

THE WORKSHOP

# User Story Mapping

A pattern for laying out the whole user experience as a left-to-right narrative and then slicing it into releases, so the team can see both the shape of the thing they're building and the thinnest honest version they can ship first.

2026-04-22

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

# Contents

Intent	3
Also Known As	3
Motivation	3
Applicability	3
Participants	4
Structure	5
Collaborations	5
Facilitator Playbook	6
Steering When It Goes Sideways	9
Consequences	10
Worked Example	10
Outputs & Follow-up	10
Related Patterns	11

---

*This is the pattern. If you want the story of User Story Mapping in action, read [User Story Mapping: Seeing the Whole](#) – the Greenbox team discovering that their flat backlog has been hiding a gap nobody noticed, and drawing the first honest release line. This post is the reference you keep open the morning of.*

---

## Intent

Lay out the full user experience as a left-to-right narrative, then slice it horizontally into releases, so the team can see the whole journey and commit to the thinnest honest version of it first.

---

## Also Known As

Often just called **story mapping**. Sometimes confused with **customer journey mapping** (journey mapping is research-led and emotional; story mapping is build-led and functional) and with **flat backlogs** (a backlog is a list; a story map is a grid). Invented and named by Jeff Patton in 2005, published as a book in 2014 that is still the canonical reference.

---

## Motivation

A team has a flat backlog of eighty-seven stories. They've been working through it for six weeks. Last week they shipped a beautiful payment form. This week they're building subscription upgrade logic. Next week they're building referrals. The product owner writes up a release announcement and notices, with a growing sense of unease, that nothing the team has built actually lets a new visitor sign up, choose a box, and receive their first delivery. The payment form doesn't connect to anything yet. The upgrade logic assumes a subscriber who can't yet exist. The referrals are for a product that has no users to refer. Every story delivered was a good story. The sum of the stories is not a product.

This is the flat-backlog failure mode. A list tells you what's next; it doesn't tell you what fits together. Priority orders stories by perceived value but loses the journey shape. Teams optimise each story locally and discover, six weeks in, that they've been building disconnected parts.

User Story Mapping exists to put the journey back. The wall is the shape. The vertical axis is priority; the horizontal axis is the user walking through the product end to end. A release is a horizontal line across the map, and the question the line forces is: *what is the thinnest version of this journey that still works as a journey?*

---

## Applicability

### Use when:

- You're planning a new product or a major new feature area
- The backlog has grown large and nobody can see the big picture
- You need to define an MVP or a first release and the arguments keep going in circles
- Different team members have different mental models of what the product does end-to-end

- You've finished Impact Mapping or Event Storming and now need to turn the insights into a buildable plan

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

- You're mapping a single, well-understood story – use Example Mapping
- You don't have a clear user or set of users to map for
- The work is purely technical with no user-facing narrative – infrastructure, internal tools, refactoring
- The scope is so broad that you're really trying to decide strategy – run Impact Mapping first

**Inputs.** One clear user persona (written on a card and pinned to the left of the wall), a rough agreed scope for this map (the full product, or one journey within it), and any existing research you can reference without putting it on the wall.

---

## Participants

**Facilitator.** Keeps the map growing in the right direction, catches backbone items that are really tasks, and runs the walk-the-map ritual.

**Product owner.** Mandatory. They're the narrator during the walk-the-map phase – the person who tells the user's story out loud while everyone else listens for gaps. They also make the final release-slicing call when the team disagrees.

**Developers.** At least two. They'll ground the map in what's actually buildable and catch tasks that look small but hide weeks of infrastructure work.

**Designers.** They think in journeys natively. A designer will reshape the backbone halfway through the session in ways a developer or product owner wouldn't have thought to.

**People who talk to real users.** Support, sales, operations, account managers. They'll add the unhappy paths the golden-path team forgot: the subscriber whose card declined, the pause that went wrong, the box that arrived damaged.

**Operations / SRE.** For products where operations are part of the user experience – on-call engineers, deployment pipelines, support agents using internal tools – the user being mapped might be *them*, and ops is the domain expert. Don't tuck them in as afterthoughts.

