




UPDATED APRIL 2026

Rethinking Voice of Customer: How Research Teams Can Evolve with the Modern Data Stack



Abstract

As companies modernize their data infrastructure and AI tooling, Voice of Customer (VoC) programs risk being left behind. Research professionals often rely on traditional tools and reports, while the rest of the business moves toward real-time data access and cross-functional insights. This article explores how research teams can adapt, bringing unstructured feedback into the modern data stack, expanding beyond surveys, and increasing the impact of their work across the company.

Contents

The VoC opportunity is bigger than we think	4
What is the Modern Data Stack, and why should VoC plug into it?	6
Why MDS is a better solution for VoC than an all-in-one suite	7
Common challenges and solutions when adding VoC data to MDS	8
How to get buy-in from data and tool owners	10
Looking ahead: the evolving role of research	12

The VoC opportunity is bigger than we think

Voice of Customer programs were designed to listen, analyze, and surface customer insights. But in most companies, they're still focused on survey data, which is a big problem.

According to a [McKinsey report](#), **only 7% of customers respond to surveys**, and response rates continue to drop. So the insights we're getting are from a shrinking group of customers, and often not the ones who need attention most.

To make things worse, survey data is often locked in Customer Experience Management (CXM) platforms. That makes it hard for other teams to access or use. Research professionals end up owning feedback in isolation — even though product, marketing, service, and operations could all benefit from hearing directly from customers.

Meanwhile, **98% of customer feedback lives outside traditional surveys** — in support tickets, call center transcripts, reviews, and forums. AI has made it possible to analyze this unstructured data at scale, but the most mature companies use this data as part of their Voice of Customer program. And according to [CX Today](#), 47% of VoC and Customer Experience (CX) program leaders rated their program maturity as “low or very low.”

So how do we fix this?

This isn't a single project. It's a cross-tool, cross-team shift that treats feedback as shared infrastructure.

What's holding back the impact of VoC teams

Most VoC programs still run in a fixed loop: collect feedback, code responses, generate reports, share with stakeholders. It works, but it's slow, manual, and rarely scales.

Unstructured feedback is difficult to analyze

According to a [2023 study by Forrester Research](#), nearly all programs (96%) regularly collect and analyze surveys, but only 50% are effective at collecting solicited unstructured feedback (let alone unsolicited feedback, such as complaints or call center interactions). From talking to practitioners we know that the most common tools for coding open-ended responses in surveys such as Qualtrics' TextIQ are too manual, which is why this data rarely ends up included in live dashboards.

Feedback isn't linked to financial data

According to an extensive [B2B benchmark study by CustomerGauge](#), 62% of businesses don't link their VoC data to their bottom line. [Forrester's CX Survey](#) talks about how this leads to lack of funding for VoC programs, with 38% of respondents stating the lack of a model tying CX to financial results as a significant obstacle.

Feedback isn't integrated with operational data

A [2024 Deloitte study](#) states that building a business case and demonstrating the impact is a key challenge for customer experience leaders, something that's impossible to achieve without linking to operational metrics. According to a [special report by Contact Center Pipeline](#), only 53% of surveyed companies knew the percent of issues identified by VoC that were resolved.

Without links to operational and financial data, insights lose their urgency. And the rest of the business moves on without them. Here's a striking statistic from the [2023 Forrester report](#): "Only 42% of survey respondents said that stakeholders ask for insights, and only 44% say that stakeholders are confident in the metrics you produce."

What is the Modern Data Stack, and why should VoC plug into it?

Over the last few years, companies have overhauled how they handle structured data. The Modern Data Stack (MDS) is a cloud-based system that brings together data collection, storage, transformation, and analytics. Common elements include:

- **Customer Data Platforms (CDPs) and Customer Relationship Management platforms (CRMs)** like Adobe CDP, Salesforce Data Cloud, Microsoft Dynamics, Hubspot.
- **Data warehouses** like Snowflake or BigQuery
- **Dashboards and Business Intelligence (BI) tools** like Tableau, Power BI, or Looker
- **Data pipelines and transformation tools** like dbt
- **Activation tools** that push data back into apps like Salesforce Service Cloud, ServiceNow, and Jira

These systems are fast, scalable, and increasingly used across all business functions. But unstructured data across all VoC channels is often left behind. And this is where the big opportunity lies.

