

AI INNOVATION FIELD GUIDE

How to rapidly test ideas with AI.

Hypothesis. Thinnest AI-assisted test. Kill date. Named owner. For innovation, product, L&D and digital leads who would rather leave with a decision than another slide.



Why this guide?

AI makes a first draft cheaper. It does not remove the need for a hypothesis, a kill date, a named owner, and real users. This guide is the method: one canvas, one rapid cycle, guardrails that keep a test honest, and a same-day hand-off so the room leaves with a decision, not another mural.

EVERY TEST LEAVES THE ROOM WITH

Hypothesis

If / then / because. One riskiest assumption you can be wrong about. Desirability, feasibility or viability. Pick the one that fails first.

Thinnest test

AI-assisted, constraint-safe, days not quarters. A checked draft, a clickable shell, a concierge script or a throwaway page. Not a platform.

Kill date + owner

A calendar day a risk lead and a CFO can both sign. A person who can move a roster, backlog or budget. Not a centre. Not a steering group.

WHO IS IT FOR?

Innovation, product, L&D and digital leads inside organisations that keep buying murals. Anyone who would rather leave with a decision than another slide.

HOW TO USE IT

Print the canvas. Fill it before you book the room. Run the rapid cycle. Sign the stopper. Send the one-pager the same day. Kill what fails. Scale what earns the next dollar.



01 The anti-theatre

AI makes a first draft cheaper. It does not remove the need for a hypothesis, a kill date, a named owner, and real users.

Before you book the room

If you cannot name the hypothesis, the kill date, the named owner and the real users before you book the room, do not book the room. Start with a calendar and a hated workflow, not a capability map.

RETIRE THESE THIS QUARTER

01 No AI strategy mural first

Start with a calendar and a hated workflow, not a capability map. An AI strategy bought off a vendor deck is theatre with better fonts.

02 No endless pilot

A pilot with no stop is a museum. Write the kill date on the same page as the hypothesis. Defaults of 14 or 30 days are useful, not religion.

03 No certainty theatre

NPV, IRR and tidy five-year forecasts are hostile to early-stage ambiguity. Fund the next learning milestone. Define the evidence that would justify more effort.

04 No parallel theatre

Parallel micro-tests are fine when each one has an owner, a metric and a kill. Parallel workstreams with no killing mechanism are a hobby with better catering.

05 No Monday funding

If nobody owns Monday's money, you ran a party. A named owner who can move a roster, backlog or budget must sit in the room.

What you leave with instead

- One framed problem
- One constraint
- One thinnest AI-assisted test
- One named owner who sat in the room
- One kill date
- One metric
- One page circulated the same day

If every idea graduates, none of them were experiments.

FURTHER READING: FUTURE SQUARED EP. 179 · THE BUSINESS CASE ALTERNATIVE



02

The AI idea test canvas

ONE PAGE. IF IT SPILLS, STOP.

Fill this on one page. If it spills to two, you are writing a business case. Stop.

FRAME 01 TO 04

TEST 05 TO 07

GOVERN 08 TO 12

01 Hated workflow

Not 'explore generative AI'.
'The claims explanation that
takes forty minutes because
we rewrite facts from three
systems.'

02 Who hurts

Role, not persona poster. Care
writer. Credit officer. Store
lead. Frontline claims officer.
Name a desk.

03 Riskiest assumption

The belief that, if wrong, kills
the idea. Desirability, feasibility
or viability. Pick the one that
fails first.

04 Constraint that kills

Privacy. Conduct. Clinical
safety. Industrial relations.
Brand and claims. Write it on
day one. Interview risk before
you mural them.

05 Thinnest AI test

AI-drafted copy a human
checks, a clickable shell, a
concierge script, a throwaway
page, or a shadow process.
Not a platform.

06 Hypothesis

If we run [test] with [cohort]
for [days], then [observable]
will move, because [reason we
can be wrong].

07 Metric (one number)

Hours returned. Cycle time.
Completion. Complaint
volume. Enquiry that is not
vanity traffic. Pick one.

