

How to get the **most out of your** customer feedback

UNIFY · UNDERSTAND · ACT · IMPROVE

Most companies already have the feedback they need. This is how to actually use it: four stages, and 90 days to run the loop once.

INSIGHTS FROM

Canva

 Notion

ElevenLabs

 Wispr

Contents

START HERE

Feedback isn't the problem. What happens next is.	03
--	----

THE RUBRIC

The Loop Maturity Model	04
-------------------------------	----

Add it up: scoring your loop	05
------------------------------------	----

THE FOUR STAGES

01 · Unify	06
------------------	----

02 · Understand	08
-----------------------	----

03 · Act	10
----------------	----

04 · Improve	12
--------------------	----

WHAT IT COMPOUNDS INTO

Two companies that look identical	14
---	----

Where loops break	15
-------------------------	----

Who owns the loop	16
-------------------------	----

Walk one problem through all four stages	17
--	----

Thirty, sixty, ninety	18
-----------------------------	----

Feedback isn't the problem. What happens next is.

Most companies are not short on customer feedback. It arrives every hour, from a dozen surfaces, in volumes no team can read. The problem is what happens next: most of it is never turned into a decision, a fix, or a change the customer can feel.

Collecting feedback and acting on it are different jobs. Most programs are built almost entirely for the first.

The stakes just rose. Agents and assistants now answer for the business at the speed of software. Feed them scattered, contradictory feedback and they will be confidently wrong in front of customers.

The constraint is the quality of your customer context, not the intelligence of the model reading it.

01

Unify

Every channel into one structured source of truth.

02

Understand

Quantify what customers need and who it affects.

03

Act

Route it to the work, then close it back to the customer.

04

Improve

Make every cycle sharpen the next one.

The Loop Maturity Model

Twelve questions about the loop as it runs this week, not the one on the roadmap. Score it on the next page.

	NO 0	PARTLY 1	YES 2
UNIFY DATA			
01 Every channel lands in one place without anyone copying it there.	☐	☐	☐
02 Feedback is tied to the account, plan and revenue behind it.	☐	☐	☐
03 Someone owns the taxonomy, and it matches today's product.	☐	☐	☐
UNDERSTAND EVIDENCE			
04 Issues are sized by distinct customers and revenue, not mentions.	☐	☐	☐
05 You can tell a representative theme from a loud minority.	☐	☐	☐
06 People outside research can answer their own feedback questions.	☐	☐	☐
ACT ACTION			
07 Insight reaches whoever can resolve it in hours, not weeks.	☐	☐	☐
08 Spikes surface without a human running a report.	☐	☐	☐
09 Customers who reported an issue hear when the fix ships.	☐	☐	☐
IMPROVE MEMORY			
10 Prior findings and ownership carry into the next cycle.	☐	☐	☐
11 Your AI front line answers from the same truth as the roadmap.	☐	☐	☐
12 You can trace a recent product change to the feedback behind it.	☐	☐	☐

Add it up

YOUR TOTAL _____ / 24

Two points for yes, one for partly, none for no.
Count honestly: a question you would need a week to answer is a no.

0-7

You are collecting feedback, not using it.

Five teams hold five partial versions of the truth, and decisions run on whoever spoke last.

Next — Get every channel into one deduplicated stream tied to accounts.

8-14

You can see the feedback but cannot move on it.

You know what customers said; you cannot say how many, how much revenue, or who owns the fix.

Next — Weight themes by distinct customers and revenue, and give the taxonomy one owner.

15-20

The loop runs because people push it.

Every step works, and every step needs a person to make it happen. Throughput is capped by their week.

Next — Automate routing and confirmation — the two steps that quietly consume the most human time.

21-24

The loop runs itself, and gets better as it runs.

Findings, ownership and context carry forward, so each cycle costs less than the last and the AI front line stays grounded.


Next — Widen it: more channels in, more teams served from the same truth.

Most teams score between 8 and 14: the data is in one place, the sizing is rough, and nothing routes itself. Whatever your score, your AI front line cannot outrun it.

Unify

One structured source of truth: every stream of feedback, deduplicated, tied to the accounts and revenue behind it, under a taxonomy that reflects how your customers actually talk.

“Beyond just keywords, we can actually understand what our users are saying across every interaction we plug in.”



Emma Auscher

Global VP of Customer Experience, Notion



Five teams hold five conflicting versions of the truth. Every later stage, and every AI system, inherits the quality of this foundation.

01 / Centralize every channel

Canva unifies 13,841 feedback themes across 34 sources, from Zendesk to Gong to Reddit to Discord — so a marketer, a PM and a support lead query the same reality.

02 / Normalize the data

A dozen spellings of "USA", timestamps as raw epochs, the same person a User ID here and an Author there. Standardising those fields is what makes segmenting possible at all.

03 / Tie feedback to revenue

BuildOps ties feedback to 128,000+ accounts carrying ARR, segment, churn risk and assigned CSM. A complaint arrives with a dollar figure and an owner attached.

