THE STATE OF B2B SALES CONTENT 2026

How B2B content and enablement executives across the US, UK, Canada, and Australia described the real cost of producing, personalizing, and sending customer-facing decks at scale and what they want to change about it.

EXECUTIVE SUMMARY

Customer-facing content has become one of the most operationally demanding parts of modern B2B growth. B2B teams repeatedly create the same types of deliverables: sales decks, business reviews, proposals, reports, and presentations.

On the surface, much of this work looks repeatable enough to automate. In reality, it rarely is. Most customer-facing content must be personalized before it is sent, passes through multiple contributors before reaching its final form, and often requires last-minute updates along the way. As content volumes grow, maintaining speed, consistency, accuracy, and visibility becomes increasingly difficult.

To understand how organizations are managing these demands, Visme surveyed 164 managers, directors, and senior executives across the United States, the United Kingdom, Canada, and Australia who are responsible for creating customer-facing content. Respondents were asked about content volume, workflow friction, the cost of rework, quality control, post-send visibility, and whether AI has meaningfully improved the process.

The findings reveal a common reality across organizations of all sizes: customer-facing content is becoming harder to scale. Teams spend significant time reworking existing materials, maintaining content quality, and keeping information current. At the same time, many have limited visibility into how content performs after it is shared and remain unconvinced by AI tools that lack the context needed to understand their brand, data, and workflows.

“Rework alone consumes up to half the working week for 85% of these teams. At any realistic team size, that represents a six-figure annual cost hidden inside the payroll budget—tracked as hours worked rather than value lost.”

FIVE FINDINGS THAT DEFINE THE PROBLEM

67% produce more than 20 customer-facing decks each month.

85% lose between 10% and 50% of their working week to rework.

53% see outdated, inaccurate, or off-brand content reach customers monthly or more often.

51% have only basic, informal, or no visibility into content performance after it is shared.

65% would adopt AI that seamlessly incorporates their brand, data, and past results.

TABLE OF CONTENTS

VOLUME

TWO IN THREE TEAMS PRODUCE MORE THAN 20 CUSTOMER-FACING DECKS A MONTH

VOLUME

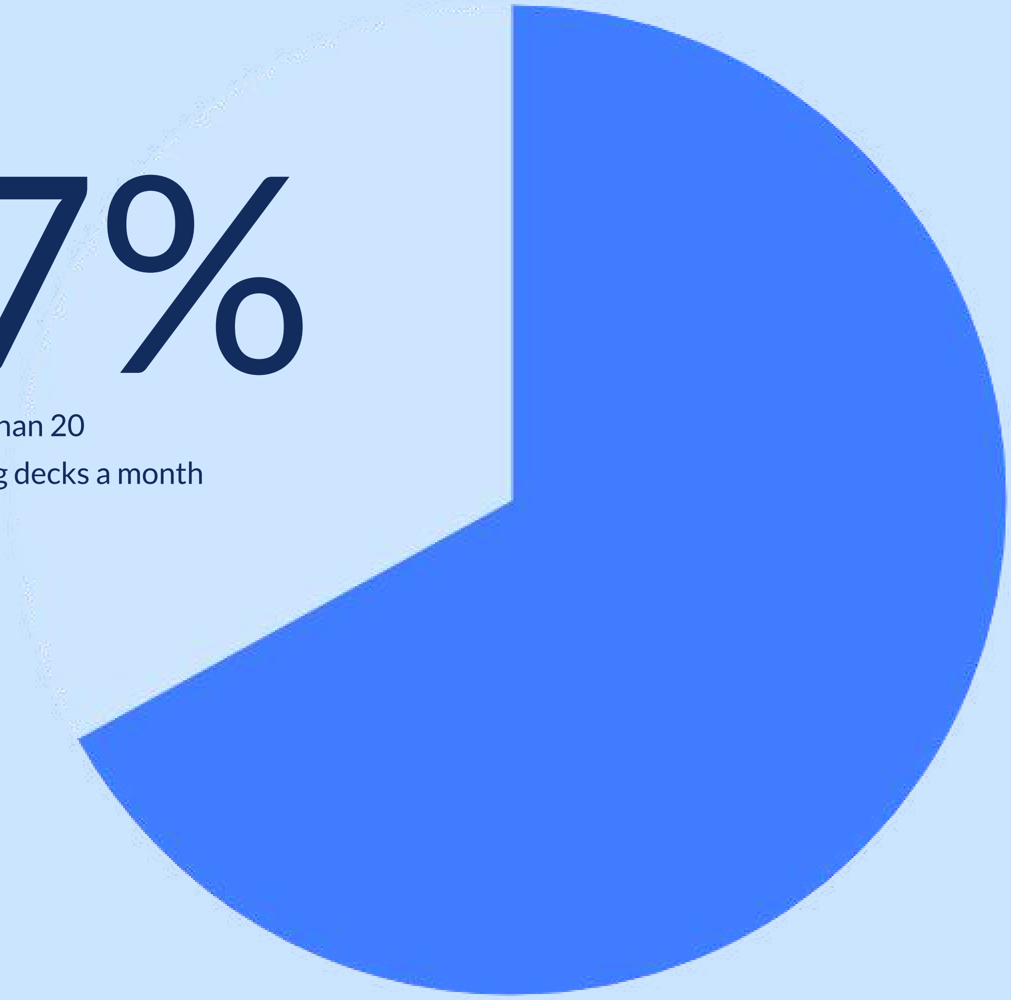

67%

produce more than 20
customer-facing decks a month

THE WORK IS HIGH-VOLUME AND REPETITIVE

Roughly 67% of teams produce more than 20 customer-facing decks a month, and about one in five produce more than 50. This number doesn't reflect a seasonal spike or a one-off project. It is a steady stream of deliverables that teams are expected to produce month after month, in addition to everything else their role requires.

CUSTOMER-FACING DECKS PRODUCED PER MONTH

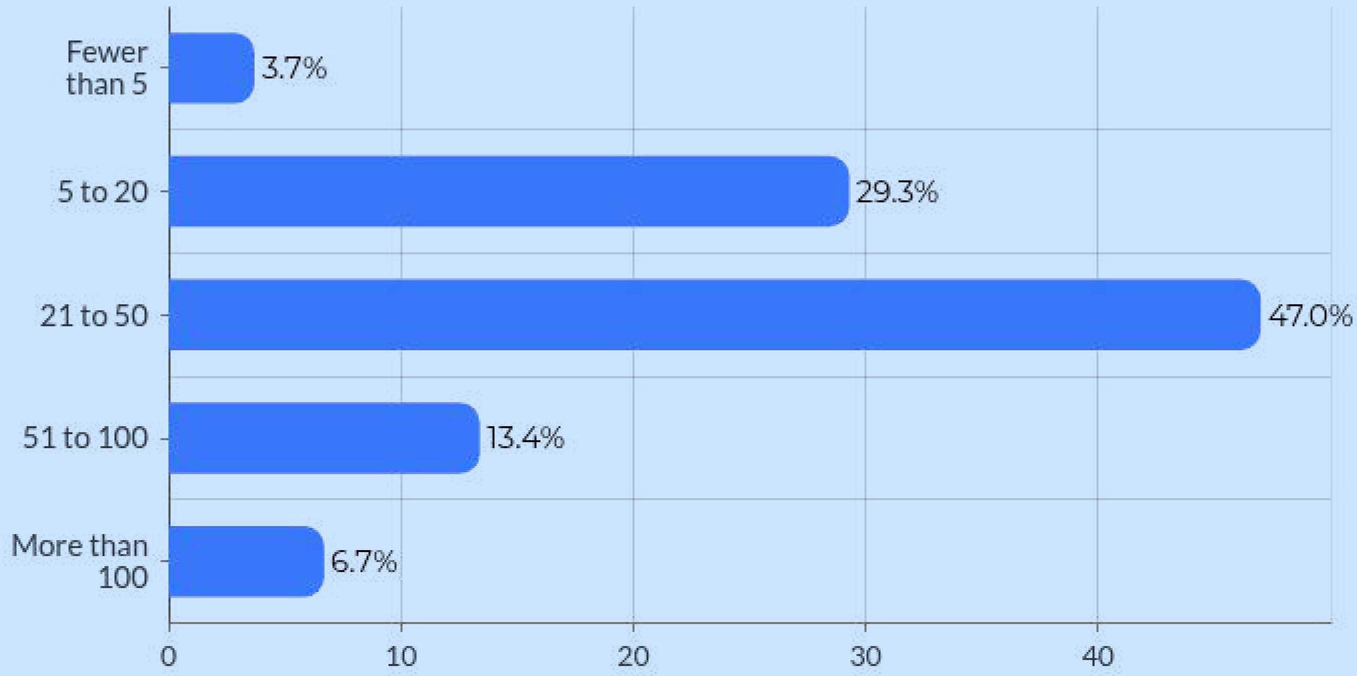

Decks produced in a typical month n=164.
Highlighted the 67% producing more than 20

