AI-Augmented Skilled Trades: Preserving Critical-Infrastructure Knowledge as the Workforce Turns Over

A policy white paper Author: Brad M. Lindsey — independent engineer · Master Electrician · Master HVAC Technician **Version:** 1.0 — 2026-06-02 **Status:** Reviewed draft, ready for placement. The calibration banner below is binding on every figure.

> **Calibration banner (binding on every figure in this document).** Macro figures are public federal or public industry data, cited to a primary source. Case-study outcomes are metrics a pilot *would measure* — not results achieved. No physical or operational validation is claimed beyond what is published in the cited public record. The author writes as an independent practitioner; no employer, contract, or specific federal facility is named or implicated.

Executive summary

The United States is losing operational control of its own critical-infrastructure knowledge. Two public, bipartisan trend lines are crossing at the worst possible time. Federal building repair backlogs more than doubled to $370 billion between fiscal years 2017 and 2024 — enough that the Government Accountability Office added federal building condition to its High-Risk List in 2025.[1] At the same time, the workforce that actually keeps those systems running is retiring faster than it can be replaced: roughly 2.1 million skilled-trades positions could go unfilled by 2030, against a labor pipeline producing about one new worker for every four jobs posted.[2]

The dollar figure is the visible problem. The invisible one is worse. Every retirement and every contract turnover walks years of machine-specific, undocumented expertise out the door — the troubleshooting instinct that tells a veteran electrician which breaker to check first, why a particular generator has always hunted on transfer, what the building did the last time it flooded. That knowledge is rarely written down. When it leaves, the next operator starts from zero, and the deferred-maintenance backlog compounds because problems that used to take an hour now take a week.

In 2026 we can finally do something about this in the field. AI is now good enough to capture an expert's reasoning, turn it into a searchable record, and give the next operator day-one fluency with a system they have never touched. But there is a catch, and it is the whole point of this paper: AI cannot be trusted near safety-critical systems unless its outputs can be governed, traced, and checked. A confident wrong answer in a hospital electrical room is not a productivity problem. It is a safety event.

[1] U.S. Government Accountability Office, *Federal Real Property: Disposing of Unneeded Facilities Could Help Reduce Maintenance Backlog*, GAO-25-108400 (2025); and opening statement of GAO Director David Marroni before the House Appropriations Committee, April 9, 2025 ("Federal building repair backlogs more than doubled to $370 billion from FYs 2017 to 2024"). gao.gov/products/gao-25-108400.

[2] JLL, skilled-trades talent research (April 2026): ~2.1 million unfilled trades positions by 2030; ~600,000 jobs posted vs. ~150,000 new entrants last year; up to $1 trillion/yr potential economic loss (citing U.S. Department of Education). jll.com/en-us/newsroom/critical-skilled-trades-shortage-threatens-economic-losses.

This paper argues that **provenance-governed AI** — AI whose every claim carries an auditable trail back to its evidence — is the trustworthy way to preserve critical-infrastructure knowledge as the workforce turns over. It proposes a low-cost federal pilot to demonstrate it. The ask is not for an appropriation. It is for **permission to demonstrate**.

---

1 — The problem: assets are deteriorating while the people who maintain them retire

Two facts, both public, both auditable.

The assets are deteriorating, and the trend is accelerating. GAO reported that combined Department of Defense and federal civilian building deferred-maintenance and repair backlogs more than doubled — from $171 billion in FY2017 to $370 billion in FY2024.[3] That is a seven-year doubling, and it prompted GAO to add federal building condition to the "Managing Federal Real Property" area of its High-Risk List in 2025.[4] The government's annual operating and maintenance cost for its 277,000 buildings exceeded $10.3 billion in FY2023.[5] The General Services Administration's own deferred-maintenance backlog, reported above $17 billion in March 2025, was assessed at roughly $50 billion a year later — more than double the agency's previous highest estimate.[6] Unless the trend reverses, GAO warns, federal assets will deteriorate to the point of premature replacement, which costs far more than maintenance done on schedule.[7]