04 / Build the taxonomy

Philo classifies 319,000+ insights into a hierarchy you navigate rather than excavate. This is the shared language every team and every model reads from.

This is maintenance, not setup. At Western Union, someone renamed a set of taxonomy keywords and every downstream report and automation built on the old mappings silently broke. No error message — the data just quietly began to mislead.

THE TEST

Pick one complaint from this morning. In under a minute, can you see everywhere else that customer has raised an issue, tied to their account and revenue?

Understand

Quantified truth: how many customers are affected, which segments, how much revenue sits behind it, and whether a theme is genuinely widespread or merely loud.

“We surface product insights in days, not weeks. Interpret aligns everyone around the top issues and requests.”



Jesse Walker

Head of Insights and User Advocacy, Canva

Canva

The default way companies prioritise is that the loudest voice wins. Measurement replaces that, fast enough that the answer arrives before opinion hardens into a decision.

01 / Quantify demand and cost

At one design platform the top request carries ~20,000 mentions, roughly twice the next item — a debate settled without opinion.

02 / Catch what instrumentation misses

Logs tell you what your software did, not which bug never threw an error. One AI coding tool surfaces memory leaks that trigger crashes; a consumer app finds crashes during photo upload.

03 / Read commercial risk early

Feedback carries churn and pricing signals before revenue does — a map of where users hit the paywall and resent it.

04 / Separate representative from vocal

Ten furious power users repeating a complaint forty times is not a theme shared by ten thousand customers. Weight by distinct customers and revenue, never by mentions alone.

Serve each team the cut they need. Roadmap evidence for product, churn signals for CS, positioning gaps for marketing, blast radius for engineering.

THE TEST

Can you state, as a number you would defend in a meeting, how many customers an issue affects and how much revenue is exposed?

Act

Change the customer can feel: insight routed to the people and systems that do the work, AI grounded on it, and the customer told when their issue is resolved.

“It elevates what our support engineers are capable of. They identify root causes and hand engineering issues that are already actionable.”



Samm Yates

Founding Head of CX, Wispr Flow



Understanding without action isn't neutral, it's corrosive. Every time a customer speaks and nothing changes, they learn that speaking up is pointless.

01 / Route it to where the work happens

One communications app links themes straight to engineering tickets, so a one-way-audio bug arrives with the complaints attached and the squad identified.

03 / Ground and check your AI front line

QA used to mean hand-reading two percent of tickets and hoping it generalised. When a machine answers, a sample that size is negligence.

02 / Act continuously, not quarterly

PMs at one creator platform get alerted to a spike in their own area within hours. A streaming service watched surcharge-driven cancellation threats rise before they hit cancellation numbers.

04 / Close the loop back to the customer

When a fix ships, Enterpret detects the theme dropping — resolved from the customer's side, not merely marked "done" — and names the accounts to follow up with.

THE TEST

Point to a change your product made last month because of customer feedback — then show the customers who reported it being told it was fixed. If you can do the first but not the second, your loop is open.

Improve

A loop with a memory: each cycle makes the next one easier and more accurate, instead of starting from zero.

“If you're not operating from the same customer understanding, within a month product, marketing and engineering are talking about completely separate things.”



Tanmay Jain

Product & Growth Lead, ElevenLabs

ElevenLabs

Most feedback work discards its own learning. A PM re-explains their product area every week; an automation re-runs the same investigation with no memory that it ran yesterday.

“Having to give it all of our ownerships and topics every single time to explain who I am as a team is a little frustrating.”

PM • EUROPEAN FINTECH

“I don't think the automation has the ability to self update itself to learn what it already checked.”

POWER USER • CONSUMER HARDWARE

A loop that does not remember does not improve. It only repeats.

01 / Keep the artifact

A new theme stays in the taxonomy. A resolution is recorded against the accounts involved.

02 / Feed it forward

Ownership, prior findings and known failure modes carry from one run to the next.

03 / Let the front line learn

Every wrong AI answer caught becomes context that prevents the next one.

Two companies that look identical

Both have dashboards, a VoC team, a weekly report, AI in support, and an executive who sincerely calls the company customer-obsessed.

FEEDBACK AS EXHAUST

Collected, sometimes analysed, rarely acted on, never closed back. Its AI answers from that same untrustworthy context with no one checking, so customers go quiet and the AI misfires at scale.

FEEDBACK AS A LOOP

One maintained taxonomy, prioritised by evidence and revenue, routed to the work, caught within hours, reported back to the customers who raised it. Its context gets richer every cycle.

In any given week the gap looks small. Then a competitor running the loop is ahead on an axis you never measured — by a distance you cannot close in a quarter.

Where loops break

None of these announce themselves. Each shows up first as something small enough to explain away.

01

The taxonomy drifts

Someone renames a set of keywords and every report built on the old mapping quietly keeps running.

YOU'D NOTICE IT AS

Two teams quote different numbers for the same theme, and both can defend theirs.