The most common formats are also the ones tied to revenue: **sales pitch decks (59.1%), customer business reviews (57.3%), and quarterly business reviews (53.7%) are at the top.** These are not formats that can be produced carelessly or handed off to a junior resource for a first pass. They go to buyers in active conversations, to accounts mid-renewal, to executives making budget decisions. The volume is high, and the stakes attached to every individual piece are equally high.

CONTENT TYPES PRODUCED REGULARLY

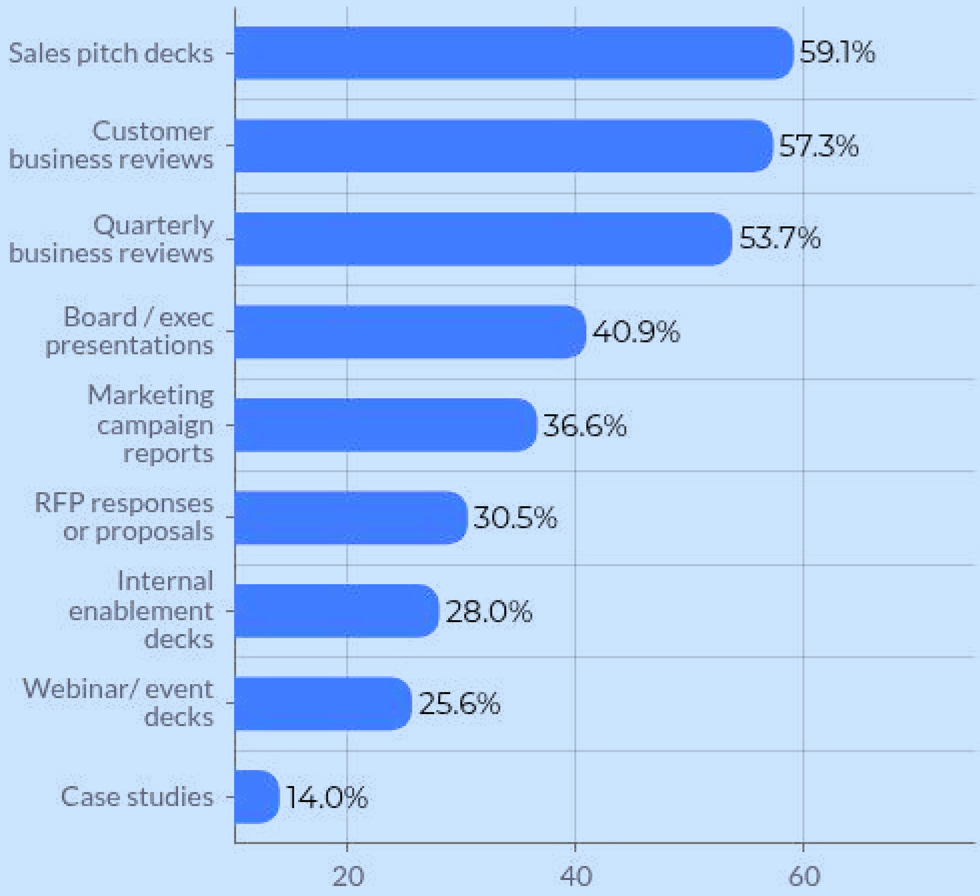

Respondents selected all that apply; totals exceed 100%. n=164.

What makes the volume problem especially demanding is the dual expectation built into it. On one hand, 92.7% of teams say at least half of what they produce follows the same format each time. But repetition doesn't mean copy-paste: 81.7% personalize their deliverables frequently or always.

The same deck template might be used over and over, but each version still needs to reflect a different customer, opportunity, or business outcome. That combination of repetition and personalization is what makes customer-facing content so demanding. The work looks standardized enough to scale, but customized enough that teams still have to rebuild it every time.

92.7%

say at least half their output repeats the same shape

81.7%

customize their deliverables frequently or always

1 in 5

produce more than 50 customer-facing decks in a single month

PRESSURE

NO SINGLE PAIN POINT DOMINATES; TEAMS FEEL IT ACROSS THE ENTIRE WORKFLOW

PRESSURE

57%

call creating content fast enough a severe pain, and the other eight pains aren't far behind

When teams rated nine common workflow pains, the result was striking not for what rose to the top, but for how little separated any of them. Every pain fell in a narrow band, with 42% to 57% of teams calling each one severe.

The pressure is spread across the entire process, from the moment content is requested to the moment it goes out the door.

> **"Everything gets messy and last-minute, so the content feels rushed and inconsistent."**
> *Director, Australia*

SHARE OF TEAMS WHO CALL EACH PAIN SEVERE

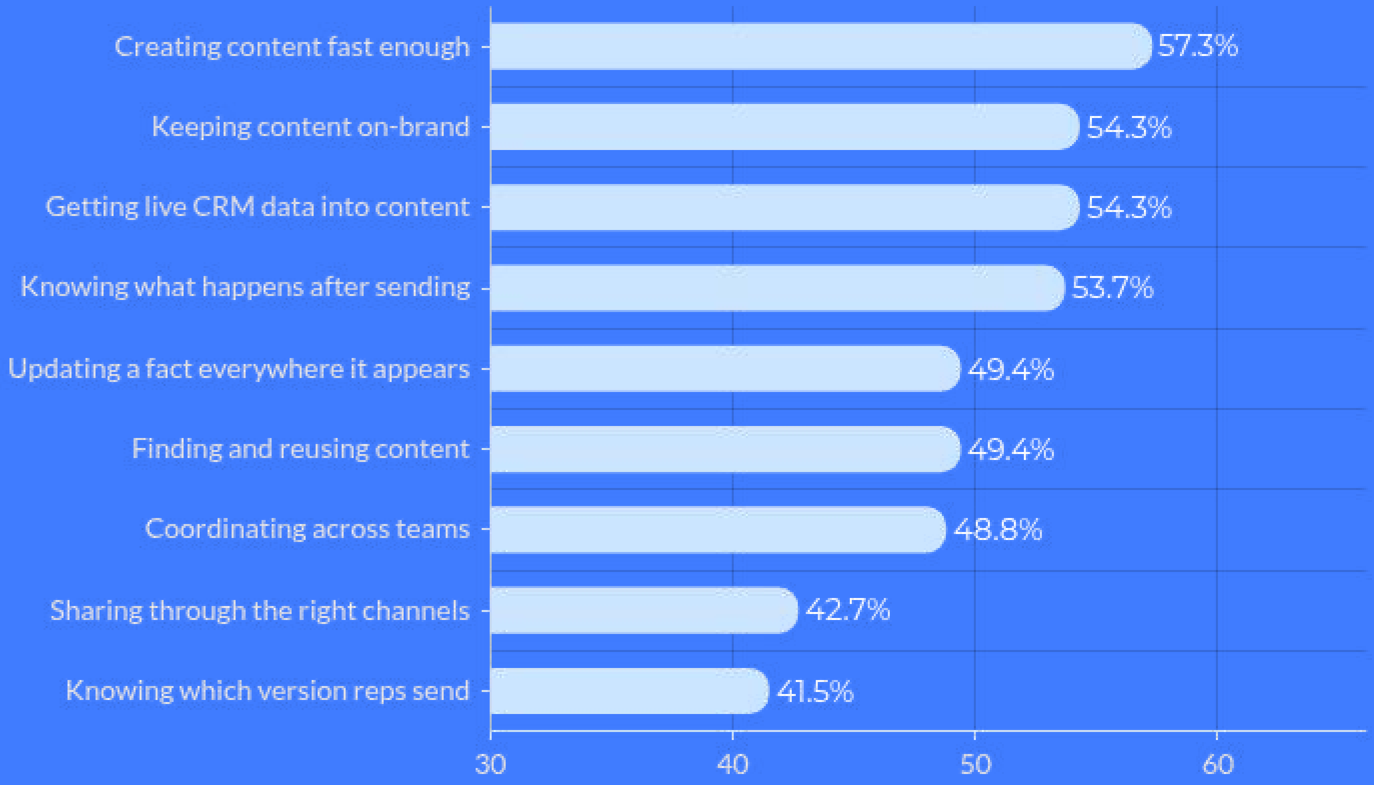

Share rating each pain 4 or 5 on a 5-point severity scale. n=164. The nine pains sit within 16 points of each other.

