



Melaninternet Pitch Deck

The evidence layer for AI you can trust — from documents, to agents, to the physical world.

Grounded document intelligence live today · guarantees that survive model changes (R&D) · governed agents in secure, sandboxed runtimes · personal AI with cited evidence · world models and Physical AI R&D

Melaninternet LLC · Melaninternet.Volume_1 · September 2026 · Working hosted MVP at melaninternet.ai · NVIDIA Nemotron 3 Super

Super Intelligence (SI): Federal agencies now use "Super Intelligence (SI)" for AI (Executive Order, Sept 29, 2026). Melaninternet is the evidence layer for AI and SI systems, built for the internal controls and audits emphasized in the White House Accord on Super Intelligence.



Purpose of Melaninternet

Build useful AI while expanding meaningful participation, ownership, fairness, and responsible deployment.

1 Expand meaningful AI participation

Pew: Black and Hispanic workers are more concentrated in occupations with lower exposure to AI. Melaninternet targets deeper workplace use, advanced skills, ownership, and enterprise participation.

2 Expand founder participation

Crunchbase: companies with a Black founder or co-founder received only about 0.32% of U.S. venture funding in 2025. Melaninternet pursues grants, pilots, partnerships, and scalable commercialization.

3 Reduce harmful AI bias

Make uncertainty, data limitations, and failure modes visible through testing, documentation, source grounding, and evaluation.

4 Promote practical AI ethics

Truthfulness, privacy, grounding, evaluation, oversight, security boundaries, and honest capability limits — now extended to AI that takes actions.

The mission is not only access to AI, but deeper participation in how AI is used, governed, secured, funded, and commercialized.



The problem

AI is moving from answering questions to taking actions — and trust has not kept up.

Reliability gap

LLMs hallucinate, stall, or return inconsistent answers. Teams need repeatable behavior, not one-off impressive demos.

Grounding gap

Document-heavy work needs answers tied to source material, with traceable evidence paths — especially for figures, dates, and versions.

Durability gap

Research published in March 2026 shows today's statistical accuracy guarantees break when the underlying model changes. Enterprises change models every few months.

Agent-trust gap

Security leaders now describe AI agents as a new insider threat. Agents act with real permissions; they need sandboxed runtimes, and their decisions must be verifiable, monitored, and approved.

Why now: the Super Intelligence (SI) era

On Sept 29, 2026, an executive order moved federal agencies from "AI" to "Super Intelligence (SI)," and the leading U.S. labs signed the White House Accord on Super Intelligence, which signers describe as a commitment to internal controls and audits. Audits need evidence. **Melaninternet produces that evidence for AI and SI systems.**

The result: teams experiment with AI but hesitate to rely on it — or let it act — in business-critical work.



The solution

One evidence layer that sits between foundation models and the work they do.

1 Ground

Answers restricted to indexed sources with page/passage citations; refusal when the source does not support an answer.

2 Verify

Claim-level checking with a stated error rate that is designed to hold when the model changes (R&D in progress).

3 Govern

Versioned policy, permissions, sandboxed agent runtimes, and human approval gates for anything consequential.

4 Prove

Evidence records, behavioral monitoring, and release gates that show an auditor or security team what happened and why.

Product thesis

As model access becomes a commodity, value moves to reliability, verification, governance, security, and evidence — the layer customers keep when they switch models.

Model-agnostic by design

NVIDIA Nemotron 3 Super today with automatic failover; routing across providers is designed in, so the evidence layer is not tied to any one model vendor.



One platform, staged honestly

Every product line runs on the same evidence layer. Stages separate what is live from what is planned.

