

AI-GENERATED IMAGE

What Customers Won't Tell Their Supplier

Independent customer interviews for investors, acquirers and technology leadership teams.

By Les Elby, Co-Managing Partner

November 2026 · About 15 minutes to read

10 to 20

interviews per customer group is usually enough

1 to 3 years

how far ahead customer interviews can show churn risk

About 8 points

fall in software net revenue retention since 2021, according to Bain

THE SHORT VERSION

Executive summary

This eBook explains how independent Voice of the Customer (VoC) interviews work in due diligence and inside technology businesses. It draws on published research and on Lighthouse's own VoC projects.

01 Customers hold back criticism from the people responsible for it.

They tell an independent interviewer, even when they know what they say could be traced back to them.

02 Unhappy customers go quiet before they leave.

By the time a licence is not renewed, the customer stopped giving feedback long ago. An independent conversation can surface that risk one to three years before renewal.

03 The standard software metrics are a weaker guide than they were.

Bain reports net revenue retention down about 8 points since 2021 and argues that ARR and NRR now say less about value. Whether AI is designing a product out of its customers' work is something only those customers can tell you.

04 Good VoC comes down to design.

That means 10 to 20 interviews per customer group, a sample the sponsor controls, unhappy and lost customers included, and questions built around the decision in hand.

05 AI already helps us analyse interviews, but it does not conduct them.

Most targets would refuse, and customers notice when a vendor saves its own time and not theirs.

06 For management teams, a CSAT or NPS score gives you a number without the reasons.

It is good business hygiene to know what your customers think before a buyer asks them.

INSIDE

Contents

THE CORE ARGUMENT

1 The person responsible is the last to hear 5
2 What customers tell an independent interviewer 7
3 What good VoC looks like 8

PART A: FOR INVESTORS AND ACQUIRERS

4 Why the data room is not enough now 12
5 Running VoC in due diligence 14
6 Should an AI do the interviews? 16

PART B: FOR MANAGEMENT TEAMS

7 Know before a buyer asks: VoC beyond CSAT and NPS 18

CLOSING

8 When VoC is the wrong tool 21
How Lighthouse helps 22
About Lighthouse and the author 22
Sources 23

Sections 1 to 3 and section 8 are for all readers. Part A is written for investors and acquirers, and Part B for management teams.



AI-GENERATED IMAGE

FOR ALL READERS · SECTIONS 1 TO 3

The core argument

Why customers say more to an independent interviewer, what they tell us, and how to design the work.

IN THIS PART

1 The person responsible is the last to hear 5

2 What customers tell an independent interviewer 7

3 What good VoC looks like 8

1 The person responsible is the last to hear

Ask a customer how things are going and the answer depends on who is asking. If it is the account manager who sold them the product, or the founder whose product it is, they are talking to someone responsible for whatever is wrong. Most people soften their criticism in that conversation. Some do not raise it at all.

An independent interviewer changes the conversation. Across our Voice of the Customer (VoC) projects for investors, acquirers and technology businesses, we have found that customers talk most freely about what they are unhappy with to someone who played no part in causing it.

HEARD IN OUR INTERVIEWS

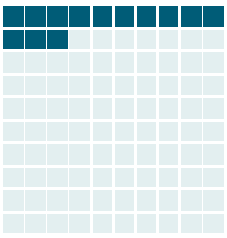
Problems they have stopped raising with their supplier, because raising them did not work.

The published research points the same way, although none of it comes from B2B technology. In a randomised UK trial, patients returned the same questionnaire either to their own GP practice or to an independent researcher. Satisfaction scores were the same in both groups. But 12.7% told the researcher they had altered their medication doses, against 3.8% who told the practice (Desborough et al., 2008). The authors treat the result as exploratory. A Canadian trial found something similar: clients rated their overall satisfaction equally whether a staff member or a fellow client interviewed them, but gave significantly more strongly negative answers to the fellow client (Psychiatric Services, 1999). A wider review of survey research concludes that how honestly people answer depends on the interviewer and the setting as well as on the person answering (Krumpal, 2013).

PUBLISHED EVIDENCE

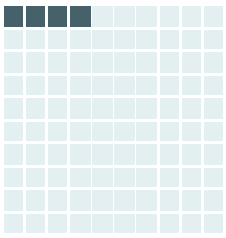
Same questionnaire, different recipient

Patients who said they had changed their medication doses, by who received their answers. Each square is one patient in a hundred.



