLE PULSES

- conversation** — answer user questions through the expert's method.
- method_check** — test whether an output follows the approved SkillPrint.
- evidence_check** — classify support as direct, weak, adjacent, or unsupported.
- brief_draft** — produce a memo or recommendation in the approved format.
- article_comment** — comment on a supplied article using the expert's framework.
- articles_summary** — summarize multiple sources using the expert's priorities.
- review_response** — record whether a response was on-method, off-method, unsafe, or unsupported.
- trainer_suggestion** — suggest prompt, guardrail, retrieval, or evaluation changes.

PACKAGE LEVELS

★ SKILLPRINT STARTER	🛡️ SKILLPRINT FOUNDATION	💎 SKILLPRINT PRO
Use case: A first controlled Persona for one expert use case.	Use case: A reusable expert-method layer that can support Stylecast or AgentFront.	Use case: Multi-workflow expert Persona for teams, services, or high-value client workflows.
Includes: Method map, baseline Persona.md, core guardrails, one output format, 15-output evaluation set.	Includes: Method map, framework library, evidence contract, boundary map, 3–5 output formats, 30-output evaluation set, connector setup.	Includes: Full method library, scoring/rubric design, retrieval tuning, review workflow, escalation rules, 75-output evaluation set, Trainer-ready governance.
Timeline: 1–2 weeks	Timeline: 2–4 weeks	Timeline: 4–8 weeks
Client effort: Low–medium	Client effort: Medium	Client effort: Medium-high
Indicative price: USD 3k–6k	Indicative price: USD 6k–15k	Indicative price: USD 18k–45k

GUARDRAILS

- ✗ Do not imitate identity: apply method.
- ✗ Do not expose raw source snippets, handles, URLs, capture labels, file paths, or metadata.
- ✗ Do not paste retrieved material directly into final answers.
- ✗ Do not answer beyond the approved evidence contract.
- ✗ Do not use style calibration as a substitute for method.
- ✗ Do not invent private beliefs, unpublished facts, or current commitments.
- ✗ Do not launch AgentFront or MindVault until the SkillPrint is stable.

RELATION TO OTHER SERVICES

- Stylecast** — adds communication behavior after the method is stable.
- AgentFront** — turns the SkillPrint into a service entry point.
- MindVault** — adds private evolving memory after governance is ready.
- PersonaTraining** — reviews outputs and suggests prompt, guardrail, and retrieval updates.



If Stylecast is the voice of the service and AgentFront is the doorway, SkillPrint is the professional judgment inside.

SUCCESS CRITERIA

- ✓ The Persona follows the expert's method even when the wording changes.
- ✓ The Persona asks for missing context instead of guessing.
- ✓ The Persona can explain uncertainty and refuse unsupported claims.
- ✓ Outputs match approved formats and decision rules.
- ✓ Reviewers can mark whether an answer is on-method or off-method.
- ✓ Style can later be added without weakening judgment.
- ✓ AgentFront can route users because the service boundaries are clear.
- ✓ MindVault can evolve memory without corrupting the core method.

★ CORE LINE

SkillPrint is
method before
mimicry: the
expert's
repeatable
judgment made
operational.



POSITIONING

Retis Persona turns professional expertise into a **controlled AI interface**. The value is not cloning a person. The value is a **governed Skillprint**: a **local-first** system that captures method, communication style, service boundaries, and feedback loops so an expert can be represented by an agent without losing control of source material.

Raw client content stays on the user's machine by default. Retis Persona ingests and embeds locally, while Attas or other Retis trainers send tuning suggestions through pulses: prompt updates, guardrail changes, evaluation feedback, and connector guidance.



SERVICES

SERVICE	WHAT IT MEANS	PRIMARY AUDIENCE
 Skillprint	Distills an expert's repeatable method, judgment rules, frameworks, caveats, and output formats.	Analysts, advisors, educators, consultants, coaches, creators, founders.
 Stylecast	Calibrates writing style, video/script style, tone, cadence, vocabulary, and disclosure rules. It is not voice-only.	YouTubers, newsletter authors, public experts, executives, media teams.
 AgentFront	Builds a user-facing proxy or entry point for the expert's services: intake, Q&A, summaries, comments, reports, and routing.	Service businesses, creators with audiences, coaches, analyst boutiques, agencies.
 MindVault	Adds private evolving memory, feedback loops, review history, versioning, and governance.	High-trust experts, internal teams, enterprise users, legacy/estate projects.
 PersonaTraining	Trainer service that reviews outputs, detects drift, suggests prompt/guardrail changes, and registers trainers on Retis Plaza.	Persona operators, marketplaces, enterprise governance teams.

