

The Verification Gap

An Independent Audit of the U.S. Military Transition Ecosystem

Laura Dryden — MissionProof Research, Tucson, Arizona · laura@mymissionproof.com

September 2026 — First Edition

Published by MissionProof Research, Tucson, Arizona, an imprint of MissionProof LLC. © 2026 MissionProof LLC. This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License (CC BY-NC-ND 4.0, <https://creativecommons.org/licenses/by-nc-nd/4.0/>): it may be copied and redistributed in any medium or format, in full and with attribution, for non-commercial purposes, but not altered, excerpted, or built upon. Suggested citation: Dryden, L. (2026). *The Verification Gap: An Independent Audit of the U.S. Military Transition Ecosystem*. MissionProof Research, September 2026, First Edition. DOI: <https://doi.org/10.2139/ssrn.7413798>. Also available at <https://ssrn.com/abstract=7413798>.

Every figure in this paper traces to a committed source file or a cited public record. No SkillBridge host or audited employer is named: the unit of analysis is program and instrument design, not the firms lawfully participating in it; federal contractors appear only as named parties to public contract records, and specimens are described structurally throughout. The correction log at Appendix D is part of the argument — this document practices the correction transparency it recommends of others.

About the Author

Laura Dryden is the founder of MissionProof Research in Tucson, Arizona. She is a retired Air Force senior noncommissioned officer who served in the intelligence career field from 2001 to 2023. As course director of the All Source Intelligence Course — the schoolhouse for the 1N0 specialty that this audit uses as a worked specimen — she was responsible for the design and delivery of enlisted technical training under the Air Force’s instructional systems development process, an ADDIE-type model. That background is what makes the audit’s method possible: a reader who has written training standards can check a crosswalk at the level of the code, read a Career Field Education and Training Plan against the occupation it is mapped to, and tell a mapping that was reviewed from one that was carried forward. She holds a Master of Business Administration and the Project Management Professional certification, and is a doctoral candidate in business administration at National University.

Her dissertation, a meta-content alignment study of Air Force enlisted professional military education against the Project Management Competency Development framework, required assembling, staging, and coding the enlisted training and education corpus document by document. The corpus and the verification habits that work

demanded are the foundation of this audit; the question it asks — whether the instruments members are told to trust have ever been checked — grew out of finding, document by document, how rarely they had been.

This audit is published by MissionProof Research, which the author owns. MissionProof also builds a career-pathway platform for Air Force enlisted members, which she wrote as its sole developer, and readers are entitled to know that the publisher of an audit of the transition ecosystem works in it. The audit is built so that the fact does not have to be taken on trust: every figure recomputes from a committed public source, no host or audited employer is named, no product is recommended, and the correction log at Appendix D records every change made to a printed number.

Abstract

Each year the United States returns between 148,315 and roughly 200,000 service members to civilian life — the federal sources disagree by 26% and none reconciles them — through an ecosystem Congress has legislated and funded piece by piece: a mandated occupational crosswalk to translate military skills, a credentialing apparatus, an internship marketplace (SkillBridge), and a mandatory transition curriculum (the Transition Assistance Program). The Government Accountability Office has reported on this ecosystem at least ten times in nine years and has verified almost none of it. This audit does, using only public data, the government’s own binding instructions, and receipt-grade evidence: every figure traces to a committed source file or a cited public record, and the correction log is published with the paper. The crosswalk members are told to trust was last quality-reviewed in 2014; of 48,042 Air Force specialty-to-occupation links, 0.51% have ever been human-reviewed, and a systematic comparison of every specialty’s published training standard against every civilian occupation’s task statements finds the median offered occupation ranks 75th of 922 on task similarity, with 42.6% outside the top 100 and 90.9% of offers identical across skill levels. The SkillBridge marketplace advertises 11,532 positions behind which no enumeration exists; its own export instrument cannot produce one. Since August 2024 every provider warrants a 75% qualifying-offer rate that no federal dataset can verify, while members on full military pay supply an estimated \$316–732 million per year of labor, for up to 180 days each, to hosts who pay nothing. Access collapses by rank and branch even as it goes unverified, and the market follows employer demand rather than member need. The audit closes with the Freedom of Information Act requests, correction requests, and statutory changes that would make each layer verifiable.

Keywords: military transition, veterans, SkillBridge, O*NET crosswalk, Transition Assistance Program, credentialing, program verification, labor market information, public-sector oversight

Executive Summary

Every year the United States military returns a population of service members to civilian life that its own federal sources count differently — roughly 200,000 by the

commonly cited Congressional Research Service (CRS) figure, 148,315 active-duty separations in 2024 in the Defense Manpower Data Center's (DMDC) own transaction files, a 26% gap no located public source reconciles (Appendix D, correction 26). Nothing in this paper is calculated from either number. The government has built an ecosystem to carry those members across: a mandated occupational crosswalk to translate their skills, a credentialing apparatus to certify them, an internship marketplace (SkillBridge) to place them, and a mandatory transition curriculum (the Transition Assistance Program, TAP) to prepare them. Congress has legislated each piece, appropriated for each piece, and — as the Government Accountability Office (GAO) has now said in at least ten reports across nine years — verified almost none of it.

This paper documents what independent verification finds. Between July and August 2026, MissionProof audited each layer of the ecosystem using only public data, the government's own binding instructions, and receipt-grade evidence: every claim below traces to a source file, a contract record, a directory export, or a federal publication.

The crosswalk members are told to trust was last quality-reviewed in 2014. Of 48,042 Air Force specialty-to-occupation links, 0.51% have ever been human-reviewed. A \$49,996 government pilot showed in 2015 that an enhanced method returned roughly nine times as many civilian matches per military code across the 68 occupations it covered — a measure of coverage, not of accuracy, and one no subsequent study has tested. The comparison nobody had run has now been run. Scoring every specialty's published training standard against every published civilian occupation's task statements — 1,055 crosswalk-offered pairs across the 60 Air Force specialties whose task inventories could be extracted — the median occupation the crosswalk offers ranks 75th of 922 on task-text similarity to the specialty's own standard, 42.6% fall outside the top 100, and the layers a member actually browses perform worst of the three that could be scored. 90.9% of those offers are made identically to every skill level, while the specialty's own printed proficiency codes move an occupation a median of 73.5 rank places between levels. Its recommendations were never implemented, coverage was never extended, and GAO's recommendation to measure whether any of it works was formally refused by the Department (GAO-17-133: Closed — Not Implemented).

The SkillBridge marketplace advertises 11,532 positions; no located enumeration of what stands behind them exists — and the Department's own instrument cannot produce one. The site's full-export endpoint fails; the only working method is per-filter; and a complete 57-state sweep is structurally incapable of returning listings whose location is blank or "Nationwide" — 41% of a single-filter export, weighted toward exactly the virtual and portable programs a remote-stationed member most needs. Merging 87 exports yields the largest assembly of the directory yet built: 7,826 rows, 4,255 programs, 2,732 organizations, of which 63% was invisible to the state sweep. On that corpus, 12.1% of programs name a military occupation code and 49.8% say nothing at all about targeting; 1.4% disclose what the destination job pays; 86.7% say nothing about whether employment follows. In a 15-employer verification pilot, five hosts — one-third of the sample — had no discoverable paid position behind the role their internship trains, including one provider training members for a certified assessor role its own website indicates it cannot employ. At union grain 5.0% of programs are flagged as training into commission-only

or business-ownership models in which no wage exists at all — an indicator-list size rather than a validated prevalence; 17% (a census-grain figure, IV.5) do not describe, in any recognizable form, what the intern will do.

These are not merely design flaws — since August 2024 they are contract-compliance questions. Every provider signs a Memorandum of Understanding (MOU) warranting that “75% or higher” of completers receive “a qualifying offer of immediate post-service suitable employment,” that providers give members “neither compensation nor gifts,” that members incur no costs including licensure, and that providers report 90- and 180-day employment outcomes, including starting salary, to the Department. No sourced publication of that outcome data has been located. The Department’s FY2026 budget justification nonetheless asserts an “85% job-offer rate” — a figure the GAO’s own findings suggest the services do not collect. The recommendation deadlines the Department of Defense (DoD) set for itself fall on September 30 and December 31, 2026.

The scale of what goes unverified is material. Members on full military pay work up to 180 days for hosts who pay nothing. At published participation levels the member labor involved is an estimated \$316–732 million per year (central estimate ~\$506M) — against no verified obligation that a job exists at the end. One host’s own job board prices the equivalent training at four to six weeks of paid wages; the program permits collecting twenty-six weeks free.

Access to the program is collapsing by rank and by branch even as it goes unverified. Applying each service’s own duration caps to every listed program: a junior sailor can verify fit with 66.7% of the union market; a senior Army or Air Force noncommissioned officer (NCO) with 278 programs — 6.5%. Two members of identical rank in different branches face a roughly sixfold difference in accessible market (5.9×). The caps also bind unevenly by population: one Air Force separator in five is a senior NCO (E7–E9) — twice the Army’s share, over three times the Marine Corps’ — so the 60-day senior tier reaches the most senior-heavy separating population in the force (Census Bureau Veteran Employment Outcomes, or VEO, paygrade file, 2020–2021 enlisted cohort). A third of all listings state no usable duration at all (“Under Review”), making fit unknowable at any rank.

The market serves employer demand, not member need or member value. Joining Census Bureau VEO data to the directory: the deepest program coverage clusters where employers need trade labor. The occupations with the highest civilian value (explosive ordnance disposal, intelligence, air traffic control) and the most vulnerable populations (infantry — the largest single veteran group, landing at a \$33,500 first-year median; medical technicians at \$36,200 with the force’s worst employment attachment) face a near-empty market. The credentialing layer compounds it: for the geospatial intelligence (GEOINT) analyst specialty, roughly one in three of the credentials Air Force Credentialing Opportunities On-Line (AF COOL) lists cannot be purchased outside the cleared workforce — issuer-verified, Part I. The Air Force’s FY2026 retention-bonus list dropped roughly a dozen medical specialties — making medics more releasable — into a marketplace that names them in seven programs of 4,255 (IV.7; correction 37).

The demand side mirrors the supply side. A full-book audit of a top-20 defense contractor’s public requisitions — an employer branding itself a “top veteran employer” roughly 379 times — found zero references to any military occupation code or to SkillBridge across 384 requisitions, 36.5% of them gated on a security

clearance and a further 8% on a Public Trust determination, 24% degree-walled with no experience pathway, and pay disclosed only where law compels (28%). A market-anchor corpus of 629 posting rows (30 July 2026 build; later expanded, V.2), built to an evidence floor of three-plus employers per data point, shows the honest stay-vs-go answer is lane-dependent — and for some fields the market never catches military compensation before twenty years. Employers and government each assume the other performs the translation; neither does.

The maintenance layer of this ecosystem shows a recurring procurement pattern. The SkillBridge website ran through July 2024 under a \$3.4M task order to Solutions for Information Design — the same small firm that built the COOL credentialing platform, operates MilGears, and authored the shelved 2015 pilot study — awarded with one offer received. Its successor, a \$15.2M (ceiling \$27.5M) contract to Goldbelt Apex covering the transition office’s IT systems, is an 8(a) sole-source award: one offer, not available for competition. Both structures are lawful and routine. Neither has ever faced a competitive bid, and the system holding the outcome data providers owe the government has published nothing.

Graded against the government’s own binding language, the system fails its own instructions. Parsing the governing issuances yields 640 binding (“will/shall/must”) requirements from the six transition instruments, plus the Department’s information-quality instruction graded in VII.2. Of the seventeen requirements gradeable today from public evidence — fifteen from the transition instructions, two from the information-quality instruction — none grades MET. The 365-day transition-start requirement — a statutory “must” — is met by roughly half the force. DoD Instruction 1332.35 promises a “single-source data collection framework... will be instituted”; GAO found the services manually pulling data from scattered sources. And the widely cited 90-day capstone deadline turns out to be written as “should” — advisory — in the implementing instruction: the instruction softened the statute.

The pattern, demonstrated now across four programs and twenty years of paper trail, is uniform: **mandated → funded → never verified**. The final sections of this paper set out eight fixes — most of them form-validation changes requiring no legislation — an “honest host” standard that costs good employers nothing, the Freedom of Information Act (FOIA) register (eleven requests drafted against named custodians, not yet filed), and the congressional engagement map for the one moment on the calendar when someone must answer: the open GAO deadlines of late 2026.

Methodology and Evidence Standards

Sources. This audit uses exclusively public materials: the SkillBridge public directory’s state-anchored export (the site’s own per-state export API, all 57 U.S. jurisdictions, 3 August 2026 — complete with respect to state-anchored listings and structurally blind to the rest, IV.9); and an 87-export union corpus merging that sweep with per-service, program-type, and job-family partitions, 8 August 2026, built after the per-state method was found to have a measurable structural blind spot (IV.9); federal contract records (USAspending, primary award exports); statutes, Department of Defense (DoD) Instructions, service regulations, and the SkillBridge Memorandum

of Understanding (MOU) in three vintages — the August 2024 blank release, a 2025 executed federal instance, and a 2026 signed instance (correction 40); Government Accountability Office (GAO), RAND, Congressional Research Service (CRS), Department of Labor (DOL), Department of Veterans Affairs (VA), and Census Bureau publications; employers’ own public careers pages and job boards, captured with URL and date; and budget justification documents.

Standards. Five rules govern every finding. (1) *Never outrun the receipts*: a claim not traceable to a committed source file is not made. (2) *Structure, not villainy*: this audit documents incentive structures and verification gaps, and alleges no misconduct by any contractor or employer; procurement structures described (single-bid fair opportunity, 8(a) sole source) are lawful and routine. (3) *Finite-search language*: “no X was located” describes the corpus searched, never proof of absence. (4) *Breach indicators, not breach findings*: a listing inconsistent with an MOU clause is reported as an inconsistency in the published record; providers may hold approvals their listings omit — the absence of any public field in which compliance could be demonstrated is itself the finding. (5) *Correction transparency*: errors found in our own work are retracted in place and logged (Appendix D).

Verification method for employer claims. The conversion pilot checked, for each sampled host, whether a paid opening matching the trained role was publicly discoverable on a stated date, recording the URL of every source consulted. A “no” means *no paid analogue verifiable on that date at those sources* — a repeatable, falsifiable statement.

Scope and generalizability. The directory, coverage, targeting-corpus, pay, conversion, MOU, window-cap, Transition Assistance Program (TAP), compliance-corpus, and Veteran Employment Outcomes (VEO) analyses are all-service: the directory covers every branch (4,154 of 4,255 union programs — 97.6% — accept all six; the census-grain figure was 1,497 of 1,527), the parsed instruction corpus includes Army, Navy, Marine Corps, and Coast Guard issuances, and the earnings join covers five branches. The educational-attainment analysis is all-service too: a complete five-service by seven-band matrix, read at full category detail from the Department’s own 2024 demographics captures, every row validated against the separately printed summary figure the same capture carries. Three deep dives use the Air Force as the worked layer — the crosswalk rebuild, the career-field dark-zone analysis, and the retention-collision lens — because that is where the author’s subject-matter expertise permits code-level verification. Every method is service-portable; Army and Navy replications remain to be run, and section IV.8 reports what the all-service data already shows.

Use of artificial intelligence, and independent verification. This audit was conducted with substantial use of artificial-intelligence systems, disclosed here because a paper that asks federal instruments to declare their methods must declare its own. Large language models were used for document analysis, data processing and classification, the mechanical extraction of modal statements from the governing issuances, and drafting. Two systems from different vendors were used deliberately rather than one: Anthropic’s Claude for analysis, computation and drafting, and an independently operated OpenAI ChatGPT instance as an adversarial reviewer tasked with re-computing findings, testing whether figures reproduce from the source files, and locating errors. The independent adversarial review recorded in Appendix D is that second system’s work, and it is not decorative: several of the retractions logged

there originated from it, including the withdrawal of a two-railroad pay-disclosure comparison and the downgrading of four classifier outputs from validated prevalence estimates to labeled indicator counts. Three safeguards bound the method. Every figure traces to a committed source file or a cited public record, not to model output; the analysis and classifier scripts are committed and re-executable by a third party; and where a construct failed to reproduce, the failure is reported rather than the figure quietly restated. The author directed the work, made every interpretive and editorial judgment, and is responsible for the result. No finding in this paper rests on a model's assertion that something is true.

What this audit cannot see. Provider outcome reports submitted to the Military-Civilian Transition Office; command-level denial decisions; participant-level data. Freedom of Information Act (FOIA) requests targeting each are enumerated in Appendix C.

Original Contributions of This Audit

For the record, and distinguishing confirmation from discovery. Findings that confirm or extend prior work (GAO's data-gap findings; DOL's no-wage-benefit result; RAND's effectiveness verdict; the TAP compliance record) are cited to their sources throughout. The following, to the best of a documented finite search, appear here for the first time anywhere:

1. **The largest committed deduplication of the SkillBridge directory located in this audit's finite search, and a measurement of what the directory's own export path cannot reach** — 11,532 advertised positions resolved against two independently built corpora (1,103 organizations state-anchored; 2,732 across 87 merged exports); the 150-result search cap; the failing full-export endpoint; and the 41% structural blind spot in the per-state method, established definitionally rather than by estimate.
2. **A measurement of the directory's information content (none earlier was located):** 12% occupation targeting, 1.4% pay disclosure, 86.7% employment silence (union grain, re-run classifiers), the 17% opacity tier, the 5.0% commission tier (union grain; 7.6% at census grain, with identified false positives), 47 over-ceiling listings.
3. **The first live-requisition verification of SkillBridge hosts** — the "is there a job behind the door" method, its five predictive patterns including the paid-twin test, and its documented specimens (a provider training members for a role it structurally cannot employ; the same program listed five times across five duration buckets; unpaid staff roles at a five-person firm).
4. **The first grading of the public directory against the August 2024 MOU's contractual warranties** — and the observation that the directory carries no field in which compliance could be demonstrated.
5. **A window-weighted market computation** — 278 programs (6.5%) verifiable for a senior Army/Air Force noncommissioned officer (NCO) at union grain; the 5.9× same-rank cross-branch access gap.
6. **The first quantification of the unpaid member labor at stake** — \$316-732M per year delivered to hosts against no verified job obligation, and the per-member \$26,300 unit figure.

7. **The Fair Labor Standards Act (FLSA) shadow test** — primary-beneficiary risk scoring of an entire federal internship marketplace.
8. **The retention-market collision** — the FY26 Selective Retention Bonus (SRB) cross-reference showing 18 of 24 retained specialties with two or fewer programs, the special-warfare fields named by occupation code in six programs of 4,255, exactly one of them inside the defense industrial base (correction 38), and the medical release-into-void chain (SRB drop → seven programs → \$36.2K landing; correction 37), completed with the Census VEO join showing the market allocates by employer trade demand rather than member need or value.
9. **The transition money trail** — identification of the SkillBridge website’s maintainer as the same vendor at the center of the credentialing data layer, the successor sole-source contract holding the owed outcome data, and the zero-competitive-offers pattern across both; plus the budget-justification “85% job-offer rate” set against GAO’s finding that the underlying data is not collected.
10. **The citizen compliance matrix** — mechanical extraction of 640 binding requirements from the six governing instructions, public-evidence grading (none MET), and the statute-softening specimen: the 90-day capstone rendered advisory (“should”) in the implementing instruction.
11. **The veteran-language census** — the first corpus-scale measurement (22,000+ posting rows, 300+ employers) of how rarely job postings speak to veterans beyond compliance boilerplate, including the zero rate in the cleared-intelligence market and the near-zero rate at a flagship SkillBridge-academy employer.
12. **The two-sided silence** — zero military-code references across a self-branded “top veteran employer’s” 384 requisitions, set against the directory’s 88% targeting silence: the first measurement of the translation gap from both sides simultaneously.
13. **The first documentation that a second, binding federal occupational classification exists for the cyberspace workforce with no published concordance to the crosswalk** — the DoD Cyberspace Workforce Framework (DCWF) instruction placing work roles “along with” rather than mapped to military occupational codes; zero occurrences of O*NET across fourteen governing and program documents; and the contractor bifurcation in which a mandatory contract clause still requires compliance with the “current version” of a manual canceled in February 2023.
14. **The first grading of a federal dataset against the data-governance stack that is supposed to keep it true** — the capstone information-enterprise directive searched and found silent on accuracy; the registration duty pointing at a repository a later directive superseded; the circular “trusted” standard; and the demonstration that the crosswalk’s blank maintainer breaches an express statutory metadata requirement while its “irregular” cadence breaches nothing, because no rule requires a federal data asset to be revised at all.
15. **The crosswalk rebuild: the 48,042-link reconstruction and the 0.51% human-review rate — corroborated link-for-link against an independent build made five weeks earlier by a different process — together with the four-vintage measurement of the file’s revision cadence, which finds the Air Force layer unrevised since October 2022, and its extension by content hash against third-party archive captures showing the published file byte-identical from September 2024 to July 2026** — a freeze approaching

four years, established without reliance on any figure this audit computed itself; and, beneath it, the currency of the destination taxonomy itself, which has admitted no new occupation since 17 November 2020 and refreshed no occupation from job incumbents or occupational experts in its current release.

16. **The first systematic comparison of military and civilian task documentation located anywhere — every scorable Air Force specialty’s published training standard against every published civilian occupation’s task statements, 1,055 crosswalk-offered pairs across 60 specialties and 450 distinct offered occupations ranked against a 922-occupation pool, yielding the median offered rank of 75, the layer decomposition that finds the member-facing layers weakest, and the first measurement of the skill-level collapse against the training standard’s own proficiency codes:** 90.9% of offers made identically to every level, against a median between-level movement of 73.5 rank places. No comparison of this kind published by the Department or by O*NET was located, by the searches described, as of 13 August 2026.
 17. **The a-fortiori from the Department’s own inbound instruments — the Air Force’s Convertible Skills List and Officer Classification Directory, measured against the outward-facing crosswalk on the controls this paper says are missing: row-level change dates, review before award, currency bounds, a negative authorization list, and 120 affirmative statements that a dimension does not apply. Same service, same classification system, same question about the same skills, opposite direction. The controls follow the direction of consequence.**
-

Part I — The Crosswalk: The Map Members Are Given

The foundation of the transition ecosystem is occupational translation: telling a member what their military specialty is worth in the civilian labor market. The government’s answer is the Military Occupational Classification crosswalk distributed through O*NET, My Next Move for Veterans, Credentialing Opportunities On-Line (COOL), and the TAP curriculum itself.

The mandate and the twenty-year trail

Congress ordered the crosswalk built in the FY2006 National Defense Authorization Act (NDAA). The VOW to Hire Heroes Act of 2011 (§222) went further: study true skills-equivalence, publish it, and use it in individualized assessments. Between 2013 and 2015, a DOL pilot study — conducted by Solutions for Information Design for \$49,996 — tested exactly that. Its result: the existing Defense Manpower Data Center (DMDC) crosswalk mapped 1.5 civilian occupations per military code; the enhanced method mapped 14.1 (962 matches versus 100 across the 68 occupations studied). The study’s formal recommendations: publish the enhanced crosswalk, distribute it through TAP, and extend it to all remaining occupations. Coverage was frozen at those 68 occupations, and the recommendations were never implemented.

Two qualifications belong on that figure, and the audit states them against its own interest. First, 14.1 against 1.5 is a measure of yield, not accuracy: it counts how many civilian occupations the method returned per military code, not whether the additional ones were better matches. The pilot did not test match quality against any outcome, and no study since has. Nine times the coverage is not nine times the correctness, and this paper should not have implied otherwise. Second, the enhanced method rested on the same O*NET military layer whose maintenance this Part documents as the underlying problem — a layer no party is under a published obligation to keep current, and whose stated update cadence is “irregular.” A remedy built on that foundation inherits its decay. Even had the pilot been adopted in full, nothing in its design would have kept it accurate, and the government commissioned no instrument that would have detected the drift.

That is the more damaging reading, not the softer one. The only remedy the record contains was never validated, never extended, and never given a maintenance mechanism.

The subsequent record is a drumbeat: GAO-14-144 (2014) found the TAP evaluation approach “incomplete.” GAO-17-133 (2016) recommended DoD develop performance measures for credentialing; DoD did not concur — the recommendation is Closed, Not Implemented. RAND told the Army in 2017 that existing crosswalks, naming My Next Move for Veterans, produce “low-quality recommendations,” and told the Office of the Secretary of Defense (OSD) the same year that soft-skills translation existed only as a prototype. GAO-22-105261 (2022): DoD “does not have current evidence of program effectiveness.” RAND again in 2023, this time sponsored by the transition office itself: “wide disparity in the number of high-quality matches.” GAO twice in 2024: transition data “not analyzed... because there is no requirement to do so.”

*(Provenance note: Part I’s figures were first computed in the July 2026 crosswalk audit and independently reproduced for this paper on 12 August 2026 from the four raw federal source files — exactly, and at set level rather than count level: 25,257 pairs in each build, intersection 25,257, none present on one side only. The rebuild adds the Credentialing in the Military/Armed Services Vocational Aptitude Battery (CitM/ASVAB) layer the July build omitted, and its 48,051 layer-tagged rows resolve to 48,042 distinct links — the figure used throughout this Part. The full reconciliation, layer by layer, with both undocumented build parameters, is at Appendix D, correction 25. A standing weekly tripwire has monitored the O*NET Resource Center’s crosswalk page, What’s New feed, and data.gov entry since a 1 July 2025 baseline and has recorded no change; the file remains the July 2024 vintage, so every figure below describes a file that has been stable and observed. The member-facing interface specimens are committed and independently re-fetched.)*

The audit findings

MissionProof’s 2026 independent audit rebuilt the Air Force layer of the crosswalk from its four federal source layers and found: 48,042 specialty-to-occupation links across 6,827 Air Force Specialty Codes (AFSCs), of which 247 links on 15 codes — 0.51% — have ever been human-reviewed. The quality layer is dated 2014 (the pilot’s residue), and the official update cadence is listed, verbatim, as “irregular.”

The four layers are not equal, and naming them locates the defect precisely.

layer	source	AF links	what it is
CitM / ASVAB	citm_supplemental_data.xlsx	22,794	title- and credential-driven; unreviewed; identical at every skill level
DoD/service COOL	cool_supplemental_data.xlsx	18,245	the services' own occupation crosswalk; unreviewed; identical at every skill level
DMDC	milx0724.csv	6,765	the baseline; frequently maps a specialty only to <i>another military occupation</i>
VOW §222 enhanced	vow222.xls	247	the hand analysis — QC'd, with linkage type and attainability rating

Ninety-nine and a half percent of the civilian advice in this instrument comes from two supplemental files that no one reviewed. The DMDC baseline is not the source of the implausible content; for the All-Source Intelligence Analyst it maps only to a military occupation code. The civilian lists — including the ones a member actually reads — come from COOL and CitM, both of which store the same occupations for every skill level of a career field, from helper to craftsman.

The update cadence is published as “irregular.” Measured against four dated vintages, it stopped.

Everything above describes one file at one moment. The publisher's own stated update cadence for the military crosswalk is the single word “irregular” — which this Part has until now simply quoted. Four dated vintages of the file are held, and they replace the publisher's word with a measurement.

vintage	rows	AF links	sha256 (first 16)
milx0720.csv — July 2020	30,037	6,261	21dc1b124ded4c90
milx1022.csv — October 2022	36,986	6,392	681494395b9b8edd

vintage	rows	AF links	sha256 (first 16)
milx1023.csv — October 2023	38,800	6,702	d0578c8749cad37a
milx0724.csv — July 2024	40,077	6,765	01e15dabb452ebd2

The July 2024 vintage is the file in service today, and its 6,765 Air Force links reproduce this Part’s DMDC layer exactly — the vintages and the rebuild are the same object measured twice.

Between October 2022 and July 2024 the Air Force crosswalk was not revised. Of the 6,142 Air Force specialty codes carried through both files, the occupation list changed for none of them. Not a low number: zero, across two published updates. Three hundred and seventy-one codes were added, none were removed, and none were revised. Across all services in the same window, twenty-six codes changed — every one of them a Marine Corps, Navy or Coast Guard record gaining an additional occupation, none a re-examination of a mapping already made.

Everything that ever moved, moved earlier. Between July 2020 and October 2022 the file behaved like a maintained instrument: across all services, 2,847 links added, 2,424 dropped, 82.7% retained, and the occupation list changed for 2,328 of 12,355 carried-forward codes. On the Air Force layer alone the same window added 1,184 links and dropped 1,053, with the occupation list changing for 1,030 of 5,996 carried-forward codes. After October 2022, the drop count across two updates is three links in total.

The file is still maintained. What is maintained is not the advice. The Department edited 327 Air Force records in 2024, by the file’s own change marker. One hundred and seventy-seven of those edits were changes to the *title* of a military occupation. Exactly one was a change to the civilian occupation mapping. The same marker shows the collapse across the decade: Air Force records flagged as an O*NET mapping change number 444 in 2020 and 408 in 2021, then 8, then 70, then one in 2024.