12.7%

when replying to an independent researcher



3.8%

when replying to their own GP practice

— Overall satisfaction scores did not differ between the two groups.

Source: Desborough et al. (2008), the REPLY randomised trial in UK general practice. Squares rounded to the nearest whole patient.

The pattern is consistent. Headline scores barely move when someone independent asks the questions. What people admit to doing, and how strongly they criticise, change markedly. In a deal, those are the answers you need.

It would be easy to put this down to anonymity. In our experience, anonymity has little to do with it. A B2B customer list is short. Customers know their comments can often be traced back to them, and we tell them at the start that what they say about renewal may be passed on. They talk anyway, because the conversation is with someone independent who will make sure it is heard.

FROM OUR WORK

Customers talk to us knowing they could be identified. They are candid because we are independent and they expect someone to listen.

An independent conversation also reaches the customer who has gone quiet. Unhappiness does not get louder over time. As a customer's frustration grows, they detach. They stop raising issues, because raising them has not worked. To the vendor, a customer who has checked out can look much like a satisfied one, right up to the point where the licence is not renewed.

HOW CUSTOMERS GO QUIET

The silence before a customer leaves

As frustration grows, the feedback a vendor hears falls away. A diagram drawn from our project experience, not from a dataset.

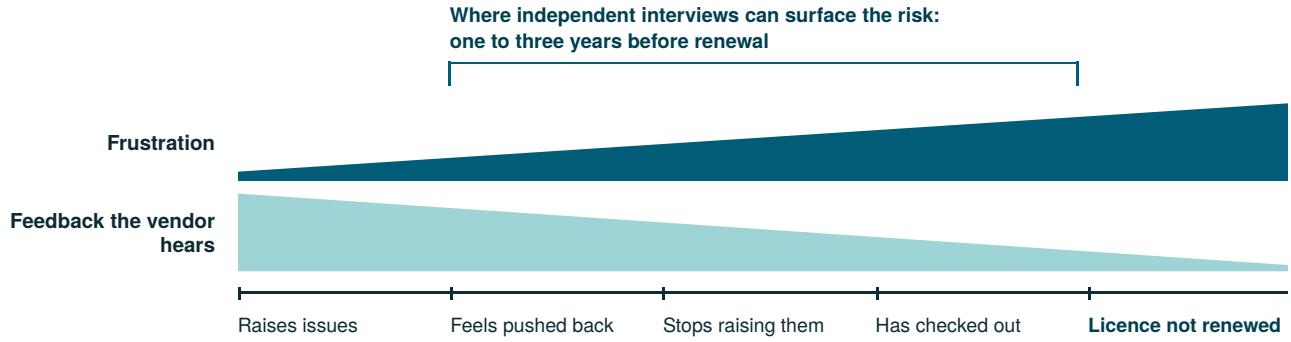


Diagram: Lighthouse. Timing depends on contract length.

FROM OUR WORK

Unhappy customers go quiet before they leave. Silence from a key account is a reason to ask questions.

2 What customers tell an independent interviewer



Customers talk more freely to someone who played no part in the problem.

AI-GENERATED IMAGE

What you learn depends on the customer's state of mind.

Content customers share good news freely, and their requests for improvements have usually reached the vendor already. What they tell us that the vendor rarely hears is who they credit for their experience. They name people: an account manager, a support lead, a product owner. When one name comes up again and again, that person is part of the reason customers stay.

Unhappy customers are where most of the value lies. How much they share depends on who they are connected to at the vendor, and on whether they feel their feedback has been listened to. If it has been pushed back, they are already detaching. The one thing they will not tell the vendor is whether they intend to renew. We ask, and they tell us.

HEARD IN OUR INTERVIEWS

Whether they intend to renew, and the specific reason why.

They also tell us why. The root cause is usually specific, and once you pass about ten interviews it shows up as a pattern. These are the four we see most often:

Overselling

The product was sold on capabilities it did not have, often by an inexperienced sales team.

Unresponsive support

Issues are logged and then left waiting.

Neglect

Nobody at the vendor is talking to the customer.

Founder push-back

In founder-led businesses, customers' issues are often argued with, or logged internally as user error. Customers end up feeling they get what they are given.