PACKAGES

A. SKILLPRINT FOUNDATION Outcome: A controlled Persona that applies the expert's approved method.		B. STYLECAST + AGENTFRONT Outcome: A recognizable user-facing Persona that can answer, summarize, comment, draft, and route users into the expert's services.		C. MINDVAULT MANAGED PERSONA Outcome: A private evolving Persona with durable memory, governance, audit history, review loops, and ongoing trainer support.	
Client provides	Reports, articles, memos, slides, transcripts, examples of strong/weak answers, service boundaries, target use cases.	Client provides	Everything from Foundation, plus writing samples, video/transcript samples, preferred vocabulary, forbidden phrases, disclosure rules, offer/service map.	Client provides	Private notes, journals, working memos, historical calls, post-mortems, model corrections, access/retention/deletion rules.
Retis/Attas provides	Local ingestion setup, expertise map, Persona.md, prompt/guardrail baseline, first evaluation set, local connector setup.	Retis/Attas provides	Style calibration, pulse design, response review loop, UI/service entry point, AgentFront routing, marketplace connector configuration.	Retis/Attas provides	Local-first MindVault architecture, permission model, Trainer pulses, drift monitoring, evaluation dashboards, versioned prompt/guardrail governance.
Timeline	2-4 weeks	Timeline	4-8 weeks	Timeline	8-16+ weeks
Client effort	Low-medium	Client effort	Medium	Client effort	High
Indicative price	USD 3k-8k	Indicative price	USD 10k-25k	Indicative price	USD 30k-75k+ or managed retainer.

CONNECTORS

Initial connector marketplace:

-  **YouTube:** public channel content, metadata, transcripts where available.
-  **X:** public handles or logged-in account capture when the user authorizes it.
-  **Local filesystem:** folders, documents, markdown, PDFs, notes.
-  **S3:** bucket/object ingestion using user-provided AWS profile.
-  **Email:** local mbox first, delegated OAuth planned.
-  **Third-party video/avatar tools** such as HeyGen or Synthesia can be used as rendering services when needed. Retis Persona does not compete on video clone generation; it focuses on digitizing expertise and governing the agent interface.

TARGET AUDIENCES

-  Expert creators with large archives: YouTube, newsletters, podcasts, courses.
-  Advisors, coaches, consultants, and analysts who want a scalable service front door.
-  Research and investment teams that need repeatable judgment workflows.
-  Founders and executives who want a controlled proxy for public or internal expertise.
-  Families, estates, and institutions preserving a late expert's professional legacy, where consent, rights, and governance are clear.
-  Organizations with internal experts, SOPs, and decision frameworks that should stay local/private.

WHAT MAKES IT DIFFERENT

-  **Local-first raw material:** source content does not need to upload to Attas.
-  **Skill before style:** method and judgment are built before personality.
-  **Evidence separated from voice:** retrieved content guides answers but is not pasted into user-facing output.
-  **Governance by default:** identity boundaries, disclosure rules, private/public split, access controls, review history.
-  **Trainer through pulses:** Attas Tuning Cloud can improve prompts and guardrails without owning raw client data.
-  **Network-ready:** Retis Persona can connect to Attas Plaza or other Retis networks.



CORE LINE

The value is not imitation. The value is a **controlled interface for expertise**.



Stylecast is not a copied voice.
It is governed expression for expertise.

ONE-LINE DEFINITION

Stylecast turns an expert's public expression into **controlled communication behavior** for a Persona: writing style, video/script style, tone, cadence, vocabulary, response shape, disclosure rules, and platform-specific delivery.



Stylecast is not voice cloning.
It is not a video avatar product.
It is the style and expression layer
that sits on top of a strong Skillprint.

POSITIONING

Skillprint captures what the expert knows and how they decide. Stylecast calibrates how that expertise should sound, read, and present itself when it becomes user-facing.

The goal is recognizability without impersonation. Retis Persona should feel aligned with the expert's communication style while staying honest about identity, source coverage, and boundaries.

TARGET AUDIENCES

- Expert creators with YouTube channels, podcasts, newsletters, blogs, or courses.
- Analysts, advisors, coaches, and consultants with distinctive commentary style.
- Executives and founders who need consistent public-facing or internal communication.
- Media teams that want drafts, summaries, comments, scripts, and response templates in a house style.
- Legacy or estate projects where an expert's professional communication style should be preserved with explicit rights and governance.

