

THE WORKSHOP

# Impact Mapping

A pattern for tracing a business goal through the people whose behaviour has to change, through the changes themselves, to the things you might build – so the team can tell the difference between work that moves the number and work that just feels productive.

2026-04-19

[barkingiguana.com/writing/the-workshop-impact-mapping/](https://barkingiguana.com/writing/the-workshop-impact-mapping/)

# Contents

Intent	3
Also Known As	3
Motivation	3
Applicability	3
Participants	4
Structure	5
Collaborations	5
Facilitator Playbook	5
Steering When It Goes Sideways	8
Consequences	9
Worked Example	10
Outputs & Follow-up	10
Related Patterns	11

---

*This is the pattern. If you want the story of Impact Mapping in action, read [Impact Mapping: Connecting Work to Goals](#) – the Greenbox team discovering that half their planned work doesn't connect to any goal they can state out loud. This post is the reference you keep open the morning of.*

---

## Intent

Trace a path from a measurable business goal, through the actors whose behaviour affects that goal, through the behaviour changes you want to cause, to the deliverables that might cause them – so the team can tell the difference between work that moves the number and work that just feels productive.

---

## Also Known As

Sometimes called **goal mapping** or **outcome mapping**, though outcome mapping has a separate formal definition in the international development world that is not quite the same thing. Sometimes confused with **story mapping** (story mapping lays out what a user does; impact mapping lays out what behaviour you want to change). Invented and named by Gojko Adzic in 2012, published as a small book that is still the best short read on the technique.

---

## Motivation

*“Which of these features, if we built them tomorrow, would move the number we're supposed to be moving this quarter?”*

Asked out loud in the right room, that question almost always lands like a flat stone hitting glass. The answer takes longer than anyone expects. Someone defends a feature by explaining why the customer asked for it. Someone else defends another by naming a competitor. A third person points to the roadmap. Nobody says, without hedging, *this one will move the number because this specific group of users will do this specific thing differently*. The question doesn't get answered; it gets re-framed until it can be.

The backlog has drifted. Each individual feature connects to *something* – a request, a rival, a hunch – but nothing connects the features to each other, and nothing connects the set to an outcome anyone is measuring. Reasonable things are infinite. Strategy is the shortlist of reasonable things that share a goal, and the shortlist has gone missing.

Impact Mapping exists to make that shortlist visible. The wall shows the goal on the left and the work on the right, with the logic connecting them drawn explicitly in between. Work that can't find a place on the wall can still be reasonable. It just doesn't belong in this quarter.

---

## Applicability

### Use when:

- You're starting a quarter, an initiative, or a new product line and need to decide what to build
- The team has a list of features but no shared story about how they connect to business outcomes
- Stakeholders are requesting features and nobody is asking *why*

- You need to say no to work and you want a defensible reason
- The organisation is measuring an outcome and the team doesn't know how their work connects to it

#### **Don't use when:**

- You already have a clear, shared understanding of the connection between goal and work
- You're planning at the sprint level – Impact Mapping is strategic; sprint planning is tactical
- Nobody in the room can articulate a measurable business goal – solve that first, separately
- The goal is imposed from above without buy-in – a mapping session with a fake goal is worse than no session

**Inputs.** One measurable, time-bound business goal, written on a card before the session starts. "Increase weekly active subscribers from 200 to 500 by end of Q3." Nothing else. If the goal isn't concrete enough to fit on a card, the session isn't ready to run.

---

## Participants

**Facilitator.** Holds the shape of the tree (goal → actors → impacts → deliverables), keeps people from jumping columns, and calls out when someone has smuggled a deliverable into the actors column.

**Product owner or business stakeholder.** Mandatory. They own the goal. They're the one who will defend it, refine it, or abandon it when the session reveals that the goal itself is the problem.

**Developers.** At least one, ideally two. They know what's cheap and what's expensive, which is what turns the deliverables column from wishful thinking into a triaged list.

**Designers.** They think about actor behaviour natively. In the impacts column – the hardest column – a designer's framing ("they would finish the sign-up flow instead of abandoning it at step three") is usually sharper than a developer's.