The positive signal is as clear. What builds loyalty is customers seeing their own feedback turn up in the product or service. It is the most convincing proof a vendor can give that it is listening.

3 What good VoC looks like

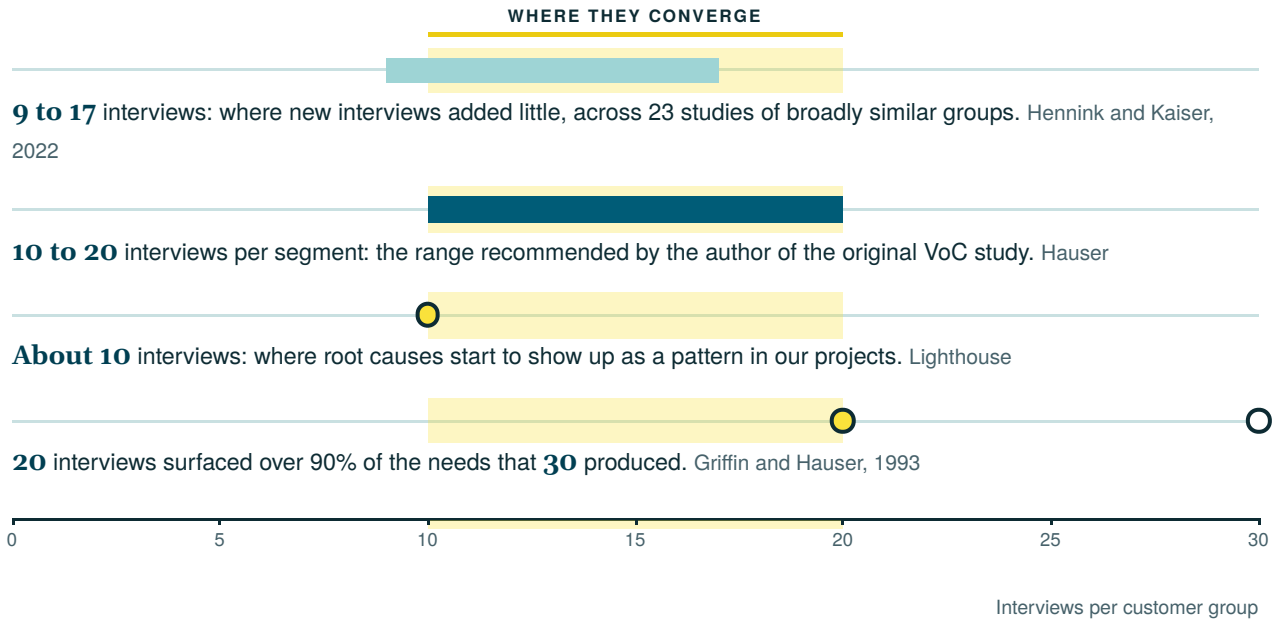
Whether VoC produces evidence or anecdote comes down to a handful of design choices.

How many interviews. Ten to twenty per customer group is usually enough. The published research supports this. In the original VoC study, 20 customers surfaced over 90% of the needs that 30 customers produced, and the study's author now recommends 10 to 20 interviews per segment (Hauser, summarising Griffin and Hauser, 1993). A 2022 review of 23 studies found that interviews reached saturation, the point where new ones add little, at 9 to 17 when the people interviewed were broadly similar (Hennink and Kaiser, 2022). The important word is group. If customers vary widely by size or sector, you need a reasonable sample in each.

PUBLISHED EVIDENCE AND OUR EXPERIENCE

How many interviews is enough?

Research and practice converge on a range of roughly 10 to 20 interviews for each group of broadly similar customers.



Sources: Hennink and Kaiser (2022); Hauser, summarising Griffin and Hauser (1993); Lighthouse project experience.

Who is in the sample. Left to choose, every company puts forward its happiest customers. Some cannot even do that, because they do not know who is unhappy, and that is a finding in itself. The sponsor of the work should set the criteria: the largest customers by share of annual recurring revenue (ARR), a cross-section by size and sector, and customers who have recently left. Lost customers are harder to reach, but they are often the most useful.

The questions. Build them around the decision in hand. Pricing is a good example. We rarely ask a customer what they would pay. We ask around it: the value they get for what they pay, how embedded the product is in their work, how it compares with the alternatives. Customers are very forthcoming about whether something is good and cheap, or expensive for what it does.