08 Named owner

A person who can move a
roster, backlog or budget.
They attend. A delegate who
cannot kill is applause.

09 Named stopper

Risk, legal, privacy, clinical,
credit or ops. Give them the
page before the room. They
mark it. They sit in the last
hour.

10 Dollar figure

Room, time, licences, checker
hours, ads. A number a CFO
can read without a pack.
Needs a business case? Not a
rapid test.

11 Kill date + criteria

A calendar day. Write the stop
in language a risk lead and a
CFO can both sign. Defaults of
14 or 30 days are useful.

12 Same-day one-pager

Who sends the one-pager
before people reach the lift.
Same day.

03

The rapid cycle

Frame. Prototype. Confront. Decide. Close with a decision, not a next-steps slide.

01 Frame

Write the hated workflow and the riskiest assumption before anyone opens a model. Use async where you can: a shared page and a short decision window beat a town hall.

02 Prototype

Speed over polish. AI is a force multiplier for thin artefacts, not a substitute for a hypothesis. Checked copy, clickable shell, concierge script, throwaway page.

03 Confront

Three operators or customers beat thirty stakeholders. Talk to the desk that hurts. If your user is only a steering committee, you tested politics, not demand.

04 Decide

Kill, iterate or scale. Publish kills in the same channel you publish wins. That is how you earn the next test.

Legitimate thin artefacts

- Copy: offer, email, care script, decision explanation (human-checked)
- UI mock or clickable shell
- Concierge script: model drafts, person delivers or decides
- Throwaway landing page
- Shadow process beside the live path, with a clock

Not prototypes

- A platform RFP
- A six-month vendor pilot with no kill
- A capability map with no desk
- A demo that never meets a customer

DECIDE

Kill

If the metric does not move enough to pay for the next window, if the stopper will not allow the paper version, or if the owner missed the heartbeat with no deputy.

Iterate

If the assumption shifted and a thinner or different test is now the honest next step. Fund the next learning milestone, not a rewrite of the five-year fiction.

Scale

Only when evidence survives a risk note and the named owner will staff the new path on live work. Then write the larger case, with evidence instead of fiction.

Publish kills in the same channel you publish wins. That is how you earn the next test.



Parallel micro-tests, not parallel theatre

When tests are cheap, run more than one. Several framed bets, short clocks, explicit kill or progress. That is aligned with evaluate-ideas discipline and with Collective Campus hack-day practice.

Rules that keep parallel work honest

- Each micro-test has its own canvas row: assumption, metric, owner, kill date.
- Shared stoppers review the set, not twelve surprise integrations in week nine.
- Cap the portfolio so a director can name every live test without a pack.
- Kill in public. Graduation is rare on purpose.

Six ideas in three hours is a framing pattern. Six unowned workstreams for six months is theatre.

DIFFERENT LOGOS. SAME CANVAS.

Worked patterns

Australian Unity PHI hackathon

SEAM

A PHI policy idea whose key assumptions have not met real feedback.

TEST

Quickly built experiments with enough signal to progress or kill.

METRIC

Evidence of appetite or rejection, not a mural photo.

KILL

The product or proposition lead who can stop owns the call.

Global FMCG hack day

SEAM

Ten cross-functional people, desire to innovate, missing methods and tools.

TEST

Six ideas framed in three hours with decision criteria on the table.

METRIC

Owners and kill dates by Monday, not a photo of sticky notes.

KILL

Any idea without an owner and a kill date dies before the week starts.

Offer test for a regulated brand

SEAM

An offer you would otherwise wait a quarter to test.

TEST

Throwaway landing page in days. AI copy OK if claims and brand sit in the chair before go-live.

METRIC

Enquiry, not vanity traffic.

KILL

If the page cannot live inside the claims rule.

Care explanation at a large AU telco

SEAM

An explanation rebuilt from a blank page every time.

TEST

A play plus a checked model draft, not a new knowledge platform.

METRIC

Minutes per explanation, then hours at day 14 and 30.