The workforce that holds the operational knowledge is shrinking and aging. JLL's 2026 skilled-trades research projects roughly 2.1 million unfilled trades positions by 2030 — electricians, HVAC technicians, plumbers, pipefitters, equipment operators, general maintenance.[8] The pipeline is not keeping up: last year, employers posted nearly 600,000 jobs across the major trades while only about 150,000 new

[3] U.S. Government Accountability Office, *Federal Real Property: Disposing of Unneeded Facilities Could Help Reduce Maintenance Backlog*, GAO-25-108400 (2025); and opening statement of GAO Director David Marroni before the House Appropriations Committee, April 9, 2025 ("Federal building repair backlogs more than doubled to $370 billion from FYs 2017 to 2024"). gao.gov/products/gao-25-108400.

[4] U.S. GAO, *High-Risk Series*, GAO-25-107743 (2025) — federal building condition added under "Managing Federal Real Property."

[5] U.S. Government Accountability Office, *Federal Real Property: Disposing of Unneeded Facilities Could Help Reduce Maintenance Backlog*, GAO-25-108400 (2025); and opening statement of GAO Director David Marroni before the House Appropriations Committee, April 9, 2025 ("Federal building repair backlogs more than doubled to $370 billion from FYs 2017 to 2024"). gao.gov/products/gao-25-108400.

[6] Public Buildings Reform Board, *The Cost of Inaction: Deferred Maintenance in GSA's Portfolio* (March 5, 2026): GSA backlog assessed at ~$50 billion, more than double the agency's previous highest estimate; GSA's own March 2025 figure exceeded $17 billion. pbrb.gov.

[7] U.S. Government Accountability Office, *Federal Real Property: Disposing of Unneeded Facilities Could Help Reduce Maintenance Backlog*, GAO-25-108400 (2025); and opening statement of GAO Director David Marroni before the House Appropriations Committee, April 9, 2025 ("Federal building repair backlogs more than doubled to $370 billion from FYs 2017 to 2024"). gao.gov/products/gao-25-108400.

[8] JLL, skilled-trades talent research (April 2026): ~2.1 million unfilled trades positions by 2030; ~600,000 jobs posted vs. ~150,000 new entrants last year; up to $1 trillion/yr potential economic loss (citing U.S. Department of Education). jll.com/en-us/newsroom/critical-skilled-trades-shortage-threatens-economic-losses.

workers entered through apprenticeship programs.[9] Demand is not cyclical — it is structural and growing. JLL's research, drawing on federal labor projections, puts electrician employment growth at 9.5% through 2034 — about three times the all-occupation average — and HVAC technicians at 8.1%; the Bureau of Labor Statistics' own Occupational Outlook rounds these to roughly 9% and 8%.[10] JLL, citing U.S. Department of Education data, puts the potential economic loss from the shortage at up to $1 trillion a year.[11]

Put the two facts together. The buildings need more maintenance than ever, and the people who know how to maintain them are walking out the door faster than replacements walk in. The gap between those lines is where infrastructure fails.

---

2 — The invisible loss: knowledge that retires with the worker

The backlog is what you can count. It is not what hurts most.

What hurts most is tacit knowledge — the expertise that lives in a person's hands and judgment and was never written down. A veteran operator recognizes subtle patterns in how equipment behaves, remembers the history of every repair decision, and carries informal troubleshooting techniques that exist in no manual.[12] In a water-treatment plant, a power plant, or a hospital electrical room, that knowledge is the difference between a thirty-minute fix and a multi-day outage. It is also, almost by definition, undocumented — which means it cannot be handed off on a checklist.