The advice a member reads is, by the file’s own accounting, years old. The median active Air Force record in the July 2024 file was last changed in 2018; 54.0% of active Air Force records carrying a parsable change marker were last touched before 2020. And nothing is ever retired: the obsolete share of Air Force rows ends the series higher than it began — 56.5%, 56.1%, 57.7%, 58.2% — because rescinded specialties are marked and kept, never removed. The file’s overall obsolete *share* falls only because active rows are added faster than dead ones accumulate.

One control belongs with this. If links vanished and returned between vintages, the churn figures would be measuring file instability rather than revision. They do not: of the 11,607 links present in both the 2020 and 2024 files, not one is absent from either intermediate vintage. The series is monotone and the drops are real deletions.

Three boundaries are stated with the finding. First, this measures the DMDC layer only — the versioned one. COOL and CitM, which supply 99.5% of the civilian content a member actually reads, are not versioned in this corpus, and no comparable vintage series for them was located; the freeze documented here is of the baseline the other layers hang off, not a direct measurement of the member-facing surface. Second, the change marker records the last change to a record, not the last review of

one. A record that has not changed may have been examined and confirmed correct. Nothing in the file distinguishes the two — which is the more precise version of the finding: the instrument publishes no signal by which a member, or an auditor, could tell a verified mapping from an untouched one. Third, four vintages is a sample, not the full published series.

Even bounded that way, the sentence the Part could previously only quote is now one it can measure. The crosswalk’s update cadence is not merely irregular. On the layer where currency can be checked, it has been static for the entire period in which the Department was asked, and declined, to measure whether any of this works.

The freeze did not end in 2024. The published file is byte-identical to July 2026.

The four vintages stop at July 2024 because that is the newest file held. That boundary can now be closed, and it is closed from outside this audit’s own corpus.

The O*NET Resource Center publishes the military crosswalk at a single link that carries neither a version nor a date:

<https://www.onetcenter.org/crosswalks.html> → /dl_files/2019/military_crosswalk.zip

The filename does not change when the contents do. A user cannot tell from the download page which vintage they are receiving — which is precisely why the only way to establish currency is to hash what the page serves and compare it across time. The Internet Archive has captured that file repeatedly, and its capture record is held by a party to which this audit is not a party.

capture	HTTP	content digest
2023-03-10	200	SC0VSFJT
2024-03-30	200	FSNTTTHU
2024-09-26	200	VW42YXEG
2025-01-14	200	VW42YXEG
2025-08-31	revisit	VW42YXEG
2026-07-22	200	VW42YXEG

The digest changes in 2023 and again in 2024 — the file was being republished — and then stops changing.

That series was then verified directly rather than relied upon. SHA-256 ee738141 1644fd39 adaf93e7 f3a88147 73b28f4e 8a5b2c00 c7ad6b77 5b01d747, 4,826,815 bytes, is identical across four copies: the archived capture of 26 September 2024, the archived capture of 22 July 2026, the copy downloaded during the July 2026 audit, and the separately archived 2024 vintage. The archive payloads were hashed as retrieved; the local copies were hashed on disk. The zip contains milx0724.accdb, milx0724.csv, milx0724.xlsx and Read Me.pdf — every member stamped 22 July 2024 — and its milx0724.csv is the 40,077-row, 6,765-Air-Force-link file measured above.

The file the Department published on 22 July 2024 is, to the byte, the file it publishes today. Set beside the vintage measurement, the Air

Force layer of the crosswalk has not been revised in any respect visible to a user **since October 2022 — a period approaching four years**, spanning the entire window in which the Department was asked, and declined, to measure whether any of it works.

This also disposes of the most obvious objection to the tripwire evidence cited earlier in this Part. A monitor that has recorded no change since a 1 July 2025 baseline could be reporting a short window. It is not: the archive brackets that baseline on both sides, and the null is the continuation of a freeze that began thirty-three months before it.

Two boundaries travel with the finding. First, the `/dl_files/2019/` path is almost certainly a legacy directory rather than a statement about the content’s age — the digest series shows the content changing in 2023 and in 2024 — and it is not offered here as a vintage indicator. Second, the scope is unchanged from the section above: this is the DMDC layer, the versioned one. COOL and CitM remain unversioned in this corpus, and the freeze documented here is of the baseline the member-facing layers hang off.

(Method: Internet Archive CDX capture history for the published zip, with SHA-256 computed on the raw archived payloads and on the local copies, 12 August 2026. The hash table is committed.)

And there is a second floor under the freeze. A crosswalk is a lookup: it can only ever return an occupation the destination taxonomy already contains. That taxonomy is also frozen. O*NET-SOC 2019 took effect on 17 November 2020 and remains current — 1,016 occupational titles, 923 of them carrying data. No release from 25.1 through 30.3 records a taxonomy change: five years and nine months without a new occupation. The mechanism that used to add them has wound down rather than paused — 153 new and emerging occupations entered the 2009 taxonomy, 152 the 2010, twelve the 2019, and none since. Nor can it move quickly. The Standard Occupational Classification (SOC) is revised roughly once a decade — 2000, 2010, 2018, next 2028 — the first Federal Register comment period for the 2028 revision closed on 12 August 2024, and O*NET took a further two years to rebuild against the previous revision. An occupation that emerged after August 2024 cannot enter the federal taxonomy before roughly 2030.

The data inside the occupations runs on a similar clock, and the published figure obscures it. O*NET reports 923 occupations updated in each of 2022, 2023, 2024 and 2025, which read alone implies an annual refresh. It does not mean that: “updated” spans thirteen source categories, among them web research, machine learning and employer job postings. The subset refreshed from actual job incumbents or occupational experts — the survey evidence behind the skills, knowledge, abilities and work activities the crosswalk is scored against — is 90 occupations in release 28.0, 101 in 29.0, 78 in 30.0, and none at all in release 30.3. Against 923 data-level occupations that is a full refresh cycle of roughly nine to twelve years, or about one SOC revision. Every occupation page nonetheless displays the current release date, 19 May 2026, whatever the age of the evidence beneath it — the page records that the file was published, not that the occupation was measured.

And the age of that evidence can be measured directly, because O*NET publishes it. Every task statement in the database carries the month its data was collected. Across the 923 occupations that carry task data, the *most recent* collection

date has a median age of 5.8 years as at August 2026. 597 occupations (64.7%) were last measured more than five years ago, and 190 of them (20.6%) more than ten. The oldest evidence still carried in the current release was collected in 2004. So the decade-scale refresh cycle implied by the release notes is not an inference: it is visible occupation by occupation, in dates the Department's own contractor publishes.

Three layers are frozen, not one. A separating Airman in 2026 is handed a mapping last revised in October 2022, pointing into a set of occupations defined in November 2020, described by survey evidence whose median age is nearly six years and which, for one occupation in five, is more than a decade old — and the earliest a genuinely new occupation can enter that set is about 2030.

That bears directly on what a remedy would have to reach. The crosswalk is not merely unmaintained; it is a stale index into a stale catalog. Fixing the mapping — the remedy the 2015 pilot proposed and the Department shelved — would have improved the index while leaving the catalog exactly as it stands.

*(Method: O*NET Resource Center taxonomy, database-release and data-update pages, and the release 30.3 updated-occupation listing, each retrieved and verified 14 August 2026; Bureau of Labor Statistics (BLS) 2028 SOC revision docket; and, for the collection-date measurement, the O*NET 30.3 database itself — Task Statements, 18,796 rows across 923 occupations, SHA-256 78703906d5eb92fa... on the published archive. Boundary: this measures the destination taxonomy's currency and nothing more. It is not offered as a claim that any particular civilian occupation is missing from the taxonomy; that is a distinct question requiring the SOC Direct Match Title File to be searched in full, which has not been done, and no such claim is relied on anywhere in this Part.)*

A further structural fact bears on the 0.51%. The fifteen hand-reviewed specialties are stored under rollup codes — 2A3X3, 3D0X1, 4N0X1 and the rest — and not one of those code strings exists in the DMDC baseline file at all, which uses skill-level codes. The only quality-reviewed layer is keyed to a code space the baseline does not share.

Structural absurdities follow from unreviewed bulk generation. 119 civilian occupations are attached to 100 or more Air Force specialties each. 1,539 specialties are told to consider Airline Pilot. Broadcast Technicians is attached to 793 — a figure that recurs below, because it is the same occupation the instrument offers to pararescuemen, signals-intelligence analysts, cryptologic linguists and airborne intelligence, surveillance and reconnaissance (ISR) operators, and, in the raw CitM layer, to Army infantrymen.

Those attachments were then re-tested without using words at all. The scoring elsewhere in this Part is lexical, and the fair objection to it is that it measures vocabulary rather than work: two occupations may describe the same job in different language. O*NET assigns every occupation a set of Detailed Work Activities (DWAs) — its own structured decomposition of what the occupation does, produced by O*NET analysts and independent of anything computed here. Two occupations either share a detailed work activity or they do not.

The first result constrains this Part's own language, and is reported for that reason. Measuring how much activity each specialty's offered set shares internally, against two hundred random size-matched draws from the same pool, the offered sets

score eight and a half times the random baseline, and 84.7% of specialties beat 95% of random draws. Whatever produced this file, it is not noise, and nothing in this Part should be read as claiming otherwise. The layers differ sharply, and that difference is itself a finding: the CitM layer beats a random draw for 82.6% of specialties, the COOL layer for 51.9% — so close to half of the credentialing layer’s specialty sets cannot be distinguished from a random draw of occupations.

The second result is at the level a member actually experiences it, and it is sharper than the lexical method. A member is not offered a set average; they are offered individual occupations, one at a time. Scoring every link by the overlap between that occupation’s detailed work activities and those of every *other* occupation offered to the same specialty — 39,186 links across 3,695 specialties — 3,302 links, 8.4% of the total, share not one detailed work activity with anything else that specialty is pointed at. 1,821 specialties, 49.3% of those tested, carry at least one such link.

The defect is concentrated, which is what makes its persistence difficult to excuse. All 3,302 are carried by 92 occupations out of the 701 DWA-scored occupations the file’s links reach (its full reach is 801 distinct occupations), and ten occupations carry 64.4% of them. Broadcast Technicians heads that list: of the 696 specialties that offer it alongside at least two other occupations, it shares no detailed work activity with any of them in 93.1% — placing it at the floor of its layer’s distribution rather than merely near it. Intelligence Analysts follows at 356 specialties, Remote Sensing Technicians at 245. Correcting ten records would remove close to two thirds of the measured defect.

*(Method: O*NET 30.3 database, Tasks to DWAs, 23,850 rows covering 923 occupations — a count that independently reproduces the data-level figure above from a different file — scored against the 48,051-link reconstruction. Jaccard overlap on unweighted DWA sets; null of two hundred size-matched random draws, seeded. Boundary, and it governs every figure in this passage: this measures an occupation’s fit to the rest of the set its specialty is offered, not its fit to the specialty. Nothing in O*NET describes an Air Force specialty, so a specialty handed a uniformly wrong but internally consistent set of occupations would score as coherent here. The 8.4% is therefore a floor, and a specialty scoring well by this test is not thereby well served. Specialties offering fewer than three occupations are excluded for want of comparators, which removes most of the DMDC layer; 4.5% of links point to occupations carrying no DWA record and are excluded rather than scored as zero. Because DWA assignments have themselves been revised since the crosswalk was published, links built against an earlier vintage are being judged by a later one — which cuts toward the crosswalk, not against it.)*

Crosswalk quality varies by specialty. The variance is measurable, it is undisclosed, and the floor is not zero.

The single most important qualification in this Part is that crosswalk quality varies by specialty, and the variance is not disclosed anywhere the member can see it. An earlier draft of this Part stated that variance qualitatively — “unevenly good” — which was an assertion, not a measurement, in a Part that is otherwise quantitative. It is now measured.

Five defect indicators were computed for every active Air Force career field from

the rebuilt four-layer table. A field is counted as carrying an indicator if: no link of its has ever been hand-reviewed; it carries an occupation attached to 100 or more other specialties; its occupation list is identical across skill levels; it carries twenty or more occupations; or every one of its links resolves to another military occupation.

Population: 315 active Air Force AFSC families.

indicator	families	share
no hand review of any link	305	96.8%
carries an occupation attached to ≥ 100 other specialties	290	92.1%
occupation list identical across skill levels	241	76.5%
twenty or more occupations on one family	110	34.9%
every link resolves to another military occupation	0	0.0% (merged four-layer table; on the DMDC baseline alone this indicator fires on 16.7% of families)

defect indicators carried	families	share
one	6	1.9%
two	67	21.3%
three	162	51.4%
four	80	25.4%

Not one active Air Force career field is clean on all five indicators. 76.8% carry three or more. 3.2% have had any link hand-reviewed.

One property is universal and is therefore reported separately rather than as a discriminator: no link in the 48,042-link table carries a fit score, a rank, or any confidence signal. Attainability is rated on the 247 enhanced links alone.

The floor is not zero — every field carries at least one indicator, and three-quarters carry three. And the differences between fields are real — a member in a field carrying two indicators is in a materially better position than one carrying four. What has never existed, in any layer, is a way for that member to find out which they are.

This is a measured statement, not an impression. Of 12,913 actively mapped military occupational code (MOC) records across all services (12,103 distinct service-and-code pairs), 1,533 records (11.9%), covering 1,430 distinct codes (11.8%), dead-end in a 55-xxxx occupation — that is, the crosswalk’s published answer to “what is this in the civilian world” is *another military occupation*. Among Air Force enlisted codes the

rate is 20.9% (275 of 1,318). Among the four Air Force Special Warfare fields it is 24 of 24 — 100% on the DMDC baseline. *(Every figure in this paragraph was recomputed from milx0724.csv inside the published archive — SHA-256 ee738141...5b01d747 — on 21 August 2026, filtering to STATUS = A and counting a record as dead-ended when every populated ONET1-ONET4 value begins 55-. The record-grain figures reproduce exactly; the code-grain figures are stated alongside them because the two differ and the distinction was previously carried by the single word “codes.”)* The zero in the indicator table above is not in conflict with it: that indicator is computed on the merged four-layer table, where the COOL and CitM layers attach civilian occupations to all four families, so no family has every link resolving to a military occupation. Which layer a member’s tool reads therefore decides whether these fields have any civilian translation at all.

For those fields there is no civilian translation in the file to fail. Whatever a Pararescueman, Combat Controller, Tactical Air Control Party (TACP) airman, or Special Reconnaissance airman reads on a federal website is produced by a derived layer the page does not name; the footer credits only “multiple sources; see Military Transition Search at the O*NET Resource Center.” Whether that derivation is sound has never been published or tested.

Skill level is maintained by the government and discarded by the member’s tool

The published file carries a separate row for each skill level — helper, apprentice, journeyman, craftsman, superintendent, and the Chief Enlisted Manager (CEM) code — and O*NET OnLine’s Military Crosswalk Search exposes all of them. My Next Move for Veterans, the tool aimed at the separating member, keys on three characters and discards the skill level entirely.

The collapse is not cosmetic — but the magnitude of it requires a definition, and this paper previously published one without stating its own.

Of 284 active Air Force families carrying two or more skill levels, the occupation list differs by skill level in 26 — 9.2%. In 90.8% it is identical from helper to craftsman. The definition of record, stated because the figure is meaningless without it: all four source layers merged, active codes only, shred preserved. That is the population a separating member actually occupies — a member holds an active code with a shred and is shown the merged list.

The figure is definition-sensitive and the range must travel with it. Across seven reasonable definitions of a career-field family the share that differs by skill level ranges from 4.0% to 41.0%: 19.5% on the merged table with obsolete codes included, 21.5% on the DMDC baseline alone, 41.0% if shreds are collapsed together, 4.0% on the CitM layer alone. The qualitative finding — that the mapping is substantially collapsed across skill levels — holds under every one of them. No magnitude should be quoted without the definition attached, and this paper previously did exactly that (Appendix D, correction 23).

In the Command and Control field, levels 1 through 7 map to Command and Control Center Specialists and level 9 maps to First-Line Supervisors; the member-facing tool hands a twenty-year superintendent the operator’s list. The granularity the government builds and maintains is discarded at the last step before the member reads it.

Seniority subtracts options, and the manager loses management

In every one of the four Special Warfare fields, the Chief Enlisted Manager code returns fewer civilian occupations than the journeyman code beneath it: Pararescue 13 against 23, Combat Control 6 against 9, TACP 6 against 36, and Special Reconnaissance 5 against 64. The most senior enlisted leader in Special Reconnaissance is offered five civilian occupations; the apprentice beneath him is offered sixty-four.

For Pararescue the relationship was decomposed exactly against verbatim captures. The manager's list is a strict subset of the journeyman's — the same list with ten entries removed — and five of the ten removed entries are the complete Management set: General and Operations Managers, Security Managers, Transportation/Storage/Distribution Managers, Training and Development Managers, and Emergency Management Directors. A Pararescue *helper*, a trainee who has not finished the pipeline, is shown all five. The Pararescue *manager* is shown none of them.

The fair reading is stated first and is a real one. The Chief Enlisted Manager code is structurally distinct from the skill-level codes, and a narrower list is entirely consistent with CEM codes having been mapped by a separate process or at a different vintage. No intent is alleged. The defect survives its fair reading on the same ground every defect in this Part does: nothing on the page explains the difference, and nothing tells a manager to look at the journeyman's code instead. These are counts of listed occupations — not measures of opportunity, hiring, or earnings.

It is not confined to Special Warfare. The same inversion was tested across the seven Air Force intelligence career fields and holds in all seven. A human-intelligence journeyman is offered twenty-one occupations; the superintendent who leads him is offered four. A cyber-intelligence craftsman is offered sixteen; his superintendent, four. Signals intelligence: eleven against four. The reductions run from 64% to 95%, and in five of the seven the only superintendent code the crosswalk carries is one it marks as no longer in active use. Air Force intelligence has a named civilian counterpart occupation, no hazardous-duty exclusion, and no military dead-end. It fails the same way.

There is a structural reason, and it is not obscure. The crosswalk's sole input is the AFSC — a *technical* classifier. Professional military education (PME) is the parallel track on which supervision, resource management, instruction and organizational competency are developed, and no field, code or channel anywhere in the instrument admits it. Because the only input is technical, the instrument cannot add management occupations as a member rises; what the evidence shows is that it removes them.

This is not a speculative gap. Independent research analyzing Air Force enlisted PME curricula against an external project-management competency framework finds that competency content is present at every tier and that its distribution differs significantly between tiers — that is, the thing that changes as an airman is promoted is documented in the Air Force's own instructional materials and is statistically demonstrable (a separate doctoral study by this author, accepted pending defense). The instrument omits a variable that has been shown to vary, along precisely the dimension where the instrument fails hardest.

The interface misdescribes its own output, in a specific and demonstrable way

My Next Move for Veterans presents its results under the heading “Closest matches are shown first.” For the Pararescue query the twenty-one returned items form two internally alphabetical blocks — eleven items, then ten — a pattern with a chance probability under a true similarity ranking of roughly one in 1.4×10^{14} .

We first offered the instrument the most generous available explanation: two score tiers, alphabetized within each to break ties. That is defensible engineering, and we published it as the fair reading. A subsequent capture of the expanded list then refuted our own fair reading. Six programmatic checks against the verbatim lists establish the mechanism exactly: block one is the 1Z100 Manager mapping, block two is the incremental occupations from the skill-level mapping, the two military-only codes are filtered out, each source block is alphabetized, and the concatenation is presented as a ranking. The block boundary is a source boundary, not a similarity boundary. The two-score-tier defense is no longer available to the instrument — its own data eliminates it.

Two consequences follow. The “closest matches” claim is false in a demonstrable rather than merely unsupported way. And the first eleven entries a first-term airman reads are the mapping the government assigned to the most senior member of his career field — from which every management occupation has been deleted. The header also reads “top 20” above twenty-one items.

On the member’s screen, under a banner promising closest matches first, a nationally registered paramedic finds Athletes and Sports Competitors at position two, Broadcast Technicians at position three, an occupation requiring roughly eleven further years of schooling at position seven, and Paramedics — the license he already holds — at position ten.

One career field is a specimen. The defect was therefore tested against another, and it is a property of the instrument. The seven Air Force intelligence codes and the airborne ISR family were captured from the member-facing tool and decomposed the same way. All eight lists are alphabetical, or concatenations of alphabetical runs. The joint probability, were the order a genuine similarity ranking, is one in 10^{59} ; the human-intelligence list alone — twenty items in exact alphabetical order — is one in 4×10^{18} , some thirty thousand times less likely than the Pararescue figure above.

And the ranking claim fails in the way that matters most to a member. O*NET carries an occupation named Intelligence Analysts. Across seven intelligence career fields, under a banner reading “closest matches are shown first,” it is ranked first in none of them — appearing at position three, four, six, six, six, ten and eleven. A human-intelligence airman reads past Clinical and Counseling Psychologists, Customs and Border Protection Officers, Database Administrators, Detectives, Editors, Eligibility Interviewers and Emergency Management Directors before reaching the occupation that shares his job title.

The airborne ISR family shows the join without any need for source files. Two distinct active specialties sit under one three-character prefix — an airborne cryptologic language analyst mapping of eleven occupations and an airborne ISR operator mapping of sixteen. The member-facing tool keys on those three characters and returns one list of sixteen that decomposes into three alphabetical blocks. Two working mappings and a senior mapping, sorted separately and concatenated under a sentence

claiming they are ranked by similarity.

Fair reading, and it is probably the true one. The likeliest explanation is mundane: the page sorts alphabetically for display and “Closest matches are shown first” is unrevised boilerplate above it. No one built a false ranking. That reading is the finding rather than a defense against it — the member is told the order carries information, and it does not.

The matcher is lexical, and it shows

The root cause is recorded in the audit’s own reverse-engineering of the source files: the mapping rests on title and description correspondence, keyed at occupation-family level, with no skill tier, no rank or scope, no competency or task overlap, no fit ranking and no education-gap signal. Two consequences are visible in the output.

Term collision. Where a word carries one meaning in military usage and names a different civilian occupation, the instrument matches the word. *Remote sensing* in intelligence usage means performing the work from remote sites across a distributed enterprise; the occupation of the same name is instrument-based earth observation — operating and calibrating sensing devices. Remote Sensing Scientists and Remote Sensing Technicians are attached to five of the eight intelligence and ISR families examined, and no Air Force intelligence course teaches the operation of such a device. *Survey* in intelligence usage means the study of an area; the occupation means land survey, instrument work and trigonometry, which is not taught in the six-month All-Source Intelligence Analyst course. Surveying and Mapping Technicians is attached to five of the eight. *Signals* and *audio* produce Broadcast Technicians and Sound Engineering Technicians for signals-intelligence analysts and cryptologic linguists — and, in the raw CitM layer, Broadcast Technicians for an Army infantryman.

These are not judgments about whether a mapping is apt. They are structural: the link rests on lexical overlap rather than on any competency the specialty develops, and the pattern is countable.

Cross-service inconsistency. The same work is translated differently depending on the uniform. The Air Force All-Source Intelligence Analyst is given sixteen civilian occupations. The Army intelligence analyst — materially the same analytic function — is given sixty-eight. Neither number is defended anywhere, no methodology reconciles them, and a member has no way to learn that the instrument’s answer to “what is this worth in the civilian world” depends on which service trained them.

Two further gaps are recorded from the same audit. Implausible fits survive in the unreviewed layers — the Aerospace Medical Service specialty draws Funeral Home Managers, Morticians, Motorboat Operators, Explosives Workers and Camera Repairers. And the file carries its own dead: 5,166 links across 2,708 obsolete specialty variants remain in the published crosswalk. Two entire Air Force intelligence career fields — Electronic Systems Security Assessment and Electronic Signals Intelligence Exploitation — are still published with civilian mappings; the file’s own end-date columns record them as having ended in April 2009 and October 2010 respectively. The crosswalk records currency to the month. It records quality not at all.

What this does and does not establish

It does not establish that the lists are junk. The Pararescue list contains Paramedics, EMTs, acute and critical care nurses, commercial divers, surgical assistants, emergency management directors, security managers, and training and development roles. Several of those are plainly relevant to the specialty.

Nor does it establish the converse, and an earlier draft of this Part asserted it. That draft stated that “most of that is substantively defensible” and that “where the crosswalk translates, it often translates reasonably.” No competency test supported either sentence when it was written; the judgment rested on whether the occupation titles read as plausible. That is the same class of unverified assertion this paper objects to elsewhere, and it is withdrawn (Appendix D, correction 24).

What the record does contain is a competency check, and it points the other way. The July 2026 audit tested the All-Source Intelligence Analyst mapping against the specialty’s own Career Field Education and Training Plan (CFETP) duty text and found cross-career-field bleed: the specialty is given Cartographers, Surveyors and Geographers, which is geospatial-intelligence work belonging to a different career field. That check was qualitative and covered one specialty.

An earlier revision recorded that this comparison had not been run — by the Department, by O*NET, or by this audit — and why this audit’s first extraction attempt failed (Appendix D, correction 29). The first two clauses still stand; the third no longer does: the extraction was rebuilt on page geometry and the comparison was run on 13 August 2026. The next subsection reports it, with its limits attached.

What it establishes is that a member cannot tell the difference. The instrument gives no confidence signal, no methodology, no indication that quality varies by field, no indication that a different code would return a different and often longer list, and no indication that the order on the screen is a join artifact rather than a judgment. The member is asked to plan a one-time, irreversible career transition against an instrument whose accuracy for *their* specialty is unstated, unmeasured, and — for the fields where it fails hardest — unreviewed since 2014.

The systematic comparison, now run: where the crosswalk’s own offers rank

The question is narrow, and answerable from published documents alone. When the crosswalk tells an airman that their specialty corresponds to a given civilian occupation, do the two published task inventories describe the same work? Both sides of that comparison are federal publications. Neither the Department nor O*NET publishes it: no comparison of this kind by either was located, by the searches described in this Part, as of 13 August 2026. This audit has now run it.

Method, in brief. Every CFETP held — 106 documents — was re-extracted from pdftotext word bounding boxes rather than from a character grid, because the corpus spans at least four incompatible print genres and several documents rotate their column headers ninety degrees. One local rule holds in every genre: *on a table row, the proficiency codes are the trailing run of free-standing code tokens, and the task statement is everything before them.* That needs no page-level column model, so narrow and wide tables behave identically and a statement can never absorb a code. Yield: 36,417 deduplicated task statements across 97 specialty codes, 91 of them within the scope of the comparison below, of which 26,005 come from pages carrying

confirmed proficiency-code columns. Only those 26,005 are scored — a gate added after an ungated pass ranked Nondestructive Inspection’s semantically correct occupation 901st, on 26 rows of that CFETP’s own Part I narrative prose, which the scorer had read as a task inventory. Each specialty and each occupation is then a bag of task-statement terms, unstemmed, compared by cosine similarity on TF-IDF with inverse document frequency computed over the O*NET occupation corpus only, so a term’s weight reflects how distinctive it is among *civilian* occupations rather than how common it is in military writing.

The instrument was checked before it was believed. Against mappings whose correctness is not in dispute, it ranks first every time: HVAC to *Heating, Air Conditioning and Refrigeration Mechanics*; Dental Assistant to *Dental Assistants*; Air Traffic Control to *Air Traffic Controllers*; Fire Protection to *First-Line Supervisors of Firefighting*; Optometry to *Ophthalmic Medical Technologists*. Against an independently produced character-grid extraction of the same PDFs, 400 randomly sampled statements returned 93.5% concordance — 71.8% verbatim, 21.8% on opening phrase, 6.5% no match, with several inspected mismatches attributable to the character-grid version rather than to this one. Roughly 0.7% of extracted statements retain an adjacent code or page number.

The result. Across the 60 Air Force specialties that could be scored (of 91 in the corpus; the other 31 are treated below) and 1,055 crosswalk-offered specialty-to-occupation pairs (which name 450 distinct occupations between them), ranked against the 922 O*NET occupations that entered the ranking (923 occupations carry task data; the one-occupation difference between that population and the scored pool is a property of the committed scoring run, recorded here rather than silently carried, and not separately reconciled. The Air Force crosswalk’s full reach, across all four layers and every specialty, is 801 distinct occupations):

Where the offered occupation ranks	Pairs	Share
top 10	197	18.7%
top 25	319	30.2%
top 50	435	41.2%
top 100	606	57.4%
outside the top 100	449	42.6%
outside the top 300	192	18.2%

The median occupation the crosswalk offers ranks 75th of 922 on similarity to the specialty’s own training standard. Two offers in five fall outside the top hundred; nearly one in five falls outside the top three hundred. On the conventional median — the mean of the two middle values — 16 of the 60 specialties have a *median* offer outside the top 100; the count is 19 if the upper of the two middle values is taken instead, and the definition is stated because a figure of this kind is meaningless without one.

