NP LAB

From Showroom Guide to Museum Interpreter

A conceptual case study examining how the existing Deckhand Dave kiosk platform could be adapted for nonprofit museums

Concept status: Research exploration. North Parade LLC has not developed, tested, or deployed this museum adaptation. No museum, collection, community, donor, or visitor participated in the preparation of this study. The proposal describes a possible reuse of an existing kiosk architecture, not a finished product or evidence of visitor outcomes.

North Parade LLC | August 2026

The problem is not a lack of information

Museums already contain information. It appears on object labels, wall panels, maps, brochures, websites, audio tours, collection databases, and the knowledge held by staff and volunteers. The harder problem is helping a particular visitor reach the right information at the moment a question occurs.

A label has limited space. A docent cannot stand beside every object. Front-desk staff may be handling admissions, memberships, school groups, accessibility needs, and building operations at the same time. Small nonprofit museums often face this problem with fewer employees, shorter public hours, older technology, and little room for a permanent digital team.

Visitors also arrive with different interests and levels of background knowledge. One person wants a short explanation. Another wants provenance, technical details, or the broader historical setting. A child may ask a question that the exhibit text never anticipated. A veteran may recognize equipment but use a term that does not appear in the catalog. A visitor with limited vision may need the same material presented differently.

The opportunity is not to replace exhibit interpretation with artificial intelligence. It is to create an optional conversational layer between the visitor and the museum's approved knowledge.

The proposed adaptation

Deckhand Dave is an existing local-first touchscreen kiosk platform designed for marine dealerships. Its current purpose is to help visitors browse site-specific information, ask questions, and request assistance without requiring a staff member to remain beside the kiosk.

The museum concept would retain the reusable platform and replace its marine-specific content, language, and workflows. The adapted system would act as a digital museum guide with a deliberately narrow role. It could help visitors explore approved collection material, locate exhibits and services, select an accessible presentation mode, and request a conversation with museum staff or a volunteer.

The museum would remain the author and authority. The kiosk would not independently decide what an artifact means, invent missing provenance, speak on behalf of a represented community, or turn speculation into fact.

Possible interaction modes could include:

- **Ask about an object.** Explore facts, context, provenance, or terminology associated with a displayed item.
- **Help me find something.** Locate galleries, restrooms, exits, elevators, quiet areas, or visitor services.

- **Give me the short version.** Receive a concise explanation in plain language.
- **Go deeper.** Access longer curator-approved material, source notes, or related objects.
- **Family mode.** Use age-appropriate prompts that encourage observation and discussion without claiming to replace an educator.
- **Accessibility options.** Increase type size, simplify navigation, read approved text aloud, or offer translated material that has been reviewed.
- **Talk to a person.** Request help from available staff or display clear directions when live assistance is unavailable.

The concept should not require a visitor to create an account, install an application, provide an email address, or surrender personal information merely to ask a question.

A fictional visitor scenario

The Harbor and Industry Museum is a fictional regional nonprofit with a small paid staff and a larger volunteer group. Its galleries cover shipbuilding, local manufacturing, immigration, military service, and the environmental history of the harbor.

During a Saturday visit, Elena and her twelve-year-old son Mateo stop at a shipyard photograph from 1943. The label identifies the location and photographer but does not explain several marks written on the workers' clothing. Mateo asks whether they were military ranks.

Elena opens the nearby kiosk and selects **Ask about this exhibit**. She enters the object number printed on the label. The system retrieves the museum's approved record and explains that the marks were most likely trade or work-group identifiers, not military rank. It clearly labels the answer as the museum's current interpretation and offers two related items: an oral-history excerpt from a welder and a photograph showing the same marking system in another building.

Mateo asks why women are visible in one photograph but not the other. The kiosk does not invent a single explanation. It summarizes the museum's reviewed material about wartime labor, notes what is known about the dates and departments shown, and identifies that the available records do not establish why the groups differ.

Elena selects **Where can we learn more?** The kiosk points them toward the labor-history gallery and offers to display a route. It also notes that a volunteer interpreter is scheduled there until 2:00 p.m.

The interaction adds depth without pretending that the machine witnessed the history, performed new research, or replaced the museum's staff.