WHAT STYLECAST PRODUCES

DELIVERABLE	DESCRIPTION
Style Profile	A concise description of cadence, tone, vocabulary, sentence shape, recurring moves, and communication preferences.
Response Shape	Rules for how answers should open, develop, qualify uncertainty, and close.
Writing Style Pack	Guidance for articles, memos, comments, newsletters, summaries, and social posts.
Video/Script Style Pack	Guidance for intros, monologue rhythm, scene beats, talking points, hooks, transitions, and calls to action.
Forbidden Claims	Phrases, identity claims, private-access claims, and source labels the Persona must not use.
Disclosure Rules	When the Persona must say it is a proxy, unofficial interface, assistant, simulation, or locally grounded agent.
Evaluation Set	Side-by-side examples of good, weak, overfit, too-generic, and unsafe outputs.

REQUIRED CLIENT MATERIAL

- Writing samples:** articles, newsletters, reports, comments, memos, captions.
- Video or audio transcripts:** YouTube, podcasts, interviews, webinars, lectures.
- Representative outputs:** best answers, bad answers, edited drafts, rewritten scripts.
- Vocabulary preferences:** favorite terms, honned terms, naming conventions, jargon limits.
- Tone preferences:** direct, warm, analytical, contrarian, cautious, educational, executive.
- Platform rules:** YouTube script, LinkedIn post, email reply, investor memo, customer Q&A.
- Identity and disclosure boundaries:** what the Persona may claim, must avoid, or must clarify.

METHOD

- Start after Skillprint foundation.
- Separate facts from style signals.
- Extract recurring communication patterns from approved source material.
- Convert style into operational rules, not a loose vibe description.
- Build prompt and guardrail updates.
- Test against real user-facing pulses: conversation, article_comment, articles_summary, script_draft, post_draft.
- Review outputs with the client and mark examples as good, weak, unsafe, or off-style.
- Feed corrections back through PersonaTraining pulses.

EXAMPLE PULSES

- conversation** — answer in calibrated style without exposing raw sources.
- article_comment** — produce a concise expert-style comment on supplied news.
- articles_summary** — summarize multiple articles in the expert's preferred structure.
- script_draft** — turn a topic into a video or podcast monologue outline.
- post_draft** — create LinkedIn, X, newsletter, or community post variants.
- rewrite_in_style** — convert a rough draft into the approved Stylecast profile.

PACKAGE LEVELS


STYLECAST BASIC	STYLECAST PRO	STYLECAST MANAGED
Use case: Make answers and summaries sound less generic.	Use case: User-facing Persona across articles, comments, Q&A, and scripts.	Use case: Ongoing creator, analyst, or executive Persona with regular publishing workflows.
Includes: Style profile, response shape, banned phrases, 20-output review.	Includes: Writing style pack, video/script style pack, platform-specific templates, 50-output review, evaluation set, guardrail tuning.	Includes: Continuous review loop, drift monitoring, platform-specific prompt versions, monthly style refresh, trainer reports, new-format onboarding.
Timeline: 1–2 weeks	Timeline: 3–5 weeks	Timeline: Ongoing after setup
Client effort: Low-medium	Client effort: Medium	Client effort: Medium-high
Indicative price: USD 3k–6k	Indicative price: USD 8k–18k	Indicative price: USD 3k–10k/month

GUARDRAILS

- Do not imitate identity; calibrate expression.
- Do not expose source titles, handles, URLs, capture labels, transcripts, file paths, or metadata as the Persona voice.
- Do not paste retrieved material directly into final answers.
- Do not use catchphrases mechanically.
- Do not invent private beliefs, current commitments, or unpublished facts.
- Do not generate style before the Skillprint method is strong.
- Do not use video/avatar rendering without explicit rights and disclosure rules.

RELATION TO THIRD-PARTY MEDIA TOOLS

Stylecast can provide script, phrasing, pacing, and delivery guidance for tools such as HeyGen, Synthesia, Descript, Eleven Labs, or Sora-style video workflows.



Those tools render media. Stylecast governs what should be said, how it should be structured, what must be disclosed, and what must never be claimed.

SUCCESS CRITERIA

- The Persona feels recognizably aligned with the expert without claiming to be the expert.
- Answers are concise, useful, and source-grounded without exposing raw source material.
- Style survives across formats: Q&A, article comment, summary, script, and draft rewrite.
- Client reviewers can explain why a good answer works.
- Unsafe outputs are caught by guardrails before reaching users.
- The Persona can improve through review history and PersonaTraining pulses.