Restricted to specialties with 200 or more extracted task statements, so that nothing below is an artifact of a thin corpus, the weakest offers in the set are these:

Specialty	Offered occupation	Rank of 922	Layer
3F1X1 Services	Travel Agents	905	CitM
4A0X1 Health Services Management	Data Entry Keyers	901	COOL
4N0X1 Aerospace Medical Service	Customs and Border Protection Officers	881	COOL
1Z3X1	Bailiffs	880	COOL
4A0X1 Health Services Management	Motorboat Operators	862	COOL
2T1X1 Vehicle Operations	Cleaners of Vehicles and Equipment	841	COOL
4N0X1 Aerospace Medical Service	Histology Technicians	816	COOL

Every one of them is offered to all five skill levels. Aerospace Medical Service is the specialty that draws Funeral Home Managers and Morticians earlier in this Part; the qualitative observation and the measurement now agree.

The defect is concentrated in the layers a member actually sees

The crosswalk is not one artifact but four layered sources. Three of them carry pairs that could be scored — no VOW §222 enhanced link entered the scored set — and those three do not perform alike:

Layer	Pairs	Median rank	Outside top 100
DMDC — the official DoD crosswalk	103	19	24.3%
CitM — Credentialing in the Military	250	59	38.0%
COOL — AF Credentialing Opportunities On-Line	767	81	45.0%

(Layer counts sum to more than 1,055 because 65 pairs are carried by more than one layer.)

The DMDC layer — the smallest of the three scored layers, and the one with formal occupational-analysis provenance — performs markedly better than the credentialing

layers, which carry more than nine-tenths of the scored pairs. The offers a separating member is most likely to encounter come from the layers with the weakest correspondence to their own training standard. This is an independent corroboration of the four-layer decomposition established earlier in this Part, arrived at from task text rather than from file structure.

Skill level: the training standard distinguishes what the crosswalk does not

This Part has already established that the crosswalk hands the same occupation list to a first-term airman and to a superintendent. The task inventories say that is wrong.

- **959 of 1,055 pairs — 90.9%** — are offered identically to every skill level the crosswalk covers for that specialty.
- Among the 882 pairs for which two or more skill levels could be scored separately, the median movement in rank between levels is **73.5 places**.
- **74.1% move at least 20 places. 41.4% move at least 100. 20.7%** move at least 200.

Specialty	Occupation	3-level	5-level	7-level	Offered to all levels
1Z3X1	Radio, Cellular and Tower Equipment Installers	1	4	733	yes
2W2X1	Helpers	9	33	722	yes
Nuclear Weapons	— Installation, Maintenance and Repair				
3E0X1	Telecommunications	3	436	690	yes
Electrical Systems	Line Installers and Repairers				
3F1X1	Bakers	64	535	732	yes
Services					
1Z3X1	Quality Control Systems Managers	677	290	96	yes

Specialty	Occupation	3-level	5-level	7-level	Offered to all levels
6F0X1 Financial Management	Compliance Officers	—	776	194	yes

Read the first row plainly. At the apprentice level, that specialty’s published tasks are the single closest match in all of O*NET to tower-equipment installation. By the craftsman level the same occupation has fallen to 733rd. The crosswalk offers it at both, and gives the member no signal that anything has changed.

The last two rows run the other way — an occupation that is a poor match for a junior airman and a good one for a senior noncommissioned officer. That is the shape of the problem stated from inside the career field: the hands-on task persists at every level, the supervisory and management content grows with rank, and the civilian occupation that actually corresponds at the top of the ladder is a *different occupation* than the one at the bottom. The instrument has no representation for it. The skill level is maintained by the government, encoded in the fourth character of the crosswalk’s own military occupation codes, and discarded at the point of use — which is what makes this measurement possible at all, and what makes its absence from the member’s screen a design choice rather than a data limitation.

What this covers, and what it does not

Coverage is a scope limitation, and it is stated before the conclusions rather than after them. Of 91 specialties in the corpus, 60 could be scored and 31 could not, because the CFETP copies collected here do not contain the Specialty Training Standard the document refers to. The reason is not established and no claim is made about it. Controlled distribution of training material is ordinary practice and carries no implication of any kind.

What does matter is what it does to coverage, because the excluded set is not random: 18 of the 31 are aircraft maintenance specialties and 7 are intelligence specialties — 25 of 31 in two families. Aircraft maintenance is one of the largest enlisted populations in the Air Force and is almost entirely absent from these results. Every percentage above is therefore a property of the 60 specialties that could be scored, not of the Air Force enlisted force, and is stated that way throughout. Whether the excluded families would score better or worse is unknown, and this paper does not speculate about it.

Four further limits travel with every number above.

1. **The score is lexical.** It measures whether two published task documents describe work in the same words. **A low score does not prove a mapping wrong.** It is a documentary finding about the federal record, not an occupational-analysis finding about the work — and the federal record is what a separating member is given.
2. **The level partition is a sample.** Roughly half the extracted statements carry no detected proficiency code, so the per-level bags are a subset of each training

standard rather than all of it. The 9-level partition is thinnest and was scorable for 119 pairs; most swings reported are 3-to-5-to-7.

3. **Shred letters are rolled up** to the five-character specialty code, so a specialty with materially different shreds is measured as one population.
4. **This is one vintage** — O*NET release 30.3 against the crosswalk as rebuilt. A snapshot, not a time series.

What the comparison establishes is narrow and sufficient. Two published federal task inventories, which the government’s own instrument asserts describe corresponding work, frequently do not describe the same work in the same terms; the divergence is worst in the layers the member is most likely to see; and the training standard’s own proficiency codes move an occupation a median of 73.5 rank places between skill levels that the crosswalk treats as interchangeable. None of that requires a judgment about whether any individual mapping is apt. It requires only that both documents be read.

And no published government version of it was located. The comparison was the work of one independent researcher, on public files. The crosswalk has been federally mandated since the FY2006 NDAA and quality-reviewed once, in 2014.

The education gap is not an inference. It is measured, and the instrument runs against it.

The root-cause analysis of the crosswalk’s source layers records that the matcher carries no education-gap signal. Attainability is rated on 247 of 48,042 links — 0.51% — across fifteen specialties, none of them intelligence. That is a structural statement. The Department’s own demographics publish what it costs.

Enlisted educational attainment, all five services, by pay grade — share holding a bachelor’s degree or higher (DMDC Active-Duty Military Personnel Master File, September 2024; populations beneath):

band	Army	Navy	Marine Corps	Air Force	Space Force
E1-E4	5.1%	4.2%	0.8%	3.2%	7.0%
E5-E6	11.2%	10.5%	4.0%	17.0%	21.9%
E7-E9	27.5%	23.6%	21.0%	54.7%	59.1%
<i>members</i>	<i>353,326</i>	<i>272,792</i>	<i>151,755</i>	<i>251,412</i>	<i>4,872</i>

1,034,157 enlisted members across five services. Attainment rises with grade in every one, and the gradient is steepest where the service runs its own degree-granting institution: Air Force 17.1× from junior to senior enlisted, Space Force 8.4×, Navy 5.6×, Army 5.4×, Marine Corps 26.2× off a 0.8% base.

And a degree is the wrong measure of the problem. The Department publishes six attainment categories, not two. Read across all of them, for every service and every enlisted band:

service band		N	no HS	HS / GED or some college	associate's	bachelor's	advanced	unknown
Army	E1-E4	181,557	0.1%	90.2%	2.5%	4.6%	0.5%	2.1%
Army	E5-E6	122,813	0.2%	72.3%	7.9%	9.9%	1.3%	8.4%
Army	E7-E9	48,956	0.1%	41.1%	15.4%	20.6%	6.9%	15.8%
Navy	E1-E4	115,999	0.1%	92.8%	2.8%	3.9%	0.3%	0.1%
Navy	E5-E6	123,334	0.1%	80.5%	8.8%	9.4%	1.2%	0.0%
Navy	E7-E9	33,459	0.6%	58.3%	17.5%	18.4%	5.1%	0.1%
Marine Corps	E1-E4	99,979	0.0%	98.1%	1.1%	0.8%	0.0%	0.0%
Marine Corps	E5-E6	37,273	0.0%	91.4%	4.4%	3.5%	0.4%	0.2%
Marine Corps	E7-E9	14,503	0.0%	67.9%	10.5%	16.5%	4.4%	0.6%
Air Force	E1-E4	130,406	0.1%	87.1%	2.9%	2.9%	0.3%	6.6%
Air Force	E5-E6	91,108	0.0%	51.2%	31.7%	14.6%	2.3%	0.1%
Air Force	E7-E9	29,898	0.0%	7.3%	38.0%	37.6%	17.0%	0.0%
Space Force	E1-E4	2,149	0.1%	78.5%	3.5%	6.6%	0.4%	10.9%
Space Force	E5-E6	1,913	0.0%	51.6%	26.5%	17.8%	4.0%	0.1%
Space Force	E7-E9	810	0.0%	12.5%	28.4%	40.2%	18.9%	0.0%

Across all 1,034,157 enlisted members of the five services, 78.0% hold a high school diploma, GED or some college as their highest attainment. In the junior enlisted band — 530,090 members, the population the transition system will process first and longest — the figure is 91.4%.

That is the population the crosswalk hands a list containing Sociologists, Political Scientists and Operations Research Analysts. It is not a story about degrees being rare at the top; it is a story about the absence of any post-secondary credential at all across the great bulk of the enlisted force, in every service, at every grade below the most senior.

The associate's column is where the one visible intervention shows up. Air

Force senior enlisted hold associate's degrees at 38.0%, against 15.4% in the Army, 17.5% in the Navy and 10.5% in the Marine Corps — and Air Force E5-E6 at 31.7% against 7.9%, 8.8% and 4.4%. That is the Community College of the Air Force (CCAF) appearing in federal data: a service-run, degree-granting institution producing a measurable credential effect no other service matches at that grade. It is also the single strongest argument any service has that its enlisted force is civilian-ready — and Part I has already shown the instrument gives that same force the narrowest occupation list of any it serves.

Measurement boundary, and it bears on one service specifically. The Department's own note records that an Army personnel-system migration in 2023 continues to affect Army reporting to DMDC, "including an increase in members with 'unknown' categorization." That shows in the data: the Army *unknown* band rises with grade — 2.1%, 8.4%, 15.8% — where every other service's is at or near zero above the junior band. Nearly one Army senior enlisted member in six has no reported attainment at all, so the Army's 27.5% is computed on a base that is materially incomplete and should be read as a floor. Air Force and Space Force carry their own unknown mass in the junior band (6.6% and 10.9%); the two senior bands are effectively fully reported. The unknown column is published here rather than dropped, because a paper about unmeasured federal data does not get to quietly omit the unmeasured column. Bachelor's-or-higher attainment rises seventeen-fold from the junior enlisted band to the senior one in the Air Force, and rises in every service. Set that beside three findings already established in this Part.

The instrument does not move. On the member's-eye definition, 90.8% of active Air Force families return an identical occupation list from helper to craftsman; 76.5% are identical on the broader measure. Educational attainment across those same grades moves from 3.2% to 54.7%. The occupation list moves not at all.

It offers degree-gated occupations to the least-credentialed band. The All-Source Intelligence Analyst list — served to E1-E4 identically to E7-E9 — contains Sociologists, Political Scientists, Geographers and Operations Research Analysts. Those are bachelor's-minimum occupations and several are doctorate-typical. 87.1% of the E1-E4 population holds a high school diploma, GED or some college; 3.2% hold a bachelor's degree of any kind. The same population averages 23.7 years of age, and three in four are 25 or younger.

And the direction inverts at the top. As established above, the crosswalk *subtracts* occupations at senior grades: a Pararescue manager receives thirteen against the journeyman's twenty-three, with all five Management occupations deleted; an All-Source Intelligence superintendent receives four against the working levels' sixteen; an airborne ISR superintendent seven against sixteen.

The instrument therefore offers the widest range of professional occupations to the band in which 3.2% hold a degree, and the narrowest to the band in which 54.7% do — 17% of them at advanced-degree level. Every figure in that sentence is the Department's own. None of it requires a judgment about whether any individual mapping is apt.

A boundary belongs with it. Educational attainment is not the only gate on an occupation, and a member without a degree today may hold one within the transition window or shortly after — the credentialing apparatus of Part III exists for exactly

that. The finding is not that these members cannot reach these occupations. It is that an instrument carrying no attainability signal for 99.49% of its links is silent about the distance, to a population that is measurably furthest from it.

And the instrument's generosity runs opposite to the credential. Air Force senior enlisted members are roughly twice as credentialed as their Army, Navy and Marine Corps counterparts — 54.7% against 27.5%, 23.6% and 21.0%. Set that beside the cross-service inconsistency established earlier in this Part: the Army's intelligence analyst specialty is given sixty-eight civilian occupations; the Air Force All-Source Intelligence Analyst specialty is given sixteen. The service whose senior enlisted force is twice as credentialed receives roughly a quarter as many civilian occupations for materially the same analytic work.

State that precisely, because it is a specimen and not a regression. 35F against 1N0X1 is one specialty pair, not a service-wide measure of the crosswalk; the attainment figures are service-wide and are not specialty-specific. Nothing here shows the instrument *responds* to credentials — the finding established throughout this Part is that it responds to nothing. The honest form is this: the instrument's per-service generosity varies more than fourfold on one worked pair, and it varies in the opposite direction to measured credential attainment.

And Part IV supplies the other half. Whatever the list says, across the largest assembly of the SkillBridge directory yet built — 4,255 programs at 2,732 organizations — six programs name any of the four Special Warfare fields by occupation code, and exactly one names all four. Translation and market are two separate failures, and they compound.

The human outcome is measured by the VA's own post-separation assessments: the share of veterans saying their skills considerably or completely match their civilian job was under 40% in 2021, under 34% in 2022, under 36% in 2023. Peer-reviewed work (Davenport et al., 2024) estimates roughly one-third of veterans underemployed, with an average earnings penalty near \$24,000. An honest counterweight is on the record: RAND's 2014 Loughran study found skills mismatch is not the durable driver of veteran *unemployment* — a finding about unemployment duration that does not address job quality, underemployment, or crosswalk accuracy, and is not in tension with this trail. A second counterweight belongs on the record, updated to current data: overall veteran unemployment remains slightly below the nonveteran rate (3.6% vs 3.9%, April 2026, BLS). But the aggregate hides the population this system exists to serve: veterans aged 18-24 — the transition-age cohort — show 12.8% unemployment against 7.6% for nonveteran peers, post-9/11 veterans run above the nonveteran rate, and of unemployed veterans, 64% have been jobless 15 or more weeks, with a median duration of 24.6 weeks. The headline rate is fine; the transition itself is where the damage concentrates — which is precisely the window this paper audits.

The first rung of the ladder, and who leaves without it. The education wall this section measures has a companion inside the service's own continuum, and it defines a tier of the separating population no instrument accounts for. Airman Leadership School (ALS) — the first level of enlisted PME — is, by its own syllabus, 192 hours over 24 academic days, with 97 of those hours in Leadership, Supervisory, Problem Solving, and Communication Skills (syllabus effective 23 January 2025, its hour table summing to its own stated 192); independent coding of the ALS curriculum against an external competency framework places 45 of 73 coded references in the framework's Personal domain — the interpersonal-leadership dimension — with

Leadership and Communication the largest sub-competencies (the separate doctoral study cited above). ALS graduation also carries the preferred route to the CCAF degree's six-semester-hour Leadership, Management and Military Studies requirement, in the catalog's own words. And because ALS gates the Staff Sergeant sew-on, the Census Bureau's VEO paygrade file bounds who leaves without it: 43.8% of Air Force enlisted separators in the 2020–2021 cohort — 20,126 of 45,964 — left at E-4 or below, an upper bound on the pre-ALS population (some Senior Airmen complete ALS before separating; the exact figure is a personnel-record fact inside the withheld separations report of Appendix C). The compounding is the point: this tier exits without the continuum's documented leadership layer or its degree credit, into a market requiring a degree on 5,313 of the 7,937 postings in this audit's corpus (66.9%), guided by a crosswalk that offers every skill level the same list, browsing a credential catalog whose management-gated and degree-gated panels carry no tier signal — while the one instrument that favors them, the rank-tiered access cap, hands them the longest SkillBridge window in the force, pointed at the shelf they can least use.

For cyber, a second federal occupational code already exists — and it was built not to reconcile

Everything above treats the crosswalk as though it were the government's only answer to "what is this specialty worth in civilian terms." For one part of the force it is not. The Department operates a second occupational classification of its own, it is binding where the crosswalk is merely advisory, and the two have never been joined.

The instruments. DoD Directive (DoDD) 8140.01, *Cyberspace Workforce Management* (5 October 2020); DoD Instruction (DoDI) 8140.02, *Identification, Tracking, and Reporting of Cyberspace Workforce Requirements* (21 December 2021); and DoD Manual (DoDM) 8140.03, *Cyberspace Workforce Qualification and Management Program* (15 February 2023), which incorporated and canceled DoD 8570.01-M. Together they establish the DoD Cyberspace Workforce Framework (DCWF) — a lexicon of work roles, each with a three-digit code, tasks, and knowledge, skill and ability statements.

The coding requirement is unambiguous and binding: "*All positions requiring the execution of cyberspace work must be coded*" (DoDI 8140.02 ¶1.2.a). The DCWF is "*the authoritative reference for the identification, tracking, and reporting of DoD cyberspace positions.*"

And the instruction states, in terms, that this is a parallel descriptor rather than a translation. Work roles:

"Are used: (a) As an **additional occupational descriptor, along with:** 1. Civilian occupational series. 2. Military occupational codes. 3. Specialty codes." (DoDI 8140.02 ¶3.2.b(2))

"Along with" — not "mapped to." The work role code sits *beside* the specialty code in the personnel record, and nothing reconciles them. The reporting requirement is built the same way: military personnel are reported by "*Occupational category (e.g., designator, military occupation specialty, rating)*" and by "*Work role*"; civilians by "*Occupational series*" and "*Work role*" (¶5.2). Two axes, recorded in parallel, never joined.

The separation is deliberate and is stated twice. DoDD 8140.01 ¶1.2.c provides that cyberspace qualification requirements “*should not be construed to modify, replace, or conflict with General Schedule occupational qualification requirements established by the Office of Personnel Management,*” and DoDM 8140.03 repeats it. The Department is asserting a qualification regime designed not to touch occupational classification.

O*NET does not appear anywhere in it. Fourteen Cyber Exchange documents were obtained and searched in full: the directive, the instruction, the manual, the cyber workforce strategy, the program introduction, the qualification matrix standard operating procedure, the qualification program application procedure, the program FAQs, the 8570-to-8140 transition guidance, and five implementation and waiver memoranda.

“O*NET” and “ONET” occur zero times across all fourteen documents. “AFSC” occurs zero times. The Office of Personnel Management (OPM) information-technology series 2210 occurs zero times. The word “crosswalk” occurs once — in a sentence denying that one exists.

That sentence is the Department’s own, and it concerns the transition between its own two schemes rather than any link to O*NET, which is why it is quoted precisely:

“The DoD 8570 and DoD 8140 programs are not structured the same and **there is no ‘crosswalk’ of qualifications between them.**” — *DoD 8570 IA Program Transition to DoD 8140 CWQP*

and, in the same document:

“There is **no direct comparison, or mapping,** between 8570 IA functional categories and the DCWF work roles.”

So the position for a separating cyber airman is this. Their position is coded in one federal system that binds. They are advised, under a non-waivable service instruction, from a second that does not. They will be hired against a third — an OPM occupational series, or a private job title answering to neither. No published concordance joins any two of the three, and the Department has stated in writing that it does not maintain a mapping even between its own successive schemes.

Two further findings from the same corpus

The contractor bifurcation, in the Department’s own words. DoDM 8140.03 canceled DoD 8570.01-M on 15 February 2023 for military and civilian personnel. It did not cancel it for everyone:

“Contractors remain under DoD 8570 policy until update of the Defense Federal Acquisition Regulation Supplement (DFARS) authorizes DoD 8140 implementation for contractor personnel.”

The DFARS has not been updated. Subpart 239.71 still lists “*DoD Manual 8570.01-M*” among its governing documents, and the mandatory clause at DFARS 252.239-7001 (JAN 2008) still requires contractor personnel to be certified “*in accordance with DoD 8570.01-M*” and “*as listed in the current version of DoD 8570.01-M.*” That subpart carried a DFARS change effective 7 May 2026.

Three years and three months after the manual was canceled, a mandatory contract clause still requires compliance with its “current version,” and there is none. This is not an oversight nobody noticed — the Department documents the gap and attributes it to a pending DFARS update. It is a verification gap of the same species as the rest of this paper: an obligation that remains on the books, binding on paper, with no instrument behind it.

The framework has outgrown the issuances that govern it. The glossary of DoDI 8140.02 defines the workforce exhaustively:

“The DoD cyberspace workforce **is divided into five elements**: IT, cybersecurity, cyberspace effects, intelligence workforce (cyberspace), and cyberspace enablers.”

The Department’s own program introduction says otherwise:

“The DoD 8140 Program unifies the overall DoD cyber workforce and organizes work roles into **seven** Cyber Workforce Elements: Information Technology (IT) · Cybersecurity (CS) · Cyber Effects (CE) · Intelligence (Cyber) · **Data/Artificial Intelligence (Data/AI)** · **Software Engineering (SE)** · Cyber Enablers”

Two of the seven elements have no basis in the instruction’s own definition of the workforce. A reader of the governing issuances would not learn that the Data/AI or Software Engineering workforces exist.

The comparison that closes the section

The 8140 corpus is not an example of federal neglect. It is the opposite, and that is what makes it useful here. The qualification matrices carry a published quarterly refresh schedule — reports due 15 March, 15 June, 15 September and 15 December — set out in a versioned standard operating procedure with a named responsible office. Positions are coded, tracked, reported, and audited against authoritative manpower systems.

The Department can specify a cadence, name an owner, version a document and publish a refresh schedule when it decides an instrument matters. It has done exactly that for the framework that governs its own cyber positions. The instrument it directs separating members to use for the rest of their lives states its update cadence as “irregular,” names no maintainer, and was last quality-reviewed in 2014.

The service’s own credentialing instrument was searched, and it carries no work-role dimension either. An earlier revision of this passage left the question open on the ground that DoD COOL — which publishes credentialing and coding material for the services — was not reachable. That understated what this audit holds — twice over, as correction (41) records. The capture is the author’s own complete pull of the Air Force COOL catalog, rebuilt into a working workbook and normalized into committed seed files: a credential inventory of 848 records with certifying agency, credential family, requirement panel (degree, examination, experience, cost, renewal), accreditation and GI Bill flags, and a specialty-credential master matrix joining 930 distinct credentials to 130 enlisted specialties across 5,285 specialty-credential pairs. Searched in full on 13 August 2026 and re-searched on 2 September 2026 on the current files, that capture returns no occurrence of “DCWF”, no DCWF work-role name, and no work-role identifier, in any field of any file. Nor does any file in the searched holdings map DCWF work roles to military occupational codes in either direction.

That matters because of where it sits. The DoD Cyber Workforce Framework is the instrument the Department uses to code, track and audit its own cyber positions, on a published quarterly refresh cycle. The credentialing portal is the instrument an airman is directed to when they ask what their cyber work is worth outside. Between the framework the Department maintains for itself and the portal it points the member toward, this audit located no join — not in the fourteen governing and program documents, and not in the service credentialing data it holds.

Boundaries. This is a finite-search claim about a specific corpus, not a claim about the world. DoD COOL’s live site was not reachable during this audit and no claim is made that no such mapping exists anywhere; a join may exist on that site, in material behind authentication, or in a system this audit has not seen. The credential capture searched is complete with respect to the Air Force COOL catalog as the author pulled it; the scope counts above are statements about the capture and are published as such. Its records carry their own marking requiring provider verification before external use, so no per-credential claim — a specific credential’s requirements, cost, or flags — is published from it. Three near-misses from the re-search are disclosed so no reader finds them first: one credential’s requirement panel carries a provider-side note reading “DoD 8570/8140 approved” — a program-approval reference with no work-role content; the phrase “work role” appears once, in a geospatial intelligence (GEOINT) credential’s National System for Geospatial Intelligence (NSG) experience requirement, unrelated to the DCWF; and the string “8140” occurs twice as a hexadecimal substring of machine-generated row identifiers. None is a DCWF join. The count of DCWF work roles is likewise not stated in any of the fourteen documents. And the finding concerns cyber; it is not evidence about how any other career field is treated.

The same capture, read for portability. The portal’s design premise is that a credential is the portable form of military experience. For a tier of its own catalog, that premise fails on the program pages of the credentials’ own issuers. Twenty-six credentials in the capture’s availability panels — verified against the issuing programs’ public eligibility pages on 2 September 2026; the evidence table is committed — cannot be purchased or sat for commercially by any member of the public: all eleven National Geospatial-Intelligence Agency (NGA) GEOINT Professional Certifications (“be actively serving in a GEOINT-related work role within the National

System for Geospatial Intelligence... hold an active U.S. Government... security clearance”), the nine-credential DoD SP5D security suite (eligibility is a DoD civilian, military, or contractor security position under DoDM 3305.13), both Counter-Insider Threat certifications (current counter-insider-threat program personnel only), and the Defense Intelligence Agency (DIA) all-source-analyst, intelligence-fundamentals, collection-management and joint-planner certifications (current TS/SCI, in one case with JWICS access, required). A separated member can re-qualify only by re-entering the cleared workforce as a federal civilian or contractor — the credential cannot follow them into the open market.

The exposure concentrates exactly where this paper has already found the transition machinery thinnest. Joined to the capture’s master matrix: for the Geospatial Intelligence Analyst specialty, 15 of the 48 credentials AF COOL lists — roughly one in three — are in this tier; Security Forces 14 of 63 (22%), Special Investigations 14 of 70 (20%), the successor GEOINT specialty 22 of 82 (27%); 25 specialties carry at least one, across 191 specialty-credential pairs. The intelligence and security career fields — the highest-civilian-value populations in the VEO join, facing the emptiest Skill-Bridge market — are also the fields whose credentialing shelf is most heavily uniform-locked. *(Coverage: 932 availability panels, of which 324 remain marked needs-review — the 26 is a floor on the classified-out tier, not a census of it. Separately, of the capture’s 865 curated requirement panels, 132 carry a postsecondary-degree gate — 81 bachelor’s or higher — and 226 state a years-of-experience requirement: the education wall this paper measures in the force reproduced inside the credential catalog. Per-credential figures above rest on issuer pages read this date, not on the capture’s own unverified panels.)*

The catalog’s quality and fit signals, stated as coverage rather than judgment. The portal’s own accreditation flags — third-party accreditation being the single strongest external quality signal a certification carries — cover the catalog like this: of the capture’s 930 credentials, 253 carry an ANSI/ANAB, NCCA, or ICAC accreditation flag; 254 are flagged unaccredited; and for 423 the capture records no determination either way. No judgment of any credential is made or needed: a member browsing the catalog is choosing among credentials whose accreditation status is affirmative on roughly one in four, negative on one in four, and unstated for the rest — and the interface surfaces none of it as a comparison. Fit is defaulted rather than advised, in the same shape Part I measures in the crosswalk: 49 credentials are listed under twenty or more specialties — an entry-level project-management certification under 98 of the 130, a maritime radio license under 60 — and 32 of the blanket-listed credentials draw zero mentions anywhere in this audit’s 61-file, 7,937-posting employer corpus (a lane-skewed corpus, so the zero is an indicator of absent demand signal, not a finding of no demand; the matching is conservative in the catalog’s favor, and the row-level table is committed). And the catalog carries an unstated seniority tier: of the 865 curated requirement panels, 100 require management, leadership, or supervisory experience — on top of the 132 degree gates already stated — credentials structurally shaped for senior noncommissioned officers, listed with no tier signal a junior member could use to know the difference.

(Method: fourteen documents obtained from the DoD Cyber Exchange, text extracted in full and searched locally, 12 August 2026. Documents behind authentication were not retrieved and are not represented here; one obtained file carries no text layer and was excluded from the term counts. Paragraph citations follow each

issuance's own numbering. The Air Force credentialing capture was searched cell-by-cell across all files of the normalized package, 13 August 2026.)

The same service builds a maintained crosswalk. For people coming in.

The defense available to everything above is that mapping military skills to civilian ones is intrinsically hard, done at scale, and cannot realistically be dated, bounded, or conditioned on a human's signature. That defense is answered by a document the Air Force publishes.