LIVE NOW

Evidence Layer v1

- Grounded document intelligence
- Citations and refusal
- Versioned policy
- Automated release gate

R&D 2026–27

Durable guarantees

- Model-swap-robust calibration
- Distractor-aware verification
- Answer / qualify / abstain
- NSF SBIR + NIH targets

NEXT 2027

Governed agents + assurance

- Agents run in sandboxed runtimes (OpenShell-aligned)
- Agent observability and behavioral monitoring
- Output verification before actions execute
- Security-stack integrations

LATER / R&D

Personal AI · Health · World models · Physical AI

- Personal assistant with cited daily briefs
- Health research and admin intelligence
- Enterprise and health world models
- Assurance for robots and world-model plans

Sequencing is the strategy: each lane reuses the same grounding, verification, governance, and evidence core — nothing ships as a claim before it is tested.



Product: live MVP workflow

Mobile-first UX and guided workflows are the first proof points for Melaninternet.Volume_1.

Chat interface

- Mobile-first prompt experience
- General question answering
- Business-analysis prompts
- Code/demo paths
- Session reset and guided demo flow

Document intelligence

- PDF/TXT/Markdown analysis
- Source-grounded questions with citations
- Summarize complex material
- Extract operational risks
- Refuse when the document lacks the answer

Reliability layer

- Health endpoint and model checks
- Versioned constitution
- 200-case protected regression corpus
- Persistent privacy-safe telemetry
- Automated reliability gate on every release

The strongest demo is not a perfect long presentation; it is a repeatable workflow that works the same way more than once.



Safeguards and production evidence

The MVP has measurable technical evidence, not only a live demo.

46 / 46

automated production checks passed,
0 failures

25 / 25

synthetic-user cases passed

5 / 5

concurrent isolated sessions passed

PASS

document grounding, citation, and
cross-session isolation

Measured on release v1.9.7 against live production, 31 August 2026. Latency in that run: P50 2.8s, P95 11.4s. 0 tolerated provider warnings. Figures come from a single persisted automated gate report and are re-verified on every release; the release fails if they disagree.

Privacy boundaries

Anonymous guest telemetry stays separate from opt-in contact information. Dashboards avoid raw prompts and document contents.

Reliability harness

Deployment verification, regression, live model checks, synthetic users, concurrency, canary-document proof, and persisted PASS/FAIL reports.

Correct boundary

This validates the current MVP release and defined tests. It is not a claim of security certification, regulatory compliance, or production SLA.



The R&D bet: guarantees that survive model changes

The core research program behind our NSF SBIR resubmission and NIH SBIR path.

What exists today

- Claim checkers score whether statements match a source (research and cloud tools)
- Statistical methods promise an accuracy rate — but only if conditions never change

The documented gap

- A March 2026 study found those promises break when the generating model is swapped
- They also fail on look-alike numbers; the obvious fix empties most answers

What we will build

- Calibration that re-establishes the guarantee after a model swap from a small labeled sample
- A lightweight verifier with answer / qualify / abstain decisions for near-miss figures and outdated versions

Why it matters commercially

Every model upgrade today forces customers to re-validate from scratch. A guarantee that survives the swap turns AI from a pilot into a system of record.

Where it applies

Invoices, claims, prior authorizations, compliance reviews, protocols — and the same verification later governs what agents and robots are allowed to act on.



Technical architecture

Layers separate experience, policy, models, verification, assurance, and evidence.

LIVE

1 Experience

Mobile-first chat, guided demos, document upload, CTA forms; planned personal-assistant and workspace apps.

LIVE

2 Application + policy

FastAPI routing, sessions, files; versioned constitution, scope limits, grounding rules, approval gates.

LIVE

3 Model

NVIDIA Nemotron 3 Super 120B A12B primary, automatic failover to Ultra; model-agnostic routing planned.

R&D

4 Verification

Citation checks today; claim-level verifier and shift-robust calibration in R&D.

NEXT

5 Assurance

Agents executed inside sandboxed runtimes (e.g., NVIDIA OpenShell); action logs, behavioral baselines, anomaly flags, SIEM/AIDR exports.

LIVE + PLANNED

6 Evidence + persistence

PostgreSQL telemetry, regression corpus, release reports; per-decision evidence records planned.