What changes from the existing Dave platform

Existing Deckhand Dave capability	Museum adaptation
Site-specific marine inventory and service information	Museum-approved collections, exhibit, event, and visitor-service content
Questions about boats, equipment, and dealership services	Questions about objects, themes, people, places, terminology, and museum logistics
Sales or service staff assistance request	Front desk, docent, educator, accessibility, or security assistance pathway
Dealer branding and operating rules	Museum identity, interpretive standards, accessibility rules, and community commitments
Local kiosk interface and local language model	Local or institution-approved model grounded in a controlled museum knowledge package
Product-oriented browsing	Curiosity-driven exploration, wayfinding, and related-object discovery

This is a meaningful adaptation, but it does not require rebuilding the entire platform. The touchscreen shell, local service structure, content retrieval pipeline, administrative controls, health checks, and managed deployment model remain useful. The largest changes would be content governance, exhibit identification, accessibility, multilingual review, visitor analytics, and the limits placed on historical interpretation.

Museum knowledge must remain controlled

A general language model knows many facts, but a museum cannot treat that broad knowledge as authoritative. Collection records can contain outdated attributions, incomplete provenance, contested terminology, donor restrictions, culturally sensitive material, or information that should not be exposed to the public. Public web sources can repeat errors or flatten disagreements.

The museum version should answer from a reviewed content package. Each content item would carry basic governance fields:

- Museum record or exhibit identifier.
- Public title and approved description.
- Source and curator or departmental owner.
- Permitted audiences and uses.
- Known uncertainties or competing interpretations.
- Cultural, legal, copyright, donor, or privacy restrictions.

- Approved related objects and visitor pathways.
- Review date, version, and expiration or recheck date.

The kiosk could use a language model to turn approved material into natural conversation, but the model should not silently expand the factual boundary. When the source package does not support an answer, the system should say so and offer a staff handoff, a related source, or a narrower question.

That limitation is especially important for Indigenous cultural material, human remains, sacred or ceremonial objects, disputed ownership, wartime experiences, racial violence, disability history, and personal records. In those areas, fluent language can make an unsupported answer sound more legitimate than it is.

Proposed technical design

The museum kiosk would not send an unrestricted prompt directly to a language model. A controlled local service would identify the exhibit context, retrieve permitted museum material, assemble the request, and check the response before display.

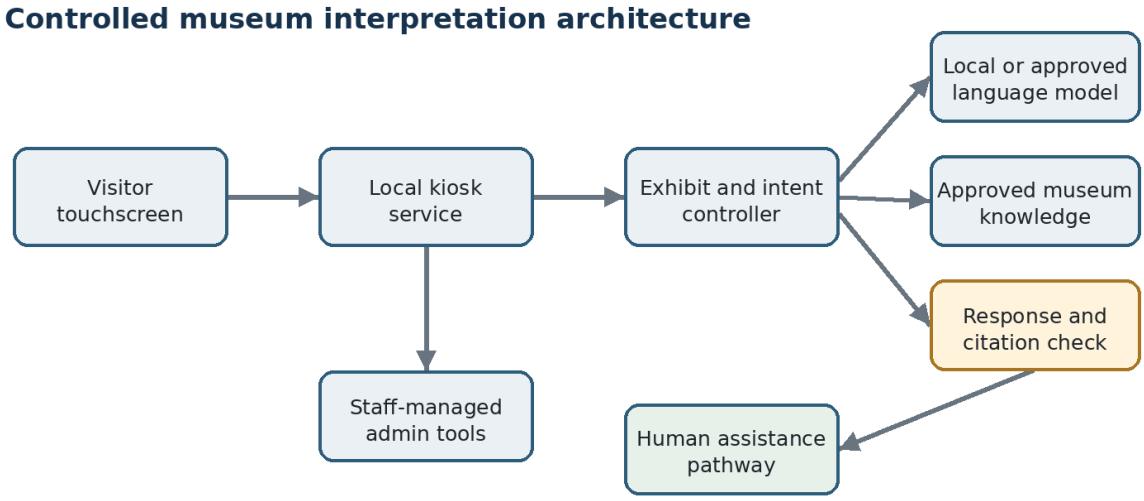

Figure 1. Controlled museum interpretation architecture

The content path should be separate from the model. Museum staff could update hours, closures, exhibit records, accessibility information, events, and approved interpretive text without retraining the model. Every response should be traceable to the content versions used to produce it.

The kiosk could operate primarily on site, allowing core exhibit questions to continue when internet service is unreliable. Features that require email, remote monitoring, cloud translation, ticketing, or external collection systems would need a network and should fail clearly rather than pretending to work.

For small museums, the practical question is maintenance. A technically impressive kiosk that requires a specialist to update one label is not sustainable. The administrative

workflow must allow authorized staff to import or revise structured content, preview likely answers, approve publication, roll back a bad version, and see whether the system is healthy.