The exposure is concentrated and immediate. Across the utility sector, reporting indicates up to roughly half the workforce may be retirement-eligible within the decade, with more than half already over 45.[13] Industry has tried for years to put a dollar figure on knowledge attrition at large organizations; the estimates are large and the trend is rising, but they are softer than the federal data above and belong here only as context. Every federal contract turnover adds a second, structural version of the

[9] JLL, skilled-trades talent research (April 2026): ~2.1 million unfilled trades positions by 2030; ~600,000 jobs posted vs. ~150,000 new entrants last year; up to $1 trillion/yr potential economic loss (citing U.S. Department of Education). jll.com/en-us/newsroom/critical-skilled-trades-shortage-threatens-economic-losses.

[10] Precise figures (electricians +9.5%, HVAC +8.1% through 2034, ~3× the all-occupation average) are from JLL's April 2026 skilled-trades research. The U.S. Bureau of Labor Statistics Occupational Outlook Handbook gives the rounded values (electricians ~9%, HVAC ~8%, 2024–2034); cite the OOH pages directly in final. bls.gov/ooh.

[11] JLL, skilled-trades talent research (April 2026): ~2.1 million unfilled trades positions by 2030; ~600,000 jobs posted vs. ~150,000 new entrants last year; up to $1 trillion/yr potential economic loss (citing U.S. Department of Education). jll.com/en-us/newsroom/critical-skilled-trades-shortage-threatens-economic-losses.

[12] Utility-workforce reporting on tacit-knowledge loss and retirement exposure (e.g., TD World; Power Engineering; ElectricEnergyOnline). Secondary industry reporting, broadly corroborated across sources; framed here as "reporting indicates." Confirm against a primary utility-association or DOE/EIA workforce study before any placement that leans on the figure.

[13] Utility-workforce reporting on tacit-knowledge loss and retirement exposure (e.g., TD World; Power Engineering; ElectricEnergyOnline). Secondary industry reporting, broadly corroborated across sources; framed here as "reporting indicates." Confirm against a primary utility-association or DOE/EIA workforce study before any placement that leans on the figure.

same loss: when a facilities or engineering contract goes out for rebid and changes hands, the incoming team rarely inherits the outgoing team's accumulated, system-specific understanding. The building stays; the knowledge of the building leaves.

I have watched this happen from the inside.

A representative case (de-identified)

I have spent more than twenty years as a Master Electrician on a large federal healthcare campus — NEC 517 essential electrical systems, isolated power, generator and automatic-transfer-switch coordination, UPS commissioning, switchgear, building-automation and variable-frequency drives, live work in occupied clinical spaces. Over that time the maintenance contract was competitively rebid and changed prime contractors multiple times. The facility never moved. The equipment never moved. But with each turnover, the institutional memory of *why the systems behave the way they do* had to be rebuilt, largely from scratch, by whoever stayed.

Some of what makes a campus like that run is not in any drawing set. It is the knowledge that a particular feeder was re-pulled after a fault years ago and now reads differently than the as-builts say; that one generator has a known behavior on transfer that is harmless if you expect it and alarming if you don't; that a given panel was relabeled during a renovation and the paperwork never caught up. None of that is exotic. All of it is the kind of thing that, when the person who knows it retires or rotates off, the next operator learns again the hard way — on an occupied clinical floor, under time pressure, with the stakes that a healthcare facility carries.

This is the representative case, not a unique one. It is the daily reality of federal facilities maintenance across the inventory. The dollar backlog is the symptom. The knowledge that never got captured is the disease.

---

3 — Why now: AI can finally do this work in the field

For most of the last two decades, the proposed answer to knowledge loss was "better documentation" — capture programs, succession binders, exit interviews. They help at the margin, and they consistently fail at the core, because the most valuable knowledge is exactly the knowledge that resists being written into a procedure. You cannot checklist judgment.

