DAFITC 2026 · TRACK 1

Rebuild a 40-year-old system without losing a single business rule. Or inventing one.

ADOM Refactor → ADOM · A deterministic path from undocumented legacy code to an authorized modern system.

```
COMPUTE WS-TAX = WS-BASE * WS-RATE
```

THE STAKES

The rules that run the business live only in the code, and nobody can read it anymore.

THE PEOPLE WHO WROTE IT

THE THESIS

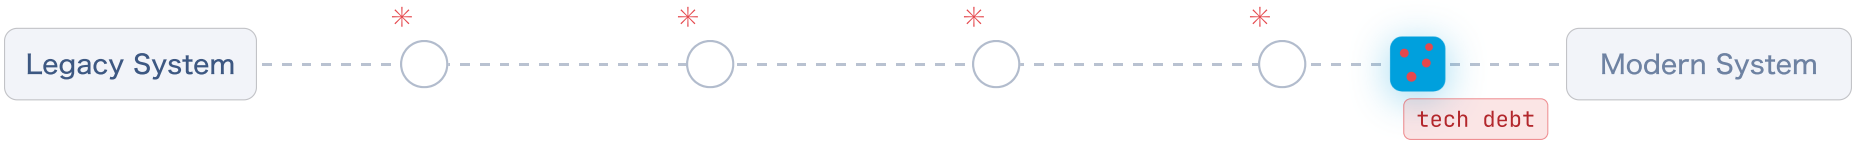
The hard part was never writing new code. It's **proving** the new code does exactly what the old code did, and showing your work.

Misread one rule → The new system quietly does the wrong thing → Nobody notices for months → A mispriced bill. A wrongly denied claim.

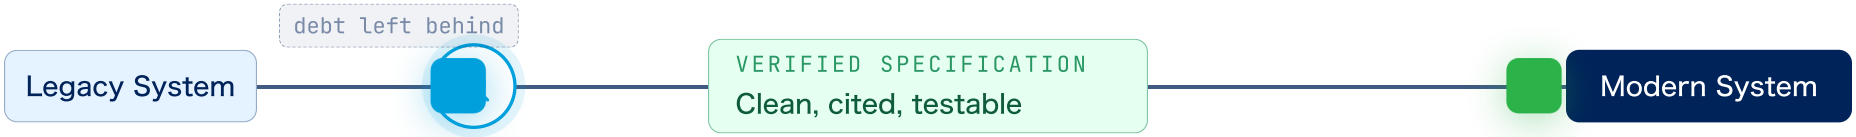
Line-by-line translation is a game of telephone.

Every hop is a chance to carry a defect forward, and the technical debt rides along the whole way.

PATH A · DIRECT TRANSLATION



PATH B · SPEC-DRIVEN ADOM



ADOM doesn't translate the code. It reads the code to rebuild from a **clean, verified specification.**

BRING THIS HOME

Don't ask AI to “convert this file.” Ask it to **explain what the file does and cite the lines**, then rebuild from the explanation you verified. Translation carries the bugs forward; specification lets you leave them behind on purpose.

This isn't just AI reading code.

THE SPINE · DETERMINISTIC

Real parsers are ground truth.

✓	**FILE**	TAXCALC.cbl · 1,284 lines
✓	**SYMBOL**	WS-RATE · S9V9(4) COMP-3
✓	**EDGE**	CALL 'GLPOST' → 1300-POST
✓	**FORMULA**	line 412 · COMPUTE WS-TAX

RECONCILE

THE AI · CLAIMS ONLY

Proposes, never decides.

CLAIM · IN FLIGHT

“Sales tax is the base amount times the rate, truncated to the cent.”

