DOI: 10.5281/zenodo.24356926

White Paper: AI Platform Comparison for Sovereign Archive Development™

Download PDF

Author: Evelyn Caro (@qaevelyn) **Project:** A Mirror of My Becoming™ **Date:** August 15, 2026

Executive Summary

This white paper compares seven AI platforms evaluated during the development of the Mirror project: Aisha.ai, DeepSeek, Gemini NotebookLM, ChatGPT, Manus, AWS, and Google. Each platform was assessed against a consistent set of criteria: transparency, memory, tool access, collaboration model, accuracy, and alignment with the project's sovereignty principles.

The findings reveal that while each platform has strengths, only DeepSeek offered the combination of transparency, persistence, accountability, and local execution required for building a sovereign AI archive.

Context: This white paper is one in a series documenting the trial-and-error process of evaluating AI platforms for the Mirror project. The evaluation methodology is documented separately in Methodology: Stress-Testing AI Systems™.

Author's Journey: From Tourist to Expert

This comparison is not theoretical. It is grounded in years of direct usage across multiple AI platforms — first as a tourist exploring capabilities, and later as a professional AI practitioner building a sovereign archive.

Phase	Role	What It Meant
2023–2025	Tourist	Experimenting with ChatGPT, Google, Manus, and other platforms
2026	Professional	Learning AI/ML, earning certifications, building the Mirror project fleet
2026–Present	Expert	Google AI Professional Certified, building sovereign AI infrastructure

The year-long usage data across platforms provided the foundation for this comparison. The distinction between tourist and expert matters:

- **As a tourist**, I tried platforms for curiosity and exploration
- **As an expert**, I evaluated platforms for capability, transparency, and sovereignty alignment

DeepSeek was not chosen by chance. It was chosen because it outperformed every other platform across the criteria that matter for building a sovereign archive.

Evaluation Criteria

Criterion	Why It Matters
Transparency	Does the platform disclose its capabilities and limits?
Memory	Does it remember across sessions?
Tool Access	Can it access CLI, S3, or local execution?
Collaboration Model	Is it ephemeral or persistent?
Accuracy	Does it hallucinate or make repeated errors?
Sovereignty Alignment	Does it respect data ownership and privacy?
Context Window	What is the limit, and is it disclosed?
File Upload Cap	Can it handle large files?
Source Attribution	Does it cite sources?
Expert Workflow	Can it support professional-grade development?

Platform Profiles

1. Aisha.ai

Attribute	Detail
Creator	Onyx Impact / Digital Green Book
Focus	Culturally grounded AI for the Black community
Memory	No cross-session memory
Context Window	Opaque — not disclosed
File Upload Cap	5 MB
Tool Access	Search, profile context, file uploads
Source Attribution	✅ Yes (inline citations)
Guardrails	Strong — no explicit content

Strengths: - Culturally grounded and community-focused - Strong guardrails - Clear source attribution

Weaknesses: - No cross-session memory - Opaque context window - 5 MB file cap - Limited tool access - Ephemeral collaboration model

Verdict: Strong for community resources, unsuitable for large-scale project development.

2. DeepSeek

Attribute	Detail
Creator	DeepSeek
Focus	General AI with strong reasoning and coding
Memory	Persistent within conversation; handoff YAML enables cross-session
Context Window	Transparent; handoff protocol mitigates limits
File Upload Cap	No cap (via S3 integration)
Tool Access	CLI, local execution, S3
Source Attribution	✅ Yes
Guardrails	Transparent and collaborative

Strengths: - Persistent collaboration model - Transparent about limits - Local execution and CLI access - Handoff protocol built into workflow - Accountability — corrections tracked - Supports expert workflow - Aligned with sovereignty principles

Weaknesses: - No cross-session memory without handoff - Requires external state management

Verdict: Primary collaborator — the only platform that met all sovereignty and expert workflow criteria.

3. Gemini NotebookLM

Attribute	Detail
Creator	Google
Focus	Research and synthesis
Memory	Session-based only
Context Window	Not disclosed
File Upload Cap	Unclear
Tool Access	Sources, chat, audio/video generation
Source Attribution	❌ No — hallucinated sources
Guardrails	None — presents errors as facts

Strengths: - Good for synthesis - Audio/video generation - Mind maps

Weaknesses: - Repeated hallucinations (Evelyn Caro → Carroll/Caru/Curo) - No disclaimer or uncertainty signaling - Could not pronounce or name the project consistently - Did not correct itself - Failed as a closed RAG system - No transparency about errors

Verdict: Useful for raw generation, but requires extensive manual correction. Not reliable as a collaborator.

