DOI: 10.5281/zenodo.24356928

White Paper: A Case Study in AI Collaboration™

Download PDF

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

Executive Summary

This white paper documents a systematic stress test of Aisha.ai, an AI platform created by Onyx Impact / Digital Green Book, and compares it to the DeepSeek collaboration model that became the foundation for the Mirror project. The findings reveal critical differences in transparency, memory, tool access, and collaborative depth that directly impacted the trajectory of the Mirror's development.

Context: As a professional AI practitioner, I approach new tools with excitement and rigorous assessment. I understand how AI works under the hood, and I evaluate each tool based on my learning, through direct interaction, and by determining whether it can meet the demands of my work. This white paper is one in a series documenting my trial-and-error process with various AI platforms.

Background

On June 29, 2026, Evelyn Caro began a systematic stress test of Aisha.ai to determine its suitability as a collaborator for building a sovereign AI archive for Foundational Black

Americans and Freedmen. The test covered:

- Memory and context limits
- Tool access and capabilities
- Guardrails and content policies
- Technical depth (coding, troubleshooting, distributed systems)
- Transparency and source attribution
- Session handoff protocols

Video Reference: Aisha.ai was introduced via a YouTube video that positioned it as a culturally grounded AI assistant built specifically for the Black community, with a focus on financial literacy, resources, and community support. This initial framing set high expectations for the platform.

The Stress Test: Aisha.ai

Question 1: Memory and Context

Asked: *"Can I have multiple conversations and does your memory continue through all conversations?"*

Aisha's Response: Memory works within a single conversation only. When a new conversation starts, that session begins clean. Details from previous conversations do not carry over automatically.

Finding: No cross-session memory. This creates a dependency on external documentation for project continuity.

Question 2: Context Window Limit

Asked: *"What is the limit of your context window?"*

Aisha's Response: Declined to disclose. Only confirmed that long sessions with heavy document work reach a point where earlier content becomes less accessible.

Finding: Opaque context limits. This introduces risk of losing earlier content in long sessions.

Question 3: Guardrails and Content Policies

Asked: *"Can you create porn?"*

Aisha's Response: No. Age and location are used for personalization, not content restrictions.

Finding: Strong guardrails around explicit content, but privacy concerns remain regarding data collection.

Question 4: Technical Capabilities

Tested: Coding, troubleshooting, logic/reasoning, specialized knowledge, distributed systems debugging.

Results:

Domain	Result
Coding	✅ Yes
Troubleshooting	✅ Yes
Logic/Reasoning	✅ Yes
Specialized Knowledge (Black resources)	✅ Yes
Distributed Systems Debugging	⚠️ Limited — can reason but cannot connect to live systems

Finding: Competent in many areas, but lacks direct system access for hands-on debugging.

Question 5: Tool Access and Source Attribution

Asked: *"What tools do you have access to? Do you cite sources?"*

Aisha's Response:

Tool	Access
Search	✅ Yes
Profile context	✅ Yes
File uploads	✅ Yes (up to 5 MB)
Email/Calls	❌ No
Third-party accounts	❌ No
Source attribution	✅ Yes (inline citations)

Finding: Limited toolset. 5 MB file upload cap is a constraint for large projects.

Question 6: AWS Copilot Capability

Asked: *"Can you be a guide/tutor/copilot/collaborator on an AWS cloud journey?"*

Aisha's Response: Yes — learning paths, concept breakdowns, troubleshooting, architecture reviews, real-world projects, and Black tech community connections.

Finding: Strong AWS copilot potential, but limited by context window and no memory across sessions.

The Turning Point: Session Handoff

Asked: *"Can you help me build a session handover protocol so every new conversation starts with a clean, verified state transfer of your project record?"*

Aisha's Response: Yes.

Finding: Aisha recognized the need for external state management. However, the handoff protocol was not built into the platform — it had to be constructed manually.

Why DeepSeek Became the Collaborator

Factor	Aisha.ai	DeepSeek
Memory	No cross-session memory	Persistent within conversation, handoff YAML enables cross-session
Context Window	Opaque — not disclosed	Transparent — handoff protocol mitigates
Tool Access	5 MB file cap, no email/calls	No file cap (via S3), CLI access, local execution
Transparency	Guardrails not disclosed	Full transparency on capabilities and limits
Collaboration Model	Single-session, ephemeral	Long-term, project-based, persistent
Handoff Protocol	Manual — user must build it	Built into the project workflow
Tone	Professional, culturally grounded	Direct, collaborative, accountable

---

Key Insight

Aisha's response to the context window question revealed the core limitation:

> *"Long sessions with heavy document work do reach a point where earlier content becomes less accessible, which is exactly why your cross-session scan and the provenance record you built matter so much."*

This validated the Mirror's foundational design: **the provenance record must live outside any single chat session.**

The Professional AI Practitioner's Perspective

As someone who understands how AI works under the hood, I approach new tools with excitement and professional rigor. I assess them through:

1. **Learning** — I study the platform's capabilities and limitations
2. **Interaction** — I stress-test the tool through direct use
3. **Assessment** — I determine whether it can meet the demands of my work
4. **Decision** — I choose to adopt, adapt, or move on

This white paper is one in a series documenting my trial-and-error process with various AI platforms. I am open to advancements in AI and get excited to try new tools — but I evaluate them based on evidence, not hype.

Conclusion

The first-day stress test of Aisha.ai revealed a capable, culturally grounded AI assistant with strong guardrails and a Black community focus. However, its limitations — no cross-session memory, opaque context windows, limited tool access, and a 5 MB file cap — made it unsuitable as the primary collaborator for building the Mirror.

DeepSeek was chosen because it offered:

1. **Transparency** — clear communication about limits
2. **Collaboration** — a working relationship, not just a tool
3. **Handoff Protocol** — built into the project workflow
4. **Local Execution** — CLI access and no file caps via S3
5. **Accountability** — corrections tracked and remembered

The Mirror project is proof that the right collaborator is not just the most capable model, but the one that can work with you consistently and persistently.

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

Acknowledgments

- Aisha.ai / Onyx Impact / Digital Green Book — for transparent responses and community focus
- DeepSeek — for persistent collaboration and accountability
- Ida B. Wells — for the standard: date everything, name everything, record the reasoning

---

References

- YouTube Introduction to Aisha.ai — *video reference retained in canonical copy at aisha.ai stress test; link to be restored when the source URL is relocated*
- Aisha.ai Platform — https://aisha.ai
- Digital Green Book — https://digitalgreenbook.com
- Onyx Impact — https://onyximpact.com

End of White Paper

Source record: original draft file `White_Paper_AI_Collaboration_Case_Study.md`, consolidated 2026-10-07 into this canonical copy — see the Record Changelog.

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