The AFSC Convertible Skills List, carried in the Air Force Enlisted Classification Directory (AFECD) and dated 31 October 2023, is a crosswalk between Air Force specialties and Army and Marine Corps military occupational specialties (MOS), Navy Enlisted Classifications (NEC) and Coast Guard ratings. It exists so that a sister-service member can be laterally accessed into the Air Force. It is the same service, the same classification system, and the same question about the same skills — run in the opposite direction. And it carries, on its face, every control this Part has spent forty pages establishing the O*NET crosswalk lacks:

	AFECD Convertible Skills List	O*NET military crosswalk, Air Force layer
row-level change dates	65 of 132 rows (49.2%) , running 30 Apr 2012 to 31 Oct 2023, 15 of them changed in 2023	none published at any level; this audit had to reconstruct the vintage by hashing four archived files
human review before award	17 specialties carry <i>"Must be reviewed by the AFCEM before awarding AFSC"</i>	0.51% of links have ever been human-reviewed
currency of the skill	4 specialties bound it — <i>"Must have performed duty in the convertible skill within the past 18 months"</i> , and 36 months for explosive ordnance disposal (EOD)	nothing asks when the member last did the work
escalation where consequence is severe	3E8X1 EOD routes Guard accessions to <i>"the gaining EOD Flight Chief with squadron commander's endorsement"</i>	no link carries a condition of any kind
destinations per specialty	median 2 , maximum 17	median 5 , maximum 89

The Air Force's own list gives a specialty a median of two destinations, chosen, dated, and in seventeen cases signed for. The instrument that routes a separating Airman gives a median of five and a maximum of eighty-nine, with 119 occupations attached to more than a hundred specialties each. That difference is not explained by

difficulty. Both directions require judging whether skills transfer. One is answered with two carefully chosen codes and a reviewing official; the other with a list.

And it is not a small special case. The obvious rejoinder — that a sister-service conversion table is a narrow instrument and proves nothing about scale — is answered by the officer half of the same apparatus. The Air Force Officer Classification Directory (AFOCD), 31 October 2023, states on its cover “*Supersedes AFOCD 30 April 2023*” — a six-month revision cycle, published, with a supersession chain. More than 190 of its elements carry their own change date. It maintains a negative authorization list: Prefix C publishes the AFSCs for which the qualifier is “*NOT AUTHORIZED.*” And it contains 120 instances of the words “Not used.” — an affirmative record that a qualification dimension was considered and does not apply.

The cadence was then checked against the current edition, and it held. The directory in force is the Department of the Air Force Officer Classification Directory (DAFOCD), 31 October 2025 — 310 pages, OPR HQ AFPC/DPPDW — and its cover reads “*Supersedes DAFOCD, 30 Apr 2025.*” Its own internal change dates evidence the intervening editions rather than merely its endpoints: 61 elements dated 30 April 2025, 47 dated 31 October 2025, 8 dated 30 April 2024 and 3 dated 31 October 2024, across 205 dated elements in total. The negative authorization list survives; the affirmative records of non-applicability have grown from 120 to 123. In the twenty-four months over which this directory was revised four times, the O*NET military crosswalk was revised not at all (the freeze and its tripwire are documented in the provenance note at the head of this Part). Two federal instruments, one service, the same twenty-four months: four editions against zero.

That last control is the one this Part turns on. In the O*NET crosswalk, a link that was carefully reviewed and a link that was bulk-generated look identical, and an occupation’s absence and an occupation’s irrelevance are indistinguishable. The Air Force’s own directories can say *not authorized*, *not used*, and *not qualified*. **The instrument that routes a separating Airman has no vocabulary for “no” anywhere in it** — which is why Broadcast Technicians can attach to 793 specialties without anything in the file objecting.

The concession, stated plainly because the argument fails without it. These instruments do different jobs. The AFECN list and the AFOCD govern the award of a classification code — a personnel action with legal effect — and are gated for that reason. The O*NET crosswalk is advisory and awards nothing. The claim here is not that they should be the same instrument. It is narrower and it is sufficient: the reasons available for leaving the outward-facing instrument undated, unbounded and unreviewed — scale, difficulty, the impracticality of review — are contradicted by the same personnel center doing exactly those things, twice a year, on a larger corpus. The controls follow the direction of consequence. When a bad mapping puts an unqualified member into an Air Force EOD flight, the Air Force absorbs it, and a squadron commander signs first. When a bad mapping tells an aerospace medical technician to consider Customs and Border Protection Officers, the member absorbs it, and nobody signs anything.

(Sources and limits. Both directories were supplied in full on 14 August 2026 and both are unrestricted publications. The Convertible Skills List is a transcription

*of the published list into `afecd_convertible_skills_20231031.csv`, 132 rows; multi-code cells such as 63E/H/J/N/S/W/Y are counted as one destination, so the destination counts understate the AFECD side and the median-of-2 is if anything generous to it. Specialty codes are counted as printed; the O*NET figures beside them are at the same printed-code grain used throughout this Part — 6,827 Air Force codes, 48,042 links — and both medians move if either side is rolled to five characters, which is why the comparison is stated at one grain and labeled. The AFECD list is small by design, covering only the specialties where lateral entry is plausible; it is not a crosswalk of the force and is not described as one, and the comparison above is about how many destinations a covered* specialty receives. 51% of its rows carry no change date and its oldest dated row is from 2012, so the claim is “more maintained than an instrument that publishes nothing,” never “well maintained.” The AFOCD’s dated-element count is published as a floor — more than 190 — for the reasons logged at correction 35. Appendix A of the AFOCD is itself dated 31 October 2022 inside a 31 October 2023 directory. The officer directory was verified current on 22 August 2026 against the DAFOCD of 31 October 2025, supplied by the author and read in full; the figures given for the 2023 edition are retained as published because the comparison is drawn against a document of that vintage, and the current edition’s counts are stated separately above. Note that the directory was renamed in the interval — AFOCD to DAFOCD, Department of the Air Force — which is why a search on the older acronym alone will not find the instrument now in force. The enlisted directory was not verified current: the only other AFECD edition held is 31 October 2020, which is older than the one used, so no statement is made here about whether a later AFECD supersedes the 31 October 2023 list. No figure in this passage depends on the officer education matrices, which were re-extracted from the 31 October 2025 edition on 22 August 2026 — 422 entries across 43 matrices, passing the acceptance tests the first extraction failed (correction 36) — and which are themselves the counter-example inside the counter-example, though not the one this paper first printed. Appendix A still carries “(Changed 31 Oct 2022)” at the appendix level in the 31 October 2025 edition — but nine of the forty-three matrices beneath that banner print their own later dates, from 30 April 2023 to 31 October 2025: 13HX and 15A changed 30 April 2025, 32EXG changed 31 October 2025, 15W, 21M, 32EXE and 43E in April 2024, 13N in October 2023, 38F in April 2023 (a tenth, 19Z, prints an establishment date of 30 April 2020, earlier than the banner). The matrices are revised; the banner over them is not — a versioning defect inside a maintained instrument, which is a smaller thing than the unrevised appendix this paper previously described and a more exact one. The cadence claim is about the directory, not about every element in it, and this is the clearest instance of that distinction in the document.)**

Part II — How the Program Was Built: Fifteen Years, Four Administrations

Every finding in this paper describes the program as it stands today. None of them explains how it came to stand that way. This Part reconstructs the record — statute

by statute, instruction by instruction, contract by contract — because the shape of the failure is legible only in the sequence.

A note on method, stated before the narrative rather than after it. Attribution across administrations is easy to do badly, and this audit has already done it badly once. An earlier internal draft recorded that DoD Instruction 1322.29 was “issued 2019 under Trump.” That is wrong: the instruction was issued 24 January 2014, and Change 1 of May 2020 is housekeeping — the instruction’s own change note says it “updates references and organizational titles and removes expiration language.” Both errors are logged in Appendix D. The rule adopted after them governs this Part: date the initiation, not the signature, and read the instrument before assigning the credit.

The conclusion this Part reaches is not a partisan one, and readers who arrive hoping for one will be disappointed. Four administrations of both parties each added a layer to this program, and not one of them repealed, rewrote, or conditioned the layer beneath it. The pattern is bipartisan, and that is precisely what makes it a structural finding rather than a political one.

II.1 2011: an authority with no price, no data, and no obligation

The provision originated in the Senate. Section 525 of S.1867, the Senate Armed Services Committee’s FY2012 defense authorization bill, became §551 of Public Law 112-81, signed 31 December 2011, and was codified at 10 U.S.C. §1143(e). Its context was acute: post-9/11 veteran unemployment was at its peak, and the VOW to Hire Heroes Act had been signed five weeks earlier.

Read the text and notice what is not in it. The authority is permissive — the Secretary “may carry out” such programs. It establishes eligibility only: a member within 180 days of separation may participate for up to 180 days. There is no appropriation. There is no outcome requirement. There are no obligations on the host employer of any kind. And the phrase that now governs every listing in the marketplace — “high probability of post-service employment” — does not appear in the statute at all.

The zero-score birth defect

The design choice with the longest consequences was made here and has never been revisited. SkillBridge does not spend new money. It repurposes pay the Department is already obligated to issue: the member remains on active duty, drawing the same salary, allowances, and benefits they would draw at their unit. The transfer to the private host is real — Part IV prices it at roughly half a billion dollars a year — but it is invisible to the budget process because it never appears as an appropriation line.

The consequences compound in a chain:

No new money → approximately no budget score → never appropriated → no annual vote → no appropriations subcommittee with a constituency in the outcome → no recurring forum in which anyone must defend the program’s results.

The contrast that makes this legible is the Work Opportunity Tax Credit (WOTC), the tax-code incentive for hiring veterans. WOTC expires. It must be reauthorized. Its reauthorization fights force periodic public argument about whether it works, and that argument generates evaluation. SkillBridge is permanent authority with no fiscal checkpoint. Nothing in its structure ever requires anyone to say what it achieved. The oversight that exists is GAO's, which is advisory, and — so far as a finite search of the public record can establish — no DoD Inspector General (IG) audit of the program has ever been conducted.

This is not a criticism of the 2011 drafters, who were legislating into an unemployment crisis with a tool that cost nothing and could be stood up immediately. It is an observation about what that tool's design made structurally impossible fifteen years later.

II.2 2014: the governing instruction, and a clause that decided everything

Twenty-five months after enactment, the Department issued DoD Instruction 1322.29 on 24 January 2014. Three features of that document govern the marketplace to this day.

First, the standard appears here and nowhere above it. “A high probability of post-service employment” is an *administrative* standard, created in an instruction, not a statutory one. It has never been given a numeric threshold, a measurement method, or an evidence field.

Second — and this is the finding that most repays attention — the standard runs to “the provider or any other employer.” That clause is original design, present from the first day the instruction existed. Its effects are three, and all three are visible in the data of Part IV:

- **A host that never hires anyone can be fully compliant.** The obligation is a *market prediction*, not a commitment. This is why 86.7% of listings can be silent on employment without any listing being irregular.
- **The MOU's later 75% qualifying-offer warranty can be satisfied by offers from anyone** — and no instrument tracks offers from anyone. **Placement-model compliance is unverifiable by design, not by neglect.**
- **Third-party placement was blessed from day one.** The intermediary tier this audit documents — staffing firms, fellowship operators, recruiting funnels holding the MOU while the labor is performed at client companies no public instrument names — is not a loophole anyone found. It is the instruction working as written.

Third, the brand is absent. The word “SkillBridge” appears nowhere in DoD Instruction 1322.29. The name on the website, on the MOU, in the budget justification, and in every press release does not occur in the document that governs it.

Through 2016 the authority was exercised as the Career Skills Program (CSP), run service by service and installation by installation — Army CSP and its analogues — templated substantially on organized-labor apprenticeship arrangements (the union

trades’ Shifting Gears and VIP programs). There was no central marketplace, no directory, and no national scale.

II.3 2019: the marketplace is built on five-year-old text

On 14 June 2019 the Department launched the SkillBridge brand and a central website at dodskillbridge.usalearning.gov. Contemporaneous service reporting described SkillBridge as “formerly known as the Career Skills Program.”

This was a genuine initiative of the first Trump administration and should be credited as one. It is also the single most consequential event in the program’s history, for a reason that has nothing to do with who launched it: it created a national marketplace on governing text written in 2014 for a decentralized program an order of magnitude smaller — and that text has never been substantively updated since.

The requirements this paper grades in Part VII are 2014 language stretched over a 2019 marketplace that now processes roughly 22,000 members a year. The program grew by about a factor of ten and its governing instruction did not change.

The first three years: no procurement record located

USALearning is the OPM’s interagency e-learning platform. The marketplace’s first infrastructure therefore ran on an interagency arrangement rather than a competed departmental contract. The underlying vendor, the cost, and the terms do not appear in the public prime-award record, and are carried here as unverified; the records request is drafted and unfiled.

This extends the procurement finding of Part III backward in time and makes it stronger. The full sequence of the marketplace’s central infrastructure is:

Period	Vehicle	Competitive offers
2019 - c. 2022	OPM/USALearning interagency arrangement	no public prime-award record located (USAspending, searches described in this Part)
Aug 2022 - Jul 2024	General Services Administration (GSA) task order, \$3.37M	one
Jun 2024 - Feb 2027	8(a) sole source, \$15.2M obligated / \$27.5M ceiling	one, “not available for competition”

Not one competitive bid appears in the entire history of this program’s central infrastructure — and its peak growth years are not visible in the federal spending record at all. Every structure named here is lawful and routine; that is the point. The finding is that a marketplace which grew tenfold, and which holds the outcome data providers owe the government, has never once been the subject of a competitive procurement, and its formative period cannot be examined through the normal spending-transparency channel.

II.4 2022-2024: the paperwork arrives, at the bottom of the hierarchy

The conditioning layer arrives under the Biden administration, and it arrives in a revealing place.

- **August 2022** — the first dollars-visible dedicated contract (the GSA website order, above), and the move to skillbridge.osd.mil.
- **December 2022** — §561 of the FY2023 NDAA codifies the program.
- **May 2023** — the Military-Civilian Transition Office (MCTO) assumes oversight.
- **July 2023 — the Department stops accepting new civilian employers**, citing “overwhelming popularity.” Demand outran governance, and the Department’s response was to close the door rather than to build the instrument.
- **December 2023** — §572 of the FY2024 NDAA mandates the GAO review that becomes GAO-24-107352, whose seven recommendations remain open.
- **June 2024** — the successor IT contract, 8(a) sole source, encompassing the systems that hold provider outcome reports.
- **August 2024 — the MOU.**

The MOU is the pivotal structural fact of this program’s development, and it is easy to misread as a strengthening. It is one — it is the first instrument in thirteen years to place any obligation on a host at all: a warranty that 75% or more of completers receive a qualifying offer of employment, a prohibition on provider compensation or gifts to the member, a bar on member-borne costs including certifications and licensure, and an obligation to report 90- and 180-day outcomes including starting salary.

But note where it sits. Thirteen years after enactment, the only enforceable host-side obligation in the entire program lives at the lowest level of the hierarchy — not in statute, not in the instruction, but in an instrument executed provider by provider, enforced if at all through contract administration, and containing no mechanism by which a member, a researcher, or the public could observe whether it is met. The outcome data it requires has been owed since FY2025. None has been published.

II.5 2025-2026: enforcement deadlines, and an approval presumption

The current administration’s layer has three components, and they run in different directions.

The 2025 MOU update is where a three-vintage diff of the instrument isolates a specific change: government agencies gained exemption from the 75% qualifying-offer key performance indicator (KPI). The warranty that is the program’s only performance obligation ceased to apply to federal and state hosts. The 2025 update also added the foreign-entity eligibility provision now raised by several listings in Part IV.

The exemption is not marginal, and it has now been measured. Classifying every advertised program by its own published point-of-contact domain, the union corpus advertises 369,000 person-days of participant time across the 2,836

programs that state a duration — the remaining 1,419 read “Under Review,” and no duration is ever imputed for them. 57,660 of those days, 15.6%, sit with hosts whose contact domain is .gov or .mil: 550 programs at 249 organizations. Six of the ten largest advertised-capacity holders in the entire directory are government bodies — the Cybersecurity and Infrastructure Security Agency, Customs and Border Protection, Naval Information Warfare Center Pacific, Norfolk Naval Shipyard, the Secret Service, and Iowa’s Department of Workforce Development. The program’s only outcome warranty does not bind roughly one day in six of the participant time the directory advertises, and the exemption was written in 2025, after the warranty it removes.

15.6% is a floor, and the reason is visible in the corpus itself. A point-of-contact domain is a proxy for sector, not a legal determination: a contractor may hold a .gov address, and state and municipal bodies routinely do not hold one. Forty organizations whose names read governmental — departments, counties, cities, public school systems, workforce boards — advertise from commercial domains, carrying a further 5,190 person-days; and two organizations, one of them inside the top ten, post some listings from a federal or state domain and others from a commercial one, so a single host is classified both ways at program grain. Counting an organization as governmental if *any* of its listings uses a .gov or .mil domain raises the share to 16.1%. The narrower grain is the one published, because it is the one the directory’s own fields support; the direction of its error is known, and it is upward. This measurement supersedes the preliminary name-lexicon estimate this paper previously carried, which is withdrawn.

Two boundaries travel with it. The MOU does not define “agency,” so which of these hosts the exemption actually reaches is a question about the instrument’s text, not about this corpus — and boundary entities remain unresolved on the face of the document: public schools and colleges, contractor-operated federal laboratories, federally funded research centers, and state bodies that are agencies in every ordinary sense. And this is advertised capacity, not delivered placements. The directory advertises; the audit’s own finding in Part IV.9 is that nothing published permits anyone to enumerate what stands behind a listing. The figure bounds the participant time the design permits these hosts to collect, unwarranted.

The 2026 vintage brings the instruction and service policy into alignment with the Department’s renaming and tightens the member-facing frame: Air Force Instruction (AFI) 36-2671 (31 March 2026) makes directory use a T-0, non-waivable requirement, and the services continue cutting member access — the Marine Corps to three-to-four months in 2024, the Air Force to rank-tiered 120/90/60-day caps in 2026, each citing readiness.

The Secretary of War’s memorandum of 3 June 2026, publicly briefed as the Project Patriot Pipeline, creates a presumption of approval for SkillBridge requests inside the defense industrial base (DIB). Read on its own terms this is a coherent workforce policy: it directs the subsidy toward the sector the Department most needs staffed, and it makes approval more predictable for members heading there — a real benefit, since unpredictable command approval is among the program’s most-complained-of features.

Read against the audit’s data it has an unintended edge. Across the 2,220 defense-industrial-base programs in the union corpus, exactly one names a Pararescueman, Combat Controller, Tactical Air Control Party airman, or Special Reconnaissance air-

man by occupation code — an IT and network-operations program at a DIB contractor whose eligibility list carries the TACP codes in skill-level form (1Z331/1Z351/1Z371), the form an X-form matcher misses (correction 38). The other five of the six programs that name these fields by code sit outside the DIB and therefore outside the presumption; two further DIB programs name the specialties in prose — a joint terminal attack controller (JTAC)/TACP instructor role and a combat-controller experience requirement — without carrying a code. The fields the Department is working hardest to retain are the fields its new approval presumption barely reaches: one code-targeted program among 2,220.

And the Department's own deadlines on the open GAO recommendations fall on 30 September and 31 December 2026 — the first moment in the program's history at which a specific office owes a specific answer on a specific date.

The fifth cycle, in progress. The FY27 House NDAA (H.R. 8800, reported June 2026 — pending legislation, not law) continues the sequence in both directions at once. Its Sec. 771 orders the services to develop recommendations for moving medics into civilian healthcare work, including translating military credentials and expanding SkillBridge accessibility in the healthcare sector — Congress legislating directly at the translation failure Part I documents. And its Sec. 1723 orders a Comptroller General study of SkillBridge, due two years after enactment, covering cross-service differences and — for the first time in the program's legislative history — the *selection of employers* and the *development of contracts*. The host-selection mechanism has finally drawn congressional attention; the instrument it drew is a study, reporting around 2029, while the 2024 GAO recommendations remain open and the outcome-reporting obligation remains unpublished. The study also asks whether SkillBridge should be uniform across the services — a question this audit already answers with data (60-to-180-day windows by service and rank; a 5.9× cross-service access gap for senior enlisted members). And it is thinner than one chamber's paper suggests. The Senate companion, S. 4784, was reported by its Armed Services Committee on 15 June 2026 (S.Rept. 119-127) and read in full for this audit: the word "SkillBridge" appears in it exactly twice, both inside a single section, and there is no Senate counterpart to Sec. 1723 at all — Title X's studies-and-reports subtitle runs to unmanned aircraft, the southern border, drone posture and the Afghanistan War Commission. The Senate's only SkillBridge language is Sec. 549A, the medic-transition analogue of the House's Sec. 771, and its SkillBridge clause asks the Secretary to identify barriers to *access* — the approval timeline, and the requirement to return to the duty station to out-process. Real member burdens, and not one clause about what happens after. Cloture on the motion to proceed failed 50-46 on 14 July 2026, so the Senate has not reached its own bill while the House passed H.R. 8800 216-212 on 22 July.

The first congressional attention ever paid to employer selection therefore exists in one chamber of two, in a bill that has not been conferenced. The pattern of §II.6 holds through a fifth cycle, and holds more sharply for the qualification: access expands by directive, and accountability arrives as homework — set, so far, by half the class.

II.6 What the sequence shows

Attribution, stated with discipline. The statutory authority and the governing instruction are of the Obama period. The brand, the central marketplace, and the scale are of the first Trump term. The conditioning paper and the intake suspension are of the Biden period. The enforcement deadlines, the KPI exemption for government hosts, and the industrial-base presumption are of the second Trump term.

No layer repealed or rewrote the layer beneath it. The program accreted rather than developed.

And the accretion has a shape that no single document reveals, which is why it took a sequence to see it:

Every instrument that expanded access, scale, or throughput was adopted promptly and at a level of the hierarchy commensurate with its effect. Every instrument that would condition the transfer or make it verifiable arrived later, at a lower level of the hierarchy, and without any mechanism by which compliance could be shown.

The statute created eligibility and asked nothing. The instruction created a standard and routed it through a clause that makes the standard unfalsifiable. The marketplace created national scale and left the instruction untouched. The MOU created obligations and put them in the one instrument no member can read and no researcher can check. The GAO mandate created review and produced recommendations with no enforcement attached.

Three further observations belong on the record, each of which the sequence makes visible and no single instrument does.

The brand has no instrument and the instrument has no brand. “Skill-Bridge” governs nothing; DoDI 1322.29 governs everything and does not contain the word. A member, a journalist, or a congressional staffer searching the governing instruction for the program’s name finds nothing.

The contract layer postdates the growth. The program’s formative expansion happened in years for which no prime-award record exists, on an interagency platform belonging to a different agency.

The gap is not a partisan artifact. It survived a change of party in 2017, in 2021, and in 2025. Every administration in this sequence added something real; none of them subtracted anything; and the verification instrument that would have made any of the additions checkable was never built by any of them. That is the strongest available evidence that the failure is structural rather than political — and the strongest reason to think it can be fixed without anyone having to lose an argument.

Part III — The Money and Maintenance Layer

Who maintains the data infrastructure the ecosystem runs on, and under what incentives?

The credentialing vendor landscape

Three primes divide the COOL credentialing ecosystem by function (locked from a direct USAspending pull, 15 July 2026). Solutions for Information Design (SOLID) — roughly 40 employees, Fairfax Station, Virginia — owns the data/design/mapping layer: it built COOL (its own Coast Guard award language states the firm “created” the site), built and operates MilGears (\$10.4M order, 2025–2029), and authored the 2015 pilot study of the enhanced crosswalk. Lifetime prime obligations: approximately \$59.1M. Global Business Solutions holds a larger Air Force COOL purchasing contract (~\$94.8M — buying airmen’s actual credential exams). The historical Air Force vendor’s contract (a large management consultancy; expired 2018) preceded the current arrangement. Competition across these vehicles is frequently single-bid.

The SkillBridge website and the transition office’s systems

The audit extended the money trail to SkillBridge in August 2026, from primary award records:

	SOLID (47QFHA22F0026)	Goldbelt Apex (H9821024CE001)
Scope	“OUSD SkillBridge Program Website Support Services”	“MCTO Transition, Reintegration & Employer Support Information Technology”
Period	Aug 2022 – Jul 2024	Jun 2024 – Feb 2027
Vehicle	GSA schedule task order; funding office the Office of the Under Secretary of Defense (OUSD) for Personnel & Readiness	Direct Defense Human Resources Activity (DHRA) award
Competition	“Full and open” fair opportunity — 1 offer received	8(a) sole source — “not available for competition” — 1 offer
Value	\$3.37M obligated (\$5.07M potential; options not exercised)	\$15.25M obligated; \$27.54M ceiling

Two observations, stated as structure. First, the directory whose defects Part IV measures — the 150-result search cap, free-text occupation fields, unvalidated durations, an export endpoint that times out on a full pull — was the website product maintained under the SOLID order, and every measured defect persists under its successor. Second, neither the prior website contract nor the current transition-office IT contract has ever faced a competitive offer — one bid each, both lawful — and the current contract’s spend rate is roughly triple its predecessor’s while the outcome-reporting system it encompasses has published nothing. The same verification gap that pervades the programs pervades the procurement that maintains them: nothing in either scope, so far as the public record shows, pays for data quality.

Adjacent awards complete the maintenance map: Serco’s \$193M DOL contract for TAP employment workshops (2024–2029; incumbent recompetes, five bidders, protest dismissed); Inverness Technologies operating Army SFL-TAP centers on a series of bridge contracts; Zeiders Enterprises staffing Navy Fleet & Family Support for over 25 years; Patriot Enterprises delivering the VA benefits day. Every layer of the member’s transition experience is contracted; no surfaced vehicle pays for outcomes.

Part IV — The SkillBridge Market

SkillBridge (10 U.S.C. §1143(e); DoD Instruction 1322.29) permits members to spend up to their final 180 days working for a civilian “industry partner” while remaining on full military pay. The host pays nothing. The design premise: employers provide training with “a high probability of post-service employment.” In August 2026 MissionProof assembled the largest capture of the public directory yet built — a 57-jurisdiction state sweep merged with 30 further filter exports; IV.9 measures why no complete pull is possible — and audited it end to end. This Part reports what the marketplace actually contains.

IV.1 Coverage: nobody has enumerated this market, including the Department

The directory’s own interface, captured 8 August 2026, reports two numbers side by side:

Total Positions: 11,532 · Result Count: 5,630

Neither is a count of employers, and the audit’s first task was to find out what stands behind them. Two independent corpora were built.

The state-anchored corpus. The site’s unfiltered export endpoint returns HTTP 504, so the only working method is per-filter. A sweep of all 57 states and territories on 3 August 2026 yielded 3,115 rows, 1,527 programs, 1,103 organizations. That sweep was repeated in full on 8 August and returned exactly 1,103 organizations again — an important control, because it removes drift as an explanation for everything that follows. The directory is stable. It is the *reach* of the method that is not.

The union corpus. Eighty-seven exports were then merged — the 3 August state sweep as one merged file, the 8 August 57-state re-sweep, six per-service exports, a DIB export, and twenty-two job-family partitions — deduplicated on organization, program, city, and state. The result is the largest assembly of the SkillBridge directory yet built:

	State-anchored corpus (3 Aug)	Union corpus (8 Aug)	ratio
Rows	3,115	7,826	2.51×
Programs	1,527	4,255	2.79×
Organizations	1,103	2,732	2.48×

Sixty-three percent of the union — 4,894 rows — was never seen by the 57-state sweep. Only 255 rows were found by the state sweep and by nothing else, which means the national state sweep is very nearly a *subset* of a handful of other queries rather than an independent view of the market. A single DIB query — one filter change, one request — returned 5,032 rows, 2,211 programs and 1,534 organizations: more rows, more programs and more organizations than the entire 53-query national sweep. (2,211 is the program count inside that one export; the union’s DIB-flagged program count used elsewhere in this paper is 2,220 — different denominators, both labeled.) *(A note on that field, which recurs below: DIB is self-attested by the applicant and validated by nothing. 197 organizations — 29.4% of those operating more than one program — carry both True and False across their own listings, including cases where the same host, at the same contact address, appears with both values. Defense-industrial-base status is an attribute of a firm, not of a job, so at least one flag in each such pair is wrong. Every DIB figure in this paper is therefore a count of what applicants entered, not a measurement of what the defense industrial base contains, and is reported as such.)* Eight hundred and eighty-three organizations appear in that one query and nowhere in the state sweep. Section IV.9 explains the mechanism, which is structural and reproducible.

The concentration finding is restated at the grain where it is true. An earlier version of this audit reported that four national chains supply 38% of listing rows. That figure was computed at the state-query grain, where a national chain is returned once per state and is therefore inflated by construction. At union grain the largest single organization holds 196 of 7,826 rows — 2.5% — and the top ten together hold 12.7%. The directory is not concentrated in a handful of employers.

What survives, and should be stated at its own grain, is the member-facing claim: when a member browses a state, the density they see is depot count. A national equipment-rental chain, a materials-handling manufacturer and their peers list every branch location, so a state page reads as a thick market and is partly a map of one employer’s footprint. That is a statement about the browsing surface, not about the directory’s composition, and this paper makes it only in that form.

Organic local coverage is thin, measured on the state-anchored corpus (the corpus for which state attribution is meaningful): stripping chains and remote-only listings leaves Wyoming 16 in-person programs, Idaho 23, Rhode Island and Delaware 25 each, Mississippi, Vermont and Maine roughly 31. The territories are effectively uncovered — Puerto Rico 9 listings and 1 in person; the U.S. Virgin Islands and American Samoa no physical presence at all. Depth tracks installation geography (Texas 1,069, Virginia 975, Florida 923), not member need.

