DAFITC 2026 · TRACK 1 · AI-DRIVEN LEGACY SYSTEM MODERNIZATION

FROM 6 WEEKS TO 6 DAYS

The agentic-AI release factory.

TRAINING SESSION · 40 MIN

COREY SOLIVAN · DATA ROOM AI / TIE LLC

THE QUESTION IS CHANGING

TWO QUESTIONS, ONE DECADE APART

YESTERDAY'S AGILE SPRINT ASKED

"What can we deliver in **six weeks**?"

THE NEW SDLC ASKS

"What can human-agent teams safely design, build, test, secure, and release in **six days**?"

The value is not that AI writes code quickly. The value is that agentic AI **compresses the translation layer** between every SDLC step. The future SDLC is not faster typing. It is **faster convergence**.

THE EVIDENCE BASE

THE FACTORY IS NOT HYPOTHETICAL

Everything in this session runs daily in our own delivery system. Measured from GitHub, Dec 01 2025 – Aug 15 2026.

28.1 PR CYCLES / DAY · LAST 30D	**225** PR CYCLES · PEAK WEEK	**94.9** COMMITS / DAY · LAST 30D	**93** REPOSITORIES IN THE FACTORY
2,539 PR CYCLES · 259 DAYS	**2,597** CODE REVIEWS	**100%** OF APPROVALS HUMAN	**11,143** REVIEW COMMENTS

SIX-DAY CYCLES ARE AN OPERATING RHYTHM, NOT A STUNT — THE TELEMETRY PANEL BEHIND THESE NUMBERS OPENS LIVE, LATER.

SECTION 01 · THE PARADIGM SHIFT

THE SIX-WEEK SPRINT IS BECOMING THE SIX-DAY RELEASE

Agile reduced waterfall waste. Agentic AI reduces agile handoff waste.

SECTION 01

THE BOTTLENECK MOVED

THE OLD CONSTRAINT

Developer capacity

Sprints exist to ration scarce build hours. Everything queues behind the keyboard.

THE NEW CONSTRAINT

Alignment speed

How quickly can intent, design, code, tests, and release evidence converge on one validated change?

Speed without validation is chaos. Speed with traceability is **advantage**.

SECTION 01

YESTERDAY: HANDOFFS, MEETINGS, WAITING, REWORK

Nine stations, eight handoffs — and context leaks at every joint.



Each step often loses context. Every handoff creates translation error, delay, and rework — that is where the six weeks live.

SECTION 01

TODAY: ONE INTELLIGENT THREAD

From user need to deployment evidence — every artifact generated in context of the last, with the codebase in view the whole way.



That chain used to be slow because humans manually translated each step. **Agentic AI compresses the translation layer** — humans stay where judgment lives: intent, priorities, gates.

SECTION 01

WHAT EACH FUNCTION BECOMES

FUNCTION	NEW AGENTIC-AI CAPABILITY
User research	Synthesizes tickets, feature requests, discrepancy reports, DX/CX trends — continuously
UI/UX design	Generates design-tool-ready concepts in minutes to hours; simulates click paths
PRD	Written with the codebase, schemas, and functions in view — code-aware from the first draft
Implementation plan	Produced as a byproduct of PRD plus technical review
EPICs / issues	Auto-decomposed with acceptance criteria, technical notes, definition of done
Development	Parallelized on file, function, and dependency awareness
Testing	Generated across schema, unit, integration, E2E, and browser user journeys
Security	Continuous code, OS, container, and dependency scanning with remediation paths
Pull requests	Draft PRs carrying integration and deployment testing evidence

SECTION 02 · THE OPERATING RHYTHM

THE SIX DAYS, ONE BY ONE

From validated need to deployment-ready change — with a human gate every day.

DAY 1 DAY 2 DAY 3 DAY 4 DAY 5 DAY 6

01 RESEARCH SYNTHESIS

Research no longer starts the cycle. It feeds every cycle.