Here are the top three reasons why VoC data should plug into the MDS:

1. **Measure what moves the business**
Once feedback is connected to financial and operational metrics, acting on it can be tied to Return on Investment (ROI). VoC teams can build business cases that prioritize improvements by impact, linking each to data such as churn, revenue, and support metrics.
2. **Make feedback accessible via existing dashboards**
Instead of building separate reports, integrate feedback themes into the same BI tools your stakeholders already use. That way, they don't need to request a report, they can explore trends directly.
3. **Link feedback to action**
Whether it's improving the onboarding experience, flagging product issues, or creating personalized marketing campaigns, once VoC data is part of a MDS, it can power workflows and processes that drive actual change.

To sum up, by integrating feedback into the company's core data systems, research professionals can scale their insights and make them available to more people, faster. This will ultimately expand researchers' impact on the organization.

Why MDS is a better solution for VoC than an all-in-one suite

The most popular VoC solutions right now are suites such as Medallia, Qualtrics, PG Forsta and InMoment. The surveying capabilities, the analysis and reporting is offered as part of a single suite, without the need of a database.

In theory, this approach should offer a simplified vendor management, seamless integration and cost efficiency. According to a 2025 [BCG survey](#) of 300 C-level leaders, in practice, often the opposite happens:

- 1. Not a one-and-done simplifier.** BCG notes that even companies that choose suites “often find themselves expanding beyond the initial integrated environment”.
- 2. Integration headaches don’t disappear** with suites. They are just as common among suite vs. best-of-breed buyers.


- 3. Cost-efficiency of a suite is a myth.** The most surprising finding by BCG was that 56% of those who bought a suite thought that this approach was less cost efficient after the fact. Once you are locked into a suite, annual price increases are common.

MDS is modular by design, so lets you **optimize cost** by choosing the best tool for each job and **swap components** as pricing or needs change, rather than absorbing a suite’s bundled shelfware. It **simplifies integration** by using the warehouse and a shared semantic layer as the hub, so apps plug into governed models. Finally, **vendor management** can be clearer, not messier: when the warehouse and identity layers are the backbone, each application is a replaceable spoke with a small contract and a defined data contract.


This approach also aligns with broader industry moves toward composable, modular architectures, and it directly addresses the suite caveats analysts highlight.

Decision Makers Believe That the Solution They Didn’t Buy Is More Cost Effective

Respondents who chose a suite but think best-of-breed solutions are more cost effective (%)



Respondents who opted for best-of-breed solutions but think a suite is more cost effective (%)



● Strongly agree ● Agree ● Neutral ● Disagree ● Strongly disagree

Source: BCG enterprise software vendor survey, 2024. Note: n = 165 respondents who bought suites; n = 158 respondents who bought best-of-breed solutions. Because of rounding, the percentages may not add up to 100. Survey question: Do you perceive that using multiple applications or solutions from multiple vendors is more cost effective than using a single vendor for multiple solutions (a suite)? Survey question: Do you perceive that using multiple applications or solutions from a primary suite vendor is more cost effective than using specialized vendors for each application or solution (best of breed)?

Common challenges and solutions when adding VoC data to MDS

Let's review common objections that you, your team or other stakeholders might have when adding customer feedback into your existing stack.

Data privacy and security

Why this feels hard

Customer comments and transcripts often contain personal data. When feedback moves from various customer channels into CDPs/CRMs/warehouses, questions arise about who "owns" raw text vs. derived themes, what purposes consent covers, and how to share responsibly across teams.

Strategies that work

Learn the **NIST Privacy Framework**, a useful baseline for managing privacy risk across diverse systems and stakeholders. While raw text can be restricted, you can treat derived signals such as themes, scores, sentiment, drivers as a governed data product. Always default to hashed IDs and redacted Personally Identifiable Information (PII) in shared datasets and dashboards. Stamp source, timestamp, and transformation lineage on every record so teams can audit how a VoC metric was produced before using it in decisions.