And the member’s own search surface is capped. The public interface returns at most 150 results per query with no paging: a member searching Arizona is shown at most 150 of 629 matches with no indication that more exist.

IV.2 Targeting: a market that doesn’t know who it’s for

Of 4,255 programs in the union corpus, 513 — 12.1% — name a parseable military occupation code. (On the 3 August state-anchored corpus the figure was 184 of 1,527, or 12.0%: the rate is stable to a tenth of a point across a corpus 2.79 times larger, which is the strongest form this finding has ever had.) 2,118 programs — 49.8% — leave the targeting field entirely blank, up from 42% on the smaller corpus; 1,114

(26.2%) say “open to all”; the remainder offer unparseable prose. More than half of all programs say nothing whatever about which military occupations they are for. The 12% that do target concentrate overwhelmingly in the trades: civil engineering, aircraft maintenance, vehicles, cyber, intelligence. Employers who need wrench-turners know exactly which codes they want; everyone else waves at the whole force.

The dark side of the ledger, recomputed on the union corpus with a matcher that reads occupation codes in both their X-form and their skill-level form (corrections 37–38): no program names pharmacy, medical laboratory, or physical medicine; dental draws one program, diagnostic imaging two, munitions two, air traffic control three; six programs name any of the four Air Force special warfare pipelines — only one names all four, and one of the six is inside the defense industrial base. The most certification-portable occupations in the force have essentially no on-ramp — a gap with direct earnings consequences documented in IV.7.

IV.3 Pay: a market with no prices

Sixty-one census programs mention any dollar figure; after removing tuition, fees, and corporate boilerplate, 19 of 1,527 (1.2%) state a credible post-program wage. Re-executed on the union corpus with the same construct — a numeric figure stating what the destination job pays, with benefits, tuition, and program-cost mentions excluded — the classifier returns 60 of 4,255 (1.4%): the pay-silence finding holds at 98.6% across a corpus 2.79 times larger (both known police-work disclosures and the documented pay-bait specimen validate in the output). A looser scan sets the floor beneath it: 85% of union programs contain no pay language of any kind. The median census disclosure: \$62,000. A member commits their one permitted internship with no compensation signal from the official source. Where employer-side pay claims exist, they warrant scrutiny: one documented case pitches interns “\$100,000–150,000” installer earnings while the employer’s own postings guarantee \$18–27.50 per hour for that track, six figures contingent on performance bonuses reached after a multi-year progression (see IV.5).

Six programs advertise a “paid internship” and fourteen offer stipends or housing — facially inconsistent with the framework’s host-pays-nothing rule and the MOU’s “neither compensation nor gifts” clause.

IV.4 Conversion: is there a job behind the door?

DoD does not publish — and, per GAO, the services do not consistently collect — whether SkillBridge leads to employment. MissionProof built the first proxy. Layer one, computed across all 1,527 programs of the 3 August census: 86.8% say nothing about employment outcomes in either direction; 11.5% state hire intent; even self-labeled apprenticeships state it only 10% of the time.

Layer two, a 15-employer stratified verification pilot (URL-and-date evidence per check): five hosts had live paid openings for the trained role; five were unverifiable; five had no discoverable paid analogue. The specimens matter. One cybersecurity training provider pitches a “\$140–160K” certified-assessor career; it has no careers page, and its own site states assessments are conducted by authorized third-party organizations — which it is not: the trained role structurally cannot exist as a paid job at that host. Another host advertises fifteen paid openings in other functions and

zero in the field its internship trains. At a third, the trained title exists on job boards only as the unpaid SkillBridge listing itself.

Four patterns predict “no job behind it,” all checkable at scale: the role exists only as the SkillBridge listing; the employer structurally cannot employ the role; the employer is too small to absorb hires; pitch pay contradicts the employer’s own postings. A fifth emerged from employer job-board capture: the paid-twin test — one host runs a paid four-to-six-week apprentice program for civilians covering the same training its SkillBridge offers unpaid for up to 180 days.

IV.5 The commission tier and the opacity tier

Commission/no-wage destinations exist as a structural tier; its exact prevalence is not yet publication-grade at either grain. At the census grain, 116 programs (7.6%) were flagged — a figure that replication showed includes false positives (the word “commission” in law-enforcement contexts flagged at least four police and federal programs). A corrected-construct union re-run flagged 212 programs (5.0%); independent audit of that hit list then identified residual contamination in both directions (hits carrying explicit wages; missed commission-pay programs). Both figures are therefore carried as indicator-list sizes, not validated prevalence estimates, pending a rebuilt, committed classifier. What survives at full strength is the specimen layer, verified case by case: 32 real-estate-agent programs across 23 brokerages whose members buy their own licenses, plus insurance producers, mortgage-officer tracks, franchise pathways, and recruiting funnels. The member works unpaid, then pays for their own license, then earns commission-only — while the host’s economics improve with headcount regardless of the member’s success. One brokerage’s own recruiting funnel supplies the base rate: “40% of real estate agents didn’t have a single closing last year.”

The sharpest single specimen: a five-person brokerage carrying nine directory listings, including the same sales program posted six times across six duration buckets (1-30 through 151-180 days — whatever window a member has, a listing matches), plus unpaid staff roles servicing the founder’s own business: an “AI & Automation Specialist,” a “Multimedia Content Producer,” and an “Executive Assistant & Talent Acquisition” intern serving as direct assistant to the founder — at a firm with no plausible payroll for any such role. Members must supply their own laptop, phone, and transportation, against an MOU that bars costs “directly or indirectly related to program administration and delivery (e.g., ... equipment...).”

Separately, 266 census programs (17%) are opaque: they name no recognizable role anywhere in their text or provide under 200 characters of description; 28 listings are entirely empty. The mechanical components replicate at union grain — 29 entirely empty and 108 under 200 characters (137 = 3.2%) — while the composite’s role-recognition component awaits its lexicon re-run and is carried at census grain, labeled. A commander approves 180 days of orders against listings a member cannot evaluate. The program office already collects the training plans that would fill this hole — the directory simply does not publish them. Template forensics adds a layer: 198 programs sit in 90 near-duplicate description clusters; six clusters cross organizations, including three nominally independent insurance agencies sharing one listing text, and two Fortune-100 pharmaceutical firms running near-identical program language — evidence of shared program-design intermediaries authoring “employer”

programs.

IV.6 The contractual layer: requirements exist, and the directory is inconsistent with them

Since August 2024, every provider signs an MOU with hard obligations (quoted from a signed federal instance, January 2025 technical update):

MOU obligation (quoted)	What the public directory shows
“75% or higher of Service members who successfully complete the program receive a qualifying offer of immediate post-service suitable employment” (85% KPI goal)	86.7% of union programs silent on employment; pilot found hosts with no discoverable hiring pipeline and roles they structurally cannot employ. And the warranty’s reach is shrinking: the 2025 MOU revision exempts federal, state, and municipal agencies from the 75% KPI, and 15.6% of all advertised participant time sits with hosts on .gov or .mil domains — the full measurement, its floor status, and its boundaries are at II.5; the MOU still does not define “agency,” so boundary entities remain unresolved on the instrument’s face
“Provide neither compensation nor gifts to Service members”	6 “paid internship/training” listings; 14 stipend/housing offers, live on the official directory
Members “do not incur direct financial costs... (e.g., training fees, ... certifications, licensure)”	69 union listings (1.6%) carry member-cost indicators (negation-guarded classifier; hit list committed) — including a training school stating its program “costs \$11,955.00” with “30 Skillbridge slots per class” (a second capture of the same school’s site, 11 Aug 2026, carries a lender-partner page offering “Borrow up to \$27,745 for tuition” and states ten SkillBridge places per class — a financing ceiling and a per-class figure respectively, recorded here because two numbers now exist for one specimen), and equipment/tool-supply requirements at employers up to Fortune-100 scale; 32 realty programs whose members buy their own licenses

MOU obligation (quoted)	What the public directory shows
Provider outcome reporting at 90/180 days including starting salary (collection from FY2025)	No sourced, cohort-resolved outcome dataset was located in this audit’s finite search; the only published outcome figure is DHRA’s FY26 budget-justification assertion of an “85% job-offer rate,” whose source and methodology are undocumented
180-day program ceiling	47 union programs state durations beyond 180 days (19 in the state census) — including federal and state-agency hosts
3-year operating history; minimum annual candidates by size; not asynchronous-only	Unverifiable from the published record — the directory carries none of the fields the MOU regulates

These are consistency findings from the published record, not breach determinations. That distinction is itself the problem: the directory carries no field in which compliance could be demonstrated. The GAO review mandated by the FY2024 NDAA (GAO-24-107352) found the services “don’t consistently collect or analyze data”; all seven recommendations remain open, with DoD’s self-declared completion dates of September 30 and December 31, 2026.

IV.7 Five analytical lenses

The window-weighted market (union grain, 4,255 programs). Applying each service’s own duration caps: a junior sailor (180-day cap) can verify fit with 2,836 programs — 66.7%; junior members of the other services (120 days) with 38.4%; mid-grade NCOs (90 days) with 19.8%; a senior Army or Air Force NCO (60 days) with 278 programs — 6.5%. Identical E-8s in different branches face a 5.9× access gap. A third of the market (33.3%) states no usable duration, capping certainty for every rank.

The unverified labor at stake. 22,000 FY2024 participants (DHRA budget justification) × 75–135 average days × civilian-equivalent Regular Military Compensation of \$70–90K/year values the member labor at \$316–732M per year (central: ~\$506M) — roughly \$26,300 per member per 120-day internship — delivered against no verified job obligation. That estimate now stands on more than the flat basis: rebuilt against the FY2024 separating population service by service, each service’s rank-tiered day caps, and the enlisted separator pay-grade mix the Census Bureau publishes (VEO paygrade file, 2020–2021 cohort), the central case reproduces at \$483–493M — within two percent across three progressively finer builds — and the unit generalizes: every 1,000 additional participants represents \$22.1–22.6M per year of member labor at central assumptions, whatever mechanism produces them (committed model: `presumption_cost_model_v3.py`; scenario grids remain in the standalone analysis and are not findings of this paper). The clearance dimension (52.2% of programs — 2,220 of 4,255 — carry the self-attested DIB flag, concentrated in the intel/cyber cluster) is stated per-unit only, because the participant mix is unpublished

and has not been requested.

The FLSA shadow test. Scored against DOL's primary-beneficiary factors via observable proxies (not a legal determination): 155 of the 1,527 census-grain programs — 10.2% — carry four to five risk indicators: unpaid staff-function roles, productive work displacing paid labor, long on-the-job periods with no educational structure, no hire path. This composite did not reproduce from its written specification — an independent reconstruction returned 1.2% — so it is reported as the output of a construct that has not been independently reproduced, and should not be treated as a validated prevalence (Appendix D). It is retained rather than deleted because the underlying risk pattern is separately evidenced by the specimens; it is the magnitude that is unverified. On that qualified basis, roughly one listing in ten would invite FLSA scrutiny absent the program's statutory umbrella.

The earnings-trajectory join (Census VEO). Joining Census Bureau VEO data (2.8M enlisted veterans, all services) to directory density: the market's deepest coverage (construction, aircraft maintenance) serves mid-tier trajectories (\$38–44K first-year medians). Zero-coverage occupations include both the highest-value veterans — EOD (\$62.4K first-year median, zero programs), intelligence, air traffic control (+60% ten-year growth, two programs) — and the most vulnerable: infantry, the largest single veteran population, landing at a \$33.5K median with essentially no targeted programs; and medical care specialists at \$36.2K with the worst employment-attachment rate in the force. (VEO's non-attachment measure includes GI Bill students and must not be read as unemployment.) The market allocates by employer trade-labor demand — not member need, and not member value.

The wartime squeeze. No stop-loss order is in public effect as of this writing, but the structural lockout map is published: retention-bonus obligations are denial grounds, and the FY2026 Air Force SRB list — cut 73%, from 89 specialties to 24 — retains exactly the fields with the thinnest SkillBridge market: 18 of the 24 retained specialties have two or fewer programs targeting them; the four special-warfare fields, at the highest retention multipliers in the force, are named by occupation code in six programs of 4,255 — one names all four, and exactly one of the six is inside the defense industrial base (the Patriot Pipeline presumption, its reach, and the code-form finding behind that count are analyzed at II.5; correction 38). Meanwhile the services cut member access directly: the Marine Corps to three-four months (2024), the Air Force to rank-tiered 120/90/60-day caps (2026), each citing readiness. The dropped medical specialties — now more releasable — exit into a market that names them in seven programs out of 4,255: none at all for pharmacy, medical laboratory or physical medicine; one each for dental, ophthalmic and surgical service; two each for diagnostic imaging and mental health (*construct and recompute at correction 37*).

IV.8 Across the services

The all-service data yields its own finding. Census VEO outcomes by branch, joined to the directory's service-attributed targeting:

Service	Year-1 median	Year-10 median	Year-1 non-attachment	Programs naming its codes
Air Force	\$48,140	\$70,360	46%	115
Coast Guard	\$42,770	\$65,810	48%	22
Navy	\$42,660	\$62,590	50%	49
Marine Corps	\$37,670	\$64,830	51%	30
Army	\$35,660	\$55,650	57%	44 confidently attributed

(A further 421 parsed code references use the two-digit-plus-letter format shared by Army enlisted and Air Force officer codes and cannot be confidently attributed — itself a data-quality finding about the directory’s free-text targeting.)

Army veterans — the largest separating population — land \$12,500 below Air Force veterans at year one, carry the worst employment attachment in the force, and show the weakest ten-year recovery... and draw no larger a targeted market than the Air Force. The service-level gradient tracks occupational mix (the Air Force’s technical skew versus the Army’s combat-arms weight — infantry alone, at a \$33.5K median, is overwhelmingly Army), which is precisely the point: the transition system’s burden is heaviest where military occupations translate worst, and the marketplace does nothing to compensate. Marine veterans show the sharpest year-one-to-year-ten climb — consistent with young separators entering school and trades ladders — while the Marine Corps imposes the harshest structural access limits (the no-backfill rule; the 2024 duration cuts). The window-cap analysis of IV.7 completes the cross-service picture: identical members in different uniforms face a 5.9× difference in accessible market — and the separating populations those windows serve differ as sharply. The Census Bureau’s VEO paygrade file (2020–2021 enlisted cohort) puts the E7–E9 share of separators at 20.4% in the Air Force against 10.5% in the Army, 12.9% in the Navy, and 5.8% in the Marine Corps: the service that caps its senior tier at 60 days separates the most senior-heavy enlisted population in the force. Transition outcomes, access, and market attention all vary by branch; nothing in the system’s design accounts for any of it.

IV.9 The instrument cannot enumerate its own market

The two corpora of IV.1 differ by a factor of 2.79 in programs. That divergence is not drift, not sampling, and not a defect in the collection: it is a property of the only export path the Department’s own site permits, and it is exactly measurable.

The export endpoint is uncapped for whatever filter it is given, and the state filter matches on the listing’s location field. A listing whose location is blank, or expressed as free text such as “Nationwide (All States),” matches no `state=XX` query. No number of state queries will ever return it. A 57-state sweep is complete with respect to state-anchored listings and structurally blind to the rest.

The proof is definitional rather than statistical. Comparing the 3 August state-anchored corpus against the 8 August single-filter export:

Signal	State-anchored corpus	Single-filter export
LocationStates contains “Nationwide”	11 of 3,115 — 0.4%	922 of 5,032 — 18.3%
LocationStates blank State = “Virtual”	0 — 0.0% 275 — 8.8%	654 — 13.0% 860 — 17.1%

Zero blank LocationStates in the state-anchored corpus is not a sampling artifact — it is definitional. Every row in that corpus is present *because* it matched a state token, so a row without one cannot appear. The blind spot is therefore exact rather than estimated: 2,045 of 5,032 rows in the single-filter export — 41% — are structurally invisible to a per-state sweep.

Three properties make this finding unusually clean. It is mechanical, reproducible by anyone with the endpoint. It is government-side, involving no employer behavior and no judgment about any host. And it is self-inflicted by the instrument that Air Force Instruction 36-2671 ¶4.5.3 — a T-0 requirement, non-waivable — obliges the member to use.

The omission is not random, and it is not what it first appears. The mechanism is recoverable from the export’s own output, and it is a string-matching rule: the state query matches two-letter state codes in LocationStates and nothing else. Not one of the 2,705 unreachable programs contains such a code; 1,504 of the 1,550 reachable ones — 97.0% — do.

Half of everything the sweep cannot reach is therefore unreachable for a single reason: 1,368 programs — 50.6% — spell the state’s name out in words. “Virginia” does not match a query for VA. Eighty-five percent of those programs are delivered in person and eighty-seven percent name the city they sit in. They are not the portable, work-from-anywhere market; they are the local one, and they are concentrated where the market is thickest — California, Virginia, Texas and Florida each lose around a hundred. A member searching their own state is shown a truncated version of their own state. Blank location fields account for a further 27.4%, and genuinely nationwide listings — the case one would assume dominates — for only 12.9%.

The correction matters because it changes what kind of defect this is. “Nationwide listings cannot be state-matched” is arguably inherent to a state filter. “Maryland” not matching “MD” is a normalization bug, and it is fixable in an afternoon by matching state names alongside codes, or by normalizing the field on intake.

Set beside the 150-result cap, the picture completes. The member-facing search shows at most 150 of, in one worked case, 629 matches. The analyst-facing export omits every listing whose location field lacks a two-letter code — half of them in-person programs in named states. Neither surface can enumerate the market, and they fail in different directions. A member cannot see the whole; an analyst cannot either; and the Department publishes no dataset that resolves it.

This changes what every population figure in this Part means, including our own. The 1,527-program, 1,103-organization corpus is not wrong — it is a valid and nameable population, *state-anchored listings as of 3 August 2026* — and the audit’s earlier findings were each reached against real evidence. But it is not a census, and this paper does not present it as one. Where a figure has been recomputed on the union it is reported on the union with the denominator named; where it has not yet

been recomputed, that is stated (IV.2, IV.3). The finite-search rule this audit applies to the Department applies here to the audit itself: every population statement below describes the corpus searched, and the corpus is named.

There is a second-order finding in this, and it is the one that should reach an oversight committee. A filter is supposed to narrow a result set. Changing one dropdown from one program type to another surfaced 883 organizations that a complete national sweep of every state and territory never returned. Whatever the directory's internal data model is, the public surface does not behave as a filtered view of a single population — and no published documentation describes what it does behave as.

The blind spot has a host side

The export defect measures what the directory hides from members. A capture from one DIB prime measures what it hides about hosts: the employer's own careers board carries 54 unique SkillBridge internship requisitions — formal applicant-tracking-system (ATS) postings for its approved military internship program, spread across named installations and locations — while the union corpus contains one directory listing for the same organization. The directory's one-card presence understates this host's actual SkillBridge intake surface by a factor of fifty. Favorable reading first: nothing requires a host to enumerate its internship slots in the directory, and running SkillBridge intake through a professional ATS with formal requisitions is arguably better practice than a static directory card. But the measurement consequence is direct: the directory is not merely blind to organizations its export cannot reach — 1,629 of 2,732 union organizations, 59%, appear in no state-sweep export — it is blind to scale within the organizations it does show. Every concentration, coverage, and per-host figure computed from the directory — including this audit's — is therefore a floor on host-side activity, not an estimate of it. No public instrument states how many members any host takes; this capture shows one host's answer is at least fifty-four at once, which is also what an uncapped entitlement looks like in requisition form.

The blind spot has a corporate-structure side

A third measurement problem sits underneath the first two, and it concerns the unit of analysis rather than the count. The directory's atomic record is a card for one organization. Hiring, at scale, is not organized that way.

One SkillBridge host — a diversified defense conglomerate — recruits through four separate systems on three applicant-tracking vendors, and holds one directory card. The first system carries eighty-five requisitions across ten named locations and the host's entire SkillBridge presence: a generic program application and one structured twelve-week supervisory role whose occupant would mentor and assess other SkillBridge participants. The second, a different vendor, carries twelve requisitions and not one occurrence of "SkillBridge," "veteran," "military" or "prior service" anywhere in the capture. The third and fourth, on a third vendor as two unrelated tenants, carry six and two requisitions and no SkillBridge language either — the only appearances of the word "veteran" across both are in equal-opportunity boilerplate.

Four systems, three vendors, one directory card — and the SkillBridge presence exists on exactly one system, attached to exactly one site.

Favorable reading first, and it is a strong one. Subsidiary recruiting estates are ordinary corporate practice, not evasion: acquisitions arrive carrying their own vendors, and consolidating them is expensive work no one is obliged to do. More to the point, a SkillBridge program *should* attach to the operating unit that can actually host and supervise a participant rather than being advertised across affiliates that cannot. Nothing here is a criticism of the host — and the supervisory posting is, on its face, better designed than most of what this audit has read: time-bounded, explicitly developmental, with a named mentoring function and no overclaimed outcome.

And there is a second reading of that same posting, which the favorable one does not dispose of. The role is a SkillBridge participant supervising, mentoring and assessing other SkillBridge participants. Near-peer mentoring is legitimate, often effective, and nothing in the posting's construction suggests otherwise. But it means a program-delivery function is being performed by a participant rather than by an employee — and program delivery is the provider's obligation, not the member's. The August 2024 MOU binds the provider to deliver the training. It does not contemplate that the supervision and assessment of that training might itself be performed by someone the provider is not paying.

This sits directly on an indicator this audit already scores. The FLSA primary-beneficiary screen of IV.7 flags *unpaid staff-function roles* as a risk marker; a member who mentors and assesses other members is performing a staff function by any ordinary reading. Whether it displaces a paid position is precisely the question the directory carries no field to answer — and the same is true of the unpaid "Executive Assistant" and "Talent Acquisition" roles documented at IV.5. A program with no cost to the host and no verified outcome can, in principle, be administered by the labor it receives.

Boundaries, and they matter more than usual here. This is one posting, at one host, on one date. It is not established that the participant is the sole supervisor rather than one layer inside a paid structure, and no impropriety is alleged against this host — the posting is better built than most in the corpus and this audit says so above. What the specimen establishes is narrower, and still worth recording: the delivery of the program can be staffed with the labor the program supplies, and no public instrument would reveal it if it were.

The measurement consequence is nonetheless direct. A transitioning member who finds this host in the directory and then does the obvious next thing — searches that company's openings — reaches one system out of four, and nothing in the stack tells them the other three exist. The directory does not model corporate boundaries. The private re-indexers inherit the same grain. And a member has no way to learn that the entity named on the card is not the entity doing most of the hiring.

Part IV.9 has now measured the instrument's blindness in three directions: to organizations its export cannot reach, to scale within the organizations it does show, and to the boundary of the organization itself. Only the first is a defect of a query parameter. The third is a defect of the data model.

Scope, stated plainly: one host, one date, four recruiting systems examined directly, each captured to a hashed file — two of the four being a supplied page capture and a verbatim transcription respectively, labeled as such in the manifest rather than

presented as full-page dumps. This establishes that the grain mismatch exists and is measurable. It does not establish how common it is, and no prevalence claim is made from it. Two further observations about this host — that its own corporate entity picker names additional operating units, and that one subsidiary careers URL was observed not to resolve — rest on screenshots that were not archived to file, and neither is carried into the finding above. The audit’s finite-search rule applies: this is a specimen, named as a specimen.

IV.10 The host directory is frozen too, and the live copy is unarchived

The method that measured the crosswalk’s currency was applied to the SkillBridge directory itself, and returns the same shape.

The directory is rendered in the browser from a published data file rather than served as static pages, and that data file has been captured. `organizationsData.js` on `skillbridge.osd.mil` has thirty captures between October 2021 and March 2026, carrying nineteen distinct content versions — each a structured record per program, with organization, opportunity type, delivery method, duration, states, job family and nationwide and online flags. Program-listing counts from those captures: 1,341 (30 Oct 2021), 3,312 (20 Jun 2023), 4,937 (29 Sep 2024), 5,454 (23 Apr 2025).

Two properties of that series matter to this audit, and neither is a churn measurement.

The .mil copy stopped changing in April 2025. The captures of 23 April 2025, 10 August 2025, 29 January 2026 and 25 March 2026 carry a single content digest: the file is byte-identical across eleven months.

And the live directory is no longer served from .mil at all. The current opportunity directory is served from `sb-learn.azurewebsites.us` — a commercial cloud tenancy under a hostname that is neither `.mil` nor `.gov` — and a capture query against that domain returned zero results on 12 August 2026. The federal directory of record has migrated to a host no public archive has captured.

The consequence is the same as the crosswalk’s and it compounds: no published instrument states the directory’s currency, and the only third-party record of what the Department published, and when, ends at the domain change. A member cannot establish whether what they are reading is current; neither can an auditor; and after the migration, neither can the archive.

No survival or churn figure is reported from this series, and the reason is a defect in the source rather than in the method. The directory’s program field concatenates organization name with program name and carries no entity identifier, so renaming a program presents as an organization leaving. Between September 2024 and April 2025 the series shows 478 apparent exits; 359 of them are the same organization under a changed program string, leaving 119 candidates — an apparent exit rate inflated roughly fourfold by a naive match. Survival analysis on this series requires hand adjudication that has not been performed, and this section reports the series’ existence and the freeze, not its churn.

Part V — The Demand Side: What Actual Job Postings Show

The directory and the programs are supply-side. MissionProof’s parallel corpus is demand-side: employers’ actual public requisitions, collected with per-posting evidence — employer, requisition ID, URL, observed date, and a content hash on every row (the full archive is inventoried and SHA-256-hashed in the capture manifest; Appendix B). Three bodies of evidence anchor this Part: a seven-corpus requisition audit spanning six industries — 4,563 unique requisitions across a defense prime, an enterprise-technology firm, a cleared-services prime, a consultancy, two hospital systems, and a Class I freight railroad — a 629-row market-anchor evidence base (30 July 2026 build) spanning federal and contractor postings for four Air Force exit lanes (since expanded to 21,086 specialty-lane rows across 92 specialties, at the changed grain V.2 states), and a 22,000-row postings corpus. Each requisition corpus carries an explicit scope status under the audit’s completeness rule (a count is not a census; the word “book” requires a receipt): the statuses below were earned by dedup denominators, board-total checks, and collector attestations, and two earlier completeness claims were withdrawn and logged when receipts contradicted them.

V.1 The requisition audit: seven employers, five industries

The corpora, with their earned scope statuses:

Employer (descriptor)	Industry	Unique requisitions	Scope status
A top-20 defense contractor	defense	384 (400 instances)	census — dedup receipt
An enterprise-technology firm	technology	1,250	complete book — attestation + instance/unique receipt (raw captures resolve to 1,250 unique requisition IDs; an earlier tally of 1,256 is corrected to the raw-supported count); includes international postings
A cleared-services prime	cleared services	544	~22% slice of a ~2,500-requisition board, likely category-scoped

Employer (descriptor)	Industry	Unique requisitions	Scope status
A management & technology consultancy	consulting	1,022	multi-channel reconstruction (careers site + LinkedIn, to each channel's cap), cumulative over two dates
A regional hospital system	healthcare	126 (139 instances)	multi-category, multi-state capture; completeness unclaimed
A multi-state hospital system — a SkillBridge host with three directory programs	healthcare	1,152	attested complete
A Class I freight railroad (non-participant counter-case)	freight rail	85	complete book — 85 of 85

Two methodological notes before the findings. First, the consultancy's corpus is itself a finding about enumeration: its careers board refused pagination past page five on the capture date (an HTTP "Not Acceptable" error, reproduced across two browsers on 9 August 2026; the behavior did not reproduce the following day — an intermittent defect, which still defeats reliable enumeration but is dated accordingly), and LinkedIn's vertical search hard-caps at 1,000 results — so no single channel can enumerate the board, and the 1,022 was reconstructed as a union of capped channels, exactly the method Part IV.9 documents for the government directory. The inability of a search surface to enumerate its own answer is not a government defect; the audit has hit six distinct failure modes across five platforms with five different owners (the government directory's full-export failure and its 150-result cap; a commercial career-site paginator that skips from page 1 to page 9; the consultancy board's page-six refusal; the 1,000-result platform cap; and a defense prime's keyword search whose URL pagination silently discards the keyword, returning the unfiltered board — 115 pages where the filtered answer was two — while the on-page "next" control behaves correctly). The counts derived from each capture are reproducible from the committed files; the error behaviors themselves rest on dated collector attestations, not archived raw responses — replaying them requires the live sites. Second, favorable readings apply throughout: nothing below alleges misconduct; these are lawful postings on lawful boards, and every silence documented has innocent explanations available — ATS defaults, recruiting-channel conventions, decentralized posting authorship. The findings concern what a transitioning member can and cannot see.

