

Staff-Maintainable Local Digital: Handoffs That Outlast a Clinic Semester

Independent research note (portfolio writeup; not peer-reviewed)

Jade Zhao

Informatics student, Luddy School (affiliation for contact only)

Indiana University Bloomington

jlvhao@iu.edu

Independent personal writeup; not IU-endorsed or lab output

Abstract

Local organizations often need digital tools their staff can maintain without a developer on call. Clinic sprints that optimize for the demo audience leave partners with undocumented ETL, inaccessible content chunks, and surprise dependencies when the semester ends. This independent writeup frames staff-maintainable local digital work as the operating requirement: accessible by default, honest about scope, and explicit about renewal ownership—including when generative AI tools enter community workflows over years rather than demo week. Approach: synthesize public-safe project summaries and digital-inclusion principles from the local-digital repository covering ServeIT Clinic builds for Bloomington-area partners (Accessibility Council, Middle Way House, Indiana Phenology, Bloomington Chamber civic reporting), plus Serve-AI constraints already stated in those notes (human review before publish; no partner PII into external models without consent and clinic policy). Findings are qualitative process patterns drawn from those notes, not partner performance metrics or invented outcome percentages. A dedicated discussion argues that long-horizon AI capability growth does not remove staff-capacity and accessibility bottlenecks; AI-touched systems still fail handoff if only a student can operate, disclose, or override them. Limits: documentation reflects clinic sprint scope as recorded in public project notes; partner-owned data and live URLs stay with handoffs; this PDF is a portfolio-style research note, not peer-reviewed research, not a longitudinal adoption study, and not official IU or clinic policy.

1 Introduction

A nonprofit website or civic dashboard that only a student developer can fix has already failed its handoff.

Situation. The concrete setting is clinic work through ServeIT at IU Luddy: community partners need sites, pipelines, and reporting surfaces their own staff can edit, refresh, and escalate after the semester ends and the student team leaves [4]. The author’s public role frame is Website Team Lead alongside related public-interest AI work on the Serve-AI (PIT-UN) track; this writeup indexes community projects and process notes, not clinic ownership claims for Serve-AI upstream materials.

Gap. Many student builds optimize for a working demo and a gradebook deadline. Accessibility becomes a pre-launch scramble. Data jobs live in a personal account. Documentation describes what the builder did, not what the next maintainer must do. Partners discover those gaps in production. WCAG 2.1 baselines and clinic accessibility guides exist [1, 2]; what this note adds is an operating frame that treats staff maintainability as the acceptance test, not a soft preference after the demo—and extends that frame to AI-touched workflows that partners may adopt slowly over years.

Claims. Stated as claims, not empirical results:

1. Local community digital work should be staff-maintainable first: accessible by default, honest about scope, and free of surprise dependencies.
2. That frame is already grounded in four public-safe ServeIT project summaries indexed in local-digital (not new field experiments).
3. As AI tools proliferate long-term, accessibility and staff capacity remain bottlenecks; capability growth is not automatic inclusion, and AI-touched clinic systems still need human review, disclosable scope, and override paths partners can operate.
4. Jade’s independent writeup and indexing notes remain separate from partner-owned content and from Serve-AI / ServeIT materials that carry their own licenses and credit rules.

Section 2 situates related practice; Section 3 describes the frame and projects; Section 4 reports qualitative patterns; Section 5 discusses long-term AI growth; Section 6 states limits; Section 7 concludes.

2 Related Work

Clinic accessibility baselines. ServeIT public extension guides target WCAG 2.1 Level AA for clinic builds, with Cascade WCMS and IU Rivet composition patterns [1, 2]. Those guides are the checklist home; local-digital indexes community projects and process notes that do not belong only in accessibility runbooks.

Civic tech and nonprofit handoffs. The project docs emphasize maintainability for communications and admin staff: documented ETL schedules, clear chart semantics, crisis paths above the fold, and chunk-based editing without developer intervention. That is adjacent to broader civic-tech conversations about sustainable public-interest software. This note stays inside documented clinic sprint scope rather than surveying the whole field.

Public-interest AI constraints. When clinic work touches generative AI (Serve-AI / PIT-UN track), the operating frame requires human review before publish, disclosure where appropriate, and no feeding partner PII into external models without consent and clinic policy sign-off [3]. Serve-AI's public page frames responsible adoption for community organizations: accessibility alongside automation. That constraint is part of digital inclusion here, not a separate ethics appendix.