**Group size:** 5–8. Fewer and you miss perspectives; more and the wall becomes a crowd. If you're forced above 10, split into two sessions with overlapping attendance and reconcile the maps afterwards.

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

- **Real end users.** Their presence warps what the insiders will say. Interview users separately and bring their words into the room as evidence, not as voices.
- **Senior leaders who will turn the session into a requirements meeting.** Story Mapping is discovery; requirements come from it, not into it.
- **Spectators.** Anyone "just observing" is absorbing attention without contributing. Either they participate or they read the output.

## Structure

Phase	Duration	Notes colour	Key question
Orient, persona, scope	15 min	–	“Who is this map for and how far does it go?”
Backbone	20 min	Blue	“What does the user do at the highest level?”
User tasks	30 min	Yellow	“What specific things do they do at each step?”
Walk the map	15 min	(review)	“Does this journey make sense end-to-end?”
Slice releases	30 min	Tape lines	“What’s the thinnest complete journey?”
Wrap-up	10 min	–	“What’s in release 1?”
<b>Total</b>	~2 hours inside a 2–3 hour block		

Budget three hours for a first session on a new product. Budget ninety minutes for a map of a single feature area inside an existing product. Do not try to map two different personas on the same wall in the same session – split them.

## Collaborations

Story Mapping is a standing-up, walking-around, wall-based ritual. Nobody sits. Notes move. The shape of the room matches the shape of the map: wide, layered, with people moving along it as the conversation moves.

The rhythm is **backbone, then vertical detail, then narrate, then slice**:

- During the **backbone** phase, one person – ideally the product owner – tells the high-level story in order. Everyone else listens and places blue activity notes. It’s a single voice telling the shape of the journey; the facilitator’s job is to catch details that slip into the backbone before they belong there.
- During the **user tasks** phase, the conversation opens up. Everyone works on multiple activities at once, writing yellow task notes and placing them vertically. People move around the wall. The facilitator circulates and catches tasks that are really implementation details or screen specs.
- The **walk-the-map** phase is a ritual interruption. Stop adding notes. One person narrates the entire journey aloud, left to right, using only the notes on the wall. Everyone else listens for gaps. Then gaps get filled.
- The **slice releases** phase is the most political. A piece of tape goes across the wall. Every “above the line” decision is someone committing to ship something and someone *not* committing to ship something else. The facilitator holds the space for that trade to happen honestly.

## Facilitator Playbook

### Phase 1 – Orient, persona, scope (15 minutes)

Before any note goes up, pin the persona card to the left end of the wall and read it aloud:

*"Today we're mapping the experience of a first-time subscriber – let's call her Anna. She's health-conscious, she's busy, and she's just heard about us from a friend. The map we build is her journey, from the moment she hears about us to..."*

Then agree the end-point explicitly:

*"...where does the journey end? First delivery? Three months of subscribing? Cancellation and win-back? Pick one. We'll map that scope and call it done."*

Scope drift is the single most common Story Mapping failure. A team starts mapping "sign up to first delivery" and an hour in discovers they're also mapping pause, substitution, and cancellation. The wall fills up and the release slice becomes impossible. Agree the end-point now, write it on a card, stick it at the right end of the wall. That's the boundary.

#### What to watch for:

- **Two personas trying to share a wall.** *"But the supplier does X..."* If the conversation keeps switching personas, you're mapping two journeys. Split them into two sessions.
- **Scope that's too broad.** *"The whole product."* That's six maps, not one. Pick the first-time subscriber journey, or the pause journey, or the renewal journey. One at a time.
- **Missing persona.** The team can't quite describe who the map is for. Pause: *"Let's spend ten minutes writing a persona card before we draw anything."*

### Phase 2 – Backbone (20 minutes)

The backbone is the user's journey at the highest level – six to twelve big activities from the start of the journey to the end. These go along the top of the wall, left to right.