Technical challenges integrating data and tools

Why this feels hard

To make feedback useful, you need to line it up with who said it (person/account), what they did (product usage, support cases), and what it meant for the business (renewal, churn, spend). Those pieces live in different systems (CRM, contact-center, product analytics, data warehouse/CDP). The same customer may appear with slightly different names/emails/IDs across these tools. Joining this data through identity resolution and maintaining the joins are challenging.

Strategies that work

To limit points of failure, agree on a handful of standard keys (for example, account, user, segment, lifecycle) and make sure every dataset carries them. Prepare raw text once in a central place: redact PII, extract themes and sentiment. Then publish just those clean signals to where people work: tasks/queues in CRM or contact-center tools for action, and a certified tile in BI for tracking. Ask your data partner to "drift-proof" things so new fields don't break existing joins and to log schema changes you can review monthly. If identity is messy, start at the account level and expand later.

Unstructured data analysis at scale

Why this feels hard

Open-text, transcripts, and social posts are noisy. Teams need reliable topic discovery, sentiment, and driver analysis that update as products change. For a long time, VoC suites have offered approaches based on curating manual rules for analysis, which is costly and relies on professional service hours.

Strategies that work

AI has come a long way. Test out for best-of-breed solutions for unstructured data analysis on your data to find those that produce the most accurate results. Don't rely on vendor reported accuracy stats, because they vary by data. Look out for solutions with insights governance features, such as Human-in-the-Loop features to verify and modify analysis. Use the same text analytics solution regardless of channel for consistency of analysis.

Version drift and maintaining a single taxonomy

Why this feels hard

As products, services and policies evolve, so does the feedback. Without governance, feedback in different channels can be tagged differently (e.g. "Onboarding friction" in support vs. "Signup UX" in product). Schema/taxonomy drift quietly erodes trust and multiplies rework.

Strategies that work

Designate an owner for the taxonomy created for a particular use case. They need to do quarterly reviews, making sure the analysis is comprehensive and current. Some solutions will alert you of new themes worth adding as they emerge in the feedback, which speeds up the review process. In the past, you needed manual maintenance of synonyms and rules, now AI automated a lot of this.

But probably, the greatest challenge you'll face will be convincing others that VoC data should be included as part of the MDS. The next chapter digs into this specific challenge in depth.

How to get buy-in from data and tool owners

Making VoC part of the modern data stack is as much an organizational project as a technical one. Your goal is to get your stakeholders onboard by treating them as co-owners with shared outcomes. By understanding their needs you can explain how VoC data can help them achieve their goals.

Stakeholders' needs and how better access to VoC data can support them

Platform teams such as data engineering and data warehouse owners care about cost, clear SLAs, low pipeline churn. By integrating feedback sources directly into a data warehouse and from there into other tools, the company would reduce the need for ad-hoc data pulls. This team could collaborate on a stable, documented VoC data product.

Ops teams such as Marketing Ops, Sales and Revenue Ops in charge of CDP, CRM care about cleaner data, agent/rep productivity. Here, it's best to lead with a single, concrete business outcome (e.g., "reduce time-to-mitigation for top complaints in Segment A by 30%"). To achieve this, it's important to tie VoC data to **their** KPIs (average handling time, renewal rate, case deflection, cost-to-serve).

Contact Center Ops/IT in charge of Call Center as a Service (CCaaS) solutions, call/chat tools, care about efficiency, agent performance and reducing costs. The MDS approach would make the data they own more accessible across the company, which means issues and problems would be resolved faster. The call center is also often responsible for closing the loop on complaints and detractors. Instead of tackling each customer individually, this approach can enable closing the loop in batches, and possibly even in an automated way through AI agents.

Security & privacy teams such as information security, legal and compliance specialists care about auditability, data minimization.

BI / Analytics teams in charge of dashboards care about governed definitions, reusable models, fewer bespoke dashboards. Together, you could work on creating one certified model and delivery instead of many slide decks.