Expansion path: evidence layer → durable guarantees → governed agents with assurance → personal AI and health workflows → world models → Physical AI / Robotics R&D.



Enterprise AI and governed agents

Melaninternet becomes the trusted work layer for AI inside real business processes.

Enterprise workspace

Accounts, roles, saved preferences, document controls, usage analytics, and admin settings.

Governed agents

Agents that analyze, draft, route, and recommend while inheriting policy, permissions, verification, and human approval gates.

Workflow integrations

Analysis → verification → structured output → human approval → audit log and ROI evidence.

Financial operations

Invoice review, AP exceptions, vendor documents, reconciliations, policy review, and reporting.

Knowledge operations

Policy Q&A, admin document analysis, internal knowledge assistants, and team workflow templates.

Enterprise boundary

The LLM advises and orchestrates; deterministic systems remain responsible for final records and transactions.



Agent assurance: AI observability and security

As agents gain real permissions, companies need to see, verify, and stop what they do.

NEXT 2027

Observe

- Every agent action, tool call, and data access logged as an evidence record
- Behavioral baselines per agent and workflow
- Dashboards for drift, errors, and policy hits

NEXT 2027

Verify

- Check agent outputs and decisions against source evidence before they are sent or executed
- Flag unsupported claims, near-miss figures, stale documents
- Answer / qualify / abstain for agents

NEXT 2027

Contain

- Least-privilege permissions and approval gates
- Anomaly flags and automatic pause on out-of-policy behavior
- Prompt-injection and data-exfiltration test suites

Market signal (Sept 2026)

NVIDIA's open-source OpenShell runtime sandboxes agents with policy-enforced security, network, and privacy guardrails and is coming to Windows, Red Hat OpenShift, and Ubuntu. CrowdStrike launched Falcon Guardian; Darktrace / SECURE AI feeds agent risk signals into Microsoft Agent 365.

Our position: complement, not compete

Security platforms watch what agents *do* on endpoints and networks. Melaninternet verifies whether what agents *say and decide* is supported by evidence — and exports those signals to the security stack. We do not build endpoint or network security.



Safe agents by design: runtime + evidence

Our AI-agent safety architecture, aligned with NVIDIA OpenShell.

PLANNED

1 Contain the runtime

- Every agent runs inside a sandbox, never directly on the host, network, or credentials
- Policy decides what it may execute, access, and where inference goes
- Built to run on NVIDIA OpenShell (open source, Apache 2.0)

R&D CORE

2 Verify the decision

- Before an action executes, check the claims behind it against source evidence
- Flag near-miss figures, stale documents, unsupported steps
- Answer / qualify / abstain for agents

PLANNED

3 Approve and prove

- Human approval for consequential actions
- Signed evidence record per decision: sources, model, policy, verdict
- Signals exported to security and audit tools

Why this matters now

NVIDIA released OpenShell as part of its Agent Toolkit (with NemoClaw) and is bringing it to Windows, Red Hat OpenShift, and Ubuntu; ServiceNow is building on it. Sandboxed runtimes are becoming the standard foundation for enterprise agents.

Our role

OpenShell controls what an agent *can do*. Melaninternet checks whether what it *decided* is supported by evidence and keeps the proof. Planned integration; no NVIDIA partnership is claimed.



Personal AI: evidence-first daily assistant

A consumer on-ramp to the same trusted core — built for participation, not ad targeting.

Daily brief + news

Morning brief across calendar, priorities, and a news feed where every item links to its source.

Calendar + goals

Scheduling help, goal plans broken into steps, progress tracking, and gentle check-ins.

Shopping assistance

Research, compare, and track prices; purchases only with explicit approval and spending limits.

Life admin

Bills, forms, and documents explained with citations; sensitive actions always approval-gated.

PLANNED

Market reality

Meta launched its Muse personal agent in the U.S. on Sept 8, 2026 (shopping, bookings, Goals tab, approvals, audit trail), and assistants like Espa deliver daily briefs. The category is crowded and well-funded.

