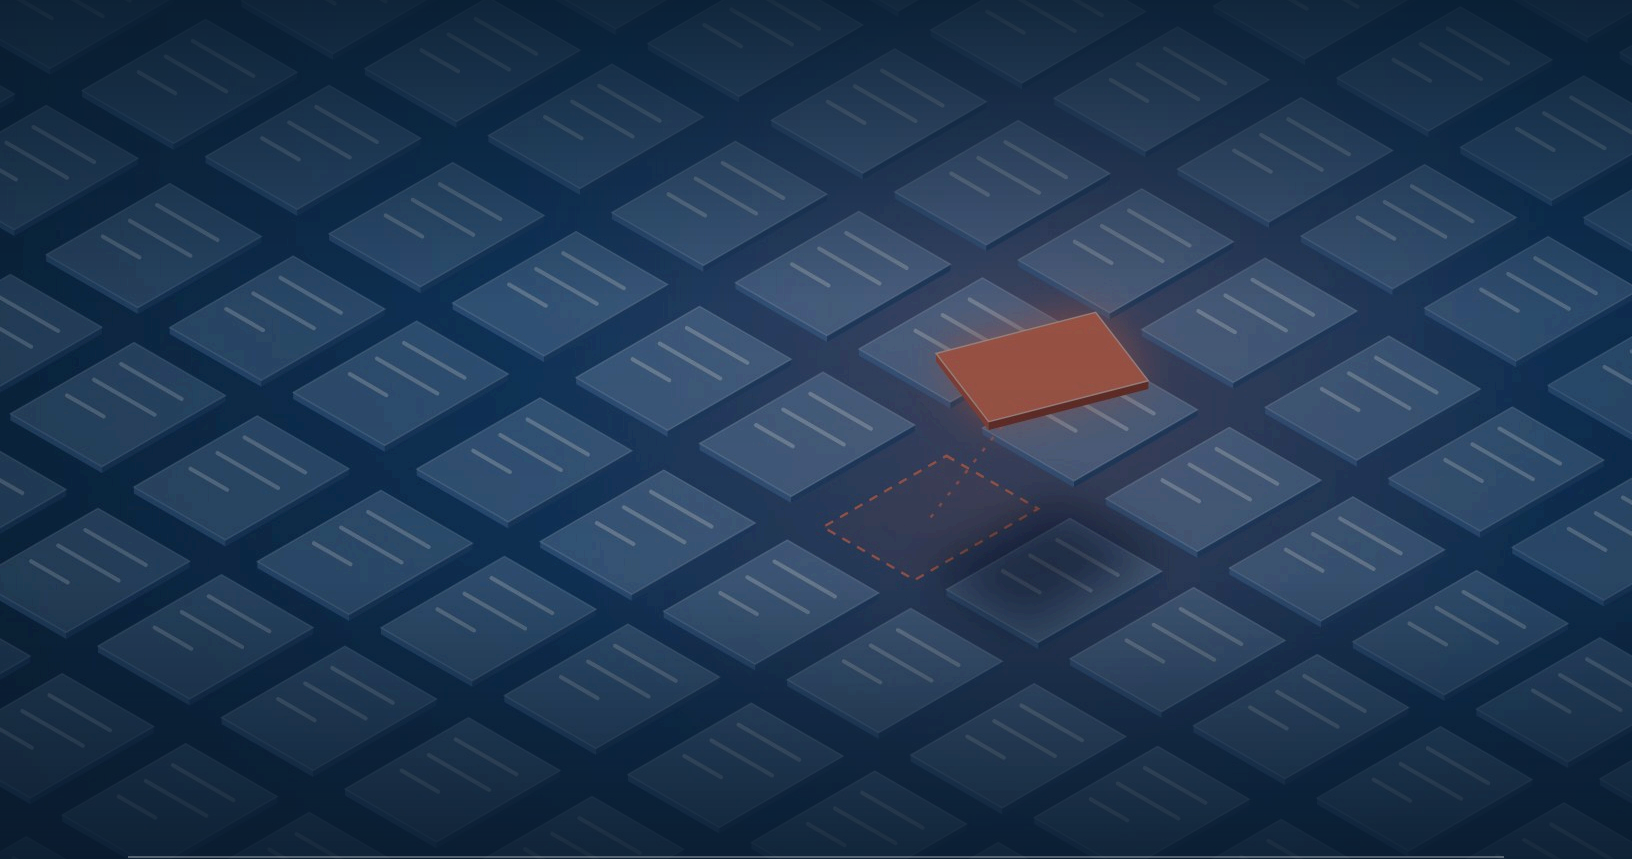
WHITE PAPER

The Cost of Content Drift

How inaccurate and outdated content creates support tickets, feeds wrong answers to AI assistants and pulls NPS down, and what correcting it returns.



