

AI PROTOTYPE TO BOARD-READY DEMO




IN 8 WEEKS (NO IT QUEUE REQUIRED)

01 SCOPE THE PROBLEM



SUMMARISE INBOUND SUPPLIER DISPUTES AND ROUTE TO THE RIGHT HANDLER.

02 8-WEEK PLAN



03 SUCCESS METRIC



04 DATA DECISION

SYNTHETIC DATA = BUILD NOW, PASS SECURITY LATER



THE CLOCK STARTS BEFORE PROCUREMENT DOES

You've got a validated AI concept, executive sponsorship, and six weeks before the next board review. You also have an IT procurement queue that moves at a different speed entirely. The two timelines don't reconcile.

This isn't an edge case. It's a common friction point for corporate innovation leaders in 2026. Internal IT and procurement queues can add months to any new initiative in large organisations – timelines vary significantly by sector and organisation size, but the drag is real and well-recognised. This kills momentum and makes the innovation team look ineffective precisely when it needs to look credible.

The good news: you don't need IT to get started. You need a tight scope, the right external tooling, and a build structure that keeps you on track without drifting into scope creep.



Here's how to go from idea to board-ready AI prototype in eight weeks, entirely outside the corporate network.





BEFORE YOU TOUCH A SINGLE TOOL: SCOPE THE PROBLEM PROPERLY

PROBLEM STATEMENT (ONE PARAGRAPH):

Summarise inbound supplier disputes and route to the right handler.



ONE INPUT, ONE OUTPUT.

The AI receives something specific and produces something specific.



A QUANTIFIABLE BASELINE.

You need to know the current state in numbers before you build.



REAL USERS YOU CAN TEST WITH.

You need three to five internal users who will give you honest feedback in weeks three through five.

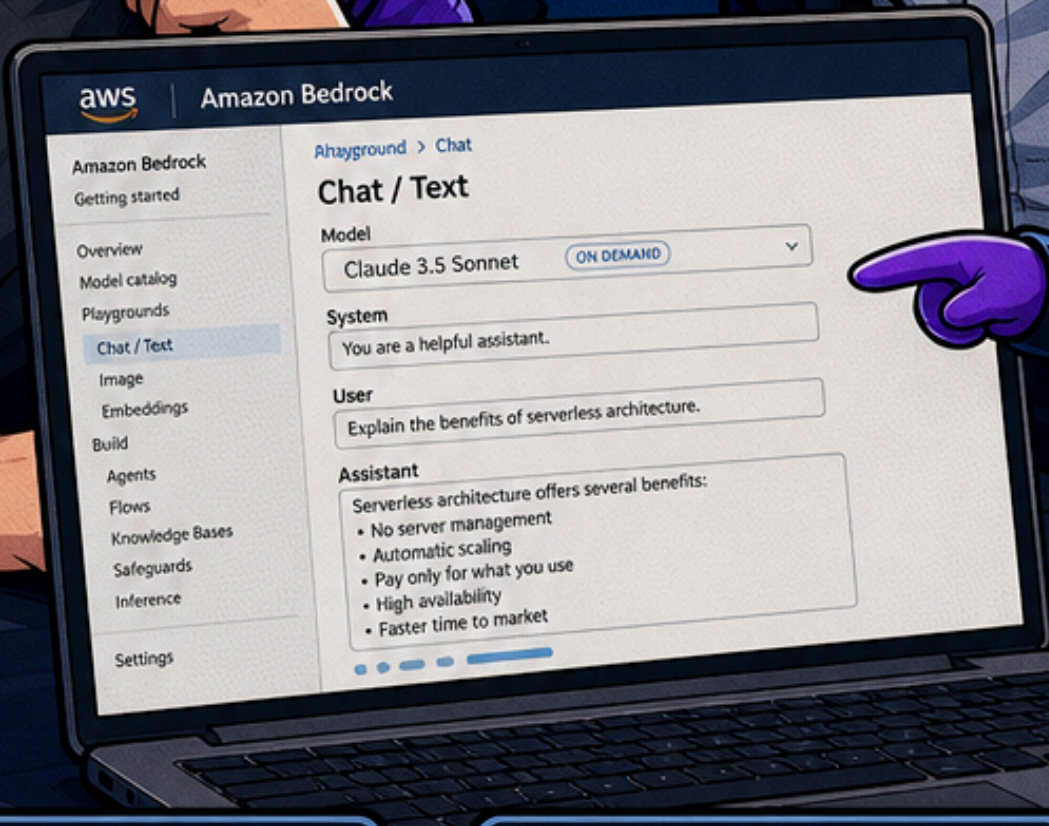


DATA YOU CAN WORK WITH IMMEDIATELY.