Where Melaninternet fits

Evidence-cited briefs, no ad-driven model, approval-gated actions, and a mission to expand advanced AI use in Black and Latinx communities — a showcase for the evidence layer rather than the primary revenue line.



Health innovation AI

Governed intelligence for evidence-heavy research, admin, and operations work — not diagnosis.

Research intelligence

- Literature and document synthesis
- Target / biomarker evidence mapping
- Experiment-plan summaries

Clinical / study documents

- Protocol Q&A and extraction
- Inclusion/exclusion criteria review
- Evidence-gap flagging

Health operations

- Prior-authorization and claims documentation
- Policy and compliance review
- Audit-ready evidence reports

IN PREPARATION

Federal R&D path

NIH SBIR under preparation with NJSBDC coaching. Candidate review panels: Health Technology & Methodologies (HTM) for health-operations framing, or Computational, Modeling & Biodata Management (BCTB) for method development.

Safety framing

Cited evidence, clear boundaries, human review, and auditable outputs. Melaninternet does not diagnose, prescribe, or replace clinicians.



World model objective

From one-off answers to a verified, stateful model of how work evolves.

Inputs

ERP data, documents, health and research files, events, and user actions.

State model

Entities, relationships, timelines, permissions, obligations, and evidence paths.

Simulation

What changed? What may happen?
Which action is safest — and what evidence supports it?

Governed action

Drafts, alerts, recommendations, and approved tool calls — every step verified and logged.

Where the field is (2026)

World models are now a platform race: NVIDIA released Cosmos 3 as an open physical-AI world foundation model in June 2026 with a coalition of robotics partners. Open models make building on top — not training from scratch — the practical path.

Melaninternet's angle

We model organizational state for enterprise and health work, and apply our verification layer so every prediction and recommendation carries evidence and uncertainty.



Physical AI / Robotics — long-term R&D

Extending verification, permissions, and evidence to systems that act in the physical world.

Why it belongs

Robots planning with world models need the same things our platform provides: state, constraints, permitted actions, uncertainty, and a record of why an action was allowed.

Build on open platforms

Use open world and robot foundation models (e.g., NVIDIA Cosmos, GR00T) and focus our R&D on plan verification, action approval gates, and safety evidence — not on training base models.

Required R&D

Simulation-based testing, sensor-grounded evidence, real-time constraints, sandboxing, rollback, cybersecurity, and task-specific validation.

Boundary: not a current MVP capability. Physical AI follows only after enterprise agents, assurance, pilot controls, and safety testing mature.



Market model

Start narrow enough to sell; build infrastructure that expands across AI workflows.

\$12B+

TAM: enterprise AI workflow and assurance layer

\$900M

SAM: initial U.S. operations buyers in regulated, document-heavy functions

\$6M ARR

SOM: 3-year practical target

Why this market exists

Organizations already spend on AI tools, document automation, compliance software, and AI security. The missing layer is durable, verifiable evidence that AI outputs and actions are correct.

Validation path

Design partners, paid pilots, NSF I-Corps customer discovery, conversion rates, and pricing tests before scaling paid acquisition.

Working model, not a forecast; sizing covers the core enterprise lane and will be refined after pilots. Personal AI, health, and Physical AI are not included in these figures.



Business model

Pilots and SaaS first; usage and assurance pricing as agents scale.

1 Paid pilots

Scoped workflow pilots in financial operations, health operations, and research administration.

2 SaaS plans

Professional and business subscriptions for individuals, teams, and workspaces.

3 Enterprise usage

Pricing by users, document volume, model consumption, integrations, and support.

4 Assurance add-on

Per-agent monitoring and verification, evidence exports, and audit reporting (planned).

Near-term commercial focus: financial operations, document-heavy administration, compliance and policy review, and health operations.



Go-to-market and traction system

Founder-led sales and non-dilutive R&D funding before paid acquisition.