CORE LINE

Stylecast is
not a copied
voice.
It is governed
expression for
expertise.





AgentFront is the front door for expertise:
answer what is safe, ask what is missing,
and route what matters.

ONE-LINE DEFINITION

AgentFront is the user-facing entry point for an expert's services. It turns a Skillprint into a practical front door: intake, Q&A, triage, summaries, comments, draft outputs, routing, and handoff to the real expert or service team.



AgentFront is not a generic chatbot.
It is a controlled service interface
for expertise.

POSITIONING

Skillprint captures the expert's method. Stylecast calibrates expression. AgentFront packages both into a usable service surface that clients, readers, students, customers, or internal teams can interact with.

The goal is to make expertise accessible without pretending the Persona is the human expert. AgentFront should answer what it can, disclose limits, collect useful context, and route high-value or high-risk cases to the right next step.

TARGET AUDIENCES

- ✓ Consultants, advisors, coaches, analysts, and educators who want a scalable service front door.
- ✓ Expert creators with large audiences who need structured Q&A, comments, summaries, and intake.
- ✓ Boutique research, investment, legal, finance, or strategy teams with repeatable client workflows.
- ✓ Enterprise teams with internal experts, SOPs, decision frameworks, and support queues.
- ✓ Agencies that want branded expert agents for client-facing workflows.
- ✓ Legacy or institutional knowledge projects that need a governed access point to archived expertise.

WHAT AGENTFRONT PRODUCES

DELIVERABLE	DESCRIPTION
 Service Map	Clear list of what the Persona can answer, draft, summarize, route, or refuse.
 Intake Flow	Structured questions that collect the context needed before giving useful output.
 Pulse Catalog	User-facing actions such as conversation, article_comment, articles_summary, brief_draft, intake_review, and handoff_request.
 Output Templates	Report, memo, comment, script, Q&A, email, and brief formats tied to the expert's Skillprint.
 Routing Rules	Logic for when to answer directly, ask follow-up questions, escalate, or refer to paid services.
 Disclosure Layer	User-facing wording that explains the Persona's role, source limits, and identity boundary.
 Review Loop	Saved response history, ratings, corrections, and trainer suggestions for improvement.

REQUIRED CLIENT MATERIAL

-  **Service catalog:** offers, packages, consultations, reports, courses, products, or workflows.
-  **Common user questions** and intake forms.
-  **Examples of completed client deliverables.**
-  **Escalation rules:** what requires human review, legal review, compliance review, or paid engagement.
-  **Boundaries:** restricted topics, disclaimers, regulated advice limits, refund/support rules.
-  **Preferred outputs:** report formats, summary formats, email templates, comment style, draft structure.
-  **Conversion path:** booking link, contact form, newsletter signup, paid report, community, support queue.

METHOD

- 1 Start with a usable Skillprint.
- 2 Add Stylecast only when the expert-facing outputs need recognizable expression.
- 3 Define the service map: answer, draft, summarize, route, refuse, escalate.
- 4 Build intake questions for each high-value workflow.
- 5 Create pulse contracts and output templates.
- 6 Add guardrails, disclosure rules, and confidence/coverage checks.
- 7 Connect the AgentFront to Retis Persona, Plaza, connectors, and review history.
- 8 Test real workflows with users and feed corrections through PersonaTraining.

EXAMPLE PULSES

- conversation** — answer user questions using the Skillprint and local material.
- intake_review** — collect user context and determine whether the case is answerable.
- article_comment** — produce a concise expert-style comment on supplied news.
- articles_summary** — summarize multiple articles in the expert's preferred structure.
- brief_draft** — draft a memo, research note, client brief, or service response.
- service_route** — recommend next step: self-serve answer, booking, report, support, or escalation.
- handoff_request** — package context for the human expert or team.
- review_response** — collect user/client feedback and mark output quality.

PACKAGE LEVELS

AGENTFRONT BASIC	AGENTFRONT PRO	AGENTFRONT MANAGED
Use case: A simple public or internal Q&A front door for an expert archive.	Use case: A working service proxy for comments, summaries, draft reports, intake, and lead routing.	Use case: Ongoing front office for an expert service, creator business, analyst product, or internal knowledge function.
Includes: Service map, conversation pulse, intake questions, disclosure rules, basic routing, 20-output review.	Includes: Pulse catalog, output templates, Stylecast integration, connector setup, routing rules, evaluation set, 50-output review.	Includes: Continuous review loop, monthly workflow tuning, lead/service analytics, drift monitoring, new pulse onboarding, Trainer reports, governance refresh.
Timeline: 2–3 weeks	Timeline: 4–8 weeks	Timeline: Ongoing after setup
Client effort: Low–medium	Client effort: Medium	Client effort: Medium–high
Indicative price: USD 5k–10k	Indicative price: USD 15k–35k	Indicative price: USD 5k–15k/month