INPUTS — ALWAYS FLOWING

- Discrepancy reports and trouble tickets
- Feature requests and support themes
- User behavior signals
- DX/CX trends and operational pain points
- Codebase reality — what the system can absorb

OUTPUT BY END OF DAY

- A ranked, evidence-backed change thesis
- What users need · cited to source tickets
- What is breaking · with frequency data
- What creates the highest mission value

HUMAN GATE · DAY 1

Product owner approves the ranked thesis — agents rank, humans decide what matters.

FROM OUR FACTORY

1,786 issues worked in 259 days — agents posted 8,295 triage and synthesis comments along the way.

DAY 1 | DAY 2 | DAY 3 | DAY 4 | DAY 5 | DAY 6

02 UX + CODE-AWARE PRD

Design in minutes. Validation in hours. Iteration without waiting.

DESIGN BECOMES RAPID SIMULATION

- Concepts and wireframes generated, compared, discarded cheaply
- Design-tool-ready UI patterns in hours, not sprints
- Click paths simulated; accessibility gaps flagged early
- Design changes linked directly to required code changes

THE PRD KNOWS THE CODE

- User need + functional requirement
- Codebase, schema, API, and UI impact · enumerated
- Test impact and security impact
- Acceptance criteria + required release evidence

HUMAN GATE · DAY 2

Design direction and PRD signed off — intent is validated before a line of build.

THE SHIFT

Requirements know the code before development begins — old PRDs described intent; new PRDs describe impact.

DAY 1 DAY 2 DAY 3 DAY 4 DAY 5 DAY 6

03 PLAN + ISSUES

When the PRD understands the code, the plan writes itself.

THE PLAN IS A BYPRODUCT

- Files to modify, functions to update, schemas to change
- Dependencies to inspect; risks and constraints
- Tests to create or update
- Rollback considerations and deployment sequence

NO MORE THIN TICKETS

- EPICs → features → stories → tasks, auto-decomposed
- Every issue: acceptance criteria, technical notes, test requirements
- Security considerations and definition of done
- Dependency chains estimated up front

HUMAN GATE · DAY 3

Tech lead approves scope, sequence, and the parallelization plan.

FROM OUR FACTORY

310 issues opened in a single week (Jul 05) — a full backlog decomposition wave, grooming compressed from days to minutes.

DAY 1 DAY 2 DAY 3 DAY 4 DAY 5 DAY 6

04 PARALLEL BUILD

Parallel work without parallel chaos.

WHY PARALLEL IS FINALLY SAFE

- File-level and function-level dependencies mapped
- Shared schemas and API contracts known
- Component ownership and test coverage visible
- Integration risk assessed before work is assigned

HOW THE DAY RUNS

- Independent workstreams fan out to human-agent pairs
- Small, traceable commits against issue-scoped branches
- Colliding streams serialized by design, not discovered in merge
- Result: fewer collisions at higher speed

HUMAN GATE · DAY 4

Increment-level code review — humans steer while agents produce.

FROM OUR FACTORY

94.9 commits/day over the last 30 days · peak day 629 commits — sustained, not spiky.

05 TEST + SECURITY

The fastest teams do not skip testing. They automate more of it.

THE GENERATED TEST STACK

- Schema, function, unit, and integration tests
- End-to-end and regression suites
- Playwright user journeys — real click-throughs in real browsers
- Accessibility checks in the same pass
- Hundreds of manual QA hours → minutes of simulation

SECURITY AS CONTINUOUS CO-PILOT

- Code, OS, container, and dependency scanning per change
- Outdated packages detected with remediation paths
- Package remediation inside the delivery flow — not a separate backlog
- Policy-as-code validation before the gate, not after

HUMAN GATE · DAY 5

Risk acceptance — a person owns every finding that ships open.

FROM OUR FACTORY

2,597 reviews across 2,539 PR cycles — ≈1.02 reviews per PR, scans on every one.

● LIVE DEMO 03 · PLAYWRIGHT JOURNEY

DAY 1 | DAY 2 | DAY 3 | DAY 4 | DAY 5 | DAY 6

06 DRAFT PR = RELEASE EVIDENCE

The pull request becomes the release control tower.