In 2026, that calculus changes. Large language models are now capable enough to do three things that documentation never could:

1. **Capture reasoning, not just facts.** An AI assistant can sit alongside an expert troubleshooting a live problem and record not only what was done but *why* — the discarded hypotheses, the tell that pointed to the real fault, the order of operations that experience dictates.
2. **Make experience searchable.** Decades of captured troubleshooting becomes a knowledge base an operator can query in plain language: *this AHU is short-cycling and the building is calling for cooling — what has caused this here before?*

3. **Give the next operator day-one fluency.** A new technician, or an incoming contractor's crew, can stand in front of an unfamiliar system and ask it questions, instead of waiting years to accumulate the instinct the last crew took with them.

This is force-multiplication, not replacement. The expert still does the work; the AI captures it and hands it forward. The new operator still turns the wrench; the AI shortens the distance between *never seen this* and *competent*. No model is being proposed as a substitute for a licensed tradesperson near energized equipment.

But capability is not the hard part anymore. Trust is.

---

4 — The trust problem: a confident wrong answer near a critical system is a safety event

Here is the question a federal facilities engineer will ask before letting any AI near a hospital electrical room, and it is the right question: **how do I know this thing isn't confidently wrong?**

Anyone who has worked with today's AI knows the failure mode. The same system that produces a correct troubleshooting sequence will, with identical fluency and confidence, produce an incorrect one. It will describe a simulated result as if it were measured. It will lose track of why it recommended something two steps ago. Left ungoverned, an AI knowledge system drifts toward whatever sounds plausible, its claims inflate beyond what the evidence supports, and the provenance of any given answer evaporates. In an office setting that is an annoyance. In NEC 517 territory — essential electrical systems in an occupied healthcare facility — a confident wrong answer is a safety event waiting for someone to trust it.

This is the reason "just deploy AI" is not a serious infrastructure proposal. Any honest plan to put AI near critical systems has to answer the trust question first, in a way an auditor and a safety officer can both verify. Most proposals don't. They lead with capability and treat trust as a disclaimer.

The answer to the trust question is the actual content of this proposal.

---

5 — The answer: provenance-governed AI

The trustworthy way to put AI near critical systems is to govern it the way a trade governs its own work — labeled, tested, and traceable back to the day it was made. What's behind the wall has to actually do what the paperwork says it does. The same standard can be enforced on an AI's outputs.

The approach is a **provenance discipline**: a set of governance mechanisms that bind every AI-produced claim to an auditable trail of evidence, so that the system cannot silently drift, cannot inflate a claim beyond its support, and cannot lose the record of *why* it said what it said. Concretely, three properties matter for a federal deployment:

- **Traceability.** Every answer the system gives can be traced to the specific captured expertise and evidence it rests on. An operator — or an inspector — can ask not just *what* but *on what basis*, and get a real answer.
- **Auditability.** The record is append-only and cryptographically anchored. What the system knew, and when, cannot be quietly rewritten after the fact. This is the same integrity property a chain of custody provides, applied to machine knowledge.
- **Calibration.** The system is built to distinguish what it has evidence for from what it is guessing, and to say so. A claim may advance only as far as its evidence supports it.

This is not a product pitch and it is not theoretical. The underlying methodology is published and open. The author has released an open-source methodology framework for governed, traceable AI-collaborative engineering — the *Lindsey Provenance Discipline*, available as the Python package `lindsey-provenance` under an MIT license — together with two companion papers deposited as citable public records: a practitioner experience report and a responsible-AI self-audit.[14][15][16] The self-audit is worth naming here, because it is the proof that the discipline is applied to its own author and not just preached: it documents a formal pass in which forty-nine of the author's own internal intellectual-property claims were checked against prior art and *zero* survived as novel — a result published rather than buried. A framework that holds its own creator to that standard is the kind that can be trusted to flag its own uncertainty in a hospital electrical room.

The point is not the specific framework. The point is that **governed AI is buildable today, with public evidence**, and that governance — not raw capability — is what makes AI safe to deploy near critical infrastructure. The capability is commodity. The discipline is the moat.

---

6 — The proposal: a low-cost federal pilot

The responsible way to establish whether this works in a federal setting is not a procurement and not a mandate. It is a small, measured demonstration with honest success criteria defined in advance.