02

Mentions stand in for customers

Volume is counted as demand, so ten power users repeating a complaint outrank a thousand quiet ones.

YOU'D NOTICE IT AS

A fix ships against the top-ranked request and the underlying volume barely moves.

03

Insight stops at the report

The analysis is produced and circulated, but never lands in a ticket, a queue or a system.

YOU'D NOTICE IT AS

A weekly readout everyone reads, and no piece of work that cites it.

04

Nobody tells the customer

The fix ships and the accounts that reported it are never told, so they learn that speaking up changes nothing.

YOU'D NOTICE IT AS

You can name last month's feedback-driven fix, but not one customer who heard about it.

05

Every cycle starts from zero

Ownership and prior findings are re-established by hand each time.

YOU'D NOTICE IT AS

Someone — or something — re-explains your product area from scratch every week.

Who owns the loop

Feedback that feeds both the roadmap and the AI front line is too consequential to sit in the gaps between teams. No new title required — the responsibilities follow the four stages.

NAME THE OWNER FOR EACH

STAGE	RESPONSIBILITY	OWNED	SHARED	NO ONE
UNIFY 01-04	Own the pipes and the taxonomy. One normalized stream, tied to accounts and revenue, under a hierarchy that matches the product.	☐	☐	☐
UNDERSTAND 01, 03	Size demand, cost and risk. Defensible numbers on what customers ask for and where revenue is exposed.	☐	☐	☐
UNDERSTAND 04	Separate representative from vocal. Keep the loud minority from shipping as the actual market.	☐	☐	☐
ACT 01-02	Route insight to where the work happens. Tickets filed with the complaints attached; spikes caught in hours.	☐	☐	☐
ACT 03	Ground and check the AI front line. Audit interactions, human and AI, at full coverage.	☐	☐	☐
ACT 04	Close the loop back to the customer. Confirm the fix from the customer's side and tell the accounts that reported it.	☐	☐	☐
IMPROVE 01-03	Carry each cycle forward. Ownership and prior findings feed the next run; the front line learns from every miss.	☐	☐	☐

If more than two rows land in **Shared** or **No one**, the loop doesn't have an owner — it has volunteers, and it will stall at the stage nobody is accountable for.

Walk one problem through all four stages

Pick one real problem your customers reported in the last ninety days.
Wherever the walk stalls is your constraint.

01 UNIFY

Can you see every customer who raised this issue, across every channel, tied to their account and revenue, in one place?

02 UNDERSTAND

Can you state how many customers and how much revenue it affects — and whether it's representative or merely loud?

03 ACT

Did the insight reach the person or system that could resolve it? Was your AI front line grounded on it? Were the customers who reported it told when it was fixed?

04 IMPROVE

Did the cycle leave your taxonomy and context sharper for the next problem, or did it start from zero?

For most teams the walk stalls at Act — the highest-return stage, and the one most competitors are also skipping.

Thirty, sixty, ninety

Every stage depends on the one before it, so the order is not negotiable. What is negotiable is scope: one product area, done properly, beats the whole company done thinly.

FIRST 30 DAYS

Unify one product area

Connect the two or three sources that carry most of your feedback, deduplicated into one stream.

Stand up the taxonomy on your own language, not a generic category list.

Tie records to accounts, ARR and an owner.

DAYS 30-60

Size it, then route it

Weight every theme by distinct customers and revenue, never by mentions.

Put monitors on the areas that matter so spikes surface in hours.

Wire the top themes to the tickets, owners and systems that resolve them.

DAYS 60-90

Close it and compound it

Audit the front line, human and AI, at full coverage rather than a sample.

Confirm fixes from the customer's side, then tell the accounts that reported them.

Carry ownership and prior findings from each cycle into the next.

EXIT TEST

Take one complaint from this morning and see everywhere else that customer has raised it, in under a minute.

EXIT TEST

Bring a number you would defend in a prioritisation meeting, and have it change the order of the work.

EXIT TEST

Name a fix, the theme that dropped after it shipped, and the customers who were told.

Ninety days is enough to run the loop once, end to end, on one problem. That is the only proof that matters.



This is the loop we build.

Interpret unifies every channel of customer feedback into one adaptive taxonomy built on your own language, tied to your accounts and revenue — then puts that context where your team already works, and where your agents can read it.

TRUSTED BY TEAMS AT

ElevenLabs

 **descript**

 **Fanatics**

 **Wispr**

Canva

 **Notion**

STRAVA

 **Apollo**

UNIFY

Every source, deduplicated and tied to accounts, in weeks not quarters.

UNDERSTAND

Your categories, quantified by customers and revenue — not a model guessing.

ACT & IMPROVE

Agent loops that route, monitor, audit and close the loop back to the customer.

See it on your own feedback.

We connect a couple of your sources and stand up your taxonomy on your real data. By the end you'll know whether this is real — on your feedback, not a demo.

BOOK A DEMO

INTERPRET.COM