Accuracy, interpretation, and institutional trust

Museums are trusted because their information is researched, contextualized, and attributable. A conversational interface changes how that authority is perceived. Visitors may assume that a confident answer is an official museum position even when the model has improvised it.

The interface should therefore make several distinctions visible:

- **Documented fact:** supported directly by approved museum records.
- **Museum interpretation:** an institutional explanation supported by cited material.
- **Uncertainty:** incomplete, disputed, or changing information.
- **Visitor reflection:** an invitation to observe or compare, not a factual conclusion.
- **Outside scope:** a question the kiosk cannot responsibly answer from approved material.

The system should not impersonate a historical person unless the museum has deliberately created and reviewed that interpretive experience. Even then, it should identify the interaction as a reconstruction. First-person roleplay can be engaging, but it can also manufacture beliefs, emotions, and statements that the historical record does not support.

Accessibility cannot be an afterthought

A museum kiosk can improve access, but a touchscreen can also introduce new barriers. The physical installation and interface would need to account for seated reach, wheelchair approach, glare, type size, contrast, motor control, hearing, vision, cognitive load, language, and the possibility that voice interaction is inappropriate in a quiet or crowded gallery.

Useful features could include keyboard or switch access, visible focus states, adjustable text, high contrast, captions, headphone support, volume control, plain-language summaries, and a consistent way to return to the beginning. Text-to-speech should read approved museum content rather than generate a different interpretation. Automated translation could assist exploration, but important public text should be professionally reviewed before the museum presents it as authoritative.

The pilot should include visitors with disabilities and accessibility specialists before public deployment. Compliance cannot be established by checking whether the screen technically responds to touch.

Privacy and visitor analytics

The safest default is to collect little. A visitor should be able to use the kiosk anonymously. Questions can reveal religion, disability, family history, political beliefs, military experience, immigration history, or other personal context even when the museum never asks for a name.

The museum may still need basic operational evidence. A privacy-conscious pilot could count sessions, selected modes, unanswered topics, response time, abandoned interactions, assistance requests, and hardware errors without retaining full conversation transcripts.

If staff need sample conversations for evaluation, the pilot should use clear notice, short retention, restricted access, and a documented review purpose. The system should not build visitor profiles, use facial recognition, infer demographics, or connect anonymous exhibit questions to ticketing, donor, membership, or surveillance records.

Local operation reduces some exposure, but it does not eliminate privacy risk. Anyone with administrative access to the kiosk, backups, logs, or remote-support tools could still reach stored information. Data minimization matters more than a broad promise that the system is private.

Why the nonprofit setting matters

Nonprofit museums vary widely. A large institution may have curators, educators, digital teams, security personnel, and formal accessibility staff. A small historical society may rely on one director and volunteers. The concept only works if its obligations fit the institution using it.

A realistic deployment would avoid assuming that staff can continuously supervise conversations or rewrite hundreds of records. It would begin with one exhibit area, a limited set of approved questions, and a maintenance schedule the museum can actually sustain.

The cost case also needs restraint. Reusing Dave's hardware and local deployment pattern could reduce new engineering work, recurring cloud use, and dependence on visitor devices. It would not remove the costs of content preparation, staff review, accessibility testing, installation, support, security updates, or replacement hardware. For many museums, those human costs would exceed the cost of model inference.

A grant, sponsor, university partnership, shared regional service, or vendor-supported pilot could make the first deployment possible. The funding arrangement should not give a sponsor control over historical interpretation or turn the kiosk into an advertising display.

Existing work supports investigation, not certainty

Conversational museum guides are not a new idea. The MAX project placed an embodied conversational agent in a public computer museum in the early 2000s. More recent studies have explored screen-based, mobile, voice, virtual, and robotic guides. Research suggests that conversational systems can support engagement, personalization, navigation, and informal learning, while also revealing problems with comprehension, response quality, authenticity, and visitor expectations.

A 2025 live-exhibition study of a voice-based retrieval system reported that a museum chatbot could support informal learning and engagement, but it was a bounded deployment rather than proof that an unrestricted guide works across institutions. Recent work on museum chatbot roles likewise emphasizes that design choices affect how visitors understand and use the interaction.

Professional museum organizations are also treating AI as a governance issue. The American Alliance of Museums has emphasized privacy, bias, cultural sensitivity, staff oversight, data quality, and the need for institution-specific protocols. ICOM and UNESCO have begun broader work on ethical guidance and current museum practice. This reinforces the central premise of the concept: the hard part is not adding a chatbot to a screen. It is preserving institutional trust while making conversation useful.

