



Agentic AI Foundation

AGNTCon + MCPCon North America

Where the **agentic stack** is being built.

The Agentic SDLC: How We Shipped AI-Native Software at a Legacy Company

Austin Brown · A Place for Mom
with the Grace engineering team

A quick show of hands

Who here has played Grand Theft Auto?

Where the **agentic stack** is being built.

Are AI coding assistants like cheating in a video game?

Or are they another tool in the toolbox?

Hi, I'm Austin.

ABOUT ME

Principal Engineer


10x Meeting Attender

paid to argue with robots

- **Founding engineer on Grace**, our digital advisor for senior care
- **I help our engineers build with AI**, and measure what it really changes
- **I like building things and am an avid boater** 

jQuery, Angular, and Yet-Another-JS-Framework

ABOUT ME



- The glory days of jQuery and jQuery UI
- Endless plugins, each promising the fastest, safest, most reliable software ever
- Each one helped with one part of the build. Sound familiar? *cough* Skills
- 2014: asked for a feature at 9 AM, in the owner's hands by 2 PM

● **But that sometimes broke things, so I became obsessed with improving and automating the SDLC**

An SDLC is people, process, tools and teams that have to stay aligned.

FRAME

People

Process

Tools

Software

Teams

How do we keep *humans* aligned?
What are *humans* driving toward?

My lived experience with an AI-Native SDLC

What worked

What failed

How much AI slop we ship

The builder and the team each need something different to excel.

FRAME

The Builder

Theory understand the tool

Proficiency be competent with it

Fluency use it swiftly, at the right times

The Team

Adoption broad, not a few champions

Alignment swimming the same direction
Solved with OKRs ★

Feedback short loops

The story, in order:

Build

Prove

Align

Iterate

WHAT WORKS

Tools

Cursor, Claude, Codex, Lovable,
OpenCode, you name it

Centralized MCP

Claude Code Actions: auto agentic review,
automations

Agent skills marketplace

Norms

Weekly demos

Encourage ALL stakeholders to use
Claude and Codex: Product, Design, Eng,
Leadership

Intent as source of truth



WHAT WORKS

Jira / Atlassian

GitHub

Observability

On-Call Systems

Cloud Vendors

Data Warehouse

Marketing

Analytics

Code Scanning

All with SSO... The IT/Security folks can sleep at night 🛌

Agentic AI Foundation
AGNTCon + MCPCon
North America

Product Design Engineering



The slowest step

The elephant in the room 🐘

Reviewing Pull Requests

Where the **agentic stack** is being built.

Pull Request Statistics

PR REVIEW

Our mission: guide caregivers and their loved ones to a place of confidence.

Grace is our digital advisor. A human advisor is always available.

Every merge changes what a family sees.

3,627

PRs opened
~15 per day

2,955

PRs merged


81%

merge rate

Jul 2025 – May 2026 (239 weekdays) · six repos, incl. backend and frontend

Bottom line: the lowest-risk PRs merge themselves when two conditions agree.

PR REVIEW



Required check: a CI job that must pass before merge. A bot does the approving.

A rule computes a minimum risk score, and the model can only raise it.

PR REVIEW

```
final score = max( model score, floor )  
floor: file paths, diff size, and always-risky areas
```

1–3

Docs, tests,
simple config

4–6

Ordinary feature code

7–8

SMS, email,
public endpoints

9–10

Auth, infra,
migrations, CI


Can auto-merge

4 and above: a person reviews

Ours auto-merges up to 3. The floor's rules come from main, so a PR can't loosen them.

On one repo, the median PR took 8.6 hours to merge instead of 27.

PR REVIEW



Aug 19

We turned on auto-merge for the lowest-risk PRs

-69%

median time to merge, frontend repo

42%

of PRs now merge themselves, in 1.6 h. PRs a person reads: 21.9 h

Frontend repo · median opened→merged · Jul 1–Aug 18 (205 PRs) vs Aug 19–Sep 23 (211) · analysis by our manager

Cool. So it's all sunshine and rainbows?

What didn't work? – AI Slop Reviews

The Flip-Flop

the codebase wins

The Phantom Bug

no trigger path, no bug

The False Duplicate

verify against source

The Scope Creep

review scope is the diff

The Implementation Leak

judge against the interface

The Verbosity Trap

triage severity first

BUSINESS VALUE GUARDRAIL

Does this PR comment solve a real-world, plausible, legitimate problem?

What didn't work? – Offloading Your Brain

THE AI SLOP CANNONS

“Did I do that?”

Not remembering a document you wrote

Engineers

sending PRs they didn't read

“What is this?”

Git blame: oh. I did that.

Product

sending PRDs they didn't review

“What does our product actually... do?”


Design

...actually, design is perfect

Models seem smarter. Do they appear intelligent, or produce output that helps you build products?

What doesn't work: Too Much, Too Fast

— Code shipped — What the team understands



Comprehension Debt

Addy Osmani's term

How the hell do you
know what you're
shipping?

Illustrative, not measured. We have proxies, not an instrument.

The pile keeps growing: 3,627 PRs in ten months, and about 78 more every week.

WHAT DIDN'T WORK

3,627


PRs opened, six repos

78

a week since the Nov 3 launch

227

in the busiest week, before launch



Cumulative PRs opened, Aug 4, 2025 – May 2026 · backend, frontend and four smaller repos

Each PR is a change someone has to understand. Volume is how comprehension debt builds.

So where does that leave us?


WHERE THAT LEAVES US

A Design Pattern Renaissance

- Accept the inevitable. Coding is not what it used to be.
- But... Design Patterns still make code more coherent, readable. Lean into them.
- Encapsulate the crap

Try this Monday

- 1 Focus more on design patterns & interfaces in your reviews
- 2 Delegate what you can measure: start with docs and tests, shrink the blast radius
- 3 **Commit to iteration: Start with a weekly demo**



Thanks for your time!

Is AI coding like cheating in a video game?

I don't think so. We move faster, but the concerns, the responsibility, and the fulfillment of building working products are the same.



Session feedback

austin.brown@aplaceformom.com