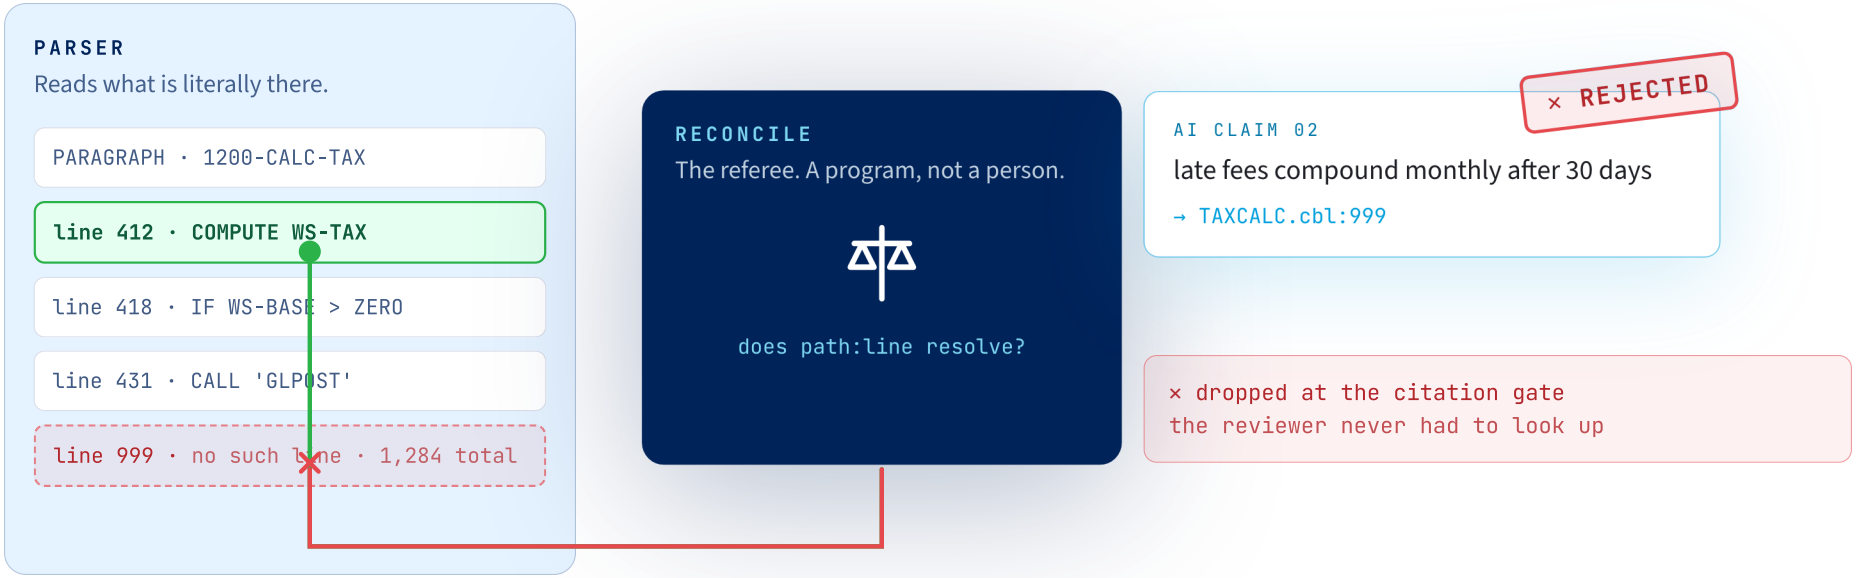
cites TAXCALC.cbl:412

It cannot cross until the seam checks it.

The parser can't guess. It reads what's on the line. **The AI proposes; the spine decides what's true.**

Every claim has to survive the citation gate.