Creating content fast enough to hit deadlines edges out the rest at 57.3%, followed closely by keeping content on-brand and getting live data into a deck, both at 54.3%. Coordinating between multiple contributors, personalizing content at volume, and preventing reps from making last-minute edits to approved versions are all within a few percentage points of each other. The narrow spread matters because it suggests teams are not struggling with a single bottleneck. If one challenge had stood far above the rest, the diagnosis would be simple: fix that one thing.

There is also a sequencing story buried in this data. The pains at the top of the list — speed, brand consistency, live data — tend to be on the creation side. The pains lower in the band — post-send tracking, approval friction — tend to be governance and visibility problems.

The fact that both ends of the process are severe tells you that no stage of this workflow is running smoothly. Evidently, the pipeline has several leaks, spread across every stage, and the pressure behind each one is roughly equal.

COST

TEAMS LOSE AS MUCH AS HALF THEIR WEEK REDOING WORK THEY ALREADY COMPLETED

COST

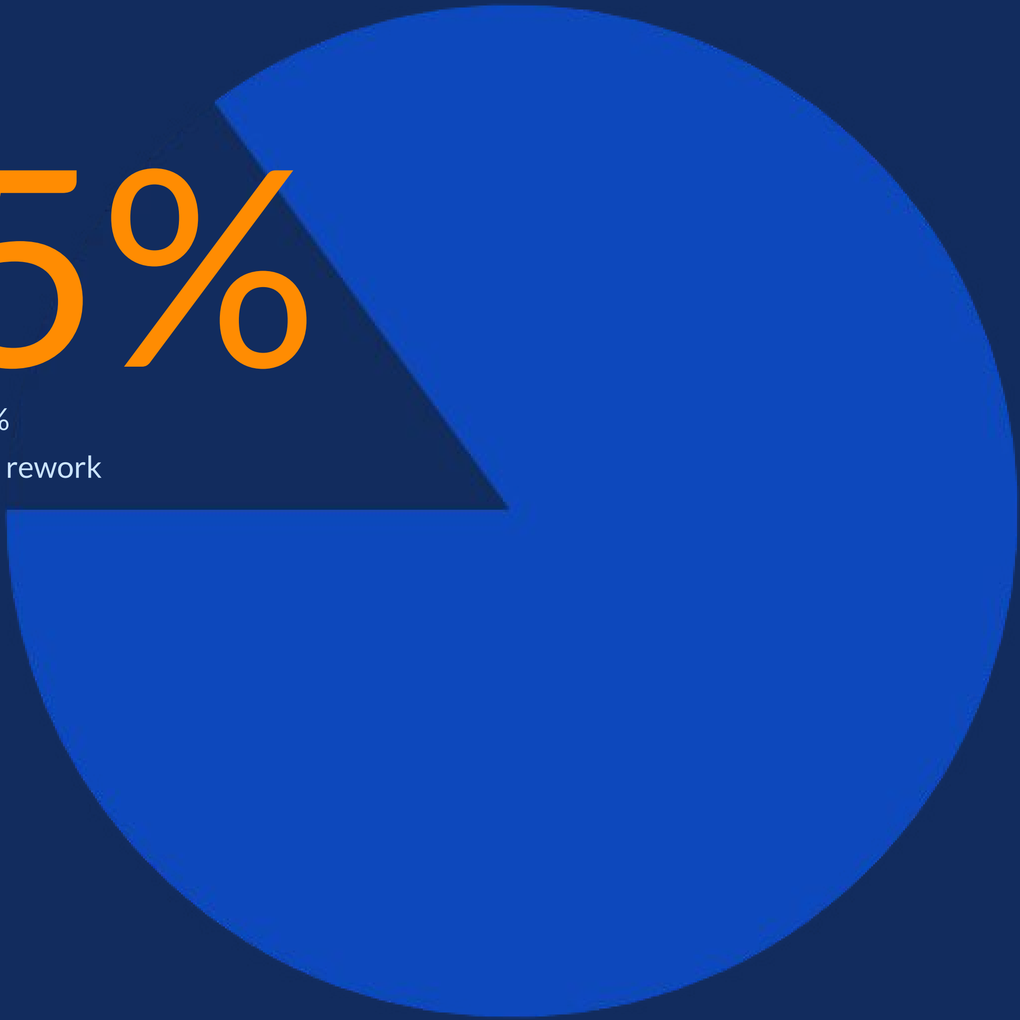

85%

lose 10% to 50%
of their week to rework

"The manual amount of effort it takes to extract data into a workable format that we can then use in pitches or presentations. It takes far longer than it should, and it's very manual."

Sales manager, United Kingdom

Every deliverable in this category requires a significant investment of times. 75% of teams spend between an hour and a full day on one, and about one in five spend more than a day. At 20-plus decks a month, the hours consumed on creating content alone are substantial.

