

THE CURRENT FRAMEWORK HAS ALREADY SOLVED HALF THE PROBLEM

Product discipline already has a validated standard for discovering what the customer needs before building. What that same discipline never solved is the inward half: what the company itself knows about its own operation, and has never asked itself.

Product discipline already has a validated standard, taught today, for discovering what the customer needs before building. What it never solved — and it is no accident — is the inward half: what the company itself knows about its own operation, and has never asked itself.

THE FRAMEWORK THAT EXISTS, HAS A NAME AND HAS AN AUTHOR

In 2021, the product researcher Teresa Torres published *Continuous Discovery Habits*, and in the years that followed the book became the most cited reference in the discipline — not because it is recent, but because it is the first product-discovery framework built to run every week, indefinitely, rather than once a quarter.

The central mechanism: a **product trio** — product manager, designer, engineer — runs a structured customer interview **every week**, without exception, and organises what it learns into an **Opportunity Solution Tree**, a visual map linking the desired outcome to the real opportunities found in conversation, and those

opportunities to possible solutions. Torres sums up the central gain in a phrase that also sums up the argument of this text: what matters is not the speed of building. It is **Time to Learn** — how long it takes between an idea and its real validation.

In a public interview, she criticised the most common and shallowest use of OKRs — the goal-setting tool created by Andy Grove at Intel in the 1970s and now standard in almost every technology company: “*we see a lot of people adopting OKRs, and the objective is ‘ship the Android app’. That’s a work result, not a business result.*” Counting delivery is not measuring impact — and Torres points to this as the most common mistake among teams who think they are measuring outcome and are, in practice, only counting output.

WHAT AI HAS ALREADY CHANGED IN THIS FRAMEWORK, THIS YEAR

In 2026, the community that follows Torres is already updating the framework for the AI era — and the most recent text on this, published in April by Perspective AI, exposes a number nobody had calculated clearly before: **a single thirty-minute interview, in the classic continuous-discovery model, costs between four and six hours of human work** — identifying the sample, pulling the list from the database, writing the invitation, scheduling, dealing with reschedules and no-shows (a no-show rate of **20% to 35%** in non-incentivised B2B research, according to the same study), running the call, transcribing, tagging excerpts, synthesising.

A trio of three people devoting six hours a week to a single interview consumes around 5% of the team's capacity — and that budget collapses at the first tight deadline, the first outage, the first urgent request from an executive. It is not a failure of the framework. It is the structural cost of doing this with people running every manual step.

The solution the product market itself is now adopting, according to the Perspective AI study of April 2026: AI-moderated conversation cuts the cost of an interview from six human hours to around **thirty minutes of trio review** — without eliminating the trio, without eliminating human judgement, only eliminating the mechanical work of recruiting, scheduling and transcribing.

WHERE THIS MEETS ZTHEX — AND WHERE IT DOES NOT

Continuous Discovery, even updated with AI, still looks **outward** from the company — at the customer, at the market, at what they want that does not exist yet. It is the external half of discovery, and it now has a validated framework, a recognised author and a proven cost solution.

The internal half never got the same treatment. Nobody has published a *Continuous Discovery Habits* for the question “what does our own operation know, and has never been asked of our own team”. Not because the problem is smaller — MIT measured, in the same year, that 95% of enterprise AI pilots fail precisely for the absence of that internal knowledge, not the external one. It is because product discipline was born looking at the customer, and the habit of looking inward never became a framework with a name.

That is the technically empty space in which Zthex operates. The same logic as Torres — structured, recurring conversation, with a defined trio, measuring time until the learning is validated — applied to the side the whole discipline forgot to name: what the maintenance technician, the service agent, the sales consultant know about their own operation, and which never reached any Opportunity Solution Tree because Torres's tree was designed to listen to customers, not colleagues.

// we do not compete with Continuous Discovery. We occupy the other half of the same problem, with the same rigour of method — only turned inward.

OKRS, AT THE RIGOUR THEIR OWN CRITIC DEMANDS

Torres's criticism of shallow OKRs applies, point by point, to what Zthex delivers too — and it is worth using as a yardstick of honesty, not only of criticism aimed elsewhere.

A badly designed objective would be: “reduce operational divergence”. That is output in disguise — it does not say what changes in the life of whoever does the work. A well-designed

objective, by the standard Torres herself defends, has to name the behaviour or result the number represents: **how many findings left the cycle with a defined destination, and how many of those destinations were, in fact, implemented** — not “how many conversations happened”, which is pure output, in exactly the way “ship the Android app” is.

That is why the Zthex Divergence Index is measured as a **baseline in the first cycle**, with the reduction assessed **between cycles**, never within the same one — because promising a drop within the same cycle would be exactly the mistake Torres names: counting the output as if it were the outcome.

THE GAP IS STILL THERE, WAITING FOR A NAME

Continuous Discovery solved, with rigour and a recognised name, the question “what does the customer want that we have not built yet”.

Nobody has yet solved, with the same rigour, its sibling question: “what does our own operation already know, and no framework has asked it yet”.

WHAT THIS TEACHES ABOUT METHOD

Three lessons, taken straight from the primary source, applicable to any discipline of discovery — customer or internal operation.

01 Measure the time to learn, not the speed of building.

Torres proved it for product discovery. The same holds for operational discovery: what matters is not how many conversations happened, it is how long it took between the operation having a problem and someone with the authority to decide knowing about it.

02 The human cost of recurring discovery is real, measurable, and falls short of zero.

Six hours become thirty minutes because AI takes over the mechanical work — recruiting, scheduling, transcribing — not the judgement. Torres's trio still decides. At Zthex the rule is the same: AILA runs the conversation; validating the finding remains always human.

03 A badly designed objective counts what was done, not what changed.

It is the simplest lesson and the most ignored — in product, in operations, in any discipline that measures itself.

BOX — THE SOURCE, IN DATES	
2016	Torres introduces the Opportunity Solution Tree, inside the framework itself
2021	Teresa Torres publishes <i>Continuous Discovery Habits</i> — the discipline's standard, still taught in 2026
4–6h	human cost of a single discovery interview in the classic model <i>(Perspective AI, April 2026)</i>
20–35%	no-show rate in non-incentivised B2B research, in the same study
~30 min	cost per interview with AI-moderated conversation and trio review, in the same study
1970s	origin of OKRs, by Andy Grove at Intel — now a market standard, with the same risk of shallowness Torres names publicly

The outward half already has a framework. The inward one is still waiting for its own.

Editorial content produced by Zthex. Facts about third parties come from public sources cited in the text.

ABOUT ZTHEX

Zthex structures the layer of knowledge no system shows. It talks to the people who do the work, removes identity before recording, and compares what was said with what policy states and what the systems record. Every divergence becomes a finding with a traceable origin.

PRESS

Signed republication and additional reporting:
zthex.com/en/imprensa
info@zthex.com



ZTHEX.COM
point your
camera