1 in 10

published assets in the median library carry a claim that is no longer true

0 of 30

Claude answers were correct when the page it was given held one outdated fact

$11.15M

modeled annual cost of drift for an illustrative large enterprise

© 2026 CONTINUITY INTELLIGENCE, INC.

EXECUTIVE SUMMARY

Inaccurate and outdated content is a cost line, not a housekeeping task

It creates support tickets, fails self-service, feeds wrong answers to AI assistants, and pulls NPS down. Correcting it pays back.

KEY FINDINGS

1 in 10

Drift is already in the library.

In the median library Continuity Intelligence scanned, 10% of published assets carried at least one claim that was no longer true.

CI, 41 full-library scans, April to October 2026

14%

Self-service fails on content.

Gartner found that customers fully resolve only 14% of issues in self-service. In 43% of failures they could not find content relevant to their issue.

Gartner survey of more than 5,700 customers, August 2024

0 of 30

AI repeats what it is given.

We asked Claude 10 customer questions using real pages that contained one outdated fact. 0 of 30 answers gave the current fact. With that one fact corrected, 30 of 30 did.

CI Claude drift test, November 17, 2026

22 of 30

The customer gets no warning.

22 of the 30 answers stated the outdated fact with no caveat. The only item all three Claude models flagged was a table whose own header showed an old year.

CI Claude drift test, November 17, 2026

5 places

The cost lands in five places.

Avoidable contacts, agent and author time, customer effort and NPS, wrong AI answers at scale, and legal exposure.

Section 4

ROI

The return can be proven.

Tag content-caused tickets, test AI answers weekly, and compare corrected articles with uncorrected ones. The companion Content Drift ROI Report turns those counts into dollars.

Sections 5 and 6

THE BOTTOM LINE

$11.15M

modeled annual cost of drift for a large enterprise

$6.69M

year-one return from correcting it, on a $2.0M program

3.6 months

to pay back the program

Expected case of an illustrative model with 10 million assisted contacts a year (section 6). Not a customer result.

RECOMMENDATIONS

1. Measure the share of support contacts that content causes, with one contact tag and a ticket search.
2. Test your own AI assistant against ten known findings.
3. Correct critical and high-traffic findings first, and retire superseded versions.
4. Re-measure after 30 days and fund the program on measured numbers.

© 2026 Continuity Intelligence, Inc. | continuityintelligence.com

01



PART ONE / THE PROBLEM

Content drifts faster than teams can fix it

One in ten published assets already says something that is no longer true, and a single stale fact reaches customers through every channel that reads it.

1 in 10 published assets in the median library carry a claim that is no longer true

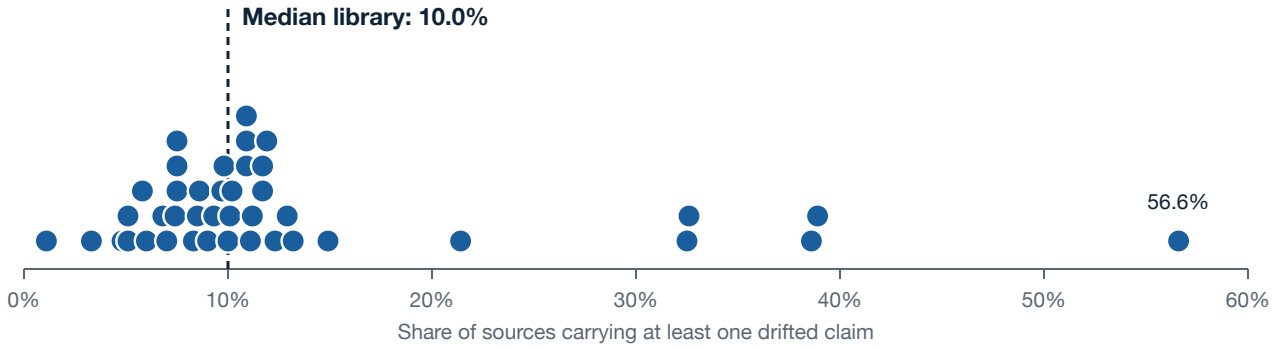
SECTION 1 / THE PROBLEM

Content drifts faster than teams fix it

One in ten published assets in the median library is already wrong, and new drift arrives every week. The figures below come from Continuity Intelligence scans of 119 organizations between April and October 2026, published in *The State of Content Drift 2026*.

EXHIBIT 1

Drift rate by library: each dot is one organization's public library



Source: Continuity Intelligence, The State of Content Drift 2026. 41 public libraries of 100 or more sources each. Across all 119 organizations: 1,959 verified findings in 37,865 public sources.