Ask the product owner:

*"Walk me through what Anna does, from the very beginning. Not in detail. Big chunks. What's the first thing that happens?"*

Write each activity on a blue note and place it. Keep the granularity high: *"Discover the service," "Sign up," "Choose a first box," "Receive first delivery," "Manage ongoing subscription," "Refer a friend."* Not *"Click the sign-up button,"* which is a task, not an activity.

Aim for 6–12 activities across the backbone. More than that and you're at the wrong granularity; fewer and you're missing stages.

#### What to watch for:

- **Starting with the system, not the user.** *"The system sends a welcome email."* Reframe: *"What does Anna do? She opens the email and reads it. That's the activity."*
- **Skipping discovery.** Teams start the backbone at *"Sign up,"* forgetting that Anna has to hear about the service first. Prompt: *"What happens before she knows we exist?"*

- **Skipping the end.** Teams end the backbone at “*First delivery,*” forgetting ongoing management, cancellation, win-back. Prompt: “*What happens after three months? After she tries to pause? After she cancels?*”
- **Tasks leaking upward.** Someone places “*Enter email address*” on the backbone. That’s a yellow task under the “*Sign up*” activity. Gently move it down: “*Great detail – let’s put it below, under ‘Sign up,’ when we get to tasks.*”
- **Operations backbones.** For an SRE-flavoured map – say, the journey of a deployment from commit to verified rollout – the backbone activities might be “*Developer pushes,*” “*CI runs,*” “*Artefact built,*” “*Staging deployed,*” “*Production rolled out,*” “*Rollback decision available.*” Same shape, different domain.

### Phase 3 – User tasks (30 minutes)

For each activity on the backbone, the team writes yellow task notes describing what the user does during that step. Tasks go vertically below their parent activity, ordered roughly top (most essential) to bottom (nice to have).

Open the phase:

*“For each blue note on the backbone, I want specific things the user does. Not UI details – intents. ‘Browse available boxes.’ ‘See what’s in each box this week.’ ‘Pick a delivery day.’ Write them on yellow, place them below the activity they belong to, and roughly stack them by importance – most essential at the top.”*

Let the conversation flow across the wall. People will jump between activities as they think of related tasks. Let them. The facilitator circulates and catches problems.

#### What to watch for:

- **UI details dressed up as tasks.** “*Click the dropdown.*” Not a user task – a UI interaction. The task is “*Pick a delivery day.*”
- **Missing unhappy paths.** The team maps the golden path. Prompt explicitly: “*What if her card is declined? What if the box she wants is sold out? What if she signs up, then immediately changes her mind?*” Unhappy paths are often where the release line is hardest to draw.
- **One activity with twenty tasks, another with two.** The dense activity probably needs splitting into two activities; the sparse one might be fine, or might be missing work.
- **Arguments about horizontal order.** Within an activity, vertical order (priority) matters more than horizontal order. If two people disagree about what comes first horizontally, there might be two valid paths – capture both.

### Phase 4 – Walk the map (15 minutes)

Stop adding notes. Everyone takes three steps back from the wall.

Ask the product owner to narrate the full journey:

*“Walk me through Anna’s experience. Left to right. Use only the notes on the wall. Tell her story as if I’d never heard it.”*

Everyone else listens. The facilitator's job is to catch the stumbles:

- *"And then she... um, signs up, and then somehow ends up with a box..."* – there's a missing activity or a missing task.
- *"She picks a box and pays..."* – wait, is the payment activity there? Or is it hiding inside "sign up"?
- *"...and she receives her first delivery."* – what happens if she doesn't? Where's the failed-delivery path?

When the narrator stumbles, pause. Ask the room:

*"What's missing here?"*

Fill the gap. Continue the walk. The walk-the-map phase catches more problems than any other single phase in the session. It is the quality check.

### What to watch for:

- **Nobody challenging the narrator.** The room is polite. Name people by the slice of reality they own: *"From what you see in deployment tickets, does this match? From what support hears on the phones, does this match?"*
- **Duplicate tasks.** The same task appearing under two activities. Is it genuinely part of both, or is one misplaced?
- **The narrator skipping sections.** *"And then all the usual stuff happens, and..."* Interrupt: *"Walk me through the usual stuff."*

### Phase 5 – Slice releases (30 minutes)

This is the most valuable and the most political phase. Take a piece of tape or draw a horizontal line across the wall. Above the line: release 1. Below the line: later.