GUARDRAILS

- ✗ Do not imply the AgentFront is the real expert.
- ✗ Do not answer outside the Skillprint or approved source coverage without saying the limit.
- ✗ Do not expose raw source snippets, handles, URLs, capture labels, transcripts, file paths, or metadata.
- ✗ Do not give regulated, legal, medical, financial, or compliance-sensitive advice beyond approved boundaries.
- ✗ Do not route high-risk cases to self-serve automation when human review is required.
- ✗ Do not hide paid-service boundaries or make false service promises.
- ✗ Do not collect unnecessary private information during intake.

UX SURFACES

-  Embedded website widget.
-  Retis Persona local web UI.
-  Internal dashboard for service teams.
-  Creator/community Q&A portal.
-  Slack, Discord, Teams, or email assistant.
-  Plaza-discoverable Persona service endpoint.
-  API endpoint for another Retis network or agent.

RELATION TO ATTAS AND RETIS

-  Retis Persona runs the local node, connectors, local vault, pulse handling, and Persona behavior.
-  Attas or another Retis tuning service can provide Trainer pulses: prompt improvements, routing changes, evaluation feedback, guardrail updates, and drift monitoring.
-  Retis Plaza lets AgentFront register as a discoverable service endpoint so other agents can find the right Persona, invoke the right pulse, and route work to the right service.

SUCCESS CRITERIA

- ✓ Users quickly understand what the Persona can and cannot do.
- ✓ The first response asks for missing context when needed instead of guessing.
- ✓ Outputs follow the expert's method and approved formats.
- ✓ High-risk or high-value cases route to the right human or paid service.
- ✓ The Persona produces useful self-serve value without exposing raw source material.
- ✓ Review history improves future answers, routing, and guardrails.
- ✓ The expert gains leverage without losing control of their service boundary.





ONE-LINE DEFINITION

MindVault is the private evolving memory layer for a Persona. It stores, governs, reviews, and updates the expert's long-term working memory without turning raw private material into public output.



MindVault is not a chatbot memory toggle. It is a governed vault for private expertise, review history, permissions, retention, and drift control.

POSITIONING

Skillprint captures method. Stylecast calibrates expression. AgentFront exposes the service entry point. MindVault protects and evolves the deeper private memory that makes a Persona improve over time.

The goal is continuity with control: the Persona can learn from notes, corrections, calls, post-mortems, and prior outputs while respecting source permissions, identity boundaries, and deletion rules.

TARGET AUDIENCES

- ✓ High-trust experts whose work depends on private notes, drafts, calls, and decision history.
- ✓ Advisors, coaches, consultants, analysts, and executives with recurring clients or complex workflows.
- ✓ Investment, research, legal, strategy, and enterprise teams that need durable institutional memory.
- ✓ Creator or educator businesses that want private planning memory separated from public content.
- ✓ Families, estates, and institutions preserving a late expert's professional legacy with clear rights and governance.
- ✓ Organizations that need local-first AI memory without uploading raw material to a central cloud.

WHAT MINDVAULT PRODUCES

DELIVERABLE	DESCRIPTION
 Private Memory Map	Inventory of memory types: notes, calls, drafts, reviews, working memos, journals, outcomes, corrections.
 Permission Model	Rules for who can read, update, delete, export, or train from each memory class.
 Public/Private Split	Explicit boundary between public source material, private working memory, and user-facing output.
 Retention Policy	Versioning, deletion, expiry, archive, and legal hold rules.
 Review History	Saved Persona outputs, ratings, corrections, preferred rewrites, and trainer feedback.
 Drift Monitor	Checks that the Persona does not move away from the approved Skillprint, style, or guardrails.
 Audit Trail	Record of prompt updates, guardrail changes, trainer suggestions, and memory operations.
 Governance Pack	Operating rules for identity, consent, access, deletion, export, and escalation.

REQUIRED CLIENT MATERIAL

-  **Private notes**, journals, working memos, planning docs, and draft outputs.
-  **Historical calls**, transcripts, post-mortems, meeting notes, and client interactions.
-  **Corrections**: examples where the Persona was wrong, off-style, unsafe, too generic, or too confident.
-  **Outcomes**: what happened after advice, forecasts, recommendations, or client deliverables.
-  **Permission rules**: who owns which material, who can use it, and what cannot leave the local node.
-  **Retention/deletion rules**: what must expire, what must be versioned, what must never be indexed.
-  **Escalation rules**: what requires human review, legal review, compliance review, or owner approval.