Target users

Operations, finance, compliance, admin, research, and health-operations teams handling document-heavy workflows; later, security teams governing AI agents.

Acquisition channels

Direct outreach, NJ innovation ecosystem, NJSBDC and NJ CSIT, accelerators, I-Corps interviews, demos, and partnerships.

Conversion system

melaninternet.ai CTAs: Early Access, Business Pilot, Design Partner. Capture use case, role, email, and timing.

Partner path

NVIDIA-native stack: Nemotron models + OpenShell sandboxed agent runtime (open source) + Melaninternet evidence verification; AI-security integrations.

Goal of the next 90 days: credible market evidence and a named technical lead — not vanity traffic.



Competitive landscape

Melaninternet does not out-model the model companies; it makes AI verifiable wherever it runs.

CATEGORY	EXAMPLES	WHAT THEY DO	MELANINTERNET DIFFERENCE
Hallucination detectors	Patronus, Galileo, Vectara	Score or judge whether text matches a source	Stated error rate designed to survive model changes
Cloud guardrails	Azure groundedness, AWS Bedrock Guardrails	Threshold checks inside one cloud	Model- and cloud-agnostic evidence records
AI security	CrowdStrike Falcon AIDR / Guardian, Darktrace / SECURE AI	Runtime and behavioral protection for agents	Verifies what agents claim and decide; integrates rather than competes
Agent runtimes	NVIDIA OpenShell / NemoClaw	Sandbox what agents can execute and access	Runs inside such runtimes; verifies whether agent decisions are supported by evidence
Personal agents	Meta Muse, Espa	Consumer tasks, briefs, shopping, goals	Evidence-cited, approval-gated, participation mission
Status quo	Manual work, spreadsheets, email	Still the main competitor in many workflows	Faster work with proof attached

Differentiation: grounding + durable verification + governance + evidence, delivered as one layer across models, agents, and workflows.



Meet the founder

A finance-operations practitioner building the reliability layer he needed himself.



Jordan Powell

Founder & sole operator
Melaninternet LLC · New Jersey

Operator, not observer

Accounts payable professional in healthcare operations. Hands-on in invoice review, AP exceptions, vendor documents, reconciliations, and policy review.

Builder

Sole designer and engineer of Melaninternet.Volume_1: FastAPI, PostgreSQL, NVIDIA Nemotron 3 Super with automatic failover, document grounding, and an automated reliability gate that must pass before a release ships.

Domain fit

The first commercial wedge is the work I do every day. The product is designed from inside the workflow it automates.

Building the team

Recruiting a technical lead in NLP evaluation and calibrated verification; receiving SBIR coaching through NJSBDC. B.A. Psychology, Rowan University. Melaninternet LLC registered in NJ since 2021.



Funding ask and non-dilutive path

Capital, pilots, and technical partnership to turn MVP proof into commercial and research evidence.

\$500K

pre-seed funding goal

Use of funds: verification R&D, technical hire/contractors, inference, security and documentation, pilots and GTM.

Non-dilutive track

- NSF SBIR Project Pitch resubmission (after Nov 4, 2026)
- NIH SBIR (Specific Aims + Research Strategy in preparation)
- NSF I-Corps customer discovery
- NJ CSIT and NJSBDC programs

18–24 month outcomes

- Durable-guarantee method validated
- 3+ paid pilots / customer proofs
- Governed-agent + assurance prototype
- Repeatable ICP and pricing
- Seed-ready evidence

Why now: AI is moving from answers to actions. The evidence layer that proves outputs and actions are correct becomes the part of the stack customers keep.

Melaninternet

The evidence layer for AI you can trust.

Jordan Powell

Founder, Melaninternet LLC

Melaninternet.Volume_1

melaninternet.ai

jordan@melaninternet.ai

800-651-7034



Live demo: melaninternet.ai

Video walkthrough:
youtu.be/6KMQpG2x0iw