3 Approach

What was analyzed: public project summaries and the digital-inclusion operating frame in the local-digital repository. What was not analyzed: private partner data, credentials, raw dumps, or participant records (explicitly kept out of the public repo).

Four documented clinic projects supply concrete context:

- City of Bloomington Accessibility Council (2024). First clinic client. Accessible page structure aligned with WCAG baselines; Cascade / Rivet composition; documentation for city staff maintainers. Framed as first proof that clinic builds could ship outside the gradebook: real residents, real maintainers, real handoff. City-owned content and deployment URLs stay with the City handoff.
- Middle Way House (2024–2025). Survivor-serving organization (domestic violence and sexual assault services in south-central Indiana). Accessible, responsive web surfaces; staff-maintainable content blocks; critical help paths kept obvious above the fold; security-conscious forms and external links; wireframes in Figma with Cascade/Rivet patterns where applicable. Live site content belongs to the partner.
- Indiana Phenology (Jan–May 2025). Citizen-science observation ingestion. Python ETL from CSV/XML into PostgreSQL; WCAG-aligned staff workflows.

Handoff bar: partner staff can run or schedule ingestion, spot-check failed rows, and pull standard reports without SSH access; failures must log clearly enough for a non-developer to escalate. Raw exports and production credentials stay with the partner. Organizational citizen-science data only—no participant records in the public repo.

- Bloomington Chamber civic reporting (Aug–Dec 2025). React / D3 dashboards with a Node.js back end for municipal service-request patterns and regional trends. Handoff bar: Chamber staff or designated partners can interpret dashboards and refresh data on a documented schedule; chart semantics and filters should read clearly without a developer in the room. Partner-owned data and live deployment URLs stay with the handoff; this note is not a Chamber product announcement.

3.1 Design Principles

Drawn from the repository's digital-inclusion frame:

- Staff-maintainable first. If only a student can publish a fix, the handoff failed. Prefer platforms and patterns the partner's communications or admin staff can edit (Cascade WCMS chunks, documented ETL scripts, clear README handoffs).
- Accessible by default. Keyboard paths, screen reader checks, mixed-language markup, and contrast are part of the build, not a separate audit phase.
- Honest about scope. Document what shipped, what is stubbed, and what needs a developer next semester. Partners and future clinic teams should not discover gaps in production.
- No surprise dependencies. Name APIs, accounts, and renewal owners in the handoff. Avoid hidden API keys, personal accounts, or one-off hosting the partner cannot renew.
- Community data stays community-scoped. Organizational data only in public docs; credentials and raw participant records stay out.
- Public-interest AI constraints. Human review before publish; disclose AI-assisted content where appropriate; do not feed partner PII into external models without explicit consent and clinic policy.

4 Findings

Observed patterns from the public project notes (qualitative; no invented metrics):

1. Handoff bar is the real acceptance test. Across projects, the stated success condition is staff operation (edit chunks, run or schedule ingestion, interpret dashboards), not merely a passing demo. The Accessibility Council sprint was framed as first proof that

clinic builds could ship outside the gradebook: real residents, real maintainers, real handoff. The same discipline—unglamorous QA, staff training, honest limitations—carries into later Chamber and Phenology sprints.

2. Accessibility and maintainability travel together. Cascade / Rivet chunk composition and WCAG baselines appear as first-class build requirements, with detailed runbooks pointed to serveit-accessibility rather than duplicated in local-digital. Accessibility is a requirement from the first paste, not a pre-launch scramble.
3. Data pipelines need human-readable failure modes. Phenology notes require clear logging so non-developers can escalate; silent ETL failure fails the inclusion frame. Staff should spot-check failed rows without SSH. That same instinct applies if future AI-assisted cleaning or labeling steps enter a pipeline: failures must remain legible to staff.
4. Crisis and civic surfaces need readable semantics under stress. Middle Way House notes keep critical help paths above the fold; Chamber notes require chart semantics and filters that read without a developer in the room. Under stress, opaque AI summaries without a reopen path would fail the same bar.
5. Public docs describe scope, not partner product launches. Each project note states that live URLs, brands, and partner-owned data remain with the handoff. This writeup inherits that boundary.