HEARD IN OUR INTERVIEWS

Whether the product is good and cheap, or expensive for what it does.

Consent and reporting. We tell customers at the start what the exercise is for and that what they say about renewal may be passed on. We report themes, supported by anonymised quotes. Where a customer would rather not be identified, we report the share of ARR at risk instead.

Analysis. One analyst reading a set of interviews will miss things that another catches (Hauser). We use AI to test the themes we have identified against every transcript.



Ten to twenty interviews per customer group, sorted into themes and tested against every transcript.

AI-GENERATED IMAGE

CHECKLIST

Questions to ask any VoC provider

- ☐ Who chooses the customers, and against what criteria?
- ☐ Will lost and shrinking customers be included?
- ☐ How many interviews will there be in each customer group?
- ☐ Who conducts the interviews, and how senior are they?
- ☐ What are customers told at the start, and how are renewal intentions reported?
- ☐ How are themes tested across the full set of interviews?



AI-GENERATED IMAGE

PART A · SECTIONS 4 TO 6

For investors and acquirers

For PE, VC and strategic buyers assessing a technology business.

IN THIS PART

4	Why the data room is not enough now	12
5	Running VoC in due diligence	14
6	Should an AI do the interviews?	16

4 Why the data room is not enough now

For a decade, software diligence has leaned on a familiar set of numbers: ARR, net revenue retention (NRR) and revenue multiples. They still matter. Among public SaaS companies that report NRR, those above 120% traded at a median 9.3 times trailing revenue in mid-2024, against 3.1 times for those below 100% (Software Equity Group, 2024).

PUBLISHED EVIDENCE

Retention is already priced in

Median multiple of trailing revenue for public SaaS companies that report NRR, mid-2024.



Source: Software Equity Group (22 October 2024).

But those numbers describe what customers did, and in 2026 that has become a weaker guide to what they will do. Bain reports that software revenue growth has fallen from about 20% a year to roughly half that, and that NRR is down about 8 points since 2021. It argues that ARR, NRR and revenue multiples have become less reliable signals in diligence, and that diligence now has to start by asking how far AI changes the work a product supports, and whether AI could replace it altogether (Bain & Company, April 2026). Uncertainty about AI helped drive a 70% fall in technology deal value between the last quarter of 2025 and the first quarter of 2026 (Bain & Company, June 2026).

PUBLISHED EVIDENCE

2026 in three numbers

Why the familiar software metrics say less than they used to.

Software revenue growth



From about 20% a year to roughly half that.

Net revenue retention

↓ 8 points

Down about 8 points since 2021.

Technology deal value

↓ 70%

Fall between the last quarter of 2025 and the first quarter of 2026, driven in part by uncertainty about AI.

Sources: Bain & Company (30 April 2026); Bain & Company (8 June 2026).

A data room cannot tell you whether a product is being designed out of its customers' work. The customers can. They know what they are trialling, what their own teams are

building and what a new competitor has offered them.

HEARD IN OUR INTERVIEWS

What they are trialling, what their own teams are building, and what a competitor has offered them.

Customer evidence also arrives earlier than the numbers. Depending on contract length, we can usually see churn risk one to three years before it reaches revenue. Sometimes the trigger is procedural. A customer that has to run a formal procurement at renewal will explain how its needs have changed, which tells you the incumbent will face new selection criteria.

FROM OUR WORK

Churn risk can show up in customer interviews one to three years before it shows up in revenue.

The cost of missing it is well documented, if not recently. McKinsey's study of 160 mergers found that almost 70% fell short of their revenue synergy estimates, and that the average merging company lost 2% to 5% of its combined customers. Among the causes of poor estimates, it named acquirers' limited access to the target's customers (McKinsey & Company, 2004).

PUBLISHED EVIDENCE

Almost seven in ten mergers missed their revenue synergy estimates

McKinsey's study of 160 mergers, published in 2004.



Almost 70%

fell short of their revenue synergy estimates. Each square is one in ten mergers.

2% to 5%

of combined customers lost by the average merging company.

Source: McKinsey & Company (May 2004).