THE DRAFT PR CARRIES

- Code changes + impact summary
- Test results and coverage movement
- Security scan results and remediations
- Deployment test output and rollback notes
- Linked issues + acceptance criteria mapping
- Human review checklist — where to look first

WHAT THE REVIEWER NOW DOES

- Audits an evidence package, not a cold diff
- Checks the traceability chain end to end
- Applies judgment where the machine flagged doubt
- Approves — or sends it back with context intact

HUMAN GATE · DAY 6

Go / no-go — release authority stays with a person, on evidence.

FROM OUR FACTORY

28.1 PR cycles/day over the last 30 days · peak week 225 — every approval human.

SCREENSHOT SET A · DRAFT PR WALKTHROUGH — SEE FACILITATOR GUIDE

SECTION 02 · UNDER THE HOOD

THE MACHINERY BENEATH THE SIX DAYS

MCP · SYSTEM AWARENESS

Model Context Protocol tools let agents interact with files, codebases, schemas, functions, browsers, dev tools, and enterprise systems in a structured, auditable way — awareness, not guesswork.

PLAYWRIGHT · SIMULATED USERS

End-to-end journeys, UI click-throughs, regression paths, and browser validation that once consumed hundreds of manual QA hours — now generated and executed per change.

CONTINUOUS SCANNING · CODE TO CONTAINER

SAST, dependency, OS, and container scans with package-update detection and remediation suggestions — wired into the same pipeline that ships features.

None of this is exotic. It is **standard tooling, orchestrated by agents, governed by gates** — which is exactly why it exports to other teams.

SECTION 03 · PROOF FROM THE FACTORY

WHAT COMPRESSION LOOKS LIKE IN TELEMETRY

Not a benchmark. Our own delivery history, week by week.

SECTION 03

THE CADENCE SHIFT, MEASURED

Pull-request cycles per week, all contributors · Dec 2025 – Aug 2026 · hover any week