PRESS R TO REPLAY

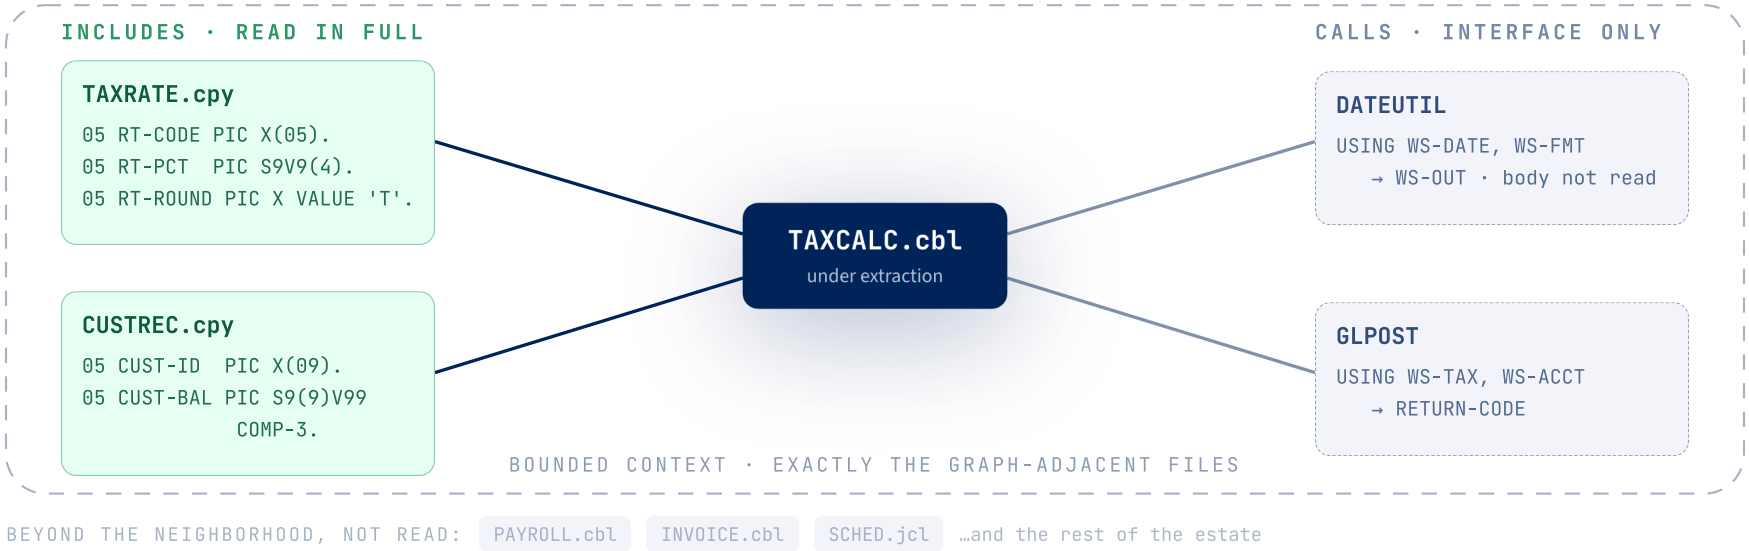


The gate is a program, not somebody's vigilance. **The AI can still be wrong; the wrongness gets caught by code,** not by hoping a tired reviewer spots it.

BRING THIS HOME

Force every AI claim to carry a `path:line` citation, then **check the citation with a script, not your eyes.** Cite-or-die kills most hallucinations before a human wastes a second on them.

Never one file alone. Never the whole repo.



BEYOND THE NEIGHBORHOOD, NOT READ: `PAYROLL.cbl` `INVOICE.cbl` `SCHED.jcl` …and the rest of the estate

Read one file alone and you miss half the logic. Read the whole repo at once and **you can't reproduce the reading twice.**

BRING THIS HOME

Give the model a **computed neighborhood**: includes in full (field widths and rounding hide there), calls as interfaces only. Graph-derived context is both more accurate and reproducible.

The half-cent that sinks migrations.

R-0412 · FIDELITY · reproduce as-is TAXCALC.cbl:412

PLAIN ENGLISH

Sales tax is the base amount times the rate, truncated to the cent, not rounded.

FORMULA

```
WS-TAX = WS-BASE * WS-RATE
```

DECIMAL SEMANTICS

PIC S9(7)V99 COMP-3 · scale 2 · ROUND: TRUNCATE

CONTROL FLOW

after 1200-VALIDATE · before 1300-POST

ERROR / RECOVERY

on overflow → 9999-ABORT (lifted as its own rule)

LAYER-1 BOUNDARY TEST

base = 1.005, rate = 0.50 → expect 0.50

BOUNDARY TEST BENCH

1.005 × 0.50

↓

TRUNCATE	ROUND
0.50	0.51
✓ matches legacy	× fails

The rounding mode rides **inside** the reproducibility hash. The grammar reads it. No model judgment involved.