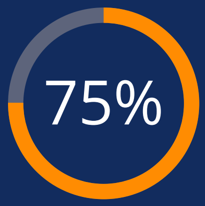

TIME TO PRODUCE ONE CUSTOMER-FACING DECK

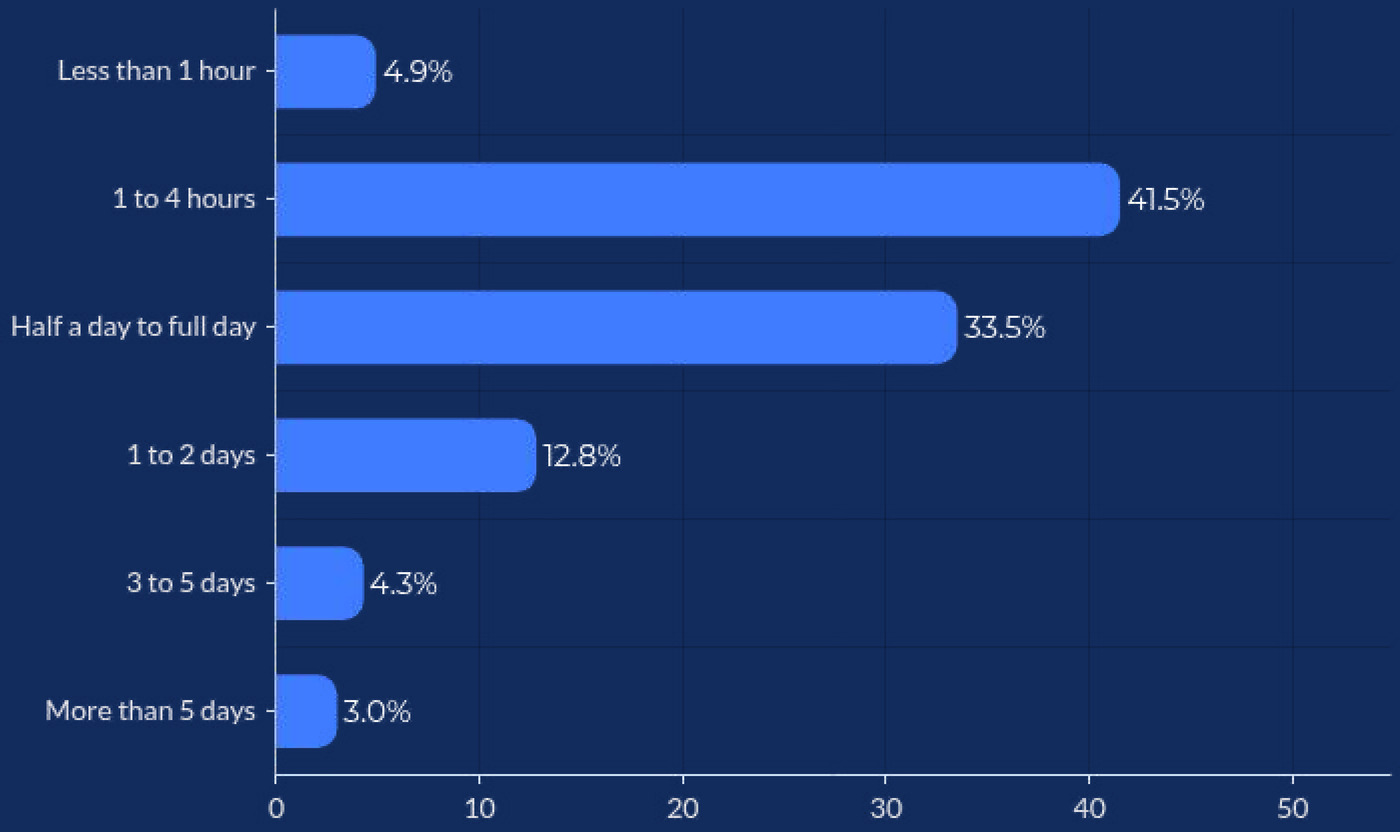

Time to produce one deck. n=164. Highlighted the 20% spending more than a day

The bigger time drain and cost is rework. 85.4% of teams spend between 10% and 50% of their week fixing and updating content they have already made. For nearly a third, that is more than a quarter of every week spent redoing finished work. In a standard 40-hour work week, that translates to at least 10 hours spent maintaining existing content rather than creating something new.

WHERE DECK DATA COMES FROM

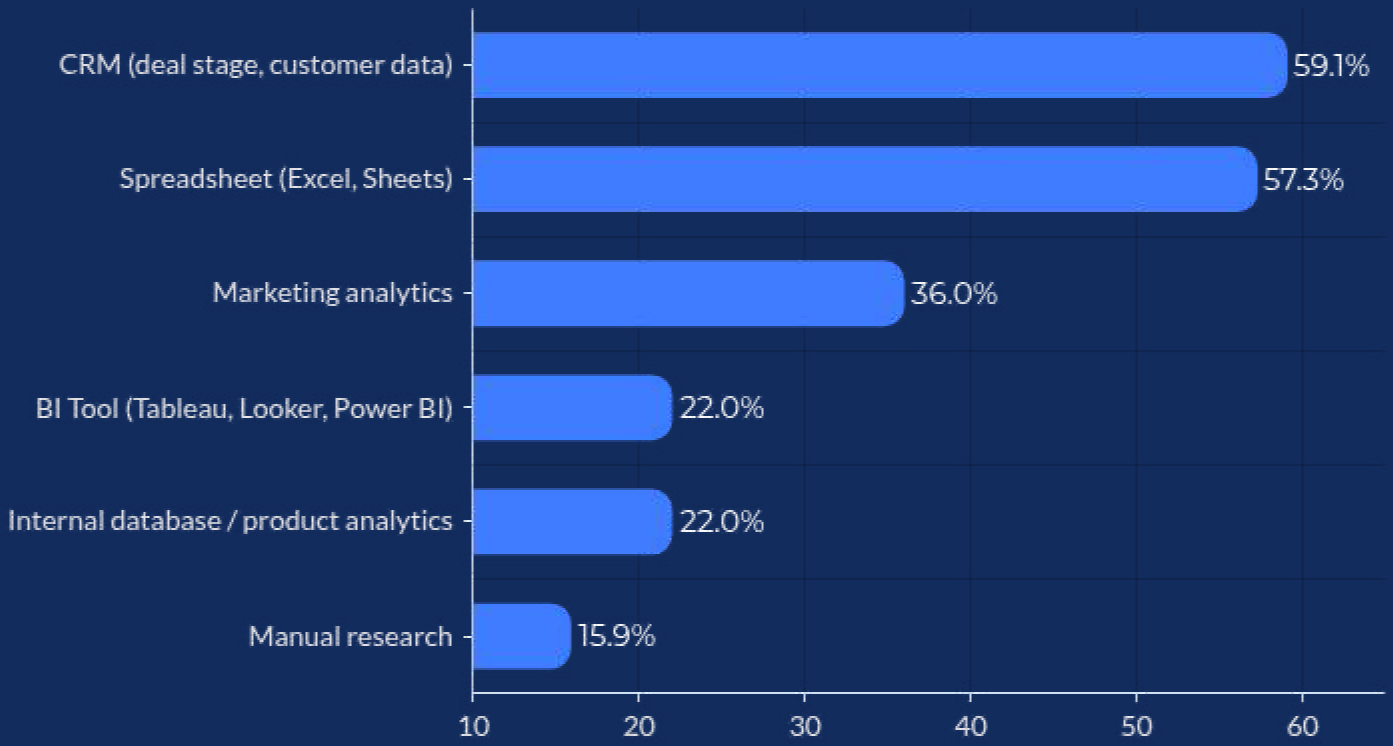

Data sources for a customer facing deck. Select all that apply n=164

Part of that time is spent on data collection and input. Most decks pull from a CRM (59.1%) and from spreadsheets (57.3%), and that information is often copied across through exported spreadsheets

SHARE OF THE WEEK SPENT ON REWORK

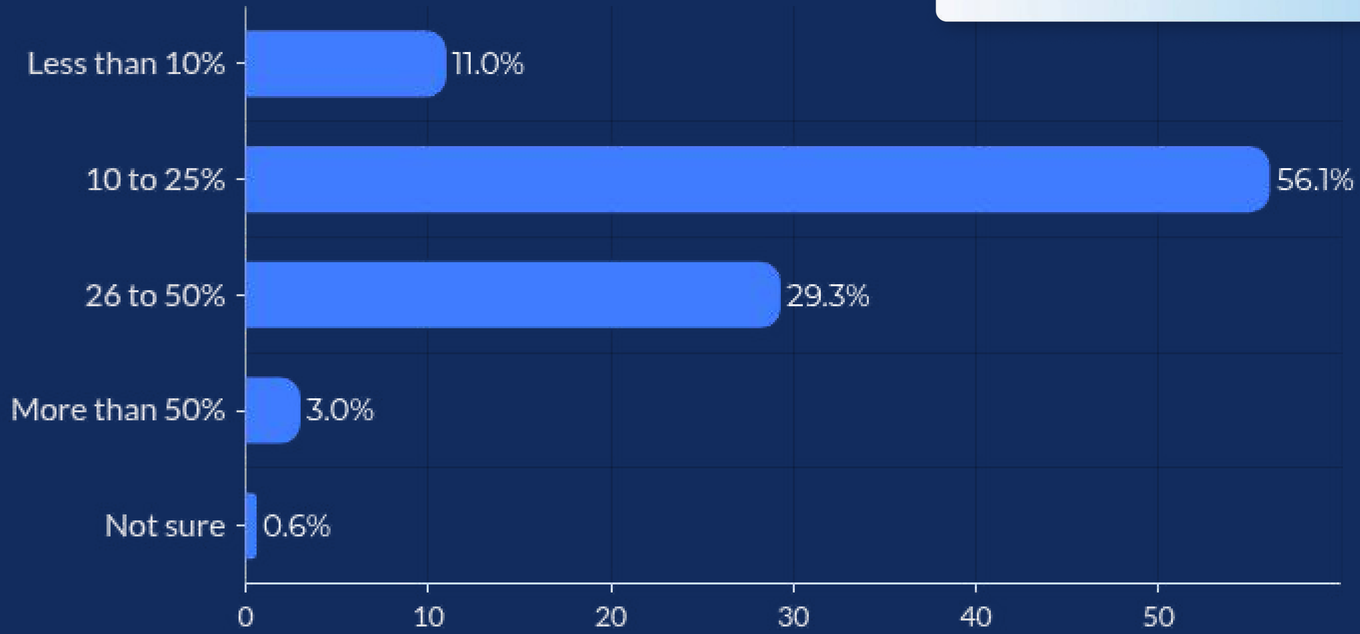

Rework versus new content in a typical week. n=164. Highlighted: the 85% at 10 to 50%

That information does not flow automatically into the deck. In most cases, it is exported and then pasted in or typed in by hand. When a deal progresses or customer data is updated, someone has to find every place that number appears and update it manually.