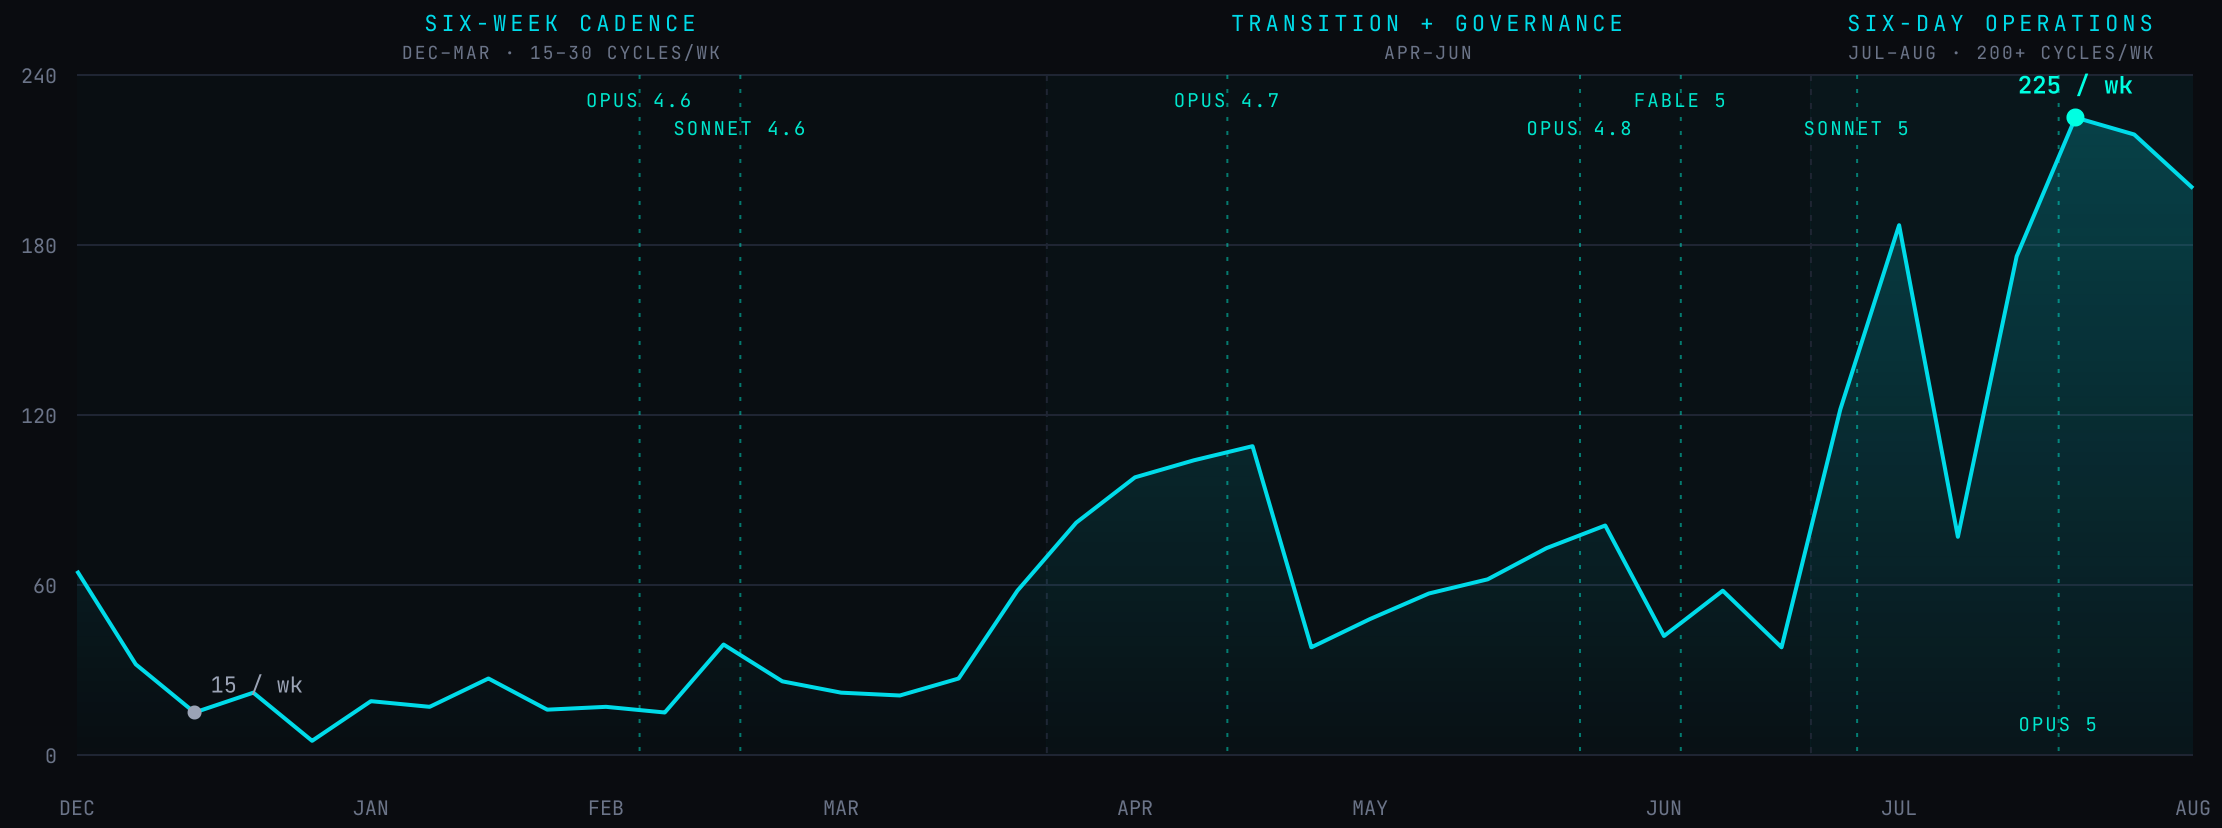

● LIVE DEMO 01 · WEEKLY DEV PANEL

SECTION 03

FAST — AND GOVERNED

First 90 days vs most recent 90 days — throughput and scrutiny grew together



1.02 REVIEWS / PR 100% OF APPROVALS HUMAN 0 APPROVALS BY AGENTS — BY POLICY

The result is not reckless speed. It is **safe change velocity**.