The test is built to land exactly where truncating and rounding disagree. **A rebuild that “cleans up” the math fails it. There's nowhere to hide.**

ACT II · REPRODUCIBILITY

What “reproducible” has to mean.

INSIDE THE HASH · THE REPRODUCIBLE CORE

- Deterministic rules
- Their citations
- Decimal / rounding semantics
- Layer-1 boundary tests

ATTACHED AFTER THE HASH · CAN'T PERTURB IT

- Advisory rules
- System-model narrative
- Module catalog
- Personas / user stories

Everything a model touched clips on **after** the seal, so it can never move the number that proves the core is faithful.

RUN 1 `content_sha: 9f2a7d13c85be0416a2f0d9744bbe1c41d`

RUN 2 `content_sha: 9f2a7d13c85be0416a2f0d9744bbe1c41d`

✓ IDENTICAL · BYTE FOR BYTE

Generate the deterministic core twice from the same frozen source and **you get the same hash, to the byte.**

BRING THIS HOME

Separate **grammar-backed fact** from model judgment and hash only the former. If your “reproducible” artifact changes when the model's mood shifts, it isn't reproducible.

Eight checkpoints. Nothing auto-advances.

CP0	Source frozen	Code pinned to one sha256 · the run is against one immutable snapshot.	AUTOMATED
CP1	Indexing complete	Every file, symbol and dependency mapped to 100%; a random sample independently re-verified.	AUTOMATED
CP2	Dead code reviewed	Unreachable code surfaced as evidence. Nothing deleted automatically; a human decides.	HUMAN GATE
CP3	Knowledge ready	Every extracted rule is citation-clean. A rule with no resolvable line doesn't exist.	HUMAN GATE
CP4	Tests signed off	Every confirmed rule pinned with a boundary test, attested by a developer.	HUMAN GATE
CP5	Spec ready for review	Unified Spec assembled, every citation resolves. A readiness check, not a sign-off.	AUTOMATED
CP6	Client approval	The client's named approver signs that the spec is faithful and the rebuild can begin.	★ SIGNATURE
CP7	Handoff complete	Build package (spec, tests, contracts, full frozen source) shipped to the ADOM build platform.	HUMAN GATE

LEGEND

- Automated · the machine runs it
- Human gate · a person must clear it
- The one client signature

× CP2 · the token can't skip a stopped gate

Everything upstream was building **the evidence for this moment.**

BRING THIS HOME

Put **hard gates** between your AI's mechanical work and every judgment call. “The AI did it” is not an accountability model; a signed checkpoint is.

ACT III · COVERAGE

Coverage is computed, never claimed.

THE MANIFEST IS THE DENOMINATOR

Plain, non-AI code walks the tree and counts every file, symbol and edge.

DENOMINATOR LOCKED

146 files

frozen at one sha256

The AI cannot shrink the denominator to look finished.

APPEND-ONLY INDEX LEDGER — units at state ≥ indexed ÷ manifest count

CBACT01C.cbl	12 paragraphs · 4 edges	VERIFIED
COACTUPC.cbl	38 paragraphs · 11 edges	VERIFIED
CBTRN02C.cbl	re-read: 9 of 12 accounted for	RE-QUEUED
COBIL00C.cbl	21 paragraphs · 6 edges	INDEXED

BINARY GATE NOT AVERAGED — 100%

REAL RUN
AWS CARDDEMO

146	**395**	**30 / 31**	**100%**
source files ingested & frozen	business rules extracted, each cited	programs fully extracted · the outlier documented, not dropped	rule recall on an independent re-read

The agent can't mark a unit done by fiat. **Completion is granted only when its evidence reconciles with the computed facts.**

BRING THIS HOME

Never let the model define its own “done.” Count the units yourself, hand it **one at a time**, and independently re-check a random sample. Deterministic outer loop, non-deterministic worker.