KILL

If the human check is skipped or the minutes do not move.

05

Guardrails

AI makes it easier to confuse 'we generated something' with 'we should build something'.

01 Data and privacy

Put the constraint on the wall in hour one. What never goes into a prompt. What stays in a private tenant. What needs a human check before a customer sees it. If the stopper cannot mark the page, you do not have a test. You have a future incident.

CHECK

Can the stopper mark the prompt boundary on page one?

02 Hallucination

Treat model output as a draft until a named checker signs. Minutes returned that skip the check are fake productivity. Human-check minutes are part of the metric, not a footnote to hide.

CHECK

Is a named checker on the canvas before any model runs?

03 Build vs generate

Generating a polished deck, a fake roadmap or a synthetic customer quote is easy now. That is not evidence. Evidence is behaviour from a real cohort against a written hypothesis. Generation with no confrontation is stationery.

CHECK

Does the thinnest test confront a real user, or only generate?

04 Production hunger

A rapid test that requires a core integration on day one is not rapid. Shadow it. Concierge it. Paper it. Earn the integration with evidence. Returned hours go to maker time, customer time, or a deleted meeting, not more theatre.

CHECK

Can the test run beside the live path without a core change?

Stopper sign-off

Give the stopper the page before the room. They mark it. They sit in the last hour.

- Risk, legal, privacy, clinical, credit or ops named on the canvas (field 09).
- Prompt boundary and customer-visible check written in stop language.
- Kill criteria a risk lead and a CFO can both sign before anyone opens a model.

SAME-DAY HAND-OFF

Before people reach the lift: one-page canvas filled; kill criteria in stop language; cohort named; human-check rule beside any model use; one-page note sent the same day; decision meeting booked on the kill date; what you will stop doing if the test lives; named person who updates the channel when you kill or scale. If any of those eight are missing, you facilitated a workshop. You did not run a test.



Selected work

COLLECTIVECAMPUS.IO/WORK

HEALTH

Australian Unity PHI hackathon

Quickly built experiments on a private health insurance policy idea so assumptions met real feedback fast enough to progress or kill, without commercialising the wrong idea.



MOBILITY

AI landing pages to test offers fast

Published Collective Campus case study: an EV and hybrid battery specialist used AI-assisted throwaway pages to test offers in days, not quarters, with claims in the chair.



ENERGY

Experiments inside a listed energy company

Lean startup practice inside AGL: framed bets, short clocks and decision criteria that survive a listed-company governance path.



SALES

2-day innovation sprint

Inkeros: a cross-functional sprint for an omnichannel sales organisation. Required hand-off: owners, kill dates and a budget line for the week ahead.

ALSO TRUSTED BY • PUBLISHED COLLECTIVE CAMPUS CLIENT WORK

AUSTRALIAN UNITY • AGL • INFINITEV • INKEROS • ALLENS • SPORTSBET

"The content and facilitation were pitched at the right level, with simple and easy-to-understand customer personas and customer journey maps helping participants to both understand fundamental concepts and give them the confidence to apply it to their work."

CLAIRE BIRCUMSHAW - CHANGE MANAGER, CITY OF MARION

"Feedback received, both formally and informally, has been overwhelmingly positive. Leaders have admitted to having their eyes opened and thinking challenged, seeing this as an opportunity to be more innovative in how we respond to and add value for our customers."

AMY HOLLINGS - ORG DEVELOPMENT PROJECT LEAD, YARRA RANGES

MEET YOUR INSTRUCTOR



Steve Glaveski

Founder of Collective Campus, author of two Wiley books (Employee to Entrepreneur, Time Rich), host of the Future Squared podcast, and regular contributor to Harvard Business Review.

GET STARTED

Leave the room with a hypothesis,
a kill date and a named owner.
Not another slide.

+61 433 261 580
STEVE@COLLECTIVECAMP.US
COLLECTIVECAMPUS.IO

INNOVATE OR DIE

COLLECTIVE CAMPUS / HOW TO RAPIDLY TEST IDEAS WITH AI

07 / 07