Not data you'll get after a six-week access request.

WRITE THE PROBLEM STATEMENT IN ONE PARAGRAPH. IF YOU CAN'T, THE SCOPE ISN'T TIGHT ENOUGH.

HOW TO BUILD ENTIRELY OUTSIDE CORPORATE IT



PUBLIC CLOUD SANDBOXES

AWS, Google Cloud, and Azure environments you can stand up in an afternoon. Pay-as-you-go. No contracts. Just a credit card.



NO-CODE & AI-NATIVE BUILDERS

Bubble, Softr, Replit, Lovable, Bolt.new, n8n, Figma AI tools. Build apps, automations, and interfaces — fast.



OPEN-SOURCE MODEL APIS

OpenAI, Anthropic, Google. Production-grade models via API. Pay-as-you-go. Inference costs are usually modest.



IMPORTANT CAVEAT: Regulated industries may require IT involvement from the outset. Check your organisation's policies before proceeding.

THE THREE-PHASE STRUCTURE THAT KEEPS YOU ON TRACK

PHASE 1

PROBLEM FRAMING
(WEEKS 1-2)



- ✓ Pressure-test problem
- ✓ Confirm success metric
- ✓ Identify test users
- ✓ Make the data decision

OUTPUT:
ONE-PAGE PROTOTYPE BRIEF



PHASE 2

ITERATIVE BUILD AND USER TESTING
(WEEKS 3-6)



- ✓ Build thin vertical slice
- ✓ Test with real users
- ✓ Get feedback
- ✓ Improve and repeat

WEEK 3: First prototype, first user session

WEEK 4: Iterate, second user session

WEEK 5: Collect performance data, third user session

WEEK 6: Feature lock

PHASE 3

DEMO HARDENING
(WEEKS 7-8)



- ✓ Make it board-ready
- ✓ Connect to business outcomes
- ✓ Prepare the narrative
- ✓ Stress-test the demo

DEMO ANSWERS:

- 1 Baseline
- 2 Prototype results
- 3 Path to production



ONE-PAGE PROTOTYPE BRIEF

	PROBLEM IN ONE PARAGRAPH	_____
	CURRENT BASELINE (NUMBERS)	_____
	SUCCESS CRITERIA (WHAT COUNTS)	_____
	TEST USERS (WHO)	_____
	DATA (WHAT WE'LL USE)	_____



THE DATA PROBLEM: SOLVE IT ON DAY ONE, NOT DAY FORTY

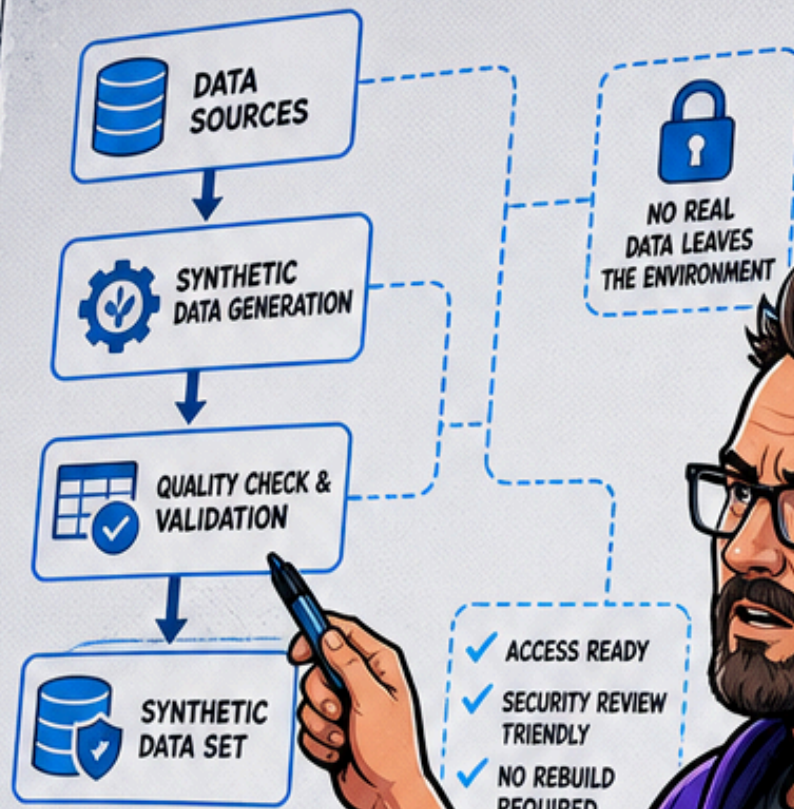
THE SINGLE FASTEST WAY TO DERAIL AN EIGHT-WEEK AI PROTOTYPE IS TO DISCOVER IN WEEK THREE THAT YOU CAN'T ACCESS THE DATA YOU NEED WITHOUT A SIX-WEEK LEGAL REVIEW.