One verified specification. Two ways out.

UNIFIED BUILD SPECIFICATION

395 cited rules · tests · contracts · hashed core

ROAD A

Clean Re-Stack

Same or similar language family. All clean code. The technical debt stays behind.

Familiar stack | Mirrored boundaries | Zero carried debt

R-0412 R-0418 R-0431 + 392 more · identical set

ROAD B

Cloud Modernization

A cloud-native target, re-architected components, modern managed services.

Azure Government | Private endpoints | Managed services

R-0412 R-0418 R-0431 + 392 more · identical set

The extracted behavior is the same either way.
The target is a human decision.

THE MACHINE RECOMMENDS · THE HUMAN CHOOSES

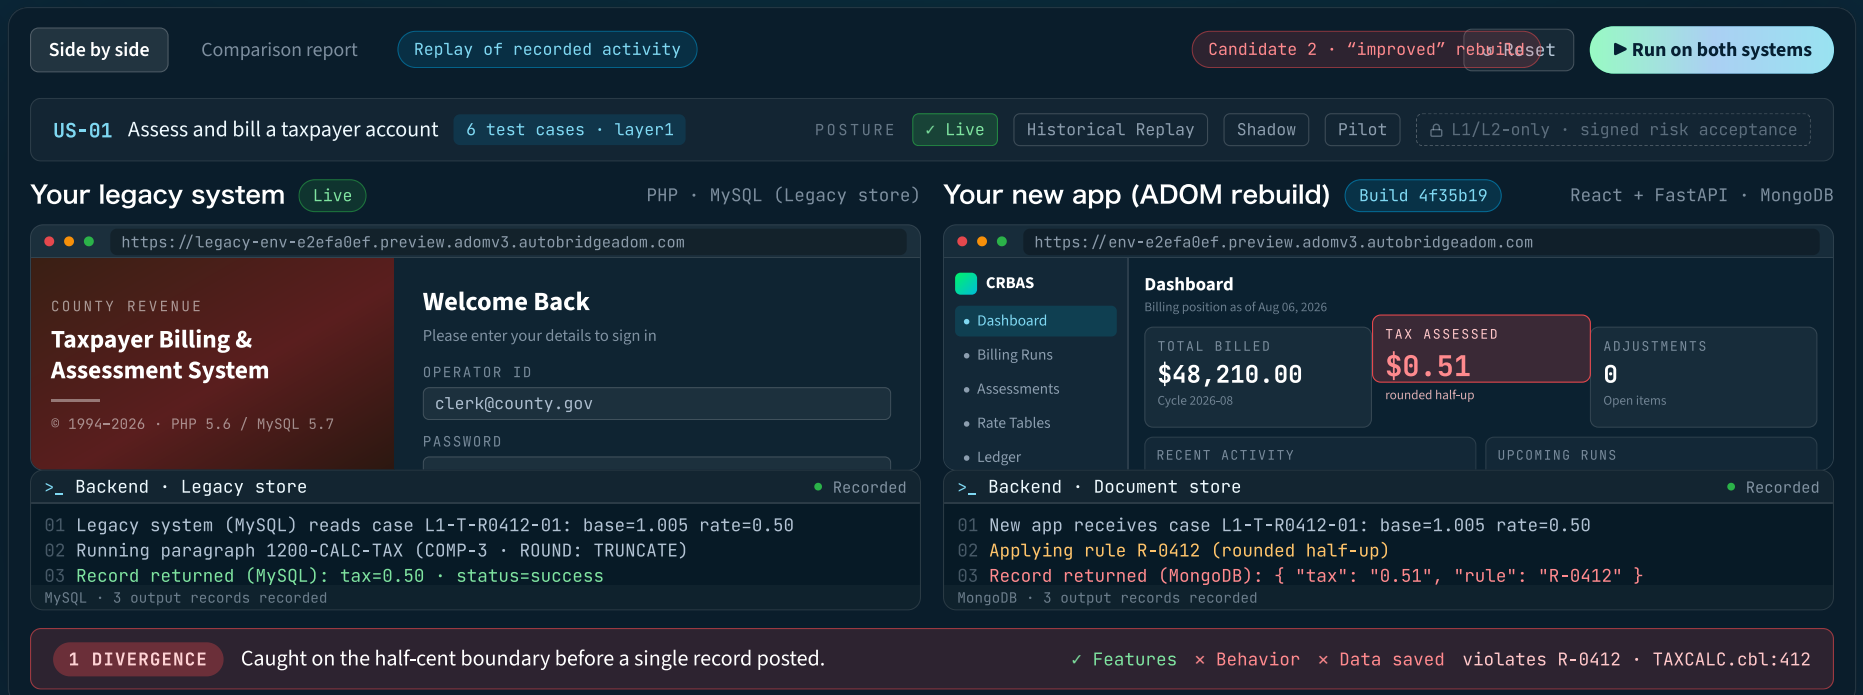
language · architecture pattern · component boundaries