The rule of the slice is simple:

*"Release 1 is the thinnest horizontal slice that still tells a complete story left to right. Anna can walk from the leftmost activity to the rightmost and achieve her goal. There will be fewer options, less polish, and more manual work – but the journey has to be whole."*

For each activity, the question is the same:

*"What's the absolute minimum version of this step that lets Anna get through it?"*

For the "Choose a box" activity, the minimum might be: *"Browse available boxes. Select a size."*

Everything else – substitutions, weekly previews, family-size recommendations, gift wrap – goes below the line.

You can draw multiple lines for multiple releases. Release 1 is the MVP. Release 2 is the next thinnest slice. And so on. The point is that every release above the line is a complete journey, not a complete activity.

### What to watch for:

- **Release 1 is "everything above the line for every activity."** The most common mistake. If release 1 is the full golden path for every step, it's not an MVP, it's the whole product. Push: *"Can a subscriber complete this step with just one option instead of five?"*

- **Cutting entire activities.** If an activity has nothing above the line, the journey has a hole. Every activity needs at least one task in release 1, even if the task is manual or minimal.
- **“We can’t launch without...”** Some things genuinely can’t be cut (payment processing, for instance). Some things feel essential but aren’t (substitution preferences for launch). Challenge each claim individually.
- **Uneven slices.** One activity has eight release-1 tasks, another has one. Sometimes that’s correct – payment really does need more than preferences – but check that the dense activity isn’t hiding overbuild.
- **Operations-flavoured slicing.** For a deployment-pipeline map, the release 1 slice might be *“Manual rollback, alerts go to one channel, health checks are basic, observability is minimal”* – a pipeline that works end-to-end but isn’t polished. Later releases add automation.

## Steering When It Goes Sideways

**The architect.** Someone keeps mapping system architecture instead of user journey. *Recovery:* “What does Anna experience at this point? We’ll figure out the technical flow later.” *Stop if:* They can’t hold the distinction after three prompts. They belong in a design session that happens after the map.

**The everything-is-essential person.** Someone argues every task is critical for release 1. *Recovery:* Impose a constraint: “We have six weeks to launch. What can Anna live without until release 2?” Constraints force prioritisation in a way that abstract discussions don’t. *Stop if:* The person won’t accept any constraint. Escalate – they need a separate conversation about scope with the product owner.

**The map gets too big.** The wall is full and the team is still adding activities. *Recovery:* Scope is too broad. Pick the most important journey (e.g., first-time subscriber sign-up to first delivery) and park the rest for separate sessions. *Stop if:* The team can’t agree on which journey to focus on. That’s a strategy problem, not a mapping problem.

**Two users on one map.** The conversation keeps switching personas. *Recovery:* “I’m hearing both subscriber tasks and supplier tasks. Those are two maps. Let’s finish the subscriber one today and schedule the supplier session for next week.” *Stop if:* The team insists the personas share a journey. Check: do they really? Or are you trying to save a session?

**The design session.** People start sketching screens. *Recovery:* “Park the screens. We’re mapping what Anna needs to do, not how the screen looks. The design comes after we agree on the journey.” *Stop if:* Screens keep creeping back in. Pair the designer with a developer to hold each other honest.

**The silent release-slicing.** The team is quietly placing tasks above or below the line without any debate. *Recovery:* Slow it down: “Before we go further, can someone explain out loud why the substitution preferences are above the line? I want to hear the reasoning.” The point of the slice is the conversation about the trade-off. *Stop if:* The team still won’t engage. Something else is going on – maybe the release date is imposed and the slice is performative. Name it.