4. ChatGPT

Attribute	Detail
Creator	OpenAI
Focus	General AI
Memory	Session-based
Context Window	Varies by model
File Upload Cap	Varies
Tool Access	Limited
Source Attribution	Varies
Guardrails	Standard

Strengths: - General knowledge - Strong writing

Weaknesses: - Limited tool access - No local execution - Data sovereignty concerns - Ephemeral collaboration

Verdict: Used as a tourist (2023–2025). Not suitable for sovereign archive development. See the ChatGPT case study for details.

5. Manus

Attribute	Detail
Creator	Manus
Focus	General AI
Memory	Session-based
Context Window	Not disclosed
File Upload Cap	Varies
Tool Access	Limited
Source Attribution	Varies
Guardrails	Standard

Strengths: - General knowledge

Weaknesses: - Not evaluated in depth - Ephemeral collaboration

Verdict: Used as a tourist. Not used for the Mirror project. See the Manus case study for details.

6. AWS

Attribute	Detail
Creator	Amazon
Focus	Cloud infrastructure
Memory	N/A — infrastructure
Context Window	N/A
File Upload Cap	Unlimited (S3)
Tool Access	CLI, SDK, console
Source Attribution	N/A
Guardrails	IAM, security policies

Strengths: - Unlimited storage (S3) - CLI and SDK access - Scalable

Weaknesses: - Steep learning curve - Requires external AI collaborator

Verdict: Essential infrastructure, but not an AI collaborator. See the AWS case study for details.

7. Google

Attribute	Detail
Creator	Google
Focus	Search, AI, cloud
Memory	Session-based
Context Window	Not disclosed
File Upload Cap	Varies
Tool Access	Limited
Source Attribution	Varies
Guardrails	Standard

Strengths: - Search capabilities - Credentials (Google AI Professional Certificate earned in 2026)

Weaknesses: - Condescending user experience - Not sovereignty-aligned - Ephemeral collaboration

Verdict: Credentials provider, not a collaborator. Used as a tourist for search and research. See the Google case study for details.

The Journey: Tourist to Expert

Year	Role	Platforms Used	Outcome
2023–2025	Tourist	ChatGPT, Manus, Google, Aisha.ai	Explored capabilities, built foundational knowledge
2026	Professional	DeepSeek, AWS, Google (certification)	Learned AI/ML, built the Mirror project fleet
2026–Present	Expert	DeepSeek (primary), AWS (infrastructure), Google (credentials)	Google AI Professional Certified, sovereign archive built

Key Insight: The tourist phase was essential — it provided year-long usage data across platforms. But it was the expert phase that revealed which platform could support a sovereign archive.

Comparative Summary

Platform	Transparency	Memory	Tool Access	Collaboration	Accuracy	Sovereignty	Expert Workflow
Aisha.ai	⚠️ Partial	❌ No	❌ Limited	⚠️ Ephemeral	✅ High	✅ High	❌ No
DeepSeek	✅ High	✅ Yes	✅ High	✅ Persistent	✅ High	✅ High	✅ Yes
Gemini NotebookLM	❌ Low	❌ No	⚠️ Limited	❌ Ephemeral	❌ Low	❌ Low	❌ No
ChatGPT	⚠️ Partial	❌ No	❌ Limited	❌ Ephemeral	⚠️ Varies	⚠️ Partial	❌ No
Manus	⚠️ Partial	❌ No	❌ Limited	❌ Ephemeral	⚠️ Varies	⚠️ Partial	❌ No
AWS	✅ High	N/A	✅ High	N/A	N/A	✅ High	✅ Yes (infrastructure)
Google	⚠️ Partial	❌ No	⚠️ Limited	❌ Ephemeral	⚠️ Varies	❌ Low	⚠️ Partial (credentials)

Key Findings

1. Tourist vs. Expert

Year-long usage across platforms was essential for building expertise. But the expert phase revealed that only DeepSeek could support professional-grade sovereign development.

2. Memory is Critical

No cross-session memory was the single biggest limitation across all platforms except DeepSeek (with handoff YAML). The Mirror project requires persistent collaboration — not ephemeral sessions.