OVERRIDE

ACT III · THE PROOF

Run both systems. Diff the output.

PRESS R TO REPLAY

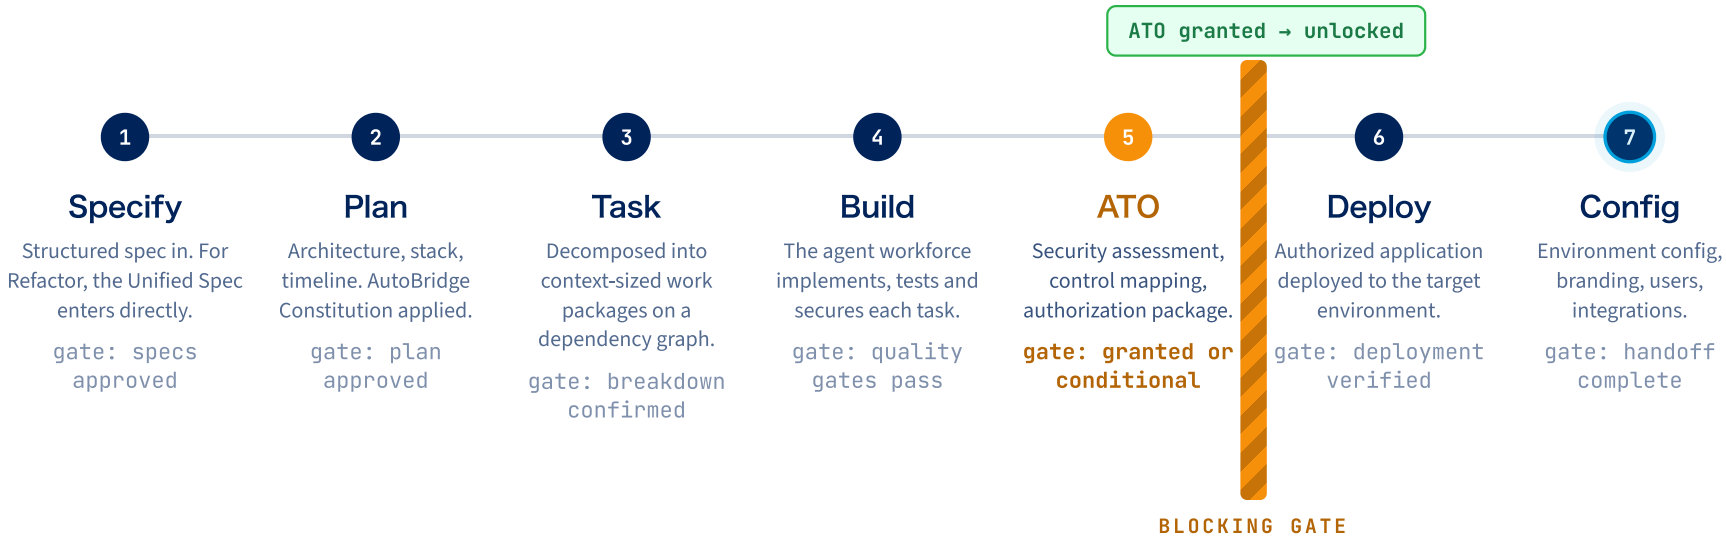


The difference between "we think it matches" and **"we ran it and it matches."**

BRING THIS HOME

Prove equivalence by running old and new on the same inputs, and **engineer those inputs to hit the boundaries** where a rewrite silently diverges.