Objective. Demonstrate that a provenance-governed AI system can (a) capture expert troubleshooting workflows from veteran tradespeople, (b) convert that field experience into a searchable, auditable knowledge base, and (c) measurably shorten the time for a less-experienced operator to diagnose and resolve real faults — without ever presenting an ungoverned or unverifiable claim to a user.

Candidate site types (generic — no specific facility implied): a military or veterans'

[14] B. M. Lindsey, *Lindsey Provenance Discipline* — open-source methodology framework; Python package `lindsey-provenance` (MIT). github.com/bradmlindsey/lindsey-provenance.
[15] B. M. Lindsey, *One Operator, Nine Trunks, Seven Weeks* — practitioner experience report (Zenodo deposit; DOI per RELEASE_SEQUENCE).
[16] B. M. Lindsey, *Zero of Forty-Nine* — responsible-AI self-audit (Zenodo deposit; DOI per RELEASE_SEQUENCE).

healthcare facility, a federal office building under GSA, or a federal civilian campus with an aging maintenance workforce and a documented backlog.

What the pilot would measure (these are the metrics, defined as targets — not results):

- Mean time to diagnose a defined set of common faults, expert vs. assisted-novice.
- Equipment downtime on covered systems, baseline vs. pilot period.
- Training time to competency for a new operator on a covered system.
- Knowledge-capture coverage: share of a retiring expert's troubleshooting workflows captured in auditable form before departure.
- Governance integrity: rate at which the system correctly flags low-confidence or unverifiable answers rather than asserting them.

Governance and guardrails, built in from day one. Every captured workflow and every answer carries its provenance trail. The deployment de-identifies sensitive facility and system specifics. No AI output is presented to an operator as authoritative without its evidence trail attached. The system is positioned as an assistant to licensed personnel, never a replacement, and never in the loop on live energized work.

Cost posture. The methodology is open-source and the supporting toolchain is lightweight. This is a demonstration measured in a single facility and a small set of systems, not an enterprise rollout. The cost to the government is the access and the measurement — not a development contract.

---

7 — The ask

This paper does not ask Congress for money. It asks for **permission to demonstrate** — a sponsored, low-cost pilot at a willing federal facility, with the metrics above measured honestly and the results published whichever way they come out. If the demonstration succeeds, the government has a low-cost, auditable way to preserve critical-infrastructure knowledge as its workforce turns over. If it does not, the measurement itself is a public good: it tells us what governed AI can and cannot yet do near critical systems, which is worth knowing before anyone deploys at scale.

The macro case is bipartisan and already on the record — a $370 billion and growing backlog, a 2.1 million-worker shortfall by 2030, and a GAO High-Risk designation. The capability to capture expert knowledge in the field now exists. The one thing standing between the problem and the tool is trust, and trust is precisely what a provenance discipline is built to earn. A small pilot is the honest next step.

I am one Master Electrician who has watched two decades of institutional knowledge cycle through a federal facility across multiple successive maintenance contracts, and who then built and published the discipline this proposal rests on. I am asking for the chance to show it works where it matters.

Notes and references

Author's note on sourcing: the figures below are public federal and public industry data. Before any external placement, each starred figure should be re-confirmed in its primary PDF; see the companion `FIGURES_VERIFIED_2026-06-02.md` for the verification log.

Author's disclosure. This document, and the methodology it advocates, are the product of LLM-collaborative work conducted under that same provenance discipline. The point of the discipline is to keep that collaboration honest and auditable; this paper was written to its own standard.

Calibration attestation. Every claim here is held to the proof-state on record as of June 2, 2026. Macro figures are cited to public primary sources and re-confirmable in the companion verification log; case-study outcomes are framed as pilot targets, not results. No fabricated figures, vendor-supplied targets, or unvalidated outcome claims appear.

Outreach — to a committee or member office, an agency modernization office, or a trades association, via an RFI or public-comment channel — is the author's relationship move, not part of this document.