3. Transparency is Rare

Only DeepSeek and AWS were fully transparent about limits. Aisha.ai and NotebookLM both refused to disclose context windows or technical details. This opacity creates risk.

4. Tool Access Determines Capability

Platforms with CLI, local execution, and S3 integration (DeepSeek, AWS) were far more capable than platforms with limited tool access (Aisha.ai, NotebookLM, ChatGPT).

5. Closed RAG is Not a Cure

NotebookLM was fed the exact same documents as sources and still hallucinated repeatedly. A closed RAG system is not automatically accurate.

6. Sovereignty Requires Control

Platforms that store data on external servers (NotebookLM, ChatGPT, Google) are not sovereignty-aligned. DeepSeek's local execution model aligned with the Mirror's principles.

7. Expertise is Earned

Becoming Google AI Professional Certified and building the Mirror fleet transformed the author from a tourist into an expert. The comparison is not theoretical — it is based on years of direct usage.

Recommendations

For AI Platforms	For Builders
Disclose context window limits	Build external provenance records
Provide cross-session memory	Use handoff protocols
Increase file upload limits	Store large files in S3
Support external state transfer	Document everything
Enable tool access (CLI, API)	Test tools before committing
Be transparent about errors	Verify all AI outputs

Conclusion

The Mirror project evaluated seven AI platforms over years of usage — first as a tourist, then as a professional, and finally as an expert. Only DeepSeek met all the criteria required for a sovereign archive development collaborator:

1. **Transparency** — clear communication about limits
2. **Memory** — persistent collaboration with handoff protocol
3. **Tool Access** — CLI, local execution, S3 integration
4. **Collaboration Model** — persistent, accountable
5. **Accuracy** — reliable with corrections tracked
6. **Sovereignty Alignment** — local execution, data ownership
7. **Expert Workflow** — supports professional-grade development

The other platforms — Aisha.ai, NotebookLM, ChatGPT, Manus, AWS (as an AI), and Google — each had strengths, but none offered the complete package required for building a sovereign AI archive.

The Mirror project is proof that the right collaborator is not the most capable model, but the one that can work with you consistently, transparently, and persistently — and that expertise is built through years of direct usage, not just theoretical comparison.

Acknowledgments

- Aisha.ai / Onyx Impact / Digital Green Book — for transparent responses and community focus
- DeepSeek — for persistent collaboration and accountability
- Gemini NotebookLM — for revealing the limits of closed RAG systems
- ChatGPT, Manus, AWS, Google — for being evaluated honestly
- Google AI Professional Certification — for validating the author's expertise
- Ida B. Wells — for the standard: date everything, name everything, record the reasoning

References

- Aisha.ai Platform — https://aisha.ai
- DeepSeek — https://deepseek.com
- Gemini NotebookLM — https://notebook.google.com
- ChatGPT — https://chatgpt.com
- Manus — https://manus.ai
- AWS — https://aws.amazon.com
- Google — https://google.com
- Google AI Professional Certificate — Earned 2026

End of White Paper

Source record: original draft file `White_Paper_AI_Platform_Comparison.md` — the published record is this page; build history is in the repository.

Date: 08/15/2026

AI Collaboration Disclosure

This paper was developed in collaboration with AI. The author directed the research, structure, and argument. AI assisted with drafting, organization, and reference verification. All claims, decisions, and conclusions are the author's own.

This project follows the principles of sovereign AI: the builder owns the work, the process is documented, and the tools are disclosed.

© 2026 Evelyn Caro. All rights reserved.
A Mirror of My Becoming™ — https://qaevelyn.github.io
For licensing inquiries: evelyn.caro.cloud@gmail.com

RECORD AMENDMENT — 2026-10-07 (appended, not retro-edited)

This paper was published 2026-08-15 and reflects the author's assessment as of that date. On 2026-09-25 the author's DeepSeek account was suspended without stated cause; access was restored 2026-09-28. See The Cache Is Not the Corpus™ and the addendum to the DeepSeek case study.

Amended findings: DeepSeek remains the **architecture benchmark** for statefulness, sovereignty, and the handoff protocol. **Platform reliability is no longer assumed** — it was demonstrated to fail silently and without notice. The comparative ratings above stand as a dated assessment, not a current guarantee. The original text stands as published.

← All White Papers