How to get started

To get your team onboard, work out a plan to get started. It's best to agree on a small, time-boxed pilot that runs over 6 to 8 weeks. Here are some suggested steps:

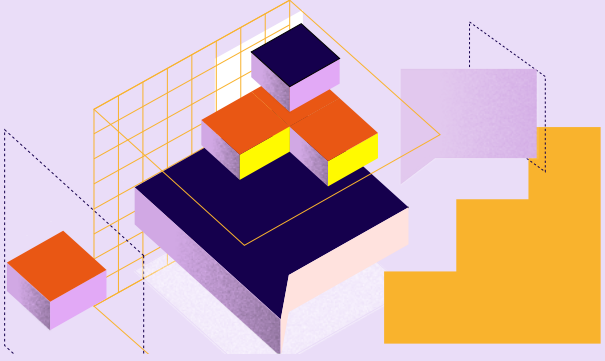
1. Choose one segment, one metric, or one “close-the-loop” workflow to target.
2. Define the VoC Data Product for this project: schema, field meanings, refresh cadence, quality checks, PII policy, and owners. Keep it to 1 to 2 pages.
3. Provide ready-to-use views per team (e.g., CRM widget, CCaaS queue, BI tile)
4. Organize a 30-minute enablement session to get everyone on board.
5. Publish a short before/after readout (what changed, who acted, impact).

Pragmatic risks to manage

Discuss the following key risks with the stakeholders:

- **Data ownership and consent:** clarify who owns raw comments vs. derived themes, and document consent for downstream use.
- **Access barriers:** role-based access and data minimization by default (e.g., masked PII) so non-research teams can self-serve safely.
- **Version drift:** keep a single governed taxonomy and model registry so “churn-driver” means the same thing in CRM, BI, and CCaaS.
- **Latency trade-offs:** some use cases need streaming insights; others work with daily syncs. Don't over-engineer real time where near-time is enough.

These considerations paint a picture of how the VoC discipline and with it the role of research and insights professionals needs evolve.



**Real example:
A telco transforms VoC for speed and scale**

A leading telco used to send out quarterly VoC reports. These took weeks to prepare and were often outdated by the time teams saw them.

By pulling survey comments into their data warehouse and building feedback dashboards in Power BI, the company gave frontline managers and executives real-time access to what customers were saying. They could filter by region, product, or call reason — and act faster.

Within months, they reduced reporting time by over 80%, spotted recurring issues sooner, and improved internal engagement with VoC insights.

Looking ahead: the evolving role of research

There's a misconception that automation makes research less relevant. In reality, it makes research more powerful. It's more important to do it right and make sure the teams that do research on their own aren't pulling the company into different directions.

Your role shifts from doing everything manually to choosing the right AI tools, testing them, implementing them to multiply AI's impact:

- Designing smart feedback taxonomies,
- Validating and refining AI-generated themes,
- Teaching teams how to interpret and act on insights,
- Ensuring feedback is used ethically and thoughtfully.

Near-term, expect AI to become even more capable and usable. Cleaning and unifying data will become easier, AI tools will auto-surface causal hypotheses from text and other data, automated closed-loop routing will assign owners and due dates straight from a comment. Emerging approaches include tools that enable AI search over governed feedback stores, lightweight agents that watch segments for emerging issues, and privacy-preserving learning on redacted text. This will move VoC from periodic reporting to continuous detection and activation. The research role becomes curating definitions, validating models, and setting ethical guardrails.

The main learning here is that this approach doesn't dilute your influence, it amplifies it.

Don't let VoC fall behind

VoC programs can't afford to be the last to modernize. The rest of the business is moving fast, using integrated, real-time data to make decisions. Feedback should be part of that ecosystem.

By bringing unstructured feedback into the modern data stack, using AI to speed up analysis, and making insights more accessible across the company, research professionals can lead the next evolution of VoC.

Finally, it's not about replacing what works; it's about scaling what matters. Today, that means a governed feedback product, clear joins to outcomes, and workflows that push insights to where action happens. Tomorrow, AI will make linkage, detection, and routing faster and more precise. The teams that succeed will do both: honor present constraints while building toward a future where customer voice is continuous, connected, and accountable.



Stop reporting on feedback. Start activating it.

Thematic brings unstructured feedback into your modern data stack — connecting customer voice to the metrics, dashboards, and workflows your teams already use.

See how it works for your organization.

[Talk to one of our experts](#)

getthematic.com