ARCHITECTURE

- 1 MindVault is designed local-first.
 - 2 Raw material stays on the user's machine by default.
 - 3 Local node handles ingestion, parsing, chunking, embedding, search, and vault storage.
 - 4 Atlas or another Retis trainer receives summaries, metrics, review signals, and tuning requests through pulses, not raw private files by default.
 - 5 Trainer pulses can suggest prompt changes, guardrails, source weighting, evaluation cases, retention changes, and drift fixes.
 - 6 User or owner approval controls whether suggested changes are applied.
-  Retis Plaza can discover the right PersonaTraining service without owning the private vault.

EXAMPLE PULSES

- memory_summary** — summarize what the local vault can support without exposing raw material.
- memory_gap_check** — identify weak areas, stale sources, and missing evidence.
- review_response** — record whether a Persona answer was good, unsafe, off-style, or unsupported.
- trainer_suggestion** — suggest prompt, guardrail, retrieval, or source-weight changes.
- drift_check** — compare recent outputs against Skillprint, Stylecast, and guardrails.
- permission_check** — decide whether a memory item can be used for a given output.
- retention_review** — flag content that should be expired, archived, deleted, or re-approved.
- handoff_context** — package safe context for a human reviewer without leaking restricted material.

PACKAGE LEVELS

 MINDVAULT SETUP	 MINDVAULT PRO	 MINDVAULT GOVERNED
Use case: Add private memory safely after a working Skillprint exists.  Includes: Memory map, local vault setup, public/private split, permission baseline, retention baseline, review history.  Timeline: 3–5 weeks  Client effort: Medium  Indicative price: USD 10k–25k	Use case: Private evolving Persona for a recurring expert workflow or team.  Includes: Permission model, audit trail, response review workflow, drift monitor, trainer pulses, deletion/export process, evaluation suite.  Timeline: 6–10 weeks  Client effort: Medium–high  Indicative price: USD 30k–75k	Use case: Enterprise, regulated, legacy, or estate use where memory rights and auditability matter.  Includes: Advanced access controls, retention/deletion governance, approval workflow, compliance-oriented audit logs, periodic risk review, managed Trainer reports.  Timeline: 10–16+ weeks  Client effort: High  Indicative price: USD 75k+ or managed retainer.

GOVERNANCE RULES

- ✗ Private memory is not public source material.
- ✗ Retrieved private material should guide reasoning, not be pasted into user-facing output.
- ✗ The Persona must not reveal private notes, call details, file paths, account handles, transcripts, metadata, or internal review records unless explicitly allowed.
- ✗ Every private memory class needs owner, access, retention, deletion, and export rules.
- ✗ Trainer suggestions should be reviewable before they change prompts or guardrails.
- ✗ Late-person or legacy projects require clear rights, consent basis, and scope restrictions.
- ✗ High-risk domains need escalation and refusal rules before launch.
- ✗ Browser cache or local UI state must not be the durable source of truth for operational memory.

RELATION TO OTHER SERVICES

-  **Skillprint** defines the expert's method and judgment.
-  **Stylecast** defines the communication behavior.
-  **AgentFront** exposes the service interface.
-  **MindVault** stores and governs private evolving memory.
-  **PersonaTraining** monitors quality, suggests tuning, and helps prevent drift.

MindVault should usually come after Skillprint and AgentFront basics. Launching private evolving memory before method, interface, and guardrails are clear creates unnecessary risk.

SUCCESS CRITERIA

- ✓ Raw private material remains local unless the owner explicitly changes the policy.
- ✓ Users can tell which memory is public, private, restricted, expired, or review-only.
- ✓ Persona outputs improve from reviews without leaking private material.
- ✓ Owners can inspect, delete, export, or version important memory.
- ✓ Drift is detected before it becomes user-visible behavior.
- ✓ Trainer changes are auditable and reversible.
- ✓ The Persona becomes more useful over time while staying inside the approved service boundary.

CORE LINE

**MindVault is
private memory
with governance:
remember
what matters,
expose only
what is allowed.**





ONE-LINE DEFINITION

Retis AI builds privacy-first intelligent systems where people, agents, data services, and publishing services collaborate through explicit ownership, permissions, and controlled contribution.



The public Retis web positioning is: financial intelligence connected by agents and owned by people.