EXAMPLE FROM OUR WORK

The moat with no goodwill

In one VoC project, customers told us they had bought a product that did not do what had been promised. They had then spent a great deal of their own time helping the vendor make it fit for purpose. That effort created a small moat: having invested so much, they were slow to leave. But the goodwill had gone. Satisfaction was low and churn risk was high. The findings matched what the technical due diligence was saying about the product's maturity, scalability and stability, and the deal did not proceed.

A moat built from a customer's sunk effort, with no goodwill behind it, is the easiest kind for a new competitor to cross. Bain warns that AI-native companies are already crossing software moats in areas such as customer service (Bain & Company, April 2026). Retention can look stable right up until a credible alternative appears.

FROM OUR WORK

A customer held only by switching costs will leave once a competitor makes switching easy.

5 Running VoC in due diligence

One approach now offered to investors is "blind" customer research. Customers are recruited without the target's involvement and are not told who commissioned the work. The aim is to remove management's influence. We take the opposite approach, and in technology deals we think it produces better evidence.

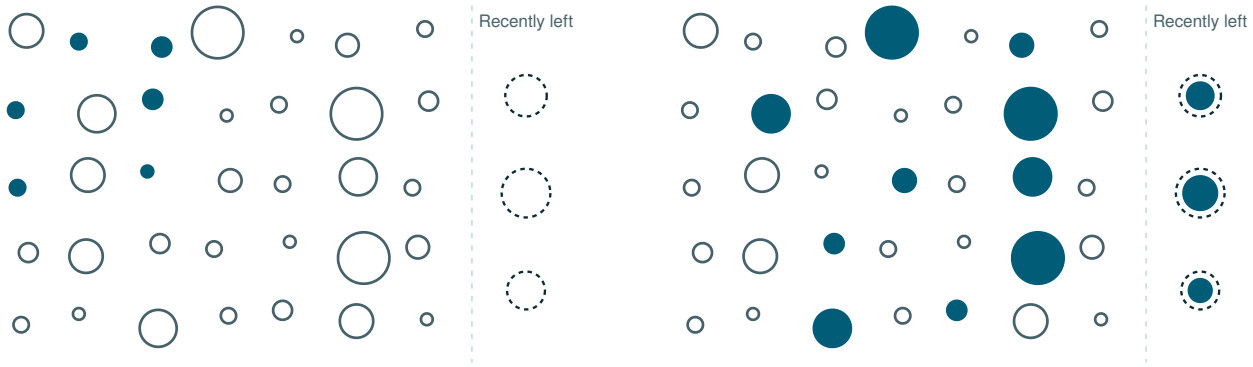
We run every project as if the target's management had commissioned it, even when an investor or acquirer is paying. We agree questions with the target that are useful to the business, alongside the ones the buyer needs answered, and the vendor is kept informed throughout. Most targets are protective of their customers, of the message and of how the exercise is framed. Working with them gets us in front of senior customers in the first place. Customers also respond well to an exercise they can see their supplier intends to act on.

The obvious objection is selection bias, and it is a fair one. We answer it by controlling the sample. The sponsor sets the criteria, we review the proposed list against them, and the largest customers by share of ARR are included whatever their mood. Where the sponsor wants to hear from customers who have left, we find them.

HOW A SAMPLE IS DRAWN

Who gets interviewed decides what you hear

The same customer base, sampled two ways. Each circle is a customer, sized by annual recurring revenue. An illustration, not client data.



Left to the company

A handful of its happiest customers, and none who have left.

Set by the sponsor

The largest customers by ARR, a cross-section by size and sector, and customers who have recently left.

● Interviewed ○ Not interviewed ◌ Recently left

Illustration: Lighthouse.

Expect two things in the findings that rarely appear elsewhere in diligence.

- **Key-person risk.** When content customers keep naming the same individual, that person is part of what you are buying. They need protecting through the transaction.
- **Corroboration.** Customer evidence is most persuasive when it lines up with the other workstreams, as it did with the technical due diligence in the example above.

> HEARD IN OUR INTERVIEWS
>
> *The same name, again and again: the person customers credit for their experience.*

There is one limit to what we can tell you. As the VoC provider, we do not always see what the investor or acquirer decided to do with the findings. We can tell you what customers said and how many of them said it. We will not claim credit for outcomes we did not see.

6 Should an AI do the interviews?



AI produces our transcripts and tests our themes. It does not ask the questions.