A responsible first pilot

The first museum deployment should be a limited interpretive pilot, not a museum-wide replacement for labels, docents, or visitor services.

A practical pilot might include:

- One nonprofit museum and one contained exhibit area.
- Fifty to one hundred approved content records.
- A fixed kiosk, model, prompt, and knowledge-package version.
- Four to eight weeks of public testing.
- Visible disclosure that responses are generated from museum-approved material.
- A curator or content owner, operational owner, and technical support contact.
- Accessibility review and testing with representative users.
- A clear staff handoff path during published hours.
- Anonymous use by default and a short, documented analytics policy.
- Daily health checks and a simple way for staff to disable the conversational feature.

The evaluation should answer modest questions:

- Can visitors understand what the kiosk is and what it is not?

- Do visitors reach relevant approved information faster?
- Which exhibit questions are not covered by the museum's source material?
- Does the kiosk encourage exploration of physical exhibits rather than screen fixation?
- Are citations, uncertainty, and interpretive boundaries understandable?
- Can visitors with different access needs use the system?
- How often does the system provide an incorrect, unsupported, or misleading answer?
- How much staff time is required to maintain the content?
- Do assistance requests reach a real person when the kiosk says they will?
- Can the museum operate the system within its actual budget and staffing capacity?

Success should not be measured by the length of conversations. A better measure is whether the kiosk helps visitors notice, understand, locate, or discuss something in the museum that they otherwise might have missed.

Major risks and stopping conditions

The primary risks include invented facts, flattened or culturally insensitive interpretations, outdated collection data, inaccessible interaction, visitor overtrust, inappropriate historical roleplay, privacy loss, sponsor influence, staff burden, hardware failure, and a screen that competes with the objects it is supposed to support.

The pilot should stop or be materially redesigned if testing shows:

- Repeated unsupported claims presented as museum facts.
- Misrepresentation of people, communities, provenance, or contested history.
- Restricted or nonpublic collection information appearing in responses.
- Visitors regularly believing the kiosk is a human, curator, or independent historical authority.
- Accessibility barriers that exclude intended visitors.
- Conversation logging beyond the museum's disclosed purpose.
- Staff unable to review content or correct errors promptly.
- Assistance requests that imply a response no person is assigned to provide.
- The kiosk drawing attention away from exhibits without a meaningful interpretive benefit.
- Maintenance demands that the nonprofit cannot sustain after the pilot.

An unsuccessful pilot could still be useful. It may show that a simpler searchable kiosk, better labels, an audio guide, or more volunteer coverage solves the problem with less risk and cost.

Conclusion

Deckhand Dave's existing kiosk architecture could be adapted for nonprofit museums without pretending that a marine dealership and a cultural institution have the same needs. The reusable value lies in the local touchscreen platform, controlled content retrieval, managed deployment, and staff assistance workflow. The museum-specific work lies in interpretation, source governance, accessibility, privacy, and trust.

The strongest version of the concept would not try to become the museum. It would help visitors reach the museum's knowledge in a form that responds to their curiosity. It would cite approved material, show uncertainty, protect anonymity, and direct people back toward objects, spaces, and human expertise.

That makes the idea technically plausible and worth testing. It does not make deployment automatically appropriate. Each museum would need to decide whether conversation genuinely advances its mission and whether it has the staff, content, and governance to support the system after the novelty wears off.

Selected sources

- American Alliance of Museums. “Artificial Intelligence in Museums: Discussing Ethics and Protocols.” 2025. AAM ethics and protocols
- American Alliance of Museums. “Transforming Museum Workflows with Artificial Intelligence (AI).” 2024. AAM museum workflows
- American Alliance of Museums. “Beyond the IT Office: Solving for Ethical AI in Cultural Institutions.” 2025. AAM ethical AI
- Chen, B., et al. “The Role of Chatbots in Shaping Interactive Experiences.” *CHI Conference on Human Factors in Computing Systems*, 2025. ACM Digital Library
- Kopp, S., et al. “A Conversational Agent as Museum Guide: Design and Evaluation of a Real-World Application.” 2005. ACM Digital Library
- Kucia, F. J., et al. “How to Make Museums More Interactive? Case Study of Voice-to-Voice RAG-Powered Chatbot.” 2025. arXiv
- International Council of Museums. “Digital Tools for Cultural Protection and Promotion.” 2025. ICOM
- Trichopoulos, G., et al. “Crafting a Museum Guide Using ChatGPT4.” *Big Data and Cognitive Computing*, 2023. MDPI