THE ANSWER IS **SYNTHETIC DATA**, AND IT HAS BECOME CONSIDERABLY MORE CAPABLE IN RECENT YEARS.

FOR A FOCUSED CORPORATE AI PROTOTYPE, SYNTHETIC DATA SERVES TWO PURPOSES. FIRST, IT LETS YOU BUILD AND TEST IMMEDIATELY WITHOUT WAITING FOR DATA ACCESS APPROVALS. SECOND, IT MEANS THE PROTOTYPE CAN PASS A SECURITY REVIEW LATER WITHOUT REQUIRING A REBUILD FROM SCRATCH, BECAUSE NO REAL PERSONAL OR COMMERCIAL DATA EVER ENTERED THE EXTERNAL ENVIRONMENT.

DOCUMENT YOUR DATA HANDLING APPROACH IN THE PHASE ONE BRIEF. WHEN SECURITY REVIEW EVENTUALLY HAPPENS, YOU WANT TO SHOW THAT GOVERNANCE WAS CONSIDERED FROM SPRINT ONE, NOT RETROFITTED UNDER PRESSURE.

SYNTHETIC DATA FLOW



SPRINT ONE: GOVERNANCE


- DATA SOURCES & ACCESS
- SYNTHETIC DATA STRATEGY
- SECURITY & PRIVACY
- COMPLIANCE & RISK
- REVIEW & APPROVAL PATH

WHAT BOARD-READY ACTUALLY MEANS

A WORKING INTERFACE IS NOT A BOARD-READY DEMO.
A BOARD-READY DEMO IS A WORKING INTERFACE
PLUS A CLEAR ANSWER TO THE QUESTION:

**"WHY SHOULD WE INVEST IN
TAKING THIS TO PRODUCTION?"**

THE SUCCESS METRIC NARRATIVE IS WHAT
MAKES THE DIFFERENCE. EVERY FEATURE YOU
DEMONSTRATE SHOULD CONNECT TO A NUMBER
FROM YOUR OWN TESTING.




- ✓ NOT "THIS SAVES TIME" BUT A SPECIFIC, TESTED REDUCTION IN HANDLING TIME WITH THE ACTUAL FIGURES FROM YOUR USER SESSIONS.

-38%
AVG. HANDLE TIME
(N=127 SESSIONS)

- ✓ NOT "THIS IMPROVES ACCURACY" BUT THE CONCRETE ERROR-DETECTION RESULTS FROM YOUR OWN TEST SET, WITH THE NUMBERS FILLED IN FROM YOUR REAL PROTOTYPE EVALUATION.

92.6%
ERROR DETECTION
RATE
(N=2,450 CASES)

 SPECIFIC NUMBERS FROM REAL USER TESTING, EVEN SMALL-SCALE TESTING, CARRY MORE WEIGHT WITH EXPERIENCED BOARD MEMBERS THAN POLISHED UI OR IMPRESSIVE AI VOCABULARY. THEY'VE SEEN ENOUGH AGENCY DECKS TO KNOW THE DIFFERENCE BETWEEN A DEMO AND A **VALIDATED RESULT**.

- ⚠ IF YOUR TEST DATA SHOWS THE PROTOTYPE ISN'T PERFORMING WELL ENOUGH TO JUSTIFY PRODUCTION INVESTMENT, SAY THAT CLEARLY TOO.

A BOARD PRESENTATION THAT SAYS "HERE'S WHAT WE BUILT, HERE'S WHAT WORKED, HERE'S WHAT DIDN'T, HERE'S WHAT WE'LL FIX" BUILDS FAR MORE TRUST.

**TRUST
>
POLISH**

EVOTRON



<https://evotronstudio.co.nz>