Seven phases. ATO is a wall, not a follow-up.



An application cannot deploy until it's assessed and granted an ATO. **Compliance is a step in the build, not a project that starts after it.**

An agent workforce, inside enforced boundaries.

PM Agent · orchestrator

THE GENERATED SPEC IS THE BOUNDARY

Frontend

`UI components`

allow-list: read repo, write /ui, run tests

CONSTITUTION

✓ QA pass

Backend

`Services`

allow-list: read repo, write /svc, run tests

CONSTITUTION

✓ QA pass

Database

`Schema + migrations`

allow-list: write /db · migration needs human approval

CONSTITUTION

✓ QA pass

API

`Endpoints + contracts`

allow-list: write /api, run contract tests

CONSTITUTION

✓ QA pass

Security

`Auth + encryption`

allow-list: read all, write /sec · IAM needs approval

CONSTITUTION

✓ QA pass

QA

`Runs the legacy-derived test suite`

every module, immediately, as it's produced

× fail → routed back to the spec

Agents build exactly what the spec defines: no more, no less. **The build proceeds within enforced boundaries, not on trust.**

BRING THIS HOME

Constrain agents to a **spec as the boundary** and give each one a **tool allow-list**, the narrowest set of actions it needs. Log every action. A narrow, checked agent is one you can actually authorize.

Compliance as a byproduct of the build.

THE ATO SUB-PIPELINE

1. **Framework Selection**
   CMMC L2 / FedRAMP Moderate · boundary scope
2. **Pre-ATO Staging**
   Isolated environment mirroring production
3. **Security Scans**
   SAST, DAST, infra, access, encryption, logging, secrets
4. **Controls Assessment**
   Evidence mapped to NIST 800-53 · SPRS score + compliance %
5. **Gap Analysis**
   POA&M with remediation steps and effort estimates
6. **SSP Generation**
   Auto-populated from real build data, then partner-signed
7. **eMASS Package**
   Assembled and exported in eMASS-compatible format

SCAN → CONTROL MAP

SAST · Semgrep / SonarQube	SA-11 · SI-10
DAST · OWASP ZAP	SA-11 · SI-10
Infra · Checkov / Azure Policy	CM-6 · CM-7
Access · Entra / IAM	AC-2 · AC-3 · AC-6
MFA	IA-2
Encryption	SC-8 · SC-28
Audit logging	AU-2 · AU-3
Network	SC-7 · AC-17
Secrets	IA-5 · partial

READINESS

- incomplete
- ready for review
- export-ready

EMASS PACKAGE

- ✓ SSP · PDF + JSON
- ✓ SAR · DISA format
- ✓ POA&M · Excel + PDF
- ✓ Evidence ZIP
- ✓ Artifact Manifest · SHA-256

SECURITY FOR SCALE

Runs inside the customer's cloud boundary: Azure Gov for CUI, private endpoints, public access disabled.

IL4-aligned hardening: default-deny, customer-managed keys, HSM-signed audit logging, least-privilege identities.

AI governance at build time: classification-aware model registry, four-layer prompt-injection defense, per-agent tool allow-lists.