## Consequences

### Benefits

- A shared picture of the whole product, end to end, that persists after the sticky notes come down
- An MVP definition the team can defend because the journey is visibly complete
- A backlog organised by both priority (vertical) and journey stage (horizontal), so new work has a place to go
- The “what about...” questions caught on the wall rather than mid-sprint
- An artefact that works as a communication tool for stakeholders who weren’t in the room

### Costs

- 10–24 person-hours for a 2–3 hour session with 5–8 people
- A long wall and a room that can accommodate people standing and moving for hours
- The discipline to protect the scope when the map wants to grow
- Coordination cost of getting the right people in one room for a long session

### Failure modes

- Scope drift: the map grows beyond the agreed journey and release slicing becomes impossible
- The map ends up being mostly one person’s mental model because the walk-the-map phase didn’t catch it
- Release 1 is “everything above the line” because the team couldn’t hold the thin-slice rule
- The map is produced and photographed but nobody updates it, so it goes stale within weeks

### Stop signals

- The team can’t agree which persona to map
- Scope has quadrupled two hours in and the wall is full but incomplete
- The walk-the-map narration reveals that nobody in the room actually understands the user

Stopping when the scope isn’t right is not failure. Producing a beautiful map of the wrong journey is.

---

## Worked Example

See [User Story Mapping: Seeing the Whole](#) for the Greenbox team’s first mapping session – including the moment the walk-the-map phase reveals a gap between “pays for the subscription” and “receives first box” that nobody had noticed, and the release-slicing conversation that saves a month of scope.

---

## Outputs & Follow-up

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

- Panoramic photographs of the full wall. Good lighting, sharp focus, enough resolution to read every note.
- Close-up shots of each activity section, so the detail is preserved even if the panorama isn’t sharp enough.

- 
- Transcribed release-1 tasks into the backlog with the activity as context, so every story knows which journey stage it belongs to.
  - A message to all participants with the photographs and the release line clearly marked.

### The product owner's week

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

- **Turn release-1 tasks into backlog items.** Each yellow note above the line becomes a backlog item, with the activity as context. The shape of the wall becomes the shape of the sprint plan over the next several iterations.
- **Protect the release-1 slice.** The single hardest follow-up work. Every time a stakeholder asks for "just one more thing in the first release," check the wall. Either it moves above the line (and something else moves below) or it waits for release 2. The map is the reason you can say that and have it be defensible.
- **Begin Example Mapping on the top release-1 tasks.** Story Mapping produces tasks; Example Mapping makes them buildable. Start on the most essential tasks first.
- **Walk the map to anyone who couldn't attend.** Their perspective may reveal gaps the original group missed, or validate the slice. Either outcome is valuable.
- **Schedule any discovery work.** Tasks on the wall that are guesses – "we think Anna will want this" – become research or experiment proposals. Don't let the guesses graduate into commitments without being tested.

### Ongoing

- Keep the map visible while the work is active. Print it, photograph it, pin it near the team's desks. Teams that can glance at the map during standups and planning make better decisions than teams working from memory.
- Update the map after each release. Move the line to the next slice. Add new tasks learned from real user feedback. Remove tasks that turned out not to matter.
- When new work is proposed, place it on the map. If it doesn't fit, either the map needs updating or the work doesn't belong. That's a valuable filter.

---

## Related Patterns

- **Event Storming** – Event Storming maps the system from the inside; Story Mapping maps the user experience from the outside. Running Event Storming first can reveal the hotspots; Story Mapping turns the insights into a release plan.
- **Impact Mapping** – Impact Mapping picks the deliverables; Story Mapping arranges those deliverables into a user journey and slices them into releases. Impact first, then story.
- **Example Mapping** – the release-1 tasks from a Story Mapping session are the input to Example Mapping. Story Mapping gives you the list of stories; Example Mapping decides whether each one is ready to build.
- **Sprint Planning** – once the release-1 slice exists and Example Mapping has run on the top stories, Sprint Planning turns the map into committed sprints.

- **Assumption Mapping** – the release-1 slice is a stack of assumptions about what users need. Assumption Mapping stress-tests the slice before you commit to building it.

## 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-user-story-mapping/](https://barkingiguana.com/writing/the-workshop-user-story-mapping/).

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