**People who talk to the actors.** Sales, support, operations, account managers. Whoever has the least-filtered view of the people whose behaviour you're trying to change. They will contradict the optimistic assumptions in the room and that is exactly what they're there to do.

**SRE / Operations.** For infrastructure or reliability initiatives, SRE is the domain expert on actors and impacts – "on-call engineers stop being paged for the billing cron" is a valid impact, and "our customers stop opening support tickets about missed deliveries" is downstream of it.

**Group size:** 4–6. Impact Mapping is a thinking-aloud exercise and the room has to stay a conversation. Above six, it becomes a meeting with a whiteboard.

#### **Who to leave out:**

- **Large groups of stakeholders.** If seven people need to shape the goal, that's a pre-session, not this session. Come to Impact Mapping with the goal agreed.
- **People who can't say no.** Someone who will accept every proposed deliverable without challenge makes the prioritisation phase impossible.
- **Pure spectators.** Impact Mapping is not a presentation; observers change the dynamic and absorb oxygen without contributing.

## Structure

Phase	Duration	Notes colour	Key question
Orient on the goal	10 min	Green (one card)	"What are we trying to move?"
Actors	15 min	Yellow	"Whose behaviour affects the goal?"
Impacts	25 min	Orange	"How would their behaviour change?"
Deliverables	20 min	Blue	"What could we do to cause that change?"
Prioritise	10 min	Marks on the map	"What's the highest-leverage path?"
Wrap-up	10 min	–	"Who owns what next?"
<b>Total</b>	<b>~90 minutes</b>		

The map grows left-to-right, one column per phase. Skipping columns is the single most common failure mode. An impact without an actor is a feature. A deliverable without an impact is a hunch. The discipline of the four columns is the technique.

## Collaborations

Impact Mapping alternates between open conversation and quiet placement. The goal is fixed; everything to the right of it is debatable.

- **Goal** is pre-agreed. The session doesn't debate whether the goal is right – it maps *from* the goal. If the goal is unstable, stop and fix the goal separately.
- **Actors** emerge from prompts. The facilitator asks, the room calls out, the notes go up. Push for specificity: not "users" but "new visitors who haven't subscribed yet," not "the team" but "the warehouse lead during Friday pack-out."
- **Impacts** are where the technique earns its cost. Each impact is a *behaviour change* attached to one actor. The facilitator's job in this phase is to intercept every "we could build X" and convert it into "what behaviour would X cause?"
- **Deliverables** come last and come quickly. Once the impacts are honest, the deliverables are almost obvious. The problem is never the supply of ideas – it's whether the ideas connect.
- **Prioritisation** is the payoff. The team traces a path: one actor, one impact, the smallest deliverable that might cause the impact. Everything else on the map is real work for a different conversation.

The key rhythm is **work backwards**. Goal before actor, actor before impact, impact before deliverable. Any deliverable that appears before its impact gets politely moved into a holding area until someone can connect it.

## Facilitator Playbook

### Phase 1 – Orient on the goal (10 minutes)

Put the goal card on the far left of the wall. Read it aloud:

*"Our goal for this quarter is to increase weekly active subscribers from 200 to 500 by end of Q3. That's on the wall. We are not here to debate whether this is the right goal. We are here to map how we could move it."*

Then make sure everyone understands it the same way:

*"Before we go any further – what does 'weekly active' mean here? What counts as a subscriber? If someone paused mid-July, are they in or out of the 500?"*

Often a five-minute clarification at this stage reveals that the goal is ambiguous, and the map would have split into three directions based on three different interpretations. Resolve it now. If you can't, the session isn't ready.