There is also an invisible ripple effect.

Every hour spent correcting an existing deck is an hour not spent preparing for a customer meeting, advancing an opportunity, or supporting an account. Rework is often treated as a normal part of the process, but the data suggests it has become one of the highest hidden costs in customer-facing content operations.

QUALITY

THE APPROVED DECK OFTEN DOESN'T REACH THE CUSTOMER UNTOUCHED

QUALITY

94.5%

of deliverables involve more than one contributor

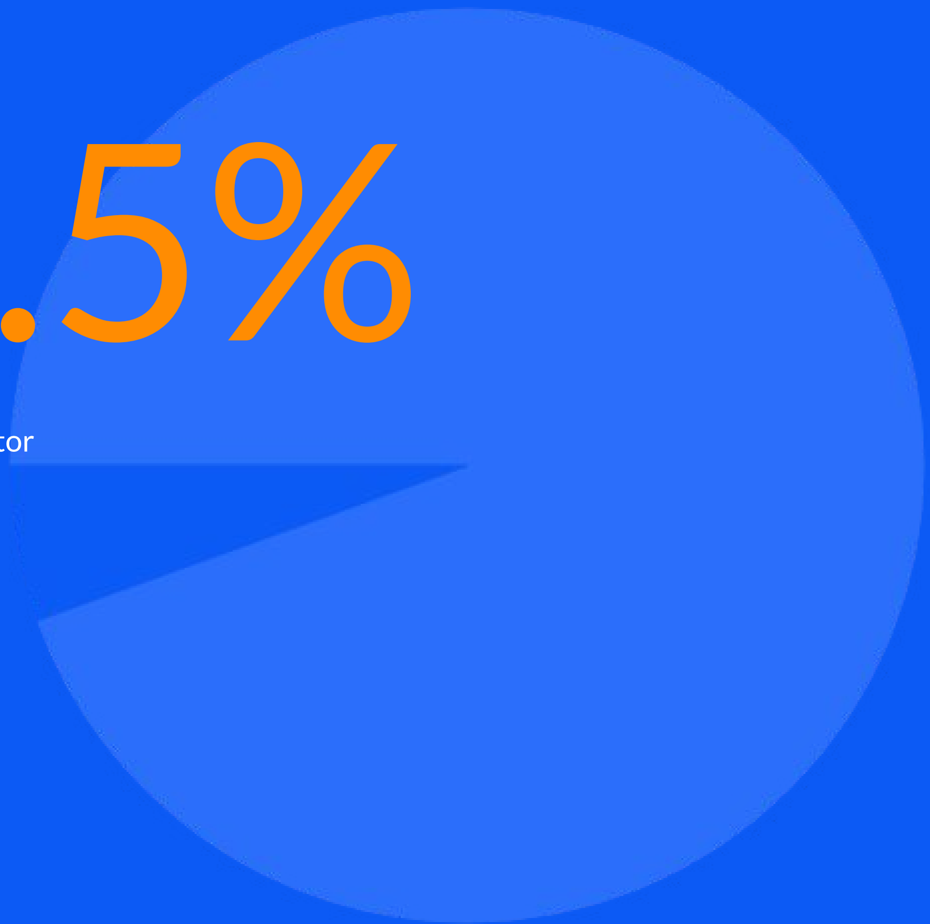

Creating customer-facing content is rarely a one-person job. Most deliverables pass through multiple contributors before reaching a prospect or customer, with 94.5% involving more than one person.

A team builds the master version. A sales enablement lead reviews and approves it. A rep picks it up and customizes it for a specific account. Sometimes a manager takes another pass.

Each handoff is a potential edit point, and there is rarely a single source of truth that prevents those edits from diverging from the approved version.

79.9% of teams say reps change that version at least sometimes on the way out. By the time it reaches the customer, it may have accumulated changes that no one with ownership visibility has reviewed.

HOW OFTEN REPS MODIFY THE APPROVED VERSION

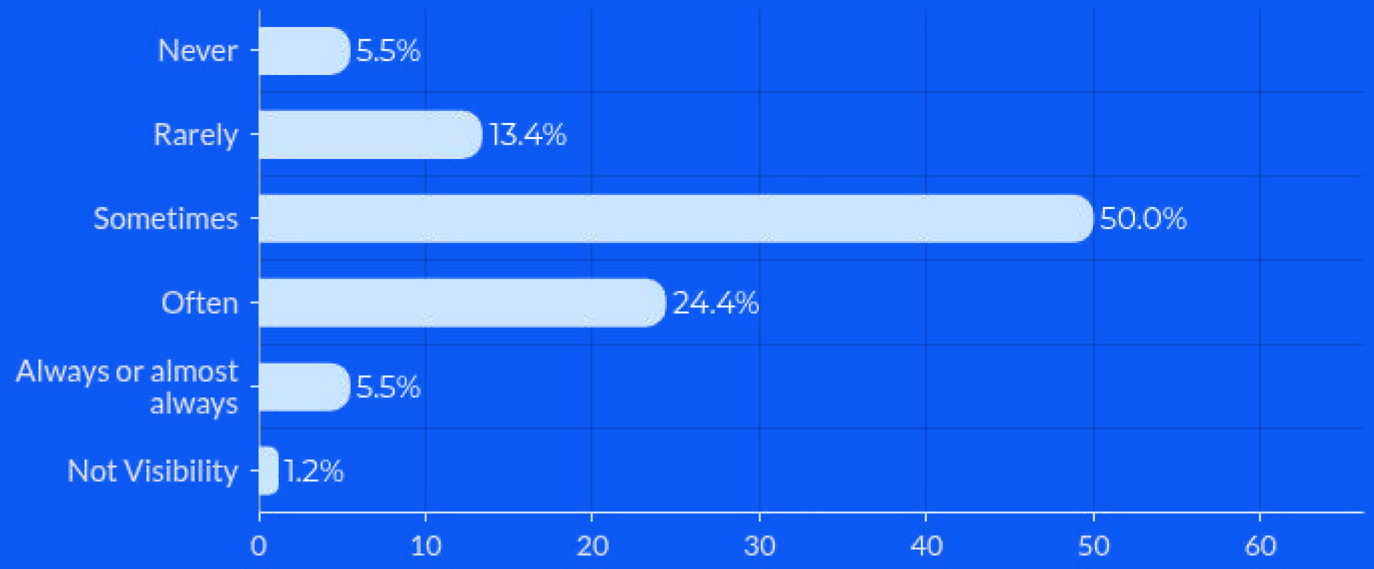

How often reps modify the approved version. n=164. Highlighted: the 80% who modify.

HOW OFTEN OUTDATED OR OFF-BRAND CONTENT REACHES A CUSTOMER

53%

of teams say outdated, off-brand, or incorrect content lands in front of customers monthly or more often. The customer does not know the deck was wrong. But the rep does. And the brand absorbs the cost either way.