SECTION 03 · THE MONTHLY ARITHMETIC

FOUR MONTHS. THREE YEARS.

The cadence normalized to one 30-day month — the same unit on both sides. Apples to apples.

x8.0	x9.2	x5.7	x7.1	x3.5
PR CYCLES	ISSUES	REVIEWS	APPROVALS	COMMITS
105 → **842** /mo	53 → **487** /mo	140 → **794** /mo	88 → **628** /mo	824 → **2,847** /mo

BASELINE = DEC 01–FEB 28 ÷ 3 — EXACTLY 3 × 30 DAYS · NOW = LAST 30 DAYS, JUL 17–AUG 15 — OPUS 5 (JUL 24) IN-WINDOW, FABLE 5 + SONNET 5 UPSTREAM

PROJECTION · SEP 01 – DEC 31 · CURRENT CADENCE HELD FLAT

4 MONTHS ≈ 3 YEARS

3,368 PR CYCLES = 32 BASELINE MONTHS (2.7 YRS) · 1,948 ISSUES = 37 (3.1 YRS)

AT THE OLD PACE, THIS WORK CLEARS IN 2029.

Assumes zero further improvement — no new model releases, no added scaling. The curve two slides back says that assumption is pessimistic.

SECTION 03 · RECAP

THE SIX-DAY SDLC ON ONE SCREEN

Click a row to relight its day.

DAY	FOCUS	PRIMARY OUTPUT
Day 1	Research synthesis	Ranked, evidence-backed change thesis
Day 2	UX + PRD	Validated design and code-aware PRD
Day 3	Planning + issues	Implementation plan, EPICs, issues with acceptance criteria
Day 4	Parallel build	Agent-assisted code changes in small, traceable increments
Day 5	Test + security	Automated test suite, scan results, remediations
Day 6	Draft PR + release evidence	Evidence package, validation output, go/no-go review

EVERY DAY ENDS AT A HUMAN GATE. EVERY ARTIFACT TRACES TO THE ONE BEFORE IT.

▼ THIS IS THE SLIDE YOU MAY WANT TO PHOTOGRAPH ▼

SECTION 04 · THE MINDSET SHIFT

STOP ASKING WHAT FITS IN THE SPRINT

SECTION 04

THE QUESTION THAT CHANGES THE ORG

OLD QUESTION

“How much can we complete in the sprint?”

NEW QUESTION

“How fast can we safely move from validated user need to tested, secure, deployment-ready change?”

The winner is not the team with the most AI.
The winner is the team with the most **disciplined AI-enabled release system**.

TAKE IT WITH YOU

RUN YOUR FIRST SIX-DAY CYCLE ON YOUR OWN BACKLOG

THE AGENTIC DELIVERY EXPORTABLE

The factory in this session ships as a deployable kit: the six-day cycle playbook, agent role definitions, gate templates, test and scan pipeline configs, draft-PR evidence formats, and the telemetry panel you watched live — stood up inside your environment, on your repos. Ask for the exportable brief.

CONTACT

Corey Solivan

corey@product-ties.com · 703-955-0326

Data Room AI · TIE LLC

ATDC · C/O Data Room AI · 75 5th St NW, Suite 2310 · Atlanta, GA 30308

CONNECT

youtube.com/@coreysolivan

linkedin.com/in/corey-solivan

product-ties.com

draics.com

draics.com/DAFITC26 · this brief

aqtx1.substack.com · AQtx

What's possible today · what's coming tomorrow · what's stopping us now.

CONTRACT

OTA · SBIR · CSO

corey@draics.com

Rapid-acquisition pathways — vehicle mapping in the exportable brief.

QUESTIONS — 5 MIN