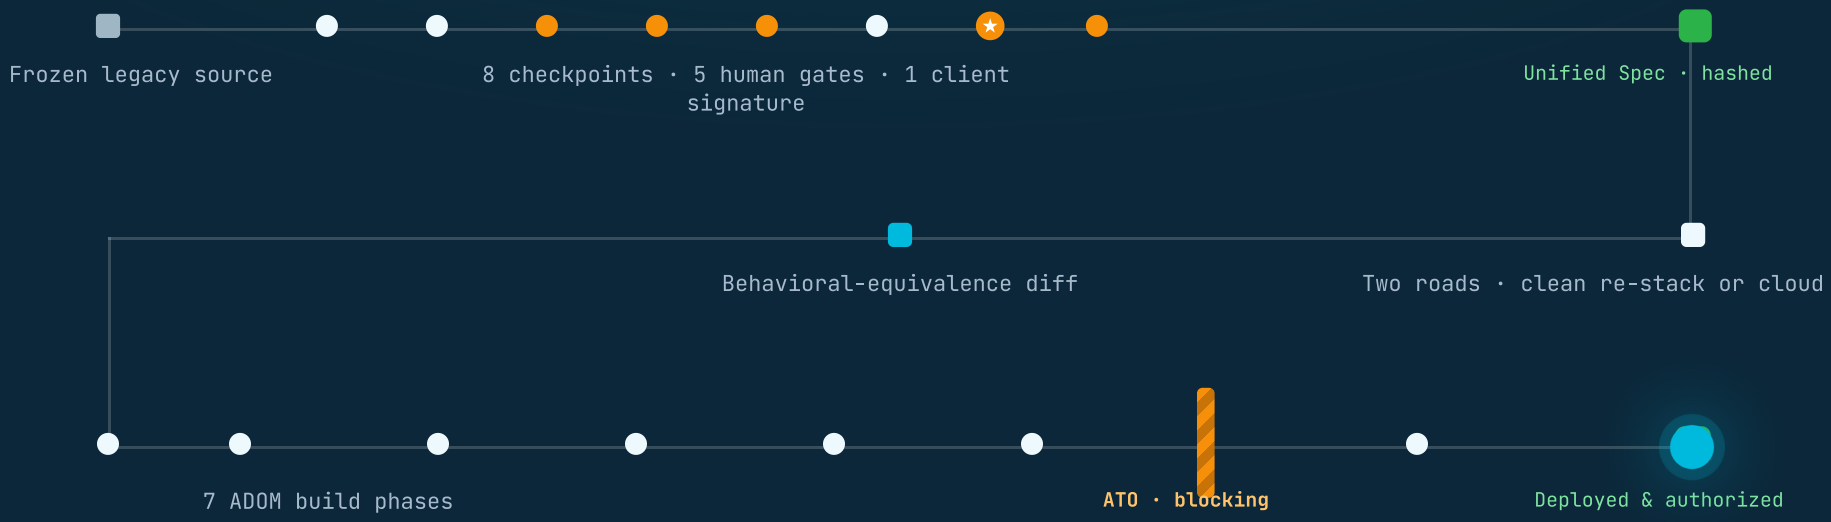
Human approval for every consequential action (deploys, IAM, schema migrations) cryptographically signed into a tamper-evident chain.

BRING THIS HOME

Generate compliance evidence **from the system you built, as you build it**: map scans to controls automatically and keep a hashed manifest. That beats compliance reconstructed by hand months later.

THE FULL ARC

Legacy line → modern artifact → authorization evidence.



THE LINE WE STARTED WITH

```
COMPUTE WS-TAX = WS-BASE * WS-RATE
```

now **R-0412** · preserved, cited, boundary-tested, and passing in the new system.

Humans in control at every gate. **Faster, without reckless.**

DAFITC 2026 · TRACK 1

Scale through automation. Control through humans.

AUTOBRIDGE SYSTEMS · ADOM REFACTOR → ADOM

THE ASK

A **bounded pilot**: one legacy application or domain, taken end-to-end from extraction through build to a generated eMASS authorization package, with behavioral equivalence verified against the legacy system.

THE PORTFOLIO A PILOT CALIBRATES

84 legacy systems

118+ applications

AF/A1 Agile Refactoring & Integration

AutoBridge Systems

CAGE 17GB7 · UEI DUY6ZAM6Z8T5 · NAICS 541512 · Small Business, Minority-Owned
micah@autobridgesystems.com · (202) 372-7829