



CUSTOMER DISCOVERY: An Iterative Process – Episode 2

TITLE:

Turning User Interviews into a Solution Sketch

DESCRIPTION:

Building on trusted frameworks like jobs-to-be-done to sketch the first version of a compelling solution through a creative process and structured guidance

[LEARN@LIFT Episode Intro]

In this episode, we're going to discuss turning all of the research you conducted with potential users during Episode 1 into a rough, first version of your product using a handful of creative structures to guide you.

The next step after finishing your user interviews and creating your job map, is moving to build out a simple problem map. A problem map connects all of the different but associated problems your users face to each other, into a set of cause-and-effect links.

Your job map may have produced a whole list of pains. The question is, what contributes to those pains? What ultimately causes them?

To begin your problem map, select a pain point that was cited most often during the course of your interviews, and one that felt central to the core job your user is trying to get done.

From there, use the "5 Why's" to dig deeper. If you can ask, "Why?" five times to explain why a problem happens, you've hit a reasonable level of depth.

As you continue to develop your problem map, it's common to the problem you're tackling spider out into different directions. So to start your solution design, you'll want to pick a point on the problem with the highest leverage — a point where a fix will end curing the pain from a number of associated, "downstream" problems.

For example, let's say you want to help struggling marketing teams deliver more leads to their sales team. But the first problem is ads are becoming more expensive for the same lead volume. The next problem here is an over-reliance on paid over organic lead generation. Which could be attributed to a lack of SEO content. And this may be related to limited bandwidth on the marketing team to build the content. You get the idea.

A good old fashioned whiteboard can be perfect for this exercise, but online tools like Miro, or LucidChart are great options as well, and come with a set of templates you can leverage for this exercise.

Once you've picked the starting place for your solution, your next step is to create a series of "How might we?" statements.

These statements suggest varying approaches to attacking the problem. It's a simple exercise that works just like it sounds. On your whiteboard, begin generating a series of sticky notes that finish the question, "How might we," with a suggested fix.

For example, if we continue on with the idea of an organic lead generation solution for B2B marketers, we might ask ourselves, "How might we..."

- ...generate first drafts of SEO content using AI?
- ...automate keyword briefs with a ready-made content outline?
- ...create content from customer conversations and support surveys?

The key here is to try to develop as many distinct approaches to solving the problem as you can. You'll want to stretch yourself to a minimum of 20 different statements, but know that some problems could have two or three times as many how-might-we options.

As you do this, you'll notice certain themes beginning to appear. You can start to group those statements to identify patterns on possible solutions that might work for your target users.

After you've completed your How Might We exercise, the next exercise is called Crazy 8's.

Crazy 8's is the process of sketching out a series of visuals for how a solution might be designed. You limit yourself to only 8 screens, to focus on outlining the basic user flow of a product — whether it's hardware or software, business or consumer.

Crazy 8's can be done by folding a sheet of paper in half, and then in quarters, and sketching on the paper manually. Or, you can use basic shapes and lines inside a whiteboarding tool like Miro.

The goal isn't artistic perfection. Rather, the goal is to create a wide and divergent group of multiple solution options, which you can then start to combine into a single approach, based on patterns and commonalities you find.

To help you generate as many ideas as possible, it can be helpful to set a timer for 5 minutes. Make sure to complete all 8 screens of your user flow inside this time, and let the ideas flow freely. Complete this exercise at least six times, and you'll have created a solid set of options to begin evaluating by comparing them to the user interview summary you created.

By the way, all of these exercises are ones that you can do alone, but it's ideal to do them as a group. You'll begin to see how other people think about approaches and implementing solutions that you may not have considered on your own.

At this point, you should have a rough idea of what your solution will look like. Let's suppose we decided to combine two of our How Might We statements from earlier into one tool that scans customer surveys for patterns and common questions, and compares this content to an SEO database that scores online search keywords, to build a blog post outline for B2B marketing teams.

Our next step is to create a low-fidelity prototype that shows potential users how they'd interact with the solution.

In the case of our marketing product, we might start with a low-cost prototyping tool, like Proto.io, Figma, or even Miro. If you're building a physical product, CAD or sketch software is going to be a fast, cheap way of communicating your idea to your customers.

If you keep in touch with the customers you interviewed previously, reach back out to them, and let them know you've made progress after several more interviews, and have a solution to run by them. Their willingness to connect with you for a second time and help design your solution is a great sign you're solving a truly meaningful problem for them.

During these conversations, you'll want to continue asking open-ended questions about how they'd attempt to interact with the product, what they'd expect it to do, and how helpful they believe the general approach is.

It's okay to stress that this is a basic version one, without the final graphics, designs, and aesthetics. For now, you're only focused on the rough function and feasibility of the solution.

From here, you'll find yourself in the middle of a build-learn-adapt loop.

You'll want to take the feedback on likes and dislikes, what's clear and what's confusing, and apply it into a second version of what you've developed. The changes you make should focus on what the product does and does not do — not how it looks, the details of how a user clicks through, or purchases the product.

Most every low-fidelity or "Lo-Fi" prototype should be black and white, static, and rather rudimentary to help you ensure you're on the right track before leaning into the visual design and user experience of your product.

This is because the goal of this approach is to spend as little capital as possible in the early days. It can be tempting to believe you've got something great without seeking multiple rounds of feedback from a fair number of people.

Product development is the most expensive part of company-building in the early days, so before you pay developers to code or manufacturers to produce your product, you've got to be as close to certain as possible.

In summary, selecting a high-leverage point in a commonly-expressed problem, then brainstorming a number of approaches through How Might We statements, paired with a low-fidelity sketch of your solution for user feedback is the best combination to ensure you're ready to begin testing your product with users.

This topic — user testing — will be the focus of our next episode.

[Insert LIFTLabs Outro]