### What to watch for:

- **The goal isn't measurable.** *"Grow the business."* The session cannot proceed. End it and schedule a goal-setting conversation.
- **The goal is actually three goals.** *"Grow subscribers and reduce churn and increase average box size."* Pick one for this session. Run the others separately.
- **Silent disagreement.** The goal is on the card but one person clearly doesn't believe it. Surface it: *"You look sceptical. Is the goal wrong or is it the number?"*

## Phase 2 – Actors (15 minutes)

Ask the room:

*"Whose behaviour, if it changed, would affect this goal? I want names or roles, not 'users.' Specific enough that we could identify them in our database or watch them at their desk."*

The team writes actors on yellow notes and places them in the second column. Actors can be external (subscribers, prospects, referrers, suppliers, journalists) or internal (support agents, warehouse staff, operations). They can be automated (the billing cron, the churn-prediction model, the weekly newsletter). Automated actors are first-class here – a scheduled job that sends the wrong email affects the goal exactly as much as a person who does.

Push for specificity:

*"'Subscribers' is too broad. Which subscribers? First-month subscribers? Subscribers who've paused once and come back? Subscribers whose delivery day has changed in the last sixty days?"*

Different slices of subscribers have different behaviour, and the map gets useful when the slices are named.

### What to watch for:

- **Jumping to deliverables.** *"We need a referral programme."* That's a blue note. Pull it back: *"Who would use the referral programme? What behaviour would change? Start from the actor."*
- **Forgetting internal actors.** Teams focus on external customers and forget that their own support team, warehouse, or on-call engineer is an actor whose behaviour affects the goal.
- **Forgetting adversarial actors.** Churning subscribers are actors. People who try the service and don't convert are actors. Don't only list the actors you want to help.

- **Too many actors.** Above eight, the map becomes unreadable. Group the similar ones or focus on the actors most central to the goal.

### Phase 3 – Impacts (25 minutes)

This is the hardest and most valuable phase. For each actor, ask:

*“How could their behaviour change in a way that helps us hit the goal? I want verbs. What would they do differently?”*

Write impacts on orange notes. Each impact sits in column three, connected to its actor. Good impacts are behaviour changes, not features:

- *“New visitors sign up on their first visit instead of leaving to think about it.”* (good)
- *“Existing subscribers tell one friend within their first month.”* (good)
- *“On-call engineers get paged fewer than twice per week for billing issues.”* (good, SRE flavour)
- *“We build a landing page with better copy.”* (not an impact – that’s a deliverable)

Then flip it:

*“Now the negative version. How could their behaviour change in a way that hurts the goal?”*

Negative impacts are where the defensive work lives and where the risks hide. They are usually where the biggest savings come from: preventing a bad behaviour is often cheaper than causing a good one.

#### What to watch for:

- **Deliverables disguised as impacts.** *“Subscribers use the mobile app”* is a deliverable dressed up as a behaviour. The impact is *“subscribers manage their subscription on the go”*; the app is one possible deliverable.
- **Vague impacts.** *“Subscribers are happier.”* Not actionable. Push: *“What would a happier subscriber do differently? Stay longer? Refer? Upgrade? Complain less?”*
- **One actor absorbing all the attention.** Time-box each actor. You can come back if needed.
- **No negative impacts.** Prompt directly: *“What could this actor do that would make the goal harder to hit?”* If the answer is nothing, the actor probably doesn’t belong on the map.

### Phase 4 – Deliverables (20 minutes)

For each impact, ask:

*“What could we build, do, write, or change to cause this behaviour? I want a list, not a single answer.”*

Write deliverables on blue notes and place them in the fourth column, connected to their impact. A good deliverables column contains multiple options per impact, ordered roughly from cheapest to most ambitious:

- Features (a referral programme, a pause flow, a rollback automation)
- Content (a welcome sequence, a runbook, a one-pager for support)
- Processes (a proactive call to at-risk subscribers, a handover checklist for on-call)
- Experiments (a landing page, a prototype, a manual concierge version of the feature)
- Changes to existing things (copy edits, configuration tweaks, prompt updates)

The most valuable column in Impact Mapping is not the widest – it's the one where cheap experiments live next to expensive builds. If every deliverable is a multi-month project, you've filled the column wrong.

