



SESSION 3 OF 9 · SUMMER 2026 PROGRAM

AI-NATIVE AGILE · HANDS-ON

The Build 2

From PRD to working software. Live.

Amr Hamdy
Engineering Manager

One discipline. **Any team size.**



Amr Hamdy

Engineering Manager

Defines how to build it; balancing architecture, scalability, delivery, and technical excellence to bring products to life.

WHAT WE COVER TODAY

Architecture

Delivery

Owning the output

Nine weeks. One story.

Every Thursday · 6:00 PM · Each session stands alone

SESSIONS 01 - 09



W1 **Surfing the Wave**

Why this moment matters

Yassin Gamal

W2 **The Build 1**

Laying the foundation for real AI-assisted products

Nada + Amr

▶ W3 **The Build 2**

How to ship: architecture, delivery, owning the output

Amr Hamdy

W4 **Reality Check**

What AI gets wrong, and why you need to know

M. Gamal + Mohab

W5 **Learning in the AI Era**

Using AI to learn faster and grow further

Fares Waheed

W6 **AI on Legacy Code**

AI-native workflows for codebases that already exist

Sherin Medhat

W7 **Can You Trust This Model?**

How to evaluate and trust the models you build with

Mo'men Hamdy

W8 **The Way We Work Now**

How AI is reshaping work, is anyone truly AI-native?

Panel · Yassin Gamal


W9 **The Wave Ahead · Open Waters**

Where we go from here, together

Closing panel + your Q&A

— WHERE WE LEFT OFF

Last session: we built **the PRD.**



Today we build it.
On screen. Live.

Four acts.



01

Ways to drive the agent



02

The Greenfield pipeline



03

Building it live



04

THE WHOLE POINT


Reviewing every step

01

— ACT 01

Pick your weapon

A few ways to drive the same agent.



What survives **the session?**

 Claude Code — direct

 Plan, then implement

 Superpowers

 BMAD

 Spec Kit

 Hybrid

Nothing

A plan. This session only.

Habits and a journal

PRD and architecture

A spec in your repo


Whatever you chose to keep

02


— ACT 02

The run map

Where we are going, before we go.



The Greenfield pipeline.



The pipeline never changes.

03

— ACT 03

Live build

Slides get out of the way now.



Recap the idea

01

From last session's PRD



WATCH FOR → The stories we'll ship today

STATION 01

DEMO

Reading the PRD

Steal the theme

02

Find the visual direction



WATCH FOR → Same style, new palette

S T A T I O N 0 2

DEMO

Finding the direction

Generate the design system



Theme to tokens, type, components

WATCH FOR → Design decided before code

STATION 03

DEMO

Tokens and components

Commit to the repo

The scaffold is now truth



WATCH FOR → Specs describe behaviour only

S T A T I O N 0 4

DEMO

The scaffold lands

Four things go in **first**.



Context7 MCP

Live docs



BMAD

Agent roles



Spec Kit

Spec to tasks



Playwright

Proof

STATION 05

DEMO

Arming the repo

The constitution

06

Ask the architect, not the intern



WATCH FOR → BMAD Architect writes our constitution prompt

STATION 06

DEMO

The constitution

Into **memory.**



Keep memory and docs always current



Use Context7 for latest docs, always



Prefer `init` and `npm` over hand-written files



Everything dockerized, nothing runs on the host


Written once. Obeyed every session.

STATION 07

DEMO

Rules into memory

Ship **one** story.



✓ You review after every single step.

NOT OPTIONAL

WATCH FOR → One story at a time. Never two.

S T A T I O N 0 8

DEMO

Ship stories with Spec Kit

Prove it with Playwright

09

Green, or it isn't done



WATCH FOR → Coverage per acceptance criterion

STATION 09

DEMO

Prove it with Playwright

Four **gates.**



01

Human read

Line by line



02

Fresh AI session

Different model



03

CodeRabbit

Review skill



04

Security scan

Before it ships

| Fix everything. Re-run.

S T A T I O N 1 0

DEMO

Four review gates

Commit and push

11

Checkpoint, then ship



WATCH FOR → The checkpoint is non-negotiable

STATION 11


DEMO

Commit, push, checkpoint


**The AI writes. You
review. That's the job.**

Where reviews happen.

✗ One big review at the end



✓ A gate at every small step



What reviewed means.



**Against
acceptance
criteria**

Not "it runs"



Fresh context

Catches different bugs



**Reject boundary
violations**

Even working ones



Update the spec

If reality diverged

Why we do it.



Quality



Maintainable



Scalable



**You still know
your system**

Small steps.
Reviewed every time.

Five **takeaways.**

01

**Specs are the
source of truth**

02

**Context is
owned, not
assumed**

03

**Pick the gear,
not the tool**

04

**Review every
small step**

05

**Checkpoint
before every
handoff**

Nine weeks. One story.

Every Thursday · 6:00 PM · Each session stands alone

SESSIONS 01 - 09



W1 **Surfing the Wave**

Why this moment matters

Yassin Gamal

W2 **The Build 1**

Laying the foundation for real AI-assisted products

Nada + Amr

▶ W3 **The Build 2**

How to ship: architecture, delivery, owning the output

Amr Hamdy

W4 **Reality Check**

What AI gets wrong, and why you need to know

M. Gamal + Mohab

W5 **Learning in the AI Era**

Using AI to learn faster and grow further

Fares Waheed

W6 **AI on Legacy Code**

AI-native workflows for codebases that already exist

Sherin Medhat

W7 **Can You Trust This Model?**

How to evaluate and trust the models you build with

Mo'men Hamdy

W8 **The Way We Work Now**

How AI is reshaping work, is anyone truly AI-native?

Panel · Yassin Gamal

W9 **The Wave Ahead · Open Waters**

Where we go from here, together

Closing panel + your Q&A

SESSION 03 · THE BUILD

Thank **you.**

Now run the loop on your own repo.

One story. End to end. Bring what breaks.

• **Next Thursday · 6:00 PM · Reality Check**