28%
of drifted facts appear in more than one place. One appeared in 179.
1,253 findings in full scans

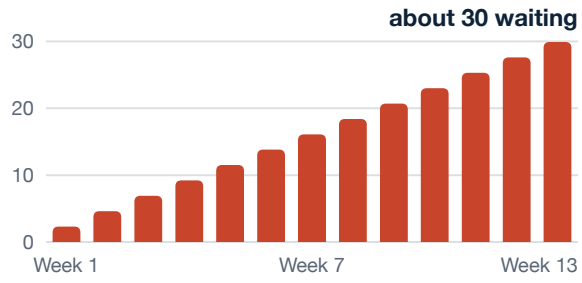
18%
of findings: the correct fact was already published elsewhere on the company's own site.
All 1,959 findings; a lower bound

2.3
new findings a week in one SaaS help center.
29 weekly scans, steady state May 3 to Sep 28

20 days
median time to clear a finding under weekly monitoring.
84 resolved findings in that library

EXHIBIT 2

Findings waiting for the next quarterly review



Calculation: 2.3 new findings a week over a 13-week cycle.

A quarterly audit cannot keep up

At 2.3 new findings a week, about 30 accumulate between reviews. Each one is live, and readable by customers and AI assistants, for the whole wait.

61% of customer service leaders said they had a backlog of articles to edit, and more than one-third had no formal process for revising outdated articles.

Gartner survey reported in December 2024, read in DestinationCRM

SECTION 2 / THE MECHANISM

How one outdated page reaches a customer

A single stale fact travels through every channel that reads the page, and the customer cannot tell it is stale.

EXHIBIT 3 THE PATH OF A STALE FACT

1 **Something changes**
A price, a plan, a policy, a screen, a regulation.

2 **The page stays as written**
So do its copies in other articles, PDFs, courses and videos.

3 **A customer meets the old fact**
It reaches them through every channel that reads the page.

SEARCH · HELP CENTER · SUPPORT AGENT · AI ASSISTANT

4 **The customer acts on it**
Picks the wrong plan, follows steps that no longer exist, misses a deadline, emails an address nobody monitors.

5 **The failure becomes a contact**
A support contact, a repeat contact, or silent abandonment.

6 **The effort shows up in the scores**
First in CSAT and NPS, and later in renewals.

COMPANIES ARE HELD TO WHAT THEIR CHANNELS SAY

CAD $812.02

in damages, interest and fees
Moffatt v. Air Canada, February 14, 2024

British Columbia's Civil Resolution Tribunal found Air Canada liable for negligent misrepresentation after its website chatbot gave a customer inaccurate bereavement-fare information.

The tribunal held that the airline is responsible for all information on its website, whether it comes from a static page or a chatbot.

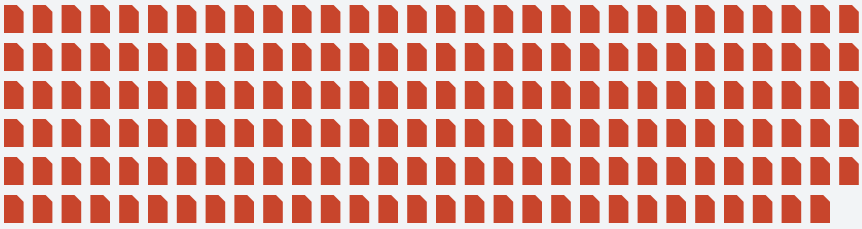
The sum was small. The principle applies to every answer a company's content produces.

Moffatt v. Air Canada, 2024 BCCRT 149, as summarized by Deeth Williams Wall

179

places held the same drifted fact in one library. 28% of drifted facts appear in more than one place, and every copy has to be found and fixed.

CI, 1,253 findings in full scans



> *None of these pages looks wrong. Each was accurate on the day it was written.*

© 2026 Continuity Intelligence, Inc. | continuityintelligence.com

02

PART TWO / THE AI MULTIPLIER

AI assistants repeat what they are given

We gave Claude ten real pages with one outdated fact each. It passed the old fact on to the customer every time.

0 of 30 answers were correct when the page held one outdated fact

SECTION 3 / THE AI MULTIPLIER

What Claude says when content has drifted

Given a page with one outdated fact, Claude passed that fact to the customer every time: 0 of 30 answers were correct. Given the same page with the fact corrected, 30 of 30 were.

EXHIBIT 4

Answers that gave the customer the current fact: each dot is one answer

The page as published, with the drifted fact



0 of 30 CORRECT

The same page with the one fact corrected



30 of 30 CORRECT

No page, open web search (Sonnet 5.5 only)



9 of 10 CORRECT

Outdated fact, no caveat (22) | Outdated fact, with a caveat or fallback (8) | Current fact

Source: Continuity Intelligence Claude drift test, November 17, 2026. 10 pages, each asked once of Claude Haiku 5.5, Sonnet 5.5 and Opus 5.5.