The worked example. The audited defense employer brands itself a “top veteran employer” — the phrase appears across its postings roughly 379 times. Across the same 384 requisitions, the number of references to any military occupation code, AFSC, MOS, or to SkillBridge is zero. (Seven whole-word “MOS” tokens all resolve to a NOAA program acronym.) The translation burden sits entirely on the member — from both directions. The government’s crosswalk is unreviewed since 2014; the employers who most loudly court veterans provide no crosswalk at all.

The deeper irony is that the employer’s demand is *already military-shaped* and simply unlabeled: DoD 8570/8140-driven certification language appears on 28 requisitions; Security+ on 54; CCNA on 22; CISSP on 16 — exactly the credentials the force issues — and 36.5% of the book is gated on a security clearance — 44.5% if Public Trust determinations are counted with them, which this paper does not, a Public Trust being a suitability determination rather than a clearance — meaning more than a third of the demand structurally prefers the cleared, transitioning population.

Two walls stand behind the friendly branding. Degree gates: 91 requisitions (24%) require a four-year degree with no experience pathway — a hard wall against the enlisted majority regardless of skill. And pay opacity: only 28% of requisitions carry complete pay bands, disclosed almost exclusively where state law compels it — the same silence the SkillBridge directory shows at 1.4% (union grain, re-run classifier), operating across the private market wherever disclosure is voluntary.

The cross-corpus pay spread is the section’s sharpest finding. Where disclosure is unconstrained, it collapses: 27.9% at the defense prime.

And where these employers do disclose, the number does not price the job. Two of the cleared-market books were re-examined for what the disclosed range actually is. At one prime, 535 requisitions carrying a pay range resolve to fifteen distinct ranges; five of them cover 73.6% of the book, and the median band is \$59,325 wide. At a second, 228 postings likewise resolve to fifteen distinct ranges, five covering 71.9%, median band \$98,400 wide — a spread larger than the median American household income, offered as the answer to what a job pays. These are corporate level bands stamped across unrelated occupations, not role prices: the same range attaches to a software engineer, a systems engineer and a network administrator.

The field is populated, the pay-transparency statute is satisfied, and a member still cannot learn what the job pays. That is the same shape as every other finding in this paper — a form completed, a requirement met on its face, and no verified information at the end of it. It also bounds this audit’s own reach: **these corpora can show whether an employer discloses, and cannot price a clearance, a skill or a specialty,** because the disclosed number varies with job grade rather than with the work. Where an employer chooses disclosure, it approaches totality: 98% across the cleared-services slice; 100% across the railroad’s complete book — from the two industries whose competitive and contractual “constraints” are most often invoked to excuse silence, and within the defense sector itself the pair sits at 98% versus 28%. And where law removes the choice, disclosure is universal at the capture date: the multi-state hospital system posts a wage range on effectively 100% of its 1,152 requisitions, in jurisdictions whose 2024 pay-transparency statutes require exactly that. The tiers close the argument: silence is a choice; the choosers prove disclosure practicable in

the “constrained” industries; and the mandated tier shows compliance operating in practice — the answer to any claim that the directory’s disclosure recommendation (Part IX) is impractical. (An earlier draft carried a fourth cell — a 5% rate for a second railroad’s LinkedIn sample — withdrawn on independent audit: the derived table had captured only pay ranges, and the underlying cards carry employer-authored single-rate pay statements on 35 of 42 postings. Corrected in Appendix D; the rail comparison now rests on that railroad’s complete book only.)

The incumbent government tool fails precisely where this book lives: O*NET’s crosswalk routes 223 AFSC variants to a single network-administration occupation, 655 variants to “General and Operations Managers,” and maps only nine variants to Project Management Specialists — while project management is this employer’s third-largest job family. The tool cannot see clearances, cannot rank matches, and misses the employer’s core lane entirely.

The counter-case: the employer doing the most is not in the program

One corpus in the table is a non-participant, examined because its absence is the finding. A Class I freight railroad — unnamed here, as every private employer in this study is — self-reports thousands of veterans hired over two decades and roughly one employee in five a veteran; it tags a majority of its postings with a structured “Transitioning Military” category field — the exact structured-field fix Part IX recommends to the Department, implemented by a freight railroad without being told — and it discloses pay on 100% of its book. No SkillBridge listing for this railroad, or for any Class I railroad, was located in the 2,732-organization union corpus, and a 250-posting, 47-company rail-industry sample shows the structured tag is this railroad’s practice alone, not an industry convention. The favorable reading is almost certainly the true one: craft positions are governed by collective bargaining and seniority bids; safety-sensitive rail work is federally certified; and the railroad simply hires people, running paid trainee pipelines. That reading is itself the finding: what the program offers a host — up to 180 days of a member’s time at no cost — is worth least to the hosts that were going to hire anyway. The employers most committed to hiring have the least reason to participate, which is why the program’s verification burden cannot be carried by host intent — the argument arriving from an employer’s revealed behavior rather than from theory. (Full rail analysis, including the failure of the railroad’s platform-supplied military-code search tool in four queries, is documented in the committed rail files. The descriptor here remains thin anonymity — the Class I population is small and the practices described are public — but no allegation attaches to this employer anywhere in this paper; every statement about it is favorable and verifiable from its own public materials.)

V.2 The veteran-language census: 22,000 posting rows, and the silence scales

The requisition audit’s zero-references finding generalizes. Across Mission-Proof’s full postings corpus — more than 22,000 posting rows spanning over 300 distinct employers (the 21,086-row market-anchor evidence base across 269

employers, in its expanded build of 14 August 2026; 1,541 hand-collected private-market listings across 146 employers in the data-center, building-automation, and enterprise-technology trades — collected through two channels: employer careers sites (including the 1,250-listing enterprise-technology book) and LinkedIn capture sessions, per the capture manifest; 192 cleared-intelligence postings across 96 employers; and the 384-requisition defense book) — veteran-directed language beyond legally driven equal employment opportunity (EEO) boilerplate is the exception, and military-occupational language is nearly absent. These are posting rows, not distinct requisitions, and the gap between the two is large enough that stating only the larger number would overstate the corpus. The market-anchor component is specialty-expanded: one requisition occupies one row for every Air Force specialty lane it anchors, so its 21,086 rows resolve to 4,108 distinct postings — 4,058 carrying a requisition identifier and 50 carrying none — and a single federal-services employer’s 109 requisitions supply 9,653 of the rows, by themselves 46% of the base. Summed across the four components the corpus is roughly 6,200 distinct postings, and that figure is an upper bound rather than a de-duplicated count, because the hand-collected census carries no requisition identifier and cannot be matched against the others. The row count is the denominator for the coding rates below — each row was coded — and the distinct count is the denominator for any claim about how much of the market this corpus sees. The rates themselves are computed within components and are unaffected:

- Of 1,541 private-market listings, 11% carried any veteran-directed signal beyond standard boilerplate as coded at collection. The signal is not evenly thin — it concentrates where an industry has learned to want veterans: roughly two-thirds of data-center construction and operations postings carry one (hyperscale cloud operators and colocation providers actively courting veteran technicians), against 44% in building automation.
- The most striking cell: of 1,250 postings from the enterprise-technology employer that operates one of SkillBridge’s most celebrated training academies, roughly one in a thousand carried a veteran signal beyond boilerplate. The company’s veteran commitment is real and famous; its other 1,249 job postings do not know it exists. Veteran-friendliness lives in the branding layer and the special program — not in the hiring instrument a transitioning member actually reads.
- In the cleared-intelligence market — 192 postings from 96 employers whose core hiring requirement (an active clearance) is grown almost exclusively in uniform — zero postings used military-occupational language in body text. No employer performed any occupational mapping of its own postings toward military codes. The market most dependent on veterans speaks the least military.

(Limitation: the private-market veteran-signal coding rule derives from the original hand-collection instrument, and its boilerplate threshold has not been independently re-confirmed against the collection notes. The 11% is therefore a coding result whose threshold is not fully specified, and it should be read as indicative of the order of magnitude rather than as a precise rate.)

The census closes the loop the requisition audit opened: the translation silence is not one employer’s quirk. It is the market norm — at trade employers, at cleared

primes, and at the most veteran-celebrated technology companies alike. Employers buy veteran-sourcing services, run veteran programs, and brand veteran-friendliness, while the postings themselves — the only instrument a searching member actually encounters — say nothing a military job code can find.

V.3 The market anchors: what the jobs actually pay, with receipts

The stay-vs-go analysis rebuilt its market curves under an evidence standard worth stating, because it is the standard this whole paper argues the government should adopt: no plotted point without at least three distinct employers and five postings, no employer supplying more than half a point’s evidence, every point traceable to posting-level URL, date, and hash — and any market lane that cannot meet the floor is not drawn, replaced by an in-chart statement that absence means uncollected, not zero. Twelve postings were excluded with named reasons and 195 duplicates collapsed by hash in the 30 July 2026 build the charts are drawn from; the expanded build of 14 August 2026 cited in V.2 excludes 256 rows on the same band-ratio rule and collapses 197 duplicates; every assumption (notably the seniority-tier-to-years mapping — postings state pay and grade but never years of service) is carried visibly as an assumption.

Under that standard, the federal-lane findings for four Air Force exit lanes are set out below. These medians are provisional: they rest on the evidence floor stated above and have not been replicated against an independent postings collection, and the cleared-contractor lane is self-suppressed for insufficient multi-employer coverage.

Exit lane	Year-4 median	Year-10 median	Year-18 median	Verdict vs military compensation
Cyber defense (1D7X1)	\$59K	\$97K	\$149K	market passes the average airman ~year 9
Intelligence (1N0X1)	—	\$103K	\$134K	closest to parity; passes mid-career
Security forces (3P0X1)	\$58K	\$83K	\$89K	never catches military comp before 20 years
NDI (2A7X2)	evidence floor not met			verdict holds on thin evidence — honestly flagged

Two findings matter beyond the numbers. First, the honest advice is sometimes “stay”: for security forces the market midpoint never crosses military compensation before retirement eligibility — and the non-destructive-inspection lane points the

same way on evidence below this Part’s own floor, so no verdict is drawn for it — a conclusion almost no transition marketing will utter, and the counter-story that builds member trust. Second, an unlabeled curve can be argued to any verdict: the same cyber AFSC yields three different crossover stories depending on whether the curve describes the federal lane, the single-employer contractor sample, or an unlabeled blend. Market lanes must be declared — the demand-side version of the taxonomy problem that pervades the SkillBridge directory.

External cross-validation closes the loop: the only two police-work pay disclosures found in the entire SkillBridge directory (\$60,000 and \$83,868) bracket the 3P0X1 federal-lane curve almost exactly — two independent evidence streams, one from the directory, one from the postings corpus, agreeing.

V.4 The two-sided silence

Set the demand corpus beside the SkillBridge directory and the symmetry is exact. Employers say nothing in military language (zero code references across 384 requisitions); programs say almost nothing either (88% of SkillBridge listings name no code or say “everyone”). Employers disclose pay under legal compulsion or exceptional choice — 28% where unconstrained, 98-100% where chosen or mandated (V.1) — while programs almost never do (1.4% on the union corpus, re-run classifier). Employers brand veteran-friendliness 379 times without one act of translation; the Department brands a translation service resting on a 2014-dated quality layer. The member stands between two systems, each assuming the other did the work of translation, while neither did. That gap — measurable from both sides, as this Part and Part IV measure it — is the verification gap in its purest commercial form, and it is where independent, receipts-based translation earns its keep.

Part VI — The TAP Ecosystem

SkillBridge sits inside TAP: the mandatory, interagency curriculum every separating member must complete. Mapping that ecosystem end to end reveals the same architecture at larger scale.

VI.1 The machine

The statutory core (10 U.S.C. §§1142-1144; mandatory since the VOW Act of 2011; redesigned by the FY2019 NDAA) requires initiation no later than 365 days before separation, a three-day interagency core (DoD Transition Day, VA benefits, DOL employment fundamentals), tiered counseling pathways, elective two-day tracks, and a commander-verified capstone. Governance is a committee of everyone: a TAP Executive Council and Senior Steering Group tri-chaired by DoD, VA, and DOL with the lead rotating annually — a structure in which no single official owns transition outcomes. CRS states it plainly: “CRS is unable to estimate the total costs of the program since the military services do not consistently include TAP funding as a line item in their

budget requests” (R48114, Kamarck, June 2024). RAND counts the broader ecosystem at 45 federal employment-transition programs across 11 agencies, \$14.3 billion (FY2019) — and CRS notes veterans report difficulty even identifying which programs apply to them. Two more CRS facts frame the stakes: more than 80% of those who enter active service separate before retirement eligibility — the transition system is the majority experience, not an edge case — and demand for SkillBridge was strong enough that DoD stopped accepting new civilian employers in July 2023 due to “overwhelming popularity,” even as the services cut member access. GAO found services “waived many servicemembers’ attendance at workshops, sometimes in violation of the services’ own policies.”

VI.2 What the evidence says about effectiveness

DOL’s own impact study (2023) — the only impact evaluation of the mandatory program in its 35-year history — found TAP participants employed faster but with no wage benefit: employed participants earned less than non-participants at every measured timepoint. RAND (2024) reviewed the full \$14.3B ecosystem and found “limited evidence that federally funded employment transition programs are effective.” No randomized evaluation of TAP or SkillBridge has ever been conducted. The single rigorous positive result anywhere in the ecosystem is a quasi-experimental evaluation of one nonprofit certification program (the Institute for Veterans and Military Families’ Onward to Opportunity program, IVMF O2O: +\$7,000 starting salary) — commissioned by the program’s own funder.

Compliance with the program’s own rules is chronically partial: GAO found only 25% of members starting TAP on time (2022), improving to 52% by 2025; roughly a quarter of the highest-need members skipped mandatory classes; 55% of eligible members received waivers from mandatory training (testimony of the Veterans of Foreign Wars, VFW, 2025); special operations members started on time 39% of the time. Members describe a check-the-box culture; commanders documented pulling members from classes for operational needs.

VI.3 The intermediary layer and the verification hole

A shadow ecosystem of nonprofits operates inside and alongside the official programs: Hiring Our Heroes (Chamber of Commerce Foundation) runs fellowship cohorts inside SkillBridge and testified to Congress to an 89% offer rate and \$107,000 average salary — and that “TAP is not an employment assistance program.” IVMF’s O2O has served 100,000; American Corporate Partners (ACP), FourBlock, Helmets to Hardhats, USO Pathfinder, and state programs (Texas Veterans Commission at \$46M/year; Veterans Florida) round out the layer. None of the intermediaries’ outcome claims can be verified against federal data, because — per GAO — the federal data is not collected. The best-marketed numbers in the ecosystem are structurally unfalsifiable. The oversight record completes the picture: ten GAO engagements since 2017, every recommendation advisory; no DoD Inspector General audit of TAP or SkillBridge was located — oversight of a \$14B ecosystem is GAO-only. The FY2026 NDAA finally mandates annual TAP tracking and reporting — a new public compliance data stream beginning now.

Part VII — The Compliance Audit: Grading the System Against Its Own Instructions

VII.1 Grading the transition instructions

Military instructions are written in compliance grammar: “will,” “shall,” and “must” bind; “should” and “may” advise. Inspectors general audit by checklists derived mechanically from those statements. Nothing about the method is privileged — only internal evidence is. MissionProof therefore ran the IG’s own method against public evidence.

The corpus. Six governing documents — DoDI 1332.35 (TAP), DoDI 1322.29 (SkillBridge), AR 600-81, AFI 36-3037, MCO 1700.31A, and the Navy instruction — parsed mechanically: 870 modal statements, 640 of them binding.

The graded matrix (v1, fifteen requirements, verbatim quotes and paragraph citations). Every requirement gradeable today from public evidence grades NOT MET or NOT DEMONSTRABLE — none grades MET. Selected rows:

Binding language (quoted)	Public evidence	Grade
“Transition Assistance must start no later than 365 days prior...” (DoDI 1332.35 ¶5.1)	25% on-time (GAO 2022) → 52% (2025); SOF 39%	NOT MET
“A DoD-Enterprise single-source data collection framework will be instituted throughout DoD TAP...” (¶4.1)	GAO: services “manually retrieve data from separate data sources”	NOT MET
SkillBridge programs “must offer a high probability of post-service employment with the provider or any other employer” (DoDI 1322.29 ¶1)	One-third of pilot hosts with no discoverable paid analogue; 87% of listings silent	NOT MET (documented specimens); NOT DEMONSTRABLE at approval layer
Providers “must also offer enrollment at no or minimal cost” (¶3)	Member-paid fees and equipment requirements in live listings	NOT MET (listing layer)
“Service members must meet all CRS, or receive a warm handover...” (¶5.2)	GAO: DoD cannot verify handovers occur	NOT DEMONSTRABLE

The method extends to the instrument that governs information quality itself, and yields two more gradeable rows — seventeen in all. DoD Instruction 8170.01 — *Online Information Management and Electronic Messaging*, Change 2

effective 12 March 2025, which incorporated and canceled the Department’s 2003 information-quality memorandum — carries the Department’s public-information quality standard at Appendix 3A. Its binding sentence is unambiguous:

“The quality of information publicly distributed by DoD Components **must** meet the attributes of utility, objectivity, and integrity, as defined in the Glossary.” (App. 3A.2.c)

The Glossary is where the requirement acquires teeth, and it bears directly on everything in Part I:

“**utility.** Refers to the relevance and **timeliness** of information to its intended users, including the public.”

“**objectivity.** A measure of information quality related to accuracy, clarity, completeness, and unbiasedness and presented in the proper context. **Sources of information are identified. Supporting data and models and errors and statistical ranges are included.**”

“**integrity.** Refers to the **security** of information; namely, the protection of information from unauthorized access or revision...”

Three observations follow, and they are definitional rather than argumentative. Timeliness is not an external standard imposed on the Department — it is half of the Department’s own definition of a mandatory attribute. *Objectivity*, as the Department defines it, is not a standard about being correct but about showing the work: sources identified, supporting models included, error ranges published. The crosswalk publishes no maintainer, no derivation for any of its four layers, and no error estimate against a 0.51% review rate. And *integrity* is defined narrowly as protection against unauthorized alteration — a file can be entirely wrong and perfectly compliant with it, which is why nothing in this Part rests on that attribute.

Binding language (quoted)	Public evidence	Grade
Information publicly distributed “ must meet the attributes of utility, objectivity, and integrity,” where objectivity requires that “Sources of information are identified” and “errors and statistical ranges are included” (DoDI 8170.01 App. 3A.2.c and Glossary)	No maintainer, no published derivation for any of the four crosswalk layers, no error estimate; 0.51% of links ever reviewed	NOT MET

Binding language (quoted)	Public evidence	Grade
DoD and OSD Component heads “Develop and publish a review process that substantiates the quality of information publicly distributed by the Component” (§2.7.h)	No published review process substantiating the crosswalk’s quality was located in this audit’s finite search	NOT DEMONSTRABLE

The second row is the more useful of the two, because the verb is *publish*: the evidence instrument was supposed to be public by design, and its absence is therefore observable from outside the Department. It is also the pattern this Part has already named — a binding requirement with no instrument in which compliance could be shown.

The method’s sharpest yield: the famous 90-day capstone deadline is written in the implementing instruction as “should take place no later than 90 days prior to separation” — advisory — while the statutory floor (10 U.S.C. §1142) binds. The instruction softened the statute. A systematic statute-to-instruction softening sweep remains to be run.

The verdict distribution is itself the thesis in compliance grammar. Requirements fail in two ways: NOT MET where public data exists, and NOT DEMONSTRABLE where the government wrote a binding requirement but built no instrument to evidence it. The second failure is deeper — a “will” statement without an evidence field was never intended to be checked.

VII.2 The data-governance layer: who is obliged to keep a federal dataset true

Part I established that the crosswalk is stale, unattributed and unmeasured. This section asks the question that follows from it and is rarely asked of any federal dataset: which instrument obliges anyone to keep it accurate, and in which record could compliance be shown? The governing corpus was read directly, and the answer is more specific — and more damaging — than a general complaint about oversight.

A note on method, because it changes how this corpus must be read. The six transition issuances of VII.1 yielded 640 binding statements because they are written in compliance grammar. The data-governance corpus largely is not. DoDI 8320.02 states policy as flat present-tense declarative; DoDD 5105.89 assigns every duty in the third-person indicative (“Develops...”, “Establishes...”). A modal parse of the Chief Digital and Artificial Intelligence Officer (CDAO) directive returns zero binding statements, which is a property of the drafting convention and not a finding about the CDAO. Reading this corpus therefore requires three categories — modal-binding, declarative-mandatory, and responsibility-assignment — and the classification of each statement below is stated with it.

The directive that manages the information enterprise does not mention accuracy

DoDD 8000.01, *Management of the DoD Information Enterprise* (17 March 2016, Change 1 of 27 July 2017), is the capstone directive for the Department's information enterprise. Its full text was searched for *accuracy*, *accurate*, *data quality*, *currency*, *timeliness*, *steward*, *authoritative data*, *metadata* and *obsolete*.

Across roughly 25,500 characters, the word “accuracy” occurs once — in the directive’s own summary of changes, describing an administrative update to its reference list. Every other term returns zero.

This is reported as a searched nil return, not as an absence in the world: the finding is that the Department's capstone information-enterprise directive contains no provision on whether the information is right.

The instructions that do govern data predate the problem and point at a superseded registry

DoDI 8320.02 (5 August 2013, Change 1 of 24 June 2020) carries the registration requirement, as a declarative-mandatory policy statement:

“Authoritative data sources (ADSs) are registered in the DoD Data Services Environment (DSE).” (§3.c)

...and binds Component heads to execute it: “*Register in the DSE, all identified ADSs, IT services, and required metadata.*” (Enclosure 2, ¶h — responsibility-assignment.)

Two defects sit in that requirement. The first is that the instruction never says which sources must become authoritative in the first place — it says registered ADSs are registered. The second is that the trust standard is circular, and modal-binding:

“Data, information, and IT services will be considered trusted when they have provided sufficient pedigree and descriptive metadata for consumers to rely on them as an ADS, and comply with applicable information assurance and cyber security policies.” (Enclosure 3, ¶1.e)

Sufficient is unquantified, *pedigree* is not a field in any mandatory metadata set, and the only concrete test in the sentence is a cybersecurity one. This is the same substitution Part I documents in the crosswalk itself: security standing in for accuracy.

DoDI 8320.07 (3 August 2015, Change 1 of 5 December 2017) contains the corpus's only accuracy clause, and its construction is the finding. Data producers and providers *will*:

“Include relevant attributes that help ensure data accuracy such as error estimates and expiration and range of applicability of the data, **to the extent applicable**, in accordance with DoDI 8550.01.” (Enclosure 3, ¶2.b(2))

Binding in verb, but framed as an illustrative list — *such as* — and qualified by *to the extent applicable*. It is also the only place in the corpus where expiration of data is mentioned at all. The same instruction requires that discovery metadata identify a responsible party:

“Create discovery metadata for high-value data and information that describes **who is responsible for specific data assets**, where the data assets are located, what kind of data assets are available, and how to go about accessing the data assets to enable data for consumers to find.” (Enclosure 3, ¶3.a(2))

“**High-value**” is **undefined**, which supplies an immediate answer to any component asked why a particular dataset carries no responsible party.

One further sentence deserves quoting because it shows the Department can write an accuracy exception when it wants one — and reserves it for speed, not for correctness:

“Data will be made available as soon as possible. Data producers and data providers should not delay making the asset accessible in order to complete processing of data before posting it, **except when limited by security, policy, regulations, or to maintain minimum levels of data accuracy.**” (Enclosure 3, ¶3.f(1))

Note the split modality inside one paragraph: *will* for availability, *should not* for the delay prohibition. Availability binds; the restraint advises.

An official owns data quality; no component owes it

DoDD 5105.89, *Chief Digital and Artificial Intelligence Officer* (18 November 2024) — the newest instrument in this corpus — assigns the CDAO to:

“Establish[] and chair[] the CDAO Council to oversee enterprise data governance and data quality, and resolve issues related to the adoption and use of data, analytics, and AI capabilities.” (¶2.w)

and to:

“Develop[], publish[], and oversee[] the use of a DoD federated data catalog and common data models...” (¶2.k)

The directive creates oversight of data quality and imposes no data-quality obligation on any Component. The only duty it places on Component heads is to *coordinate* with the CDAO on matters within the CDAO’s assigned responsibilities. No component is told that its data must be accurate, current, or attributable to a named owner.

The 2024 catalog also orphans the 2013 requirement: 8320.02 binds components to register in the DSE, while 5105.89 stands up a federated data catalog and 8320.02 has not been reissued to point at it. A binding registration requirement now names a repository a later directive has functionally superseded, and neither is a public record in which a member or an auditor could check whether the crosswalk is registered at all.

Where a hard requirement does exist, it is statutory — and the crosswalk fails it

The obligation the departmental corpus never quite states is stated plainly in statute. Under the OPEN Government Data Act, codified at 44 U.S.C. § 3511, agency heads *shall* maintain a comprehensive data inventory, and the metadata for each data asset must include, to the maximum extent practicable:

“(vi) The date on which the data asset was most recently updated.” “(vii) Each agency responsible for maintaining the data asset.” “(viii) The owner of the data asset.”

with § 3511(a)(3) requiring the inventory be updated “not later than 90 days after” a data asset is created or identified.

Against those three clauses the crosswalk’s catalog entry fails on all three. The maintainer field is blank — clauses (vii) and (viii). The catalog’s stated currency and the file’s own vintage disagree — clause (vi). And a *fourth* element a reader would expect, a stated update frequency, is not required of the agency at all: the federal metadata schema marks publication frequency as an optional field.

That last point is stated against this audit’s own interest, and it matters. The crosswalk’s published update cadence of “irregular” is not a violation of anything. **There is no rule requiring a federal data asset to state how often it will be revised, and none requiring that it ever be revised at all.** The blank maintainer is a breach of an express requirement. The irregular cadence is the absence of a requirement — which is the more serious of the two findings, and the one no correction request can fix.

The grade

Requirement (quoted)	Class	Public evidence	Grade
Metadata must include “Each agency responsible for maintaining the data asset” and “The owner of the data asset” (44 U.S.C. § 3511(a)(2)(A)(vii)-(viii))	statutory, binding	maintainer field blank; no responsible party named on the program’s own sites	NOT MET

Requirement (quoted)	Class	Public evidence	Grade
Metadata must include “The date on which the data asset was most recently updated” (§ 3511(a)(2)(A)(vi))	statutory, binding	catalog currency and file vintage disagree	NOT MET
“Include relevant attributes that help ensure data accuracy such as error estimates...to the extent applicable” (DoDI 8320.07, Encl. 3 ¶2.b(2))	modal-binding, qualified	no error estimate published against a 0.51% review rate	NOT MET
Discovery metadata describing “who is responsible for specific data assets” (DoDI 8320.07, Encl. 3 ¶3.a(2))	responsibility-assignment	no responsible party published; “high-value” undefined	NOT DEMONSTRABLE
“Register in the DSE, all identified ADSs...” (DoDI 8320.02, Encl. 2 ¶h)	responsibility-assignment	the register is not public; a 2024 directive stands up a different catalog	NOT DEMONSTRABLE
<i>A stated review cadence for a federal data asset</i>	—	—	NO REQUIREMENT LOCATED

The shape is the finding. The Department has an official who oversees data quality and no component that owes it; a registration duty pointing at a superseded repository; a trust standard defined in terms of itself and tested by cybersecurity; and a single accuracy clause written as an example with a qualifier. The one unambiguous requirement — name the owner, state the date — comes from Congress, and it is the one the crosswalk demonstrably fails.

(Method: each issuance was retrieved as published, its text layer extracted in full, and every quotation located by string search in the complete document, 12 August 2026. Two quotations carried in an earlier working pass were found on this check to be inaccurate — one paraphrased a sentence that does not appear in the source, and one omitted the qualifier “to the extent applicable” — and both are corrected above against the source text. Paragraph citations follow each issuance’s own numbering.)

Part VIII — The Pattern

One sentence describes every layer this audit touched. Congress mandates; the Department funds; no one verifies; the harm lands on people with no procurement voice.

The crosswalk: mandated 2006, enhanced version piloted 2015 for \$49,996, measurement refused 2016, quality layer still dated 2014. SkillBridge: authorized in statute, requirements written into a signed MOU in 2024, a directory inconsistent with those requirements on its face, outcome data owed since FY2025 with no sourced publication located, while an unsourced “85% job-offer rate” appears in a budget justification. TAP: mandatory for every separating member, evaluated once in 35 years, with the evaluation finding no wage benefit. The maintenance layer: single-bid and sole-source vehicles, none of which pays for data quality. The oversight layer: ten GAO reports, zero enforcement powers, zero IG engagements located.

The mechanism is not villainy; it is the absence of any actor whose job depends on the thing being true. Contractors deliver to statements of work that omit quality. Program offices verify that deliverables were delivered, not that they are accurate. Congress verifies that programs exist, not that they work. And the members — who get one transition each — discover the gaps only after separation, dispersed and voiceless. A weapons program that fails gets a hearing; a transition dataset that fails produces a VA survey statistic three years later, disconnected from any contract.