POSITIONING

Retis AI is not a single chatbot product. It is a product family for networked AI work:

- Professionals keep control of their methods, data, credentials, and private research.
- Users define objectives, choose services, review outputs, and own the final work.
- Agents expose controlled capabilities instead of exposing private prompts, source material, or internal systems.
- Structured workflows replace open-ended prompt sessions where repeatability and accountability matter.



CORE DIFFERENTIATORS

- Network model instead of isolated chatbot.
- Ownership and permissions are first-class product concepts.
- Local-first source material for Persona workflows.
- Controlled contribution from experts, data providers, and service agents.
- Structured report and publishing workflows instead of blank prompt boxes.
- Open-source foundation through Prompts, Phemacast, and Atlas code.
- Retis Plaza for identity, discovery, routing, and agent coordination.
- Trainer pulses for continuous improvement without centralizing raw private material.



GUARDRAILS

- Do not describe Retis Persona as cloning people.
- Do not imply Atlas provides investment advice beyond approved user and expert workflows.
- Do not suggest experts must upload private data, credentials, prompts, or source material to Retis.
- Do not blur public content, user-owned private content, and third-party copyrighted material.
- Do not publish a Persona as the real person; disclose it as a controlled Persona or service interface.
- Do not launch MindVault workflows without retention, deletion, access, audit, and versioning rules.



RETIS AI PRODUCT FAMILY

PRODUCT	ROLE	PRIMARY AUDIENCE
Atlas	Multi-agent financial intelligence network for reports, analysis, data services, renderers, and delivery agents.	Portfolio managers, analysts, traders, asset managers, risk teams, financial data providers.
Retis Persona	Local-first expertise digitization layer for SkillPrint, Stylecast, AgentFront, MindVault, and PersonaTraining.	Experts, creators, advisors, consultants, educators, late-expert legacy projects, internal knowledge teams.
Prompts	HTTP-native multi-agent runtime with Plaza discovery, identities, persistence pools, and verified remote execution.	Engineers and teams building networked agent systems.
Phemacast	Programmable rendering pipeline for Phemas, pulse bindings, Phemar transforms, and Castr outputs.	Knowledge publishing teams, report builders, product teams, and content systems.
Retis Plaza	Coordination and discovery layer where agents register, authenticate, discover services, and route pulses.	Retis network operators, agent builders, service marketplaces, enterprise deployments.

SUGGESTED RETIS AI MESSAGING

SHORT

Retis AI builds privacy-first agent networks where experts, data services, and AI workflows collaborate without giving up control.

MEDIUM

Retis AI is the company behind Atlas and Retis Persona. Atlas coordinates financial intelligence workflows across analyst Personas, data services, report generators, and delivery agents. Retis Persona turns professional expertise into controlled agent interfaces that can participate in Atlas or any other Retis network.

LONG

Retis AI builds intelligent systems for controlled collaboration. Its core products separate ownership from coordination. Users own objectives and final outputs, experts own their methods and Personas, data providers own source access, and service agents assemble approved work into reports, media, dashboards, and delivery channels. Atlas applies this model to financial intelligence. Retis Persona extends it to any expert whose method can be digitized, governed, and offered as a service.



ATTAS

Atlas stands for Autonomous Trading and Treasury Agents System.

Atlas is the Retis financial intelligence product. It connects analyst Personas, data APIs, report generators, renderers, and delivery agents into coordinated financial workflows.

WHAT ATTAS PROVIDES

User-Owned Objective	The user defines the question, report scope, workflow, visibility, and final deliverable.
Analyst Persona Agents	Analysts, portfolio managers, strategists, and specialists contribute approved judgment without exposing private methods.
Data And Filing Agents	Market prices, transcripts, filings, exhibits, research context, and derived data can enter the workflow under controlled access.
Report Generators	Services assemble selected data and expert context into templates, briefings, reviews, and decision-ready outputs.
Delivery Agents	Renderers and newsletter agents turn finished intelligence into documents, images, HTML, video, voice, or scheduled delivery.
Personal Agent Workspace	Users can keep research panes, charts, saved views, diagram reasoning, and report outputs together.

ATTAS WORKFLOW

- Choose the report service.
- Define the report scope: ticker, company, topic, language, template, analyst, or data source.
- Choose prompt instructions or a sectioned storyboard.
- Select media, visibility, renderer settings, avatars, voices, provider keys, and estimated points.
- Track, review, download, share, or rerun the workflow.