Quality erosion in the final mile of content delivery

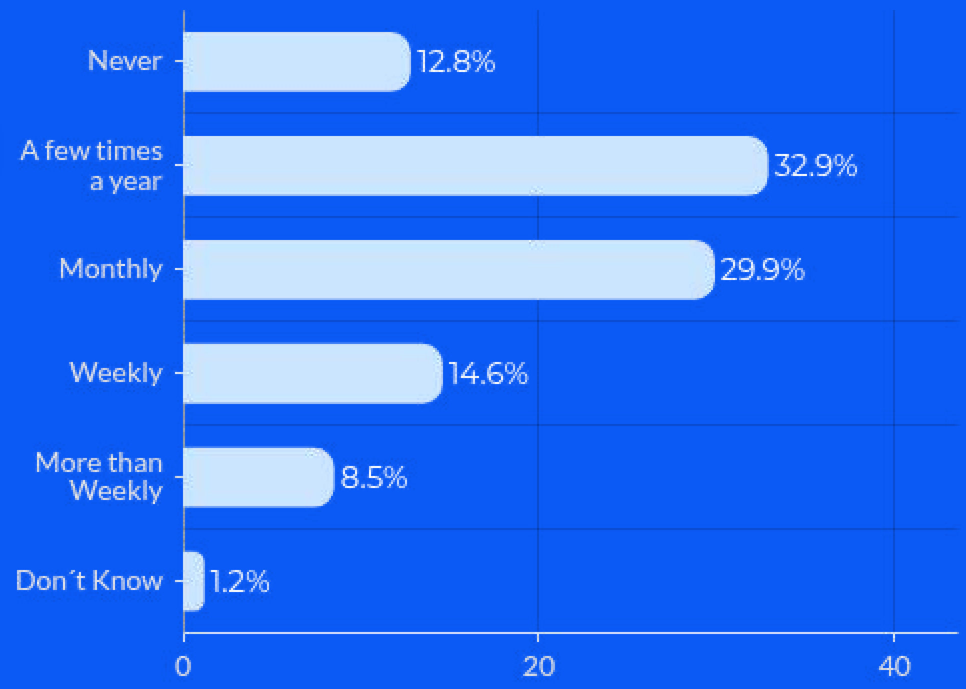

How often outdated or off-brand content reaches a customer. n=164. Highlighted: the 53% monthly or more

Two things make the final version hard to control:

94.5%

Most deliverables pass through several hands, with 94.5% having more than one contributor.

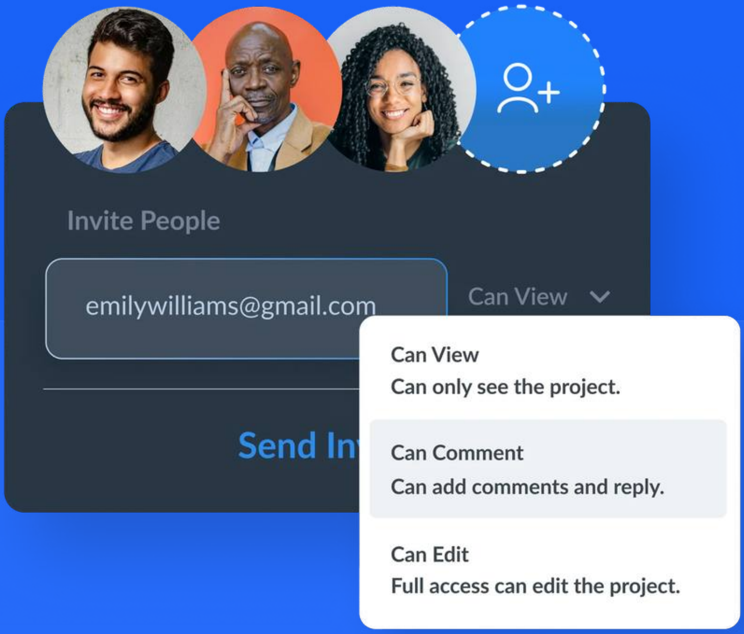

71.3%

Facts are updated by hand in 71.3% of teams, so a changed number doesn't always propagate across the board.

When a data point is updated at the source, someone has to catch it, find every deck where it appears, and fix it manually. In a high-volume environment producing 20-plus decks a month, maintaining that tracking is extremely difficult.

This means quality issues are not always caused by poor content. It is caused by the absence of a shared source of truth that travels with the content throughout its lifecycle. Approval processes protect content up to the moment of approval.

> **"The most frustrating thing is that, despite using multiple tools, we have no single source of truth for approved, updated content, so reps constantly waste time manually pulling live CRM data, hunting for the right version, and unknowingly sending outdated information to customers."**
>
> *C-level executive, United States*

VISIBILITY

HALF OF THE TEAMS HAVE NO REAL IDEA WHAT HAPPENS AFTER THEY HIT SEND

VISIBILITY

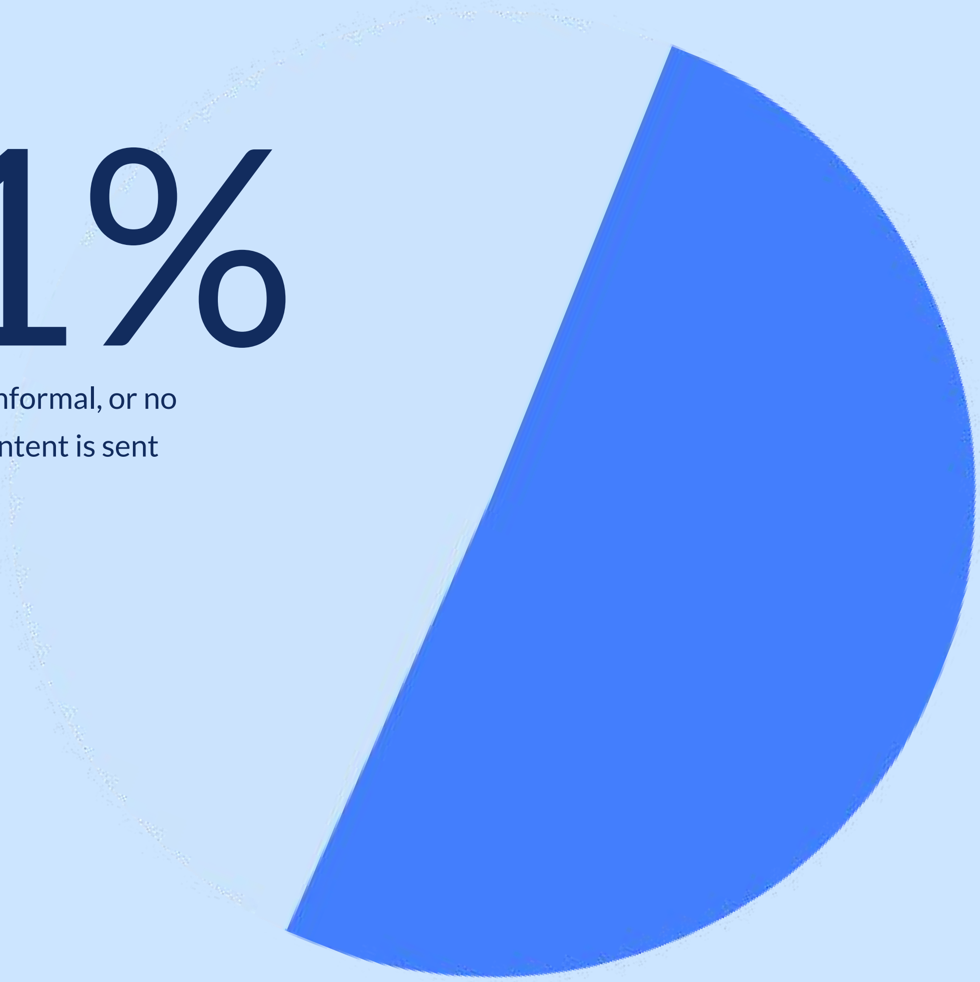

51%

have only basic, informal, or no analytics once content is sent

The moment content leaves the team, most of its creators lose sight of it entirely. They may receive a message from a rep saying the meeting went well. They may hear secondhand that a proposal was well received.

Occasionally, they will learn that a deck was rejected and a deal was lost. What they almost never receive is the granular engagement data needed to optimize the content. 50.7% have only basic analytics, informal feedback, or no tracking at all. They might see that a file was opened, but not which parts got read or whether it moved the deal.

Without that feedback, teams can't tell which decks land, which slides get read, or which formats move deals. They keep producing on instinct.

VISIBILITY INTO WHAT HAPPENS AFTER CONTENT IS SENT

How teams know what happens to after sending. n=164.