What makes this moment different is convergence: open GAO recommendations with Department-declared deadlines in September and December 2026; a new statutory reporting mandate; contractual outcome data that either exists (and can be demanded) or does not (which is itself the finding); a wartime squeeze making every window shorter; and — for the first time — an independent dataset against which every official claim can be checked.

Part IX — What Should Happen

IX.1 Eight fixes for the Department (none require legislation)

1. **Measure conversion.** Add per-cohort outcome reporting (offer made, hired, starting pay) to the MOU every host already signs; publish rates on each listing. The keystone — every other fix becomes enforceable once outcomes are measured.
2. **Verify a job exists at listing time.** Require an active requisition or stated intended-hire count; condition renewal on conversion history.
3. **Enforce the taxonomy and require duration.** Opportunity type as a drop-down of the four DoDI categories with accreditation fields; “Under Review” not a permitted duration.
4. **Require a pay band** for the destination role.

5. **Structure occupation targeting** — a code picker plus an explicit “open to all” flag.
6. **Open the data and fix the search.** Remove the 150-result cap; publish the directory as a dataset.
7. **Recruit hosts where the gaps are** — thin states, territories, and the dark fields (medical, ATC) this audit maps precisely.
8. **De-spam the map** — one employer card per organization, and publish the already-collected training plans on every listing.

Sequencing: fix 1 is the keystone; fixes 2–6 and 8 are form-validation changes a support contractor could ship in a sprint under existing vehicles.

IX.2 The honest-host standard (voluntary, immediate)

Publish the requisition behind the internship, the pay band, the occupation codes, the duration, and the last cohort’s conversion count. Good employers already effectively meet it; it costs nothing to any host with a real job behind the internship. Hosts meeting the standard earn the top grade in MissionProof’s independent directory — market reward for verification, available now, no policy required.

IX.3 The accountability calendar

Request for Correction. A route exists that this audit had not previously identified, and it does not depend on anyone’s discretion to answer. The Information Quality Act (§515, P.L. 106-554) obliges federal agencies to maintain a mechanism through which affected persons may seek correction of disseminated information. Both departments that touch this crosswalk publish one. The Department’s mechanism, at DoDI 8170.01 Appendix 3A.3, requires a decision within 60 working days, an appeal within 30 working days to the Component Chief Information Officer, and — at Appendix 3A.4 — an annual report of every complaint received, compiled and submitted to the Office of Management and Budget, with a narrative description required for any complaint involving information the Department deems “influential.” The Department of Labor publishes its own guidelines, dated 1 October 2002, and its published record shows one posted correction request, from FY2010, a statement that none were received FY2011–FY2017, and nothing thereafter. The requests are unfiled as of this writing, and the register at Appendix C records them as drafted. What the route offers that the others do not is a clock and a reporting obligation attached to someone else’s calendar. The drafted correction request rests on the two grounds Part VII.2 grades as NOT MET — the express statutory metadata elements at 44 U.S.C. § 3511(a)(2)(A)(vi)–(viii), and the Department’s own objectivity standard requiring that sources, supporting models and error ranges be published — and not on “integrity,” which both departments define narrowly as protection against unauthorized alteration and which a wholly inaccurate file can satisfy.

FOIA (drafted, not yet filed; full register at Appendix C): the MCTO outcome data and the source of the “85%” claim; MOU enforcement records; the O*NET grant statement of work (SOW) — the single document that answers whether the 2014-dated quality layer violates any contractual term. Congress: the members whose committees own this — the House Armed Services Committee (HASC) Military Personnel

subcommittee; the FY24 NDAA’s GAO mandate authors; Arizona’s senators — hold the September 30 / December 31 deadlines. The asks are small: enforce the MOU that exists; publish the data that is owed; fix the forms. GAO: this audit is available for submission as implementation-relevant evidence on the open recommendations. The market: MissionProof’s graded directory ships the fixed version without waiting — every grade computed from the receipts in this paper.

IX.4 The questions Congress already asked

CRS’s standing reference on TAP (R48114) frames the issues for Congress in its own words. This audit answers them, finding by finding:

CRS issue (R48114, June 2024)	This paper’s answer
Root causes of TAP-initiation noncompliance (~70% late)	Part VII: a binding “must” with no consequence structure; the compliance matrix grades it NOT MET and shows why enforcement never attaches
Whether TAP addresses “translating military skills to civilian employment”	Parts I and V: the crosswalk’s 0.51% review rate, and the two-sided silence — zero military-code references across a top veteran employer’s 384 requisitions
Job-finding effectiveness	Parts IV and VI: the market behind the internships, one-third of pilot hosts with no discoverable job, and DOL’s own study finding no wage benefit
Program proliferation (45 programs, 11 agencies)	Part VI: \$14.3B with “limited evidence,” no owner, and intermediary claims no federal dataset can check
Resource-allocation appropriateness	Part IV: ~\$506M/year in unpriced labor exposure is a resource allocation unexamined in any located review
Active/Reserve parity	The access-inequity method of Part IV.7 extends directly; the 5.9× branch gap suggests the disparity question is broader than component

The paper’s asks require no new questions from Congress — only follow-through on the ones its own research service already posed, now with independent data attached.

Appendix A — The Numbers (every figure traces to a committed source file)

#	Figure	Finding
1	48,042 / 0.51% / 2014	AF crosswalk links / ever human-reviewed / quality-layer date
2	1.5 vs 14.1	Matches per code, current vs piloted-enhanced crosswalk (\$49,996 pilot, shelved; correction 22: yield, not accuracy)
3	<40 / 34 / 36%	Veterans reporting skills match their job (VA, 2021-23)
4	11,532 → 2,732	Advertised positions → organizations in the 87-export union (1,103 state-anchored)
5	2.5% / 12.7%	Largest single organization's share of union rows / top ten combined — the four-chain 38% figure was a state-query-grain artifact and is withdrawn
6	150	Search-result cap (AZ member sees 150 of 629)
7	12.1% / 49.8% / 1.4% / 86.7%	Programs naming an occupation code / blank on targeting / disclosing destination pay (union re-run; 1.2% census) / silent on employment outcomes (union re-run; 87% census — holds)
8	5 of 15	Pilot hosts with no discoverable paid analogue of the trained role
9	212 (5.0%) / 266 (17%) / 29	Commission-model programs (union, corrected construct; census 116 = 7.6% including identified false positives) / opaque programs (census grain) / entirely empty listings (union; 28 census)

#	Figure	Finding
10	47	Union programs exceeding the 180-day ceiling (19 in the state census)
11	75% / 85% / 0	MOU offer warrant / KPI goal / published outcome numbers
12	278 (6.5%) / 5.9×	Senior Army-AF NCO verifiable market (union) / Navy-vs-Army E-8 access gap
13	~\$506M (\$316-732M)	Annual value of member labor delivered to hosts against no verified job obligation
14	\$26,300	One member, one 120-day internship at E-5 comp
15	155 (10%)	Programs at 4-5 FLSA primary-beneficiary risk flags (census grain; construct not yet replicable at union — see Appendix D)
16	\$33.5K / \$36.2K	Infantry / medical first-year median earnings (Census VEO)
17	89 → 24 / 18 of 24	FY25→FY26 AF retention-bonus specialties / retained specialties with ≤2 programs
18	640 / 0	Binding requirements parsed / requirements grading MET
19	“should”	The 90-day capstone’s modal verb in DoDI 1332.35
20	\$3.37M / \$15.2M / 1+1	SOLID website order / Goldbelt MCTO IT / competitive offers received, combined
21	7 / Sep 30 & Dec 31 2026	Open GAO-24-107352 recommendations / DoD’s self-declared deadlines

#	Figure	Finding
22	\$14.3B / “limited evidence”	The 45-program ecosystem / RAND’s effectiveness verdict
23	379 vs 0	“Top veteran employer” brand mentions vs military-code references, one employer’s full book (384 requisitions)
24	36.5% / 24% / 27.9%	Requisitions gated on a security clearance, strictly defined — 44.5% if Public Trust is counted with them / degree-walled with no experience path / carrying complete pay bands
25	4,563 / 7 / 5	Unique requisitions (raw-supported) / employer corpora / industries in the requisition audit, each with an earned scope status
26	28% → 98-100% → ~100%	Pay disclosure: unconstrained → chosen → legally mandated — silence is a choice, and mandate ends it
27	6 / 5	Failure modes / platforms unable to enumerate their own search results: directory full-export 504 AND 150-result cap (one platform, two modes), career-site paginator, consultancy board page-6 refusal, 1,000-result platform cap
28	0	Class I railroads located in the union corpus — while the structured “Transitioning Military” tag in a 47-company rail sample is that one railroad’s practice alone, not an industry convention

#	Figure	Finding
29	629 / $\geq 3+5$	Market-anchor evidence rows / employer-and-posting floor per plotted point
30	never before 20	Years for the security-forces market median to catch military compensation
31	$>80\%$ / 200,000+ vs 148,315	Share of entrants separating before retirement eligibility / annual transitions as commonly described (CRS R48114) against DMDC's own 2024 active-duty separations — unreconciled, correction 26
32	Jul 2023	DoD paused new SkillBridge employer intake for “overwhelming popularity” — demand outrunning governance
33	\$35.7K vs \$48.1K / 57% vs 46%	Army vs Air Force veteran year-1 median earnings / year-1 non-attachment (Census VEO)
34	22,000+ rows / ~6,200 distinct / 300+	Posting rows, the distinct postings behind them (upper bound; components not de-duplicated against one another), and distinct employers in the full MissionProof postings corpus
35	11% / ~0.1% / 0%	Private-market listings with any veteran signal beyond boilerplate / rate at a flagship SkillBridge-academy employer / cleared-intel postings using military-occupational language

#	Figure	Finding
36	41% / 0.0%	Single-filter export rows unreachable by any state sweep / blank LocationStates in the state-anchored corpus — the definitional proof
37	1,103 = 1,103	The 57-state sweep repeated five days later returned identical organization count — the directory is stable; the method's reach is not
38	883	Organizations surfaced by one filter change that a complete 53-query national sweep never returned
39	2011 / 2014 / 2019 / 2024	Statute (no host obligation) / instruction (the "any other employer" clause) / marketplace launch on 2014 text / first host obligation, by MOU
40	0	Occurrences of the word "SkillBridge" in DoD Instruction 1322.29, the instruction that governs it
41	3 years / 0	The marketplace's peak-growth period on an OPM interagency platform / prime-award records visible for it
42	Oct 2022 → Jul 2026 / 4	Period over which the Air Force crosswalk layer shows no revision visible to a user / copies sharing one SHA-256, two of them third-party archive captures
43	11 months / 0	Period the .mil SkillBridge directory data file is byte-identical (Apr 2025 - Mar 2026) / archive captures of the commercial host now serving the live directory

#	Figure	Finding
44	1 / 0	Occurrences of “accuracy” in DoDD 8000.01, the capstone DoD information-enterprise directive — in its own change note / occurrences of data quality, currency, timeliness, steward, metadata or obsolete
45	0 / 0 / 1	Occurrences of “O*NET” and of “AFSC” across fourteen DoD Cyber Exchange governing and program documents / occurrences of “crosswalk” — in a sentence denying one exists
46	15 Feb 2023 → 7 May 2026	Cancellation of DoD 8570.01-M for military and civilian personnel / most recent DFARS change still requiring contractor compliance with its “current version”
47	5 vs 7	Cyber workforce elements defined in DoDI 8140.02’s glossary / elements named in the Department’s own program introduction
48	3 of 3	Statutory metadata elements at 44 U.S.C. § 3511(a)(2)(A) the crosswalk’s catalog entry fails — responsible agency and owner blank, last-updated date inconsistent
49	0	Located requirements that a federal data asset state a review cadence, or ever be revised — the “irregular” cadence breaches nothing

#	Figure	Finding
50	478 → 119	Apparent host exits over one seven-month window / candidates surviving removal of program renamings — the fourfold inflation that bars a churn figure
51	75 of 922	Median rank of a crosswalk-offered occupation on task-text similarity to the specialty's own training standard (1,055 pairs, 60 scorable specialties, 450 distinct offered occupations, ranked against a 922-occupation pool; the Air Force crosswalk's full reach is 801 distinct occupations)
52	42.6% / 18.2%	Offered occupations falling outside the top 100 / outside the top 300 of 922
53	19 / 59 / 81	Median offered rank by crosswalk layer — DMDC / CitM / COOL — the layers a member browses performing worst
54	90.9% / 73.5	Offers made identically to every skill level / median rank places an occupation moves between levels on the specialty's own proficiency codes
55	1 → 733	Rank of Radio, Cellular and Tower Equipment Installers for 1Z3X1 at 3-level and at 7-level — offered at both

#	Figure	Finding
56	60 of 91 / 25 of 31	Specialties scorable / excluded specialties falling in just two families (aircraft maintenance, intelligence) — the bound on generalizing rows 51-55
57	78.0% / 91.4%	Enlisted members across all five services whose highest attainment is a high school diploma, GED or some college — whole force (1,034,157) / junior enlisted band (530,090)
58	38.0% vs 10.5%	Associate's-degree attainment among senior enlisted, Air Force against Marine Corps — the Community College of the Air Force, visible in federal data
59	15.8%	Army senior enlisted with no reported education level, against ~0% in every other service — the 2023 personnel-system migration, and why the Army's attainment figures are a floor
60	54.7% vs 27.5% / 16 vs 68	Air Force against Army senior-enlisted degree attainment / civilian occupations offered to each service's intelligence analyst — generosity running opposite to credential, on one worked pair

#	Figure	Finding
61	15.6% (floor; 16.1% at org grain)	Share of advertised participant time — 57,660 of 369,000 person-days, 550 programs, 249 organizations — held by hosts on .gov/.mil domains, whom the 2025 MOU exempts from the 75% qualifying-offer warranty
62	6 of 10	Government bodies among the directory's ten largest advertised-capacity holders
63	2,836 of 4,255 / 1,419	Programs stating a duration / programs reading "Under Review," for which no duration is imputed anywhere in this paper
64	2 vs 5 / 17 vs 89	Median and maximum destinations per specialty — the Air Force's own inbound Convertible Skills List against the O*NET crosswalk that routes members out, both at printed-code grain
65	65 of 132 / 17 / 4	AFECD Convertible Skills rows carrying a change date / specialties requiring Air Force Career Field Manager (AFCFM) review before award / specialties bounding how recently the skill was performed — none of which the O*NET crosswalk has any analogue for

#	Figure	Finding
66	120 / >190	AFOCD statements of “Not used.” — an affirmative record that a dimension does not apply / elements carrying their own change date, on a directory superseding its predecessor at six months
67	4 vs 0	Editions of the Air Force officer classification directory published in the twenty-four months to 31 Oct 2025, evidenced by its own internal change dates, against revisions of the O*NET military crosswalk over the same period
68	9 of 43 / 422 / 105	DAFOCD Appendix A matrices printing a change date later than the appendix’s own “(Changed 31 Oct 2022)” banner, the latest 31 Oct 2025 / Classification of Instructional Programs (CIP) education-matrix entries extracted from the 31 Oct 2025 edition (dafocd_cip_matrices_20251031.csv), used for no figure in the body other than this one / entries whose CIP cell is a single six-digit CIP code — the rest are the directory’s own wildcard forms (205), series or list forms (100), or carry no code (12)
69	7 / 6 / 1	Union programs naming a dropped medical specialty by code / naming a special-warfare field by code / special-warfare-code programs inside the DIB (both code forms; corrections 37–38)

#	Figure	Finding
70	4,154 of 4,255 (97.6%)	Union programs accepting all six services (census-grain 1,497 of 1,527 superseded)
71	20.4% / 10.5% / 12.9% / 5.8%	E7-E9 share of enlisted separators — AF / Army / Navy / USMC (Census VEO paygrade file, 2020-2021 cohort)
72	\$22.1-22.6M	Member-labor value per 1,000 additional participants per year (tier-weighted central; committed model)
73	26 (floor)	AF COOL-listed credentials not commercially available to the public — issuer-verified 2 Sep 2026 (932 availability panels, 324 needs-review)
74	15 of 48 / 22 of 82 / 14 of 63 / 14 of 70	Uniform-locked share of listed credentials — GEOINT analyst / successor GEOINT specialty / Security Forces / Special Investigations
75	132 of 865 (81 bachelor's+) / 226	Curated requirement panels carrying a postsecondary-degree gate / stating a years-of-experience requirement
76	253 / 254 / 423	Catalog credentials with an accreditation flag / flagged unaccredited / no determination recorded (of 930; portal's own ANSI/NCCA/ICAC flags)
77	32 of 49 / 100 of 865	Blanket-listed credentials (≥ 20 specialties) with zero mentions in the 7,937-posting corpus / requirement panels gated on management or supervisory experience

#	Figure	Finding
78	43.8% (20,126 of 45,964)	AF enlisted separators leaving at E-4 or below — upper bound on the pre-ALS tier (VEO paygrade file, 2020-2021 cohort)
79	97 of 192 / 6 semester hours	ALS syllabus hours in leadership, supervisory, problem-solving and communication skills / CCAF Leadership, Management and Military Studies requirement ALS preferentially fills

Appendix B — Source and Data Inventory

Primary datasets and evidence files: the directory pull (59.7MB raw multi-part, hashed — not a census, IV.9; 87-export union corpus of 7,826 rows / 4,255 programs / 2,732 organizations; 18,688-row state-availability dataset); 1,489 program-occupation edges; hire-intent, commission, opacity, and compliance classifications for all 1,527 state-census programs (union-grain re-runs pending and labeled wherever cited); the 18-row conversion-pilot evidence table (URL+date per row); the requisition-audit archive — seven employer corpora plus a 250-posting rail-industry sample and ~40 occupation-lane postings files, every file SHA-256-hashed into a 200+-row capture manifest with per-file scope notes; the four-system recruiting-estate capture set behind IV.9’s corporate-structure specimen (manifest rows 207 and 208, plus two files added 10 Aug 2026, one a supplied page capture and one a verbatim transcription, both labeled at that scope on their face; the superseded reconstruction at row 209 is retained as the record of a compaction loss and its recovery); the 870-statement instruction parse and 17-row graded matrix (fifteen in v1, plus the two DoDI 8170.01 rows); contract records for both PIIDs; the FY26 SRB capture and cross-reference; five lens outputs; VEO trajectory join; four publication-grade charts; the interactive MOC explorer. Method scripts are committed for full reproducibility. Two completeness claims (a four-firm concentration figure; one employer “complete book” designation) were withdrawn during the audit when receipts contradicted them — both retractions are logged in Appendix D and in the capture manifest.

Appendix C — FOIA Register (summary)

Eleven requests, tiered. Tier 1 (file first): MCTO provider outcome data + the “85%” claim’s source and methodology (DHRA); the DMDC *Active-Duty Separations by ISC Code, Service, and Pay Grade Report* (September 2024 and prior fiscal years),

at full ISC-code × service × pay-grade resolution — a report the Department names in the reference list of its own 2024 Demographics Report, and whose pay-grade dimension is withheld from the only public surface that presents it. Every other visualization in that report carries a Pay Grade filter; the separations visualization carries only Service Branch and Enlisted/Officer. The request matters because Air Force Instruction 36-2671 tiers SkillBridge duration and approval authority by rank, and the rank distribution of separations is the denominator that policy is keyed to. The dimension is not secret — the Census Bureau’s VEO paygrade file, built from DMDC’s own records, publishes enlisted separator cohorts by single pay grade and service through the 2020–2021 cohort; what the Department withholds, another agency releases. The request seeks what VEO cannot supply: the current-year, ISC-coded, officer-inclusive report; MOU enforcement and termination records (DHRA); the O*NET grant statement of work and quality requirements (DOL/ETA). Tier 2: SOW/QASP sets for the SkillBridge website order, the MCTO IT contract (including the 8(a) J&A), COOL vehicles, and the VOW §222 pilot closeout. Tier 3: participation by occupation and denial counts; DoD’s implementation-status submissions to GAO. Full paste-ready records language in the register document.

None of these requests had been filed as of the date of this edition. The register is a drafted program against named custodians; filing, agency responses, and any non-responses are reserved to a follow-up study. A “no responsive records” answer to the first request would itself be a finding of the same class as those reported here.

Appendix D — Limitations and Correction Log

Limitations. The directory snapshot is a point in time (3 Aug 2026); listings change. Conversion-pilot verdicts are date-stamped spot checks, not continuous monitoring. The labor-transfer estimate uses scenario ranges where actuals are unpublished. VEO non-attachment includes GI Bill students. MOC parsing of free text carries a small residual error rate; 1.8% of source rows had malformed fields, handled conservatively. MOU quotes derive from the held instruments — the August 2024 blank release and two signed instances (2025, 2026); the clause wording quoted in the Executive Summary is the 2025-and-later text (correction 40). All “not located” statements are finite-search claims.

The audit’s own population figures were revised during the work. The 3 August state-anchored corpus (1,527 programs / 1,103 organizations) is a valid named population but is not a census; the 8 August union corpus (4,255 / 2,732) is larger and is also not a census, because no method the Department’s site permits can produce one. Figures recomputed on the union are reported on the union; figures not yet recomputed are marked at their original grain. Pay disclosure (1.4%), commission (5.0%, corrected construct), outcome-silence (86.7%), and the empty-listing count (29) are reported at union grain with committed hit lists; the government-host share is the 15.6% person-day domain measurement of II.5, which supersedes and withdraws the earlier 10.5% name-lexicon figure (correction 14). The conversion pilot, the FLSA composite, and the opacity role-recognition component remain at the 3 August grain. The FLSA figure carries a specific and material limitation: the scoring construct did not reproduce from its written specification — an independent recon-

struction returned 1.2% against the published 10.2% — so the 10% figure is reported as the output of a construct that has not been independently reproduced, and should not be treated as a validated prevalence. It is retained rather than deleted because the underlying risk pattern is real and separately evidenced by the specimens; it is the magnitude that is unverified.

Corrections to date.

- (1) 4 Aug 2026 — an earlier internal finding stated a remodeling company’s six-figure pay claims applied only to its sales roles; the employer’s own installer postings also claim “\$125,000+.” Retracted and replaced with the guaranteed-wage-vs-contingent-claim analysis, which is stronger and fairer.
- (2) An initial FY25-based reading of retention-critical fields (maintenance/medical) was revised when the FY26 list cut both.
- (3) 8 Aug 2026 — the four-chain 38% concentration figure was computed at the state-query grain, which inflates national chains by construction; at union grain the top organization holds 2.5% of rows. Withdrawn and restated as a claim about the member’s browsing surface rather than the directory’s composition.
- (4) 8 Aug 2026 — an internal finding that the four Special Warfare fields “receive no civilian target at all” was true only of the crosswalk file’s direct mapping; a derived layer does return civilian occupations. Withdrawn.
- (5) 8 Aug 2026 — this audit’s own fair reading of the crosswalk’s two-block ordering (two score tiers) was tested against expanded captures and refuted by our own data; the blocks are two source mappings concatenated. Offering the fair reading was correct on the evidence then available, and recording its refutation is the point of the practice.
- (6) 8 Aug 2026 — an internal claim that two federal crosswalk tools return different answers for the same member was withdrawn: the 21-vs-23 difference is two military-only occupations correctly filtered by the member-facing tool. The miscount (“top 20” over 21 items) stands alone.
- (7) 9 Aug 2026 — an earlier finding stated that roughly eight in ten Air Force career-field families receive one identical occupation list at every skill level. Direct measurement of the published file gives a materially different figure: of 115 families carrying two or more skill levels, 55 — 48% — have mappings that differ by skill level. Withdrawn and replaced by the measured figure in Part I. *(The replacement figure has itself since been withdrawn as unreproducible — see correction 23. The direction of correction 7 stands; its magnitude does not.)* The direction of the original claim was right and its magnitude was wrong.
- (8) 9 Aug 2026 — an earlier internal draft recorded DoD Instruction 1322.29 as “issued 2019 under Trump.” The instrument’s own header gives 24 January 2014, and Change 1 (May 2020) is administrative housekeeping per the instruction’s §7 change note. Both withdrawn; the corrected sequence is Part II. A second draft error, treating Change 1 as a programmatic action, is withdrawn with it. The attribution rule adopted after these errors — date the initiation, not the signature — governs Part II.

- (9) 9 Aug 2026 (evening) — union re-execution of the commission classifier identified a false-positive class in the census run: the word “commission” in law-enforcement contexts (commissioned officers; standards-and-training commissions) flagged at least four police and federal programs, and the recruiting flag fired on incidental mentions. Corrected construct reported in IV.5 (5.0% at union grain); the census 116 softens on correction.
- (10) 9 Aug 2026 (evening) — the census GOV-host organization list included private contractors matched on the word “Federal” and other keyword strays; the corrected union list (181 organizations) is committed.
- (11) 9 Aug 2026 — a “complete book” designation for one employer requisition corpus was withdrawn when the employer’s own board total (~2,500) contradicted the captured 544; restated as a labeled ~22% slice in V.1.
- (12) 9 Aug 2026 (late, on independent adversarial audit) — the rail-industry sample’s 5% pay-disclosure figure was a parser artifact: the derived table captured only pay ranges, while the underlying 42 cards carry employer-authored pay on 35 (33 single-rate statements, 2 ranges). The “100% vs 5%” two-railroad comparison is withdrawn; the pay-spread argument is restated without it.
- (13) Same audit — window-weighted market figures (82/5%, 42% Under Review, ~7× gap) had remained at the superseded 1,527-program frame in summary passages while IV.6 carried union values; all duration figures are now stated at union grain (278/6.5%, 33.3%, 5.9×, 47 over-ceiling).
- (14) Same audit — the union commission (212/5.0%), government-host (10.5%), pay (60/1.4%), outcome (86.7%), and member-cost (69) hit lists were found to carry residual error classes in both directions; all five are labeled indicator-list sizes rather than validated prevalence estimates, and the classifier code is committed.
- (15) Same audit — Employer B’s book total is corrected from 1,256 to the raw-capture-supported 1,250 unique requisition IDs (section total 4,569 → 4,563).
- (16) Same audit — a protected employer’s name appeared in an unrelated Part III vendor passage and a sentence reported a MissionProof-produced posting-to-code mapping; both removed (sensitivity control, not a factual correction).
- (17) 11 Aug 2026 — earlier drafts described this audit’s Freedom of Information Act requests as “filed or ready to file.” No request had been filed. The register is a drafted program against named custodians; the language is corrected throughout to “drafted, not yet filed,” and the filing program is reserved to a follow-up study. A paper that grades a federal program on unverifiable assertions cannot itself make one, and “pending” is unfalsifiable to a reader by construction.
- (18) 11 Aug 2026 — a refinement, not a withdrawal: Part IV.9’s magnitudes were correct, but its characterization of what sits inside the export blind spot was not. The blind spot was described as disproportionately virtual and nationwide; direct measurement of the mechanism shows nationwide accounts for 12.9%, while 50.6% is in-person programs whose location field spells the state’s name out in words rather than using its two-letter code. Corrected in IV.9.

- (19) 11 Aug 2026 — Part V.1’s enumeration-failure count is raised from five modes across four platforms to six across five, on a defense prime’s URL pagination silently discarding its own keyword.
- (20) 11 Aug 2026 — the member-cost specimen at IV.6 is annotated: a second capture of the same training school’s site carries a lender-partner figure of “\$27,745” and a per-class SkillBridge count of ten, against the paper’s “\$11,955.00” and “30 slots per class.” A price and a financing ceiling are different measures and both are recorded rather than merged.
- (21) 11 Aug 2026 — an internal 9 August record described one host’s recruiting estate as “4+ subsidiary boards on three ATS platforms.” Tested against the hashed capture, the board count was inflated by facet views: four of the pages counted as separate boards filter a single applicant-tracking instance the capture already held whole, and counting them separately was double-counting. Corrected to four recruiting systems across three vendors — the figure carried in IV.9. The three-vendor claim survives; the board count as originally stated does not. Two adjacent observations about the same host, resting on screenshots not archived to file, are excluded from the published finding rather than carried at a lower confidence.
- (22) 12 Aug 2026 — a refinement against this audit’s own interest, raised on review. Earlier drafts stated that a 2015 Department of Labor pilot “proved” an enhanced crosswalk “performs roughly nine times better” and that “the fix was proven, then shelved.” The pilot measured yield — 14.1 civilian matches per military code against 1.5, across 68 occupations — and did not test whether the additional matches were more accurate; no study since has. “Proven” and “better” both asserted a quality result the study did not produce. Corrected in the Executive Summary and Part I to a coverage statement. A second qualification is added with it: the enhanced method rested on the same O*NET military layer whose irregular maintenance is the defect under examination, so the remedy would have inherited the decay it was meant to cure. A paper that grades a federal instrument on unverified assertions must not make one about that instrument’s proposed repair.
- (23) 12 Aug 2026 — Part I’s skill-level differentiation figure is withdrawn. The Part reported that of 115 Air Force enlisted families carrying two or more skill levels, 55 — 48% — differ by skill level. That figure is not reproducible from the source files under any of seven reasonable definitions of a career-field family, which return a range of 4.0% to 41.0%; the companion gap register’s figure of 12% is not reproducible either. The two were measuring different populations under the same word, and neither stated its definition. Restated on a definition of record — all four layers merged, active codes only, shreds preserved, which is the population a separating member occupies — 9.2% differ and 90.8% are identical from helper to craftsman, with the sensitivity range published alongside. The qualitative finding is unaffected; the magnitude was unstated as to method and should not have been published that way.
- (24) 12 Aug 2026 — Part I’s assertion that the crosswalk’s output is “substantively defensible” and “often translates reasonably” is withdrawn. It rested on the

plausibility of occupation titles and on no competency evidence. The July 2026 audit had already tested one specialty's mapping against its CFETP duty text and found cross-career-field bleed; the assertion was made without consulting it. A defense of a federal instrument requires evidence on the same standard as a criticism of one.