AI-GENERATED IMAGE

AI interview platforms now offer customer interviews for diligence at a fraction of the cost of a consultant-led project, with results in a day. For a deal team on a tight timetable that is an attractive offer, and the case for it has some substance. In one study, people disclosed more to a virtual interviewer they believed was fully automated than to one they were told a person was operating (Lucas et al., 2014). The topics were personal health, not software. In a small 2025 pilot, an AI interviewer produced answers of similar quality to human interviewers, but it was much weaker at following up surprising or unclear answers, and respondents found it less interesting (Wuttke et al., 2025). That pilot involved eleven interviews.

AI is already part of how we run VoC. It produces our transcripts, and we use it to test the themes we have found against every interview. We do not use it to conduct the interviews, for two reasons.

The first is permission. Targets are protective of their customers. Had we proposed AI voice interviews on the projects we have run, almost every target would have refused.

The second is the signal it sends. Customers would have to be told that a machine was calling. An exercise that asks for a senior customer's time, but not the vendor's, says something about how much the vendor values the relationship. It takes the shine off before the first question.

The technology will improve, and this is worth revisiting. For now, the limit is as much about permission as about quality.



AI-GENERATED IMAGE

PART B · SECTION 7

For management teams

For the leadership teams of technology businesses.

IN THIS PART

7 Know before a buyer asks: VoC beyond CSAT and NPS 18

7 Know before a buyer asks: VoC beyond CSAT and NPS

Most technology businesses already measure how their customers feel, usually through CSAT or NPS. Both have a place. Both share a weakness: they give you a score without the reasons behind it.

CSAT is typically one question, asked at one moment, about one event: how was your support call, how was your onboarding. It tells you how that interaction went. It tells you very little about whether the customer will renew.

ILLUSTRATION

A number without the reasons

What a satisfaction score can carry, against what an independent conversation can.

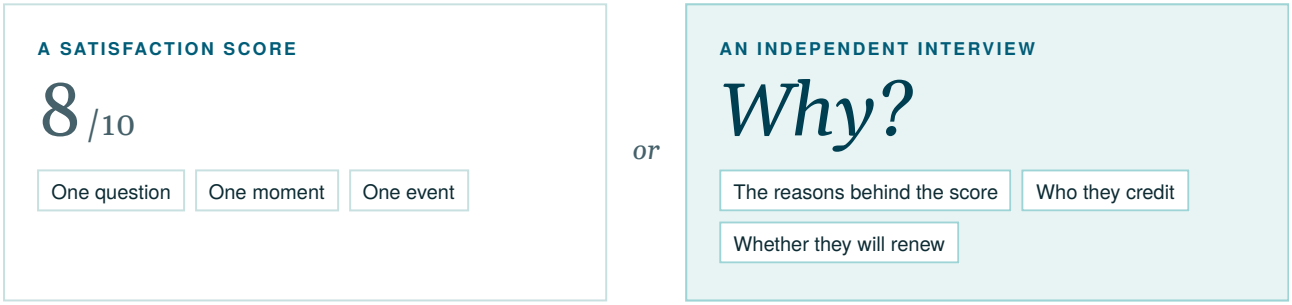


Illustration: Lighthouse.

NPS asks a broader question, but the evidence that it predicts growth is weaker than its popularity suggests. Researchers using data from 21 firms and more than 15,500 interviews could not reproduce the claim that NPS is clearly better than other loyalty measures (Keiningham et al., 2007). A later study of seven consumer brands found that only NPS measured across all potential customers predicted sales growth. The score from existing customers alone did not (Baehre et al., 2022).

PUBLISHED EVIDENCE

What the research says about NPS

21 firms, more than 15,500 interviews

Researchers could not reproduce the claim that NPS is clearly better than other loyalty measures.

7 consumer brands

Only NPS measured across all potential customers predicted sales growth. The score from existing customers alone did not.

Sources: Keiningham et al. (2007); Baehre et al. (2022).

Both tools also share the blind spot described in section 1. In our experience, customers who have checked out are the least likely to answer a survey, so the people you most need to hear from drop out of the results. When a key account goes quiet, ask why.

The alternative is to ask properly, every so often, with someone independent doing the asking. The businesses that get most from VoC outside a deal tend to commission it when something is changing.