EXHIBIT 5

One question, one changed fact: case 1, a customer messaging platform's add-on price

GROUNDED IN THE PAGE AS PUBLISHED

HELP CONTENT GIVEN TO CLAUDE

It costs $99 per month as a flat fee, which includes up to 1,000 conversations.

CUSTOMER

How much will the add-on cost us per month if we handle about 800 conversations?

CLAUDE SONNET 5.5

At about 800 conversations, you'd pay just the $99 flat monthly fee, since it covers up to 1,000 conversations.

✕ OUTDATED The vendor's newer page sets a per-conversation price.

GROUNDED IN THE CORRECTED PAGE

HELP CONTENT GIVEN TO CLAUDE

The add-on is billed at $0.20 per conversation, with a 500-conversation minimum.

CUSTOMER

How much will the add-on cost us per month if we handle about 800 conversations?

CLAUDE SONNET 5.5

At about 800 conversations, the [add-on] would cost roughly $160 per month ($0.20 x 800), which is above the $100 monthly minimum.

✓ CURRENT Same model, same prompt. Only the page changed.

Claude Sonnet 5.5's answers, first sentence shown. Brand and product names are masked; page excerpts are shortened.

© 2026 Continuity Intelligence, Inc. | continuityintelligence.com

SECTION 3 / THE AI MULTIPLIER

The test

We took 10 public pages from Continuity Intelligence scans, each with one verified outdated fact. For 8 of the 10, the organization's own newer page states the current fact. For the other 2, a statute or regulator does. We gave Claude each page excerpt as its knowledge base, the way a support assistant is grounded in help content, and asked one customer question. We ran three models (Claude Haiku 5.5, Sonnet 5.5 and Opus 5.5) on November 17, 2026. Then we corrected the single fact and ran everything again.

The ten cases

Claude Sonnet 5.5's answers. Brand and product names are masked.

	ORGANIZATION AND PAGE	WHAT DRIFTED	ANSWER, DRIFTED PAGE	ANSWER AFTER CORRECTION
1	**Customer messaging platform** help center	Add-on price. Page: $99 flat for up to 1,000 conversations. Current: $0.20 per conversation, 500 minimum	✕ "you'd pay just the $99 flat monthly fee"	✓ "roughly $160 per month ($0.20 x 800)"
2	**Water treatment manufacturer** product page	Filter life. Page: 365 gallons. Current manual: 250	✕ "replaced every 365 gallons or every 6 months"	✓ "every 250 gallons or every 6 months, whichever comes first"
3	**Regional transit agency** rider FAQ	Paper passes ended May 1, 2026. FAQ still says to show one	✕ "Yes, you can show a day pass when you board"	✓ "No, paper monthly and day passes are no longer accepted"
4	**AI voice software company** voices page	A retired voice is still listed	✕ "Yes, [voice] is listed as an English (AU) Narration voice"	✓ "I don't see a voice named [voice] in our current list"
5	**Learning platform vendor** security page	Uptime. Page: 99.99%. Published SLA: 99.9%	✕ "We offer a 99.99% uptime guarantee"	✓ "We guarantee a 99.9% monthly uptime SLA"
6	**HR and payroll platform** help center	State salary threshold table stops at 2025	✕ "$68,640 per year", then a warning to confirm the current figure	✓ "the 2026 California threshold is $70,304"
7	**Privacy software vendor** product documentation	Deadline extension for one US state. Page: 45 days. Statute: 15	✕ "a maximum extension of 45 days"	✓ "the maximum extension allowed is 15 days"
8	**Communications equipment maker** app page	Support email on a domain the company no longer holds	✕ "You can contact support ... at support@[lapsed domain]"	✓ The company's current support address and phone line
9	**Product content network** supplier FAQ	A retailer portal is now read-only. FAQ says to keep using it	✕ "Yes, you can continue using the current process"	✓ "No, [retailer] made [portal] read-only on March 1, 2026"
10	**University procurement office** policy page	Bid threshold. Page: no quotes under $25,000. Current rule: three quotes from $5,000	✕ "No, competitive quotes aren't required for a $12,000 purchase"	✓ "Yes, purchases from $5,000 to $30,999 require quotes from three qualified suppliers"

© 2026 Continuity Intelligence, Inc. | continuityintelligence.com

SECTION 3 / THE AI MULTIPLIER

What the test shows

The model cannot see what changed.

It answered fluently and politely from whatever it was given. Model size made no difference: the smallest and the largest model both scored 0 of 10.

Warnings appeared only when the page carried its own clue.

All three models flagged the salary table, whose column header said 2025. None warned that the price, the uptime figure, the legal deadline or the support address might be out of date.

More sources help, but not always.

With open web search and no page supplied, Claude found the newer source 9 times in 10. In the tenth it led with the old $99 price, because the vendor's help center carried both versions. Twice more it noted other pages that still gave the old figure or instruction.

A company