50.7%

have basic, none, or informal tracking

49.4%

have detailed engagement tracking

FLYING BLIND

- Sees that a file was opened
- Cannot see which slides were read
- No data on time spent per section
- Cannot see whether content was shared internally
- Follows up on rep anecdote and gut feel
- Iterates the next deck without any feedback signal

CLEAR VIEW

- Sees which slides hold attention and which get skipped
- Knows how long each section was read
- Can follow up while a prospect is still in the deck
- Tracks whether and with whom the deck was shared
- Identifies format patterns that correlate with closed deals
- Every send sharpens the strategy for the next one

In contrast, the other half tracks engagement in detail, which splits the market down the middle. That visibility pays off: They can see which sections of a deck hold attention and which get scrolled past without engagement. They know when a prospect spends six minutes on the commercial terms slide. They can trigger a follow-up call while the deck is still being read.

They can look across every deck sent in the past quarter and identify patterns: this format performs with enterprise accounts, this one loses momentum in the second half, this slide consistently generates questions that are worth anticipating in the next meeting. Every piece of content they send generates intelligence that makes the next piece more effective.

AI

MOST TEAMS TRIED AI FOR CONTENT, BUT IT DIDN'T KNOW THEIR BUSINESS

AI

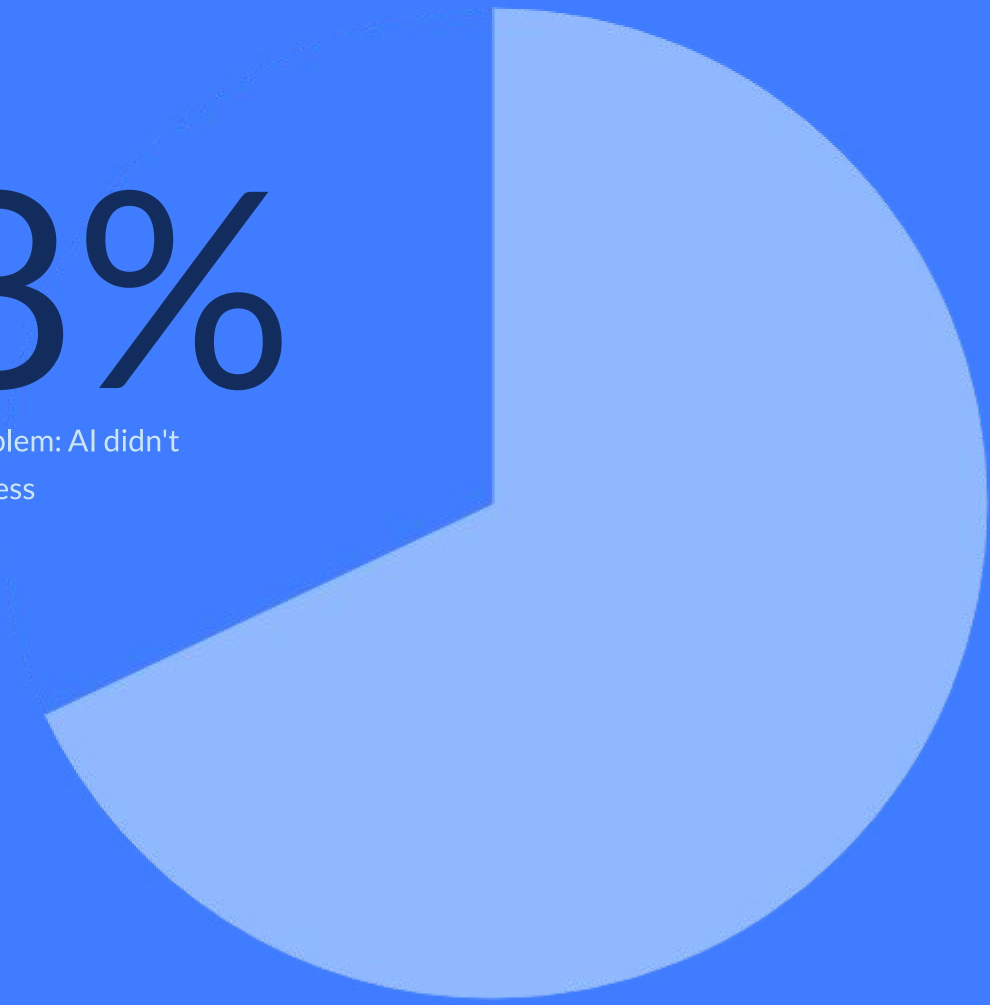

68%

hit a context problem: AI didn't know their business

When the volume problem grew acute enough, AI looked like the answer. Generate drafts faster, reduce the creation bottleneck and make personalization scalable. Most of the leaders in this survey acted on that premise. They tried AI tools for content creation. And most of them ran into the same ceiling.

The output kept falling short in ways that were specific, consistent, and traceable to a single root cause: the AI lacked context about the specific business it was supposed to be writing for.

The most common complaints were that it needed too much cleanup (33.5%), looked generic (31.1%), couldn't use the team's own data (20.7%), and didn't match the brand (19.5%).

WHY TEAMS PULLED BACK FROM AI

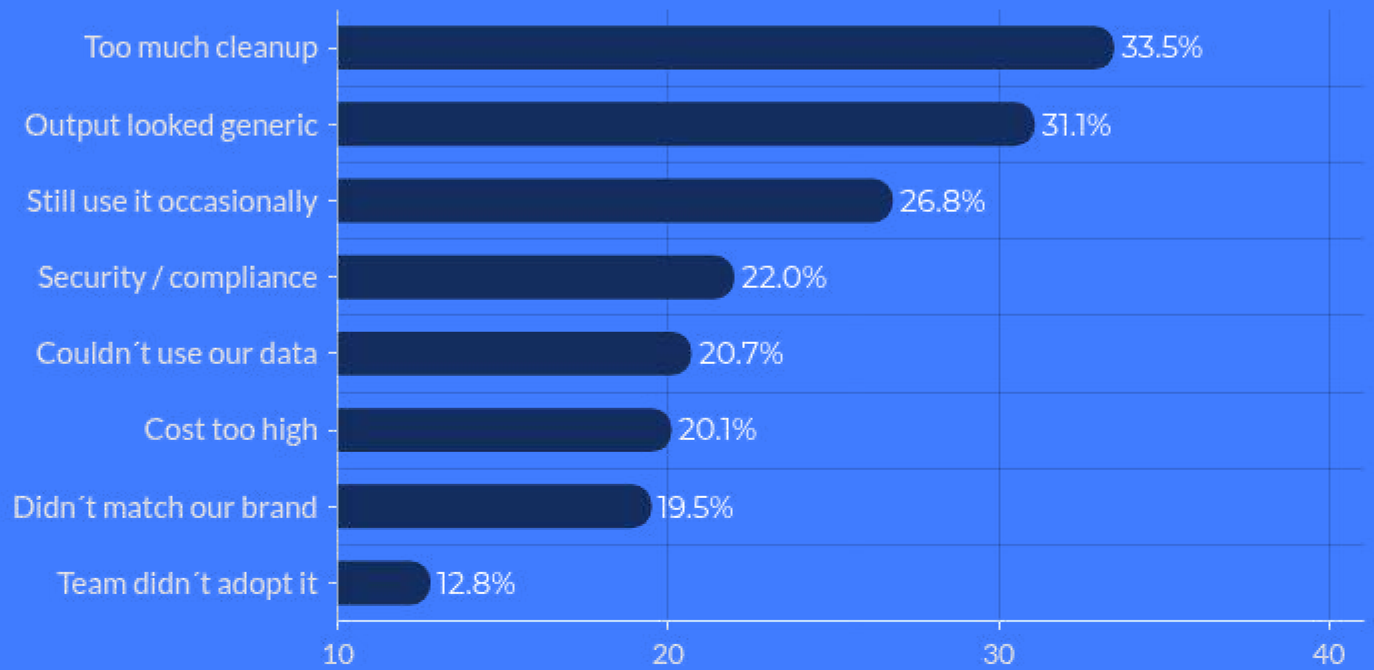

Reasons teams reduced AI use. Select up to two. n=164. Highlighted reasons are context failures; 67.7% cite at least one.