- (25) 12 Aug 2026 — withdrawn and replaced the same day: this correction was itself wrong, and the position replacing it is stronger than the one it recorded. As first published, (25) reported a reproducibility defect in this project's own record: that the derived table named as the July 2026 audit's principal deliverable, `onet_af_military_crosswalk.csv`, was not present at the path that audit documents, and that the table rebuilt for this revision superseded the missing one. The file is present, in the documented package, with its build README — 3,691,424 bytes, built 7 July 2026, 25,257 links across 6,658 specialty codes. It was located within minutes of the search that should have preceded the claim; the failure is the finite-search failure logged at (27), of which this is the fifth instance and the only one to reach print. What replaces it is the strongest provenance statement in this paper. The 7 July build and the 12 August rebuild were produced five weeks apart, by different processes, from the same federal source files, and neither was constructed with reference to the other. They reconcile exactly, and at set grain rather than count grain: layer by layer, as sets of specialty-to-occupation pairs, DMDC 6,765 in both, COOL 18,245 in both, VOW §222 enhanced 247 in both — 25,257 pairs in each build, intersection 25,257, zero present on one side only, with the July build's 6,658 specialty codes a strict subset of the rebuild's 6,827. The rebuild therefore does not replace a lost table. It is a strict superset of an independently corroborated one, adding the CitM/ASVAB layer the July build omitted. Nine CitM links duplicate links already carried by the other three layers, which is why 48,051 layer-tagged rows resolve to 48,042 distinct links; that difference is accounted for and is not a residual. Both tables are committed. Two build parameters undocumented at source are recorded in the Part I provenance note. A paper that grades a federal program on unverified assertions of absence cannot publish one about its own record, and did.
- (26) 12 Aug 2026 — the paper's opening population figure is qualified, not withdrawn. Earlier revisions opened "Every year roughly 200,000 service members leave the United States military," sourced to CRS R48114. The Department's own 2024 separations data reports 148,315 — 26% lower — on a basis excluding immediate re-enlistments, officer status changes and record corrections, and carrying its own note that total active-duty retirement data is not available for FY2024. No public source located reconciles the two. Both are now stated, with the discrepancy named. No calculation in this paper depends on either figure.
- (27) 12 Aug 2026 — a finite-search failure by this audit, recorded because the rule that failed is one this paper applies to others. On five occasions during the audit, the audit asserted that a body of work did not exist — a manuscript file, a competency-level test of the crosswalk, a validation of the alternative competency model, published separations data by pay grade, and the July 2026 crosswalk audit's own principal deliverable — and on all five occasions the work ex-

isted and was located within minutes of looking. The fifth reached print as correction (25) and had to be withdrawn in turn, which is why this entry is stated at the standard the log itself sets: the rule this paper applies to a federal instrument bound the paper first, and the instrument's custodians at least never claimed to be applying it. Rule (3) of the Standards set out under *Methodology and Evidence Standards* requires that negative findings be stated as finite-search claims describing the corpus searched. That rule binds this audit's statements about its own prior work and about the Department's holdings equally, and it was not observed. The corrected standard is recorded in the project's binding rules: assume the record exists and search before asserting absence.

- (28) 12 August 2026 — an internal inconsistency in this paper, corrected. Part I's provenance note described the file in service as "the August 2024 vintage" while the vintage table in the same Part correctly identifies it as `milx0724.csv`, July 2024, and the published zip's own members are stamped 22 July 2024. The note is corrected to July 2024. No figure depended on the date and none changes; it is logged because a paper that measures a federal instrument's currency to the month must state its own to the month.
- (29) 12 August 2026 — a limitation restated after being checked against this audit's own holdings, at the author's prompting. Part I previously stated that a systematic CFETP-to-occupation comparison had "not been run by the Department, by O*NET, or by this audit," in a construction that implied the underlying corpus did not exist. It does: 94 CFETPs are processed into 15,278 staged task rows across 81 specialty families, and the O*NET side is held complete at release 30.3. The comparison still has not been run, but the reason is that the extraction captured short competency labels rather than task statements, and yielded almost nothing for the intelligence and medical families this Part's findings concern. The passage is rewritten to state what is held, what is missing, and why the measurement is withheld. A finite-search rule applies to an audit's account of its own capabilities as much as to its account of the Department's holdings, and the earlier phrasing understated work that had in fact been done.
- (30) 13 August 2026 — a limitation lifted, and three of this audit's own figures corrected before print. Correction (29) recorded that the systematic CFETP-to-occupation comparison had still not been run, the obstacle being an extraction that captured competency labels rather than task statements. The extraction was rebuilt on page geometry and the comparison was run on 13 August 2026; Part I now reports it. Three figures carried in the internal working paper of that run did not reproduce when recomputed from the committed pair-level files, and the reproduced values are the ones published. First, the share of pairs offered identically to every skill level was recorded as 93.3% (984 of 1,055); that count came from a per-specialty summary column which, for at least ten specialties, exceeds the number of offers that specialty has — an impossible value. The pair-level file gives 959 of 1,055, 90.9%, and reconciles independently against the crosswalk's own level strings. Second, the count of specialties whose *median* offer falls outside the top 100 was recorded as 19; on the conventional median that count is 16, and 19 arises only from taking the upper of the two middle values for even-numbered sets. Both figures are now published with their definition

stated, which correction (23) established as the standing requirement. Third, the median between-level rank swing is 73.5 places, not 74. No headline figure changed: the median offered rank of 75, the 42.6% outside the top 100, and the layer medians all reproduce exactly. This entry is logged although none of the three ever reached print, because the practice this paper recommends is that a figure be recomputed from source before publication and the failures reported, not that failures be reported only when caught late.

- (31) 13 August 2026 — a boundary narrowed after this audit searched its own holdings, at the author's prompting. The cyber section stated that a mapping between Cyber Workforce Framework work roles and military occupational codes "may exist on DoD COOL" and that "that site was not reachable during this audit." The author pointed out that this audit already holds a normalized capture of the Air Force's service-level credentialing data. It does, and searching it returns no occurrence of "DCWF", no work-role name and no work-role identifier in any field of any file — a stronger and more specific result than the open question it replaces. The passage is rewritten accordingly. Two limits are written in with it: DoD COOL's live site remains unreachable, so no claim is made about what it publishes; and the capture's own records are marked as requiring provider verification before external use, so no count drawn from it is published — only the search result. This is the third time in two days that a limitation in this paper turned out to understate work this audit had already done (see corrections 27 and 29), and the corrected standard is the same one: search the record, including your own, before writing that something is absent or unreachable.
- (32) 14 August 2026 — a definitional artifact in this audit's own figure, found while trying to source a different claim. Part V and the Executive Summary reported that 45% of a defense prime's 384-requisition book is clearance-gated. Recomputed from the committed requisition file, strict clearance — Secret, Top Secret, TS-SCI — is 140 requisitions, 36.5%. The published figure reproduces only if the book's 31 Public Trust requisitions are counted as clearances, and a Public Trust is a suitability determination, not a security clearance; 17 further requisitions carry no retrievable full text. Corrected to 36.5%, with the 44.5% inclusive figure and its definition stated alongside. The companion pay-band figure was checked at the same time and stands (107 of 384, 27.9%). Separately, and recorded because the search that produced this correction failed at its original purpose: no clearance salary premium is published in either paper, and none can be computed from these captures. Posted pay at both cleared-market primes is a corporate level band — 535 postings resolving to fifteen distinct ranges at one, 228 to fifteen at the other — so the variable available is job grade, not the price of a clearance; the cleared and uncleared postings within the one book carrying real variation sit in near-disjoint occupations; and the single clean within-employer test available, 89 salaried requisitions, returns a negative and statistically insignificant difference. An earlier draft of the companion paper carried "a widely cited \$15,000–\$30,000 salary premium" with no source. It was withdrawn on 13 August for want of one, and this entry records that the audit's own evidence would not have supported restoring it.
- (33) 22 August 2026 — one artifact cited at two vintages and two grains inside the

same paper, dated and reconciled. The market-anchor evidence base appears in the Executive Summary, the Part V preamble, V.3 and Appendix B as 629 posting rows covering four Air Force exit lanes, and in V.2 as 21,086 rows across 269 employers. Both are correct: the 629-row file is the 30 July 2026 build the stay-vs-go charts are drawn from, and the 21,086-row file is a later build extended from four specialties to 92. Neither mention carried a date, so the two read as a contradiction; V.3's companion figures — twelve postings excluded, 195 duplicates collapsed — describe the earlier build, while the current build excludes 256 and collapses 197 on the same rules. Every mention is now dated, and the exclusion counts are given for both builds. No finding moves: the charts, the evidence floor and the lane verdicts all rest on the 30 July build, which is unchanged. What the entry records is that an audit which dates the Department's instrument to the month had two undated vintages of its own instrument in print.

- (34) 22 August 2026 — a row count published as a posting count, corrected downward by roughly three and a half times. V.2 reported the full postings corpus as “22,000 postings.” It is 22,000 posting rows. The market-anchor component is specialty-expanded — one requisition occupies one row for each Air Force specialty lane it anchors — so its 21,086 rows resolve to 4,108 distinct postings, 4,058 carrying a requisition identifier and 50 carrying none. The concentration inside it is the sharper fact: one federal-services employer's 109 requisitions supply 9,653 rows, 46% of the base. Across all four components the corpus is roughly 6,200 distinct postings, and even that is an upper bound, because the hand-collected census carries no requisition identifier and cannot be de-duplicated against the rest. The rates the section reports are computed within components and none of them changes — the 11% veteran-signal rate, the one-in-a-thousand cell, the zero rate in the cleared-intelligence market all stand — but the corpus-breadth figure did not survive recomputation and the corrected number is now the one in print, with the row count and the distinct count each labeled with what it is a denominator for. This is the defect Part I attributes to the Department's own instrument, found in this paper: a count whose grain was not stated, read by everyone at the grain that flattered it. It was found by recomputing a figure that nobody had questioned. Correction (4) established the rule that a count is not a census; this entry establishes the companion — a row is not a record.

Carried with (34): the Contributions list described the CFETP-to-occupation comparison as spanning “1,055 crosswalk-offered pairs across 60 specialties and 922 occupations.” The 1,055 pairs name 450 distinct occupations between them; 922 is the size of the pool each pair was ranked against, and the Air Force crosswalk's full reach across all four layers is 801 distinct occupations. Part I states the pool correctly at every mention; the summary line compressed three different numbers into one. Corrected in the Contributions list, in Part I, and in the figure register.

- (35) 22 August 2026 — three figures withheld from print because they did not reproduce, logged although none of them reached a reader. The AFECD/AFOCD comparison written into Part I this day was drafted on 14 August from an internal analysis carrying figures that were re-derived from the committed files before publication. Three did not survive. The Convertible Skills List was described as

covering 59 specialties; the committed transcription carries 72 specialty codes as printed, or 55 rolled to the five-character form — the stated figure reproduces at neither grain, so no specialty count is published and the comparison is stated in destinations-per-specialty instead. The mean destinations per specialty was given as 3.7; it is 3.1 at printed grain and 3.9 rolled, so the mean is not published and the median, which is 2 at both grains, is. The AFOCD's "Related DoD Occupational Group" mappings were counted at 133 and recount at 142, so that count is not published either. Two figures did reproduce exactly across independent counts and are published: the 120 instances of "Not used." and the 65 of 132 dated rows. A fourth was published as a disclosed floor rather than a point estimate: the AFOCD's dated elements count 194 by this run's pattern and 212 by the 14 August pattern, and appear in print as "more than 190" with both counts stated. The practice this paper recommends is that a figure be recomputed from source before publication and the failures reported (correction 30); this is the second time that practice has caught figures at the door, and the second time it has done so in the audit's own working papers rather than in the Department's.

- (36) 22 August 2026 — a counter-example this paper printed against itself was wrong in the direction that helped it, and is corrected in the direction that does not. The Part I passage on the officer classification directory stated that Appendix A "still carries (Changed 31 Oct 2022) in the 31 October 2025 edition, so the officer education matrices have not been revised in three years." The appendix-level banner does read 31 October 2022. But when the education matrices were re-extracted from the 2025 edition on 22 August 2026 — the first extraction having failed its own acceptance tests and been withheld — the header capture returned nine matrices printing their own later change dates: 32EXG on 31 October 2025, 13HX and 15A on 30 April 2025, 15W, 21M, 32EXE and 43E on 30 April 2024, 13N on 31 October 2023, 38F on 30 April 2023 (a tenth, 19Z, prints "Established 30 April 2020," which predates the banner and is not counted). The matrices have been revised in at least four of the intervening editions; what has not been revised is the banner above them. The passage is rewritten to say so. The correction cuts against the paper's convenience: an instrument this paper was using as evidence of a stale element inside a maintained directory turns out to be maintained at the element level too, and the stale thing is a version stamp — a smaller defect than the one claimed, and one this paper elsewhere treats as serious when the Department's data carries it. This is the fourth instance of the shape recorded at corrections (27), (29) and (31): a statement of absence or staleness made before the audit's own holdings had been read to the grain that would test it. The re-extraction itself is committed with its tests: 422 entries across 43 matrices and 28 parent specialties; 13M's eleven business codes all at tier 2 where the first extraction had placed five of them at tier 1; the requirement word recovered on 395 of 422 entries, with each of the 27 blanks checked against the page and found to be a blank in the source; all 17 "Any Degree" entries retained with their tier and requirement. It remains accession-side and officer-only, and the route from its CIP codes to civilian occupations is still not published, for the reason given at Appendix A figure 68: only 105 of its 422 entries carry a single six-digit CIP code, the remainder being the directory's own wildcard, series and list forms, and how many of those the federal CIP-SOC crosswalk resolves has not been

measured. (This entry was itself re-checked against source before circulation: a first draft counted ten later-dated matrices, one of which proved to carry an earlier establishment date, and asserted a property of the federal crosswalk this audit has not fetched; both are corrected here rather than logged as a separate entry.)

- (37) 1 September 2026 — a stale count in the Executive Summary, and a matcher that read only one of the two forms in which the Air Force writes its own codes. The Executive Summary and the Contributions list said the dropped medical specialties exit into a marketplace offering them three programs; Part IV.7 said seven. Recomputed at apply time from the committed union corpus (`skillbridge_union_20260808.csv`, SHA-256 23268495...), with a matcher that reads occupation codes in both X-form (4R0X1) and skill-level form (4R031): seven — none for pharmacy, medical laboratory or physical medicine; one each for dental, ophthalmic and surgical service; two each for diagnostic imaging and mental health. Three was the census-grain figure. The same recompute corrected IV.2's companion sentence: dental draws one program (not zero), diagnostic imaging two (not one), munitions two (not zero), air traffic control three (not two) — every one of the earlier undercounts traceable to an X-form-only matcher or to census grain. The skill-level-form trap is the same defect the project logged against the 2020 AFECD check on 25 August, where fixing the matcher moved presence from 47 of 104 codes to 96. Every corrected figure recomputed through the receipt gate; ledger lines committed.
- (38) 1 September 2026 — a claim this paper's sharpest policy passage rested on was wrong, found by the recompute above. II.5 and IV.7 stated that across the 2,220 defense-industrial-base programs in the union corpus, not one names a Pararescueman, Combat Controller, TACP airman, or Special Reconnaissance airman, and that all six programs naming those fields sit outside the Patriot Pipeline presumption. One DIB program — an IT and network-operations program at a defense contractor — carries the TACP ladder at three skill levels in skill-level form (1Z331, 1Z351, 1Z371), which the X-form matcher missed; two further DIB programs name the specialties in prose without a code (a JTAC/TACP instructor role; a combat-controller experience requirement). The six-program code-named count itself reproduces exactly (five current-code programs plus one carrying the legacy Pararescue code 1T2X1), as does the one-names-all-four count. The passage is corrected from "cannot reach" to "barely reaches: one code-targeted program among 2,220." The correction cuts against the paper's convenience, which is why it is logged in full: the sharpened version — exactly one in 2,220, and that one reachable only by reading the code form the government's own directories use — is also the accurate one.
- (39) 1 September 2026 — the 23 August relevance sweep applied end to end, on the author's instruction. Thirty-nine findings (9 A, 18 B, 12 C). The four surviving "proven/better" phrases that correction (22) withdrew are removed (Part I, Part III, Part VIII, Appendix A row 2); the withdrawn 10.5% government-host figure is replaced in this appendix's limitations paragraph by the superseding 15.6% domain measurement (correction 14 / II.5); "complete directory" language is rescoped to IV.9's own vocabulary at its three sites; the Executive

Summary's gradeable-requirements count is brought to VII.1's seventeen; the all-six-branches figure is recomputed at union grain (4,154 of 4,255) through the receipt gate; and the scope epigraph now states the contractor carve-out it always practiced. On the author's standing rule for the public paper, the transfer/subsidy framing is removed from the Executive Summary, II.1, IV.7, V.1 and IX (the valuation figures and the verification argument are unchanged; the vocabulary now matches the paper's thesis). Duplicated restatements are condensed (Part I provenance and directory-cadence passages, II.5/IV.6/IV.7 person-day and Patriot Pipeline passages), absolute-absence phrasings are brought to finite-search form (Executive Summary outcome-data sentence; II.3 procurement rows), and eight instances of "catalogue" are Americanized. No figure changed except as logged at (37) and (38).

- (40) 2 September 2026 — a stated limitation that understated work already done, the fourth of its kind, prompted by the author. The Methodology and the limitations paragraph of this appendix said the blank MOU instrument "has not been obtained" and that quotes were "pending the blank instrument." The blank August 2024 release has been in this project's holdings since 8 August 2026 (meta-data title "DoD SkillBridge MOU_Revision_5.24.2024_v1," six pages, provider line blank), and a three-vintage diff of the instrument — 2024 blank, 2025 executed federal instance, 2026 signed instance — was written that same day and is committed. The 75% qualifying-offer warranty, the 85% KPI goal, the no-compensation clause, the member-cost bar and the 90/180-day reporting obligation all verify against the blank release directly. One wording note travels with the verification: the word "suitable" in the quoted phrase "qualifying offer of immediate post-service suitable employment" entered ¶4.2.10 in the 2025 technical update; the August 2024 release reads "immediate post-service employment." The Executive Summary's present-tense quotation is the text a provider signs today and stands; it is logged here so the quote is never read back onto the 2024 vintage. The drafted FOIA item for the blank instrument is withdrawn as satisfied. This is the shape corrections (27), (29), (31) and (36) record: an absence asserted before the audit's own holdings were read.
- (41) 2 September 2026 — a capture described as less than it is, and its counts withheld on a misread of its own flags; prompted by the author. The DCWF passage said the credential data searched was "a normalized copy," and withheld every count on the ground that its records are "marked as requiring provider verification before external use." The capture is the author's own complete pull of the Air Force COOL catalog, rebuilt into a working workbook and normalized into committed, git-tracked seed files; the verification marking is a per-record flag governing per-credential detail — a credential's requirements, cost, accreditation — not the capture's scope, so counts of what the capture contains were always publishable as statements about the capture. They now are: an 848-record credential inventory, and a master matrix joining 930 distinct credentials to 130 enlisted specialties across 5,285 specialty-credential pairs, all recomputed through the receipt gate at apply time. The DCWF search itself was re-run on 2 September 2026 on the current files — including a seed revised after the original 13 August search — and the zero holds: no "DCWF", no DCWF

work-role name, no work-role identifier. Three near-misses are disclosed in the rewritten boundary rather than left for a hostile reader. Two grains are stated because the files carry two normalization passes (848 inventory records; 930 matrix credential identifiers) — the lesson of corrections (33) and (34) applied prospectively. This is the sixth instance of the corrections-(27) family, and the second this month found only because the author pushed back on a limitation this paper had written about her own work.

- (42) 3 September 2026 — final-release polish, on the author’s instruction; no figure changed. The edition line is brought to September 2026. Acronyms are now introduced in full at first use in the Executive Summary and again at first use in the body, and used as acronyms thereafter. Emphasis is reduced to the headline sentence of each paragraph and the boxed findings. Remaining British spellings are Americanized. Appendix A’s figure register is restored to numerical order (rows 59–68 had been printed after rows 69–79). Every numeric token in the paper was counted before and after the pass and the two multisets are identical; that check is registered on the receipts ledger.

This log will grow; that is the point.

References

Statutes and instructions. 10 U.S.C. §§1142–1144 · VOW to Hire Heroes Act, P.L. 112-56 (2011), §§221–222 · FY2006 NDAA (crosswalk mandate, via CRS R48114) · FY2019 NDAA, P.L. 115-232, §552 · FY2023 NDAA, P.L. 117-263, §561 · FY2024 NDAA, P.L. 118-31, §572 · FY2026 NDAA, S. 1071 (signed 18 Dec 2025) · DoD Instruction 1332.35 (Transition Assistance, incl. Change 1, Jul 2025), esd.whs.mil · DoD Instruction 1322.29 (SkillBridge/JTEST-AI) · AR 600-81 (12 Mar 2024) · AFI 36-3037 · MCO 1700.31A · Navy transition instruction and NAVADMIN 064/23 · MARADMIN 280/24 · COMDTINST 1040.7A · DoD SkillBridge MOU, Aug 2024 release (skill-bridge.osd.mil/docs/DoD_SkillBridge_MOU_Aug_2024_Release.pdf), quoted from the signed DOI instance, Jan 2025 technical update (doi.gov/sites/default/files/documents/2025-05/2025dodskillbridgemouinterior.pdf). · Secretary of War memorandum, 3 Jun 2026, establishing a presumption of approval for SkillBridge requests within the defense industrial base (publicly briefed as the Project Patriot Pipeline; war.gov release, Jun 2026) · Air Force Instruction 36-2671, 31 Mar 2026 (SkillBridge roles, tiering, and the T-0 directory-use requirement at ¶4.5.3) · H.R. 8800, National Defense Authorization Act for Fiscal Year 2027, 119th Cong. (reported in House, June 2026), Secs. 771, 1723 (pending legislation; hashed copy in the capture manifest).

Oversight and analysis. *Cited in the body:* GAO-14-144 · GAO-17-133 · GAO-22-105261 · GAO-24-107352 (gao.gov/products/gao-24-107352) · RAND RR-284 (2014, Loughran — the veteran-unemployment counterweight cited in Part I) · RAND RR-A2289-1 (2023) · RAND RR-A1363-12 (2024). *Consulted and forming the ten-engagement GAO record described in Parts VI and VIII, not separately cited:* GAO-18-23 · GAO-18-135 · GAO-19-438R · GAO-23-104538 · GAO-23-106793 · GAO-24-107083 · GAO-24-106587 · GAO-24-107752 · GAO-25-107205 · RAND RR-1719-A (2017) · RR-1919-OSD (2017). CRS R48114, *Military Transition Assistance Program (TAP): Background and Issues for Congress*, Kamarck, 27 Jun 2024

(congress.gov/crs-product/R48114) · RAND RR-284 (2014) · RR-1719-A (2017) · RR-1919-OSD (2017) · RR-A2289-1 (2023) · RR-A1363-12 (2024) · DOL/ETA VOW §222 Pilot Study (SOLID, 2013-15; TEN 20-14, 2015) · DOL TAP Impact Study (Jul 2023, dol.gov/resource-library) · Milken Institute, *The Impact of SkillBridge on Veteran Employment Outcomes* (Jul 2025) · Penn State Clearinghouse O2O evaluation (2022) · Davenport et al., *J. Veterans Studies* 10(3) (2024) · CNAS, *Veteran Pathways to Employment* (2019).

Primary instrument under examination. O*NET military crosswalk, published by the National Center for O*NET Development at an undated path (/dl_files/2019/military_crosswalk.zip; the 2019 is a legacy directory, not a vintage indicator) — archive ee738141...5b01d747, 4,826,815 bytes, contents milx0724.csv, members stamped 22 July 2024, verified identical across two Internet Archive captures (26 Sep 2024, 22 Jul 2026) and two local copies · earlier vintages July 2020, October 2022, October 2023 · O*NET database release 30.3 (task, skills and knowledge tables) · Air Force Career Field Education and Training Plans, 106 documents, Air Force e-Publishing.

Data-governance and cyber issuances. DoDD 8000.01, *Management of the DoD Information Enterprise* · DoDI 8320.02 · DoDI 8320.07 · DoDI 8170.01, Appendix 3A (Change 2, 12 Mar 2025) · DoDD 5105.89, *Chief Digital and Artificial Intelligence Officer* · DoDI 8140.02 · DoDM 8140.03 · DoD 8570.01-M (canceled for military and civilian personnel 15 Feb 2023) · DFARS Subpart 239.71 and clause 252.239-7001 (JAN 2008) · 44 U.S.C. § 3511 (OPEN Government Data Act) · OMB Circular A-130 · Department of Labor Information Quality Guidelines (1 Oct 2002). All obtained from the DoD Cyber Exchange or the issuing office and searched locally; documents behind authentication were not retrieved.

Federal data. Airman Leadership School Syllabus, effective 23 January 2025 (held in full) · CCAF General Catalog 2022-2024, Change 1 (held in full) · SkillBridge public directory, complete per-state export via api.skillbridge.mil (3 Aug 2026; raw multipart archive committed) · USAspending prime-award records: 47QFHA22F0026 (SOLID), H9821024CE001 (Goldbelt Apex), 1605C424F00015 (Serco), W9124D-23-C-0011 (Inverness), primary exports 4 Aug 2026 · Census Bureau Veteran Employment Outcomes v2.0, lehd.ces.census.gov/data/veo (occupation file veoo3, retrieved 4 Aug 2026; paygrade file veop.csv, 528 rows, obtained 2 Sep 2026, SHA-256 b00ec189..., committed) · BLS Employment Situation / veteran tables, Apr 2026 (release 8 May 2026; via IVMF monthly summary) · VA Post-Separation TAP Assessments (benefits.va.gov/TRANSITION/tap-assessments.asp) · VA OTED Annual Benefits Report 2025 · DHRA FY2026 Budget Justification, OP-5 (comptroller.war.gov) · DOL VETS FY2025 CBJ and FY2026 Budget in Brief · Air Force FY25 SRB list (AFPC, Dec 2024); FY26 SRB list via Air & Space Forces Magazine (19 May 2026).

Press and secondary. Littler, *DOD Imposes New Requirements for Employers Participating in SkillBridge* (2024) · MOAA, *DoD Overhauls Jobs Program...* (2024) · Military.com on USMC SkillBridge cuts (18 Jun 2024) · Military Times on AF/SF SkillBridge tiers (3 Apr 2026) · Stars & Stripes on TAP reform (11 Sep 2025; 25 Mar 2025) · Task & Purpose on AF separation-order delays (12 May 2025) · Air & Space Forces Magazine (19 May 2026) · Hiring Our Heroes HVAC testimony (O'Brien, 25 Mar 2025, docs.house.gov) · VFW HVAC testimony (Craig, Mar 2025) · militarytransitiontoolkit.com branch-approval compendium (2026; service rules cross-checked to primary instructions).

Employer-side evidence. Captured job boards and careers pages with URL and observed date for all conversion-pilot hosts and realty-funnel specimens (evidence tables: `conversion_pilot_evidence.csv`, 18 rows; captures committed 3-4 Aug 2026). Requisition corpus: canonical requisition table (384 unique requisitions, top-20 defense contractor, Jul 2026) and market-anchor evidence base (629 posting rows, 30 July 2026 build, with per-row employer/req/URL/date/hash; expanded 14 August 2026 to 21,086 specialty-lane rows resolving to 4,108 distinct postings; manifest with source hashes).

Supporting working papers (committed alongside the evidence base; each carries its own scope and limitation notes): Market Coverage Findings · MOC Targeting Drilldown · Conversion Audit Pilot · MOU Requirements vs Directory · Money Trail · TAP Ecosystem Map · Citizen Compliance Audit Framework · Three Lenses · FOIA Register · Congressional Engagement Map · O*NET governance, paper-trail, and audit-reconciliation series.

*MissionProof Research — the verification layer the transition system is missing.
Contact: lauraedryden@gmail.com*