### What to watch for:

- **Pet features appearing.** Someone places a deliverable they've wanted to build but can't connect to an impact. Challenge gently: *"Which impact does this serve? If we built it, whose behaviour would change?"* If they can't answer, park it.
- **Only big deliverables.** Push for cheap ones: *"What's the smallest thing we could do this week that would tell us whether the impact is real?"*
- **Duplicate deliverables.** The same deliverable might serve multiple impacts. That's a signal: draw lines to both. High-leverage deliverables are the ones that show up in multiple places.

### Phase 5 – Prioritise (10 minutes)

Step back. Look at the whole map. You now have a visual argument from goal to work.

Use it to pick the first path. Ask three questions in order:

*"Which actor has the most influence on this goal? Not the most numerous – the most influential."*

*"Which impact is the highest-leverage one for that actor? If we only caused one behaviour change, which one would move the number most?"*

*"Which deliverable is the cheapest way to test whether we can actually cause that impact?"*

Mark that path on the map – dot stickers, a pen line, a photograph with it circled. This is your first commitment. Everything else on the map is the second commitment, the third commitment, or "not this quarter."

### What to watch for:

- **Prioritising by excitement.** The team gravitates toward the interesting technical deliverable rather than the high-leverage one. Redirect to the goal: *"Which of these moves the number most?"*
- **Trying to do everything.** The map has twenty deliverables; the team wants to do all of them. Hold firm: *"Pick the top three. If they work, we come back for more."*
- **Ignoring the map after voting.** Someone argues for a deliverable that isn't on the map. Either put it on the map properly (with actor and impact) or park it.

## Steering When It Goes Sideways

**The goal debate.** The team starts arguing about whether the goal is right. *Recovery:* *"We can't map a goal we don't agree on. Let's pause the session, fix the goal with leadership in the next forty-eight hours, and reconvene."* *Stop if:* The disagreement is political. Mapping through a fake goal produces a map nobody will use.

**The solution-first thinker.** Someone keeps proposing deliverables without connecting them to impacts. *Recovery:* Give them a specific job: *"For every deliverable you think of, I need a yellow note and an orange note first. Who does it affect? What behaviour changes?"* *Stop if:* They can't hold the shape after three prompts. Pair them with a designer for the rest of the session.

**Analysis paralysis.** The team is stuck debating whether something is an actor, an impact, or a deliverable. *Recovery:* “The columns exist to structure thinking, not to be perfectly taxonomic. Put it in the best-fit column and move on.” *Stop if:* The same argument happens on a third note. The team is avoiding the harder conversation; name it.

**The sceptic.** Someone thinks the exercise is pointless because “we already know what we’re building.” *Recovery:* Ask them to place their planned work on the map. “Take your top three items. Which actor? Which impact? Place them.” If they can’t connect the work to the goal through an actor and an impact, the exercise has just earned its keep and they usually become the most engaged participant in the room. *Stop if:* They refuse to engage. They’re not blocking the session, just their own learning. Carry on without them.

**The everything-is-high-impact problem.** Every impact the team writes feels like it moves the goal equally. *Recovery:* Force a ranking: “If we could only cause one of these impacts, which one? Now pretend I’ve taken that one away – which of the rest?” *Stop if:* The team genuinely can’t distinguish. The goal is probably too abstract to map against; sharpen it before continuing.

**The absent stakeholder.** Halfway through, the team realises an actor is owned by someone not in the room. *Recovery:* Put a pink note on that actor: “Need to talk to [person] before we map this.” Carry on with the actors you can map. *Stop if:* More than half the map depends on people who aren’t in the room. You’re storming without the right participants.