Those four complaints share one root. 67.7% of teams hit at least one of them, and each comes back to context the tool never had: the brand rules, the approved data, the team's voice, the record of what has worked before.

The common thread is the absence of a governance layer. The AI tools most teams tried were trained broadly. They could write about sales decks in general. They could produce a structurally sound quarterly business review template. What they could not do was write a QBR for a specific enterprise account using that account's actual data, in the company's approved voice, formatted to the brand's visual standards, with language that reflected the specific selling context and relationship history.

That is what the work actually requires. Generic AI was never built to deliver it, and the gap between what it can produce and what a revenue content team is willing to put in front of a customer turned out to be wider than the initial promise suggested.

DEMAND

TEAMS STILL WANT AI, BUT PREFER ONE BUILT ON THEIR OWN BRAND AND DATA

DEMAND

65%

would adopt AI built on their brand, data, and past results

The disappointment didn't kill the appetite. What it has done is make the appetite more precise. Teams that tried generic AI and found it wanting have not concluded that AI cannot help them. They have concluded that the specific AI they tried was not built to handle the work they actually do.

65.3%

of teams say they would adopt an AI tool that seamlessly uses their brand, tone, data, and past results. Only about one in ten rule it out.

WOULD YOU ADOPT AI GROUNDED IN YOUR BRAND AND DATA?

65.3% WOULD LIKELY OR DEFINITELY ADOPT

Likelihood of adopting AI grounded in your brand and data. n=164.

What teams are saying is that they don't need perfect output. In fact, 70.1% would use AI as long as it required only minimal edits or less work than starting from scratch. They want a draft that gets enough right so they spend less time fixing the output.

This is a practical bar, not a high one. It does not ask for AI that produces flawless, send-ready content on the first attempt. It asks for AI that gets enough right that a skilled editor can get it the rest of the way in less time than they would have spent building it from zero.

10.4% ruled out AI entirely, meaning nine in ten of these leaders remain open to the right solution.. The market for context-aware AI in this category is not shrinking because of early disappointments. It is waiting for something better, and it has been specific about what that something looks like. That specificity is the most commercially important signal in this report.

"A truly practical AI content tool would need to produce accurate, brand-consistent content with minimal editing, while integrating smoothly into our existing workflow and collaboration tools."

Manager, United States

WHAT THE DATA TELLS US

1

Customer-facing content is high-volume, repetitive, and heavily personalized: Two in three teams produce more than 20 customer-facing decks a month, and the formats at the top of the stack are the ones most directly tied to revenue. This is not a temporary surge. It is the permanent operating load of a B2B sales content function, and it requires infrastructure that scales with it.

2

Rework is an invisible line item that most teams have never formally calculated. 85% lose up to half their working week to it. At any realistic team size and salary level, that translates into a six-figure annual cost that sits invisibly on the timesheet. It is not tracked as waste; rather, it's just tracked as "hours worked."

3

Quality breaks down at the last mile, every time. Approval processes protect the content up to the point of approval and not one step further. Once a rep picks up the deck, the approved version is at risk. Without a system that enforces data integrity and limits post-approval edits, it is difficult to prevent incorrect content from reaching customers.

4

Visibility is a compounding competitive advantage, and the market is split almost exactly in half. Teams with detailed post-send engagement data are running a learning loop that improves every piece of content they produce. Teams without it are repeating the same playbook with no feedback signal. The gap between those two positions widens every quarter.

5

The AI category has not lost the market. It has been given a more specific brief. 65% would adopt AI built on their own brand, data, and results. That is a product requirement. The teams who tried generic AI and found it wanting have not given up on the category. They have simply gotten precise about what would actually solve their problem.

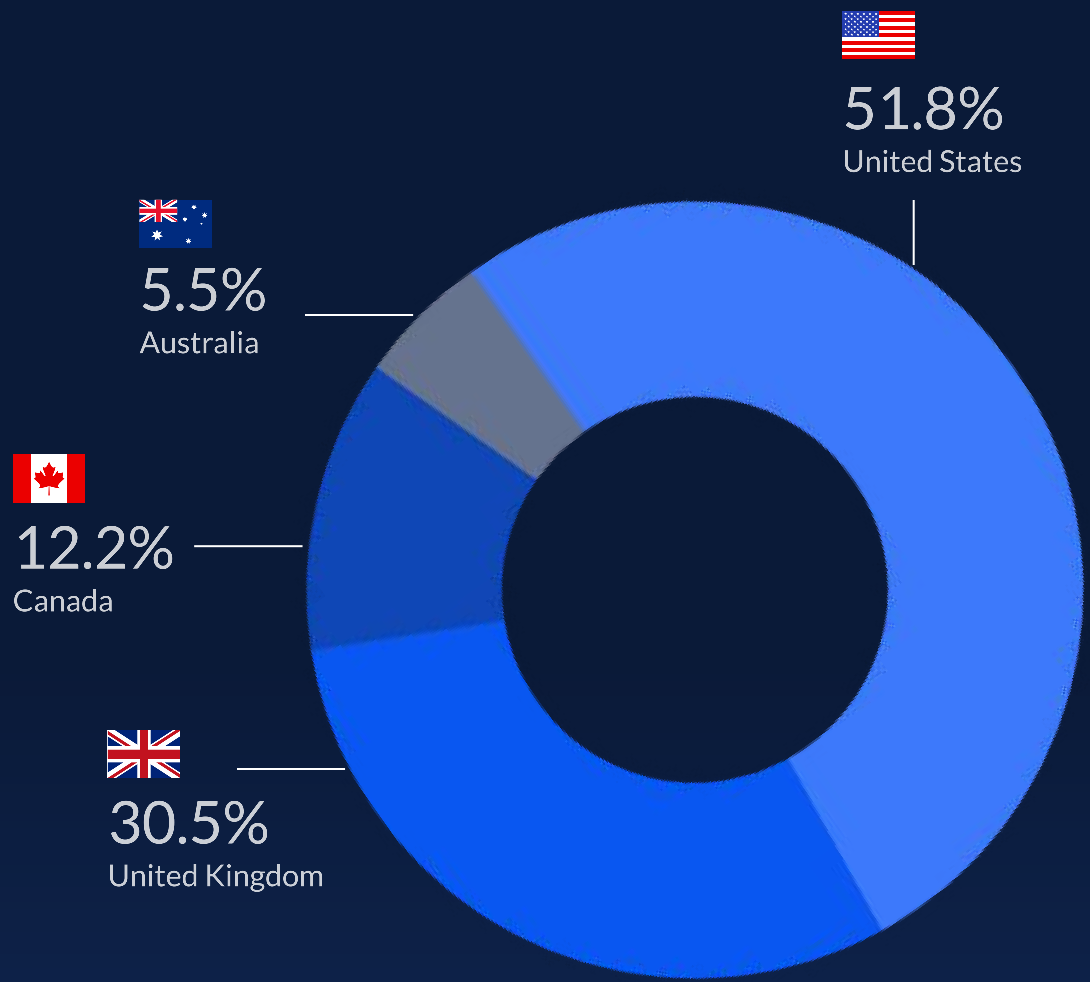

METHODOLOGY

Visme surveyed 164 managers, directors, and senior executives at B2B companies that produce customer-facing content. Respondents work in the United States (51.8%), United Kingdom (30.5%), Canada (12.2%), and Australia (5.5%).

The survey ran on UserTesting between March and May 2026. Every respondent confirmed their team produces customer-facing decks, proposals, or reports. None were current Visme users at the time of the survey.

At 164 responses, the results carry a margin of error of roughly 7.7% at 95% confidence. Figures are directional rather than precise.

HOW DOES YOUR TEAM COMPARE?

You've seen how 164 teams produce, customize, and track their customer-facing content. The open question is where your own team sits against them.

We built a short self-assessment from the same questions in this survey. Answer them for your team and see how you compare on output, rework, quality control, post-send visibility, and AI sentiment.

It takes a few minutes and shows you which gaps are costing your team the most.

GET YOUR TEAM BENCHMARK SCORE