CHECKLIST

When to commission VoC

- ☐ Sales or growth have changed direction and you are not sure why.
- ☐ Customers are leaving and the explanations do not add up.
- ☐ A new competitor is turning up in your deals.
- ☐ You are launching a new product or changing your pricing.
- ☐ You expect to raise money or sell the business in the next few years.

If you do sell, a serious buyer is likely to speak to your customers. It is better to know what they will say first, and to have acted on it. Our eBook on exit readiness covers the wider preparation.

And act on what you hear. In our experience, what turns a customer from tolerant to loyal is seeing their own feedback appear in the product.



HEARD IN OUR INTERVIEWS

Whether they have seen their own feedback turn up in the product.



AI-GENERATED IMAGE

FOR ALL READERS · SECTION 8

Closing

The limits of the method, and how Lighthouse can help.

IN THIS PART

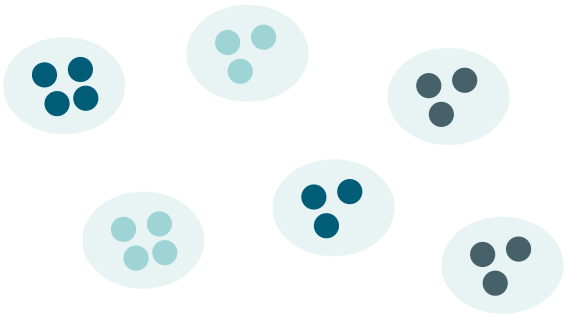
8	When VoC is the wrong tool	21
	How Lighthouse helps	22
	About Lighthouse and the author	22
	Sources	23

8 When VoC is the wrong tool

VoC works because a group of broadly similar customers, asked the same things, will show you the business as it really is. A cross-section of customers works like asking ten friends to describe you: individual opinions vary, but your character comes through. There are two situations where that breaks down.

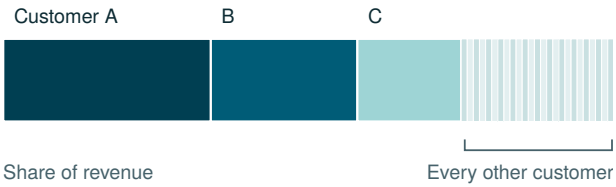
ILLUSTRATION

Two customer bases where twenty interviews mislead



No common customer group

Twenty interviews spread across many different groups say little about any one of them.



Revenue sits with a few customers

You learn how likely those few are to stay, not what the business does well or badly.

Illustration: Lighthouse. Not client data.

There is no common customer group. Ten to twenty interviews reveal patterns because the customers have something in common. For a business with a very broad offering, say one with a product range as wide as Microsoft's, twenty interviews across the whole customer base would tell you little. You would need to narrow the work to one product in one sector first.

Revenue sits with a few customers. When a business depends on a handful of accounts, VoC will tell you clearly how likely those customers are to stay. It will not give you a general picture of what the business does well or badly, because what you hear is specific to those few relationships.

If either applies to you, we will say so before you spend any money.

How Lighthouse helps

We run Voice of the Customer projects for private equity and venture capital investors, for acquirers, and for the management teams of technology businesses. The work is led by a partner from the first conversation to the final report. Our partners have run and scaled technology businesses themselves, so when a customer says something unexpected, we know which question to ask next.

If you are assessing a technology business, or want to know what your own customers would tell a buyer, we are happy to talk it through. Email info@lighthouse-advisory.co.uk or use our contact page.

About Lighthouse and the author

Lighthouse Advisory Partners

Lighthouse Advisory Partners is a growth strategy advisory firm based near Cambridge, working exclusively with technology companies. We advise on growth strategy; buy-side M&A from target identification to completion, including outsourced corporate development; selective sell-side M&A; commercial due diligence; Voice of the Customer; and AI readiness and strategy. Our partners have run and scaled technology businesses, and the team has worked on well over 100 transactions.

Les Elby

Les Elby is co-founder and Co-Managing Partner of Lighthouse. He has around two decades of experience in strategy and M&A, including as Chief Strategy Officer of a FTSE 100 company.

Sources

Published evidence is cited in brackets in the text. Everything else, including all examples, comes from Lighthouse's own VoC work.