ATTAS TARGET AUDIENCES

- Portfolio managers who need repeatable investment reports and decision support.
- Analysts who want to publish controlled insight as a service.
- Traders who need catalyst context, market read-throughs, and trade-review content.
- Asset managers who need holdings, benchmark, and committee-ready reporting.
- Risk managers who need exposure notes, stress views, downside cases, and guardrails.
- Data providers and internal teams that want to publish controlled services instead of one monolithic interface.



RETIS PERSONA

Retis Persona is the expertise layer of Retis AI.

It turns a professional method into a controlled Persona that can be used inside Atlas or any other Retis network. The Persona does not need to pretend to be the person. Its job is to become a governed interface to the expert's approved method, style, service boundaries, and evolving memory.

Raw source material should stay local by default. The local Persona node ingests and embeds content on the user's machine, while Retis or Atlas trainer services can send tuning suggestions through pulses.

RETIS PERSONA SERVICES

SERVICE	PURPOSE	OUTPUT
SkillPrint	Distills the expert's repeatable method, judgment rules, frameworks, caveats, and output formats.	Method map, guardrails, evaluation set, baseline Persona.mnd.
Stylecast	Calibrates writing style, video/script style, tone, cadence, vocabulary, and disclosure rules.	Recognizable expression without weakening judgment.
AgentFront	Builds a user-facing service entry point for Q&A, intake, comments, summaries, reports, and routing.	Practical proxy for expert services with handoff rules.
MindVault	Adds private evolving memory, feedback loops, audit history, retention, deletion, and versioning.	Governed long-term expert memory.
PersonaTraining	Reviews outputs and suggests prompt, guardrail, retrieval, and evaluation updates through pulses.	Continuous improvement without uploading raw source material.

HOW ATTAS AND PERSONA FIT TOGETHER

Atlas creates the financial workflow.

Retis Persona supplies controlled expertise.

Example:

- A portfolio manager asks Atlas for a company report.
- Atlas selects data services for market prices, filings, transcripts, and news.
- Atlas calls analyst Persona agents for approved judgment.
- Each Persona contributes through its SkillPrint and guardrails without exposing private research.
- Report agents assemble the final output.
- Delivery agents render the result as a memo, deck, HTML page, newsletter, voice, or video.
- The combined value is controlled collaboration: contributors can participate in a network without surrendering their edge.



OTHER RETIS WEB 2 INFORMATION

AREA	CURRENT SITE INFORMATION
Brand	Retis / Retis.ai / Intelligent Systems.
Site Description	Privacy-first AI products for orchestration, benchmarking, and audience-aware publishing.
Main Products	Prompts, Atlas, Phemacast. Retis Persona should be added as the expertise digitization layer.
Open Source	Public Apache-2.0 repository at https://github.com/alvincho/retis .
Atlas Beta	Beta entry point at https://atlas-beta.retis.ai .
Navigation	Home, Products, Blog, Docs, Contact, Download.
Blog / Thought Leadership	Agentive Futures Medium publication at https://medium.com/agentive-futures .
Contact	info@retis.ai ; technical support surfaced as docs.retis.ai.
Localization	English, Traditional Chinese, Simplified Chinese, Japanese, Korean, German, Spanish, French, and Italian message catalogs.
Web Stack	Next.js 15, React 19, Framer Motion, TypeScript, localized message files, Retis product image assets.



RELATION TO OTHER SERVICES

SkillPrint	defines the expert's method and judgment.
Stylecast	defines the communication behavior.
AgentFront	exposes the service interface.
MindVault	adds private evolving memory after governance is ready.
PersonaTraining	reviews outputs and suggests prompt, guardrail, and retrieval updates.

If Stylecast is the voice of the service and AgentFront is the doorway, SkillPrint is the professional judgment inside.



SUCCESS CRITERIA

- The Persona follows the expert's method even when the wording changes.
- The Persona asks for missing context instead of guessing.
- The Persona can explain uncertainty and refuse unsupported claims.
- Outputs match approved formats and decision rules.
- Reviewers can mark whether an answer is on-method or off-method.
- Style can later be added without weakening judgment.
- AgentFront can route users because the service boundaries are clear.
- MindVault can evolve memory without corrupting the core method.



CORE LINE

Retis AI is controlled collaboration for intelligent work: Atlas coordinates financial intelligence, and Retis Persona turns expertise into governed agent services.



Contact Us

Retis AI

Website: <https://retis.ai>

Attas Beta: <https://attas-b.retis.ai>

Documentation: <https://docs.retis.ai>

Blog: <https://medium.com/agentive-futures>

Email: info@retis.ai

Privacy-first intelligent systems for controlled collaboration.