Evaluation here is documentary: consistency of the operating frame against the four summaries. No A/B tests, partner satisfaction percentages, or traffic claims appear in the source docs, so none are claimed here.

5 Long-term AI Growth

This section is reflective synthesis from existing public notes, not a longitudinal study of AI adoption and not a forecast of market size or capability timelines.

AI-touched systems still fail on handoff. Over years, community organizations will keep encountering generative tools—vendor defaults, clinic experiments, and small-business navigation of “AI features” attached to existing products. Serve-AI’s public framing treats that as responsible adoption work: accessibility alongside automation, not automation instead of access [3]. A more capable model does not remove the semester-end problem. If only a student can prompt, review, disclose, renew an API account, or reopen a case after an automated path fails, the handoff still failed.

Staff capacity and accessibility as bottlenecks. Growth of AI capability is not automatic inclusion. Keyboard paths, contrast, mixed-language markup, plain-language content, and staff-editable chunks remain con-

straints. Partners without a developer on call cannot absorb opaque model updates the way a product team can. The digital-inclusion frame already requires: human review before publish; disclosure where appropriate; no partner PII into external models without consent and clinic policy. Those rules matter more as tools proliferate, because the cost of a quiet data leak or an unreviewed generated page compounds across semesters.

Small orgs adopt over years, not demo week. Clinic project notes describe multi-month sprints with explicit handoff bars (edit chunks; schedule ingestion; interpret dashboards). That tempo is the realistic unit for public-interest AI navigation: documentation, renewal ownership, and override paths that survive staff turnover. Demo-week prototypes that impress a sponsor but leave undocumented prompts, personal API keys, or silent failure modes violate the same principles that make Phenology ETL and Chamber dashboards staff-maintainable.

What this does not claim. This note does not claim measured AI adoption rates among Bloomington partners, does not invent ROI or efficiency percentages, and does not assert that Serve-AI owns or endorses this PDF. It argues only that the staff-maintainable frame already in local-digital extends cleanly to long-horizon AI use: treat maintainability, accessibility, and human override as first-class requirements before automation scales.

6 Discussion and Limitations

Readers should not conclude that these clinic builds are comprehensively evaluated outcomes research, that Indiana University or ServeIT Clinic endorse this PDF as official research, or that partner organizations have approved this writeup as a product announcement. Limitations: single-author synthesis of already-public project notes; partner deployment details intentionally omitted; clinic sprint periods bound what “shipped” means; ServeIT / Serve-AI institutional materials remain under their own credit and license rules (CC BY-SA for Serve-AI-related guides). Visual consistency with IU Rivet on the personal site does not transfer ownership of this content to IU. ProHealth / lab clinical work is out of scope and not discussed. The long-term AI section is interpretive, not empirical longitudinal evidence.

7 Conclusion

Local digital work that outlasts a semester is a maintenance problem before it is a features problem—and that remains true as AI tools enter community workflows over years. This independent writeup restates a staff-maintainable operating frame, anchors it in four public ServeIT project summaries, extends the frame to long-horizon public-interest AI constraints already present in clinic notes, and keeps partner data and institutional licensing boundaries intact. The contribution is process

clarity for builders and future clinic teams, not peer-reviewed findings and not IU research output.

Acknowledgments

Independent personal writeup synthesizing public notes in local-digital. Drafting and structure were AI-assisted; accountability remains with Jade Zhao. Clinic context credits ServeIT Clinic (IU Luddy). Where accessibility guides derive from Serve-AI / ServeIT practice, those materials are CC BY-SA 4.0 (credit Serve-AI / ServeIT Clinic). This PDF does not claim sole ownership of Serve-AI upstream materials and is not official IU research.

References

- [1] W3C Web Accessibility Initiative, Web Content Accessibility Guidelines (WCAG) 2.1, W3C Recommendation, 5 June 2018. <https://www.w3.org/TR/WCAG21/>.
- [2] J. Zhao, ServeIT Accessibility (public clinic guides), GitHub repository zhao-langxi/serveit-accessibility. <https://github.com/zhao-langxi/serveit-accessibility>.
- [3] Serve-AI project page (ServeIT Clinic / PIT-UN track). <https://serveit.luddy.indiana.edu/serve-ai/index.html>.
- [4] ServeIT Clinic, Luddy School of Informatics, Computing, and Engineering, Indiana University. <https://serveit.luddy.indiana.edu/>.