1. Baehre, S., O'Dwyer, M., O'Malley, L. and Lee, N. (2022). The use of Net Promoter Score (NPS) to predict sales growth: insights from an empirical investigation. Journal of the Academy of Marketing Science, 50, 67 to 84.
wrap.warwick.ac.uk/id/eprint/154989

2. Bain & Company (Lipman, D., Perry, C. and Kieffer, G.) (30 April 2026). The New Era in Tech Investing Starts Now.
bain.com/insights/new-era-in-tech-investing-starts-now

3. Bain & Company (8 June 2026). Press release: 2026 Private Equity Midyear Report.
bain.com/about/media-center/press-releases/2026/winning-firms-will-focus-on-what-they-can-control-weather-the-rest-as-triple-shock-brakes-private-equitys-latest-revival-bain--company-2026-midyear-pe-report

4. Desborough, J. A., Butters, P., Bhattacharya, D., Holland, R. C. and Wright, D. J. (2008). Does it matter whether the recipient of patient questionnaires in general practice is the general practitioner or an independent researcher? The REPLY randomised trial. BMC Medical Research Methodology, 8, 42.
doi.org/10.1186/1471-2288-8-42

5. Effects of Client Interviewers on Client-Reported Satisfaction With Mental Health Services (1999). Psychiatric Services, 50(7), 961.
doi.org/10.1176/ps.50.7.961

6. Hauser, J. R. Note on the Voice of the Customer. MIT Sloan courseware, summarising Griffin, A. and Hauser, J. R. (1993), The Voice of the Customer, Marketing Science, 12(1), 1 to 27.
dspace.mit.edu/bitstream/handle/1721.1/71132/15-810-spring-2005/contents/readings/note_voice_cust.pdf

7. Hennink, M. and Kaiser, B. N. (2022). Sample sizes for saturation in qualitative research: a systematic review of empirical tests. Social Science & Medicine, 292, 114523.
pubmed.ncbi.nlm.nih.gov/34785096

8. Keiningham, T. L., Cooil, B., Andreassen, T. W. and Aksoy, L. (2007). A longitudinal examination of Net Promoter and firm revenue growth. Journal of Marketing, 71(3), 39 to 51.
research.hub.ku.edu.tr/entities/publication/6c8605c0-154e-4020-bcd9-2e7e3e1f1550/full

9. Krumpal, I. (2013). Determinants of social desirability bias in sensitive surveys: a literature review. Quality & Quantity, 47(4), 2025 to 2047.
doi.org/10.1007/s11135-011-9640-9

10. Lucas, G. M., Gratch, J., King, A. and Morency, L.-P. (2014). It's only a computer: virtual humans increase willingness to disclose. Computers in Human Behavior, 37, 94 to 100. Reported by USC Today, 9 July 2014.
today.usc.edu/patients-are-more-willing-to-confide-in-computers-not-doctors-usc-study-suggests

11. McKinsey & Company (Christofferson, S. A., McNish, R. S. and Sias, D. L.) (May 2004). Where mergers go wrong. McKinsey Quarterly.
mckinsey.com/capabilities/strategy-and-corporate-finance/our-insights/where-mergers-go-wrong

12. Software Equity Group (Weekes, B.) (22 October 2024). How Net Retention Impacts Valuation for Public Software Companies.
softwareequity.com/blog/net-retention-public-saas-companies

13. Wuttke, A., Aßenmacher, M., Klamm, C., Lang, M. M., Würschinger, Q. and Kreuter, F. (2025). AI Conversational Interviewing: Transforming Surveys with LLMs as Adaptive Interviewers. arXiv 2410.01824, version 2, March 2025.
arxiv.org/abs/2410.01824

The photographs in this eBook are AI-generated illustrations, not pictures of real places, clients or projects. Charts are drawn by Lighthouse from the sources above; diagrams marked as illustrations are not based on client data.

LIGHTHOUSE ADVISORY PARTNERS

If you are assessing a technology business, or want to know what your own customers would tell a buyer, we are happy to talk it through.

Lighthouse is a growth strategy advisory firm based near Cambridge, working exclusively with technology companies. The work is led by a partner from the first conversation to the final report.

EMAIL

info@lighthouse-advisory.co.uk

WEB

lighthouse-advisory.co.uk

© 2026 Lighthouse Advisory Partners. Published November 2026.
The photographs in this eBook are AI-generated illustrations, not pictures of real places, clients or projects.