## Consequences

### Benefits

- Work connected to outcomes instead of to opinions
- A defensible answer to “why are we building this” for every item in the quarter
- A visible short list – three deliverables picked out of twenty – that the team commits to first
- A “not now” list that is just as valuable: the work you’ve deliberately agreed not to do yet
- A shared mental model of the business strategy that developers, designers, and product all recognise

### Costs

- 6–9 person-hours per session with 4–6 people
- A pre-session goal-setting conversation (sometimes the real work)
- Political cost when the map reveals that work people wanted to do doesn’t connect to any goal
- Quarterly recurrence – the map goes stale as the goal, the actors, and the learnings move

### Failure modes

- The map gets produced, photographed, and then ignored because the backlog keeps being the source of truth
- The deliverables column is full of big builds and no cheap experiments
- One actor absorbs the entire conversation and the rest of the map is thin
- The team confuses “we wrote it down” with “we agreed on it” – absent stakeholders discover the map later and veto half of it

- Impacts drift into features and nobody catches it

### Stop signals

- The goal can't survive five minutes of scrutiny
- Every impact the team writes is a feature in disguise
- Key stakeholders aren't in the room and the map depends on their actors

Stopping and fixing the goal is not failure. Running a session that produces an elegant map of the wrong thing is.

---

## Worked Example

See [Impact Mapping: Connecting Work to Goals](#) for the Greenbox team's first mapping session – including the moment they realise a feature they've been planning for six weeks doesn't connect to any impact on the wall, and the relief of deciding not to build it.

---

## Outputs & Follow-up

### Facilitator's close-out (same day)

- Panoramic photographs of the map. Good lighting, readable notes, one shot per column in close-up.
- Transcribed map in a shared document or a digital mind-mapping tool – goal, actors, impacts, deliverables, and the lines between them.
- A one-page summary message to participants and stakeholders: here's the goal, here's the prioritised path, here's what we're deferring.

### The product owner's week

This is where the pattern earns its cost, and the work is mostly the product owner's.

- **Turn the priority path into backlog items.** Each deliverable on the priority path becomes a backlog item – but with the actor and impact captured in the description. When someone later asks "why are we building this," the backlog item contains the answer.
- **Park the deferred deliverables explicitly.** A "not this quarter" list is as valuable as the priority list. Put it somewhere visible. When new work gets proposed, the first question should be "does this displace something on the not-now list?"
- **Schedule discovery for the experiments.** Cheap experiments from the deliverables column – landing pages, manual concierge runs, interview scripts – need to be started within a week of the session. If they sit, the map's value decays fast.
- **Walk the map to absent stakeholders.** Anyone who should have been in the room but wasn't gets a walk-through. Their challenges either strengthen the map or reveal a problem you need to fix before committing.
- **Use the map to say no.** This is the hardest and most important week-after task. The map gives you a defensible reason to refuse work that doesn't connect. Use it.

## Ongoing

- Review the map at the start of each quarter or planning cycle. Has the goal changed? Have you learned which impacts actually work? Update and re-prioritise.
- When someone proposes new work, ask them to place it on the map. If it doesn't trace back to the goal through an actor and an impact, it probably isn't worth doing – or the map needs to grow.
- Keep the photographed map visible where the team works. It's the reference that prevents the slow drift back into feature-list thinking.

---

## Related Patterns

- **Event Storming** – Event Storming describes how things happen now; Impact Mapping describes what you want to change. Run Impact Mapping first when you're choosing what to build; run Event Storming first when you're understanding what already exists.
- **User Story Mapping** – once Impact Mapping has chosen the deliverables, User Story Mapping lays out the user journey through them and slices it into releases.
- **Assumption Mapping** – every impact on the map is an assumption. Assumption Mapping stress-tests which of those assumptions deserve to be tested before the team commits to the deliverable.
- **Business Model Canvas** – when the goal itself is unclear or the business model isn't coherent, the Canvas is the session to run before Impact Mapping. The Canvas sets the strategy; Impact Mapping executes against it.
- **Wardley Mapping** – Impact Mapping tells you what to change; Wardley Mapping tells you where each component sits in its evolution and therefore how to change it. They compose well for strategic initiatives.

## About this playbook

This playbook is part of *The Workshop*, a reference series of facilitator playbooks published at [barkingiguana.com](https://barkingiguana.com). The canonical, up-to-date version lives at [barkingiguana.com/writing/the-workshop-impact-mapping/](https://barkingiguana.com/writing/the-workshop-impact-mapping/).

*These posts are LLM-aided. Backbone, original writing, and structure by Craig. Research and editing by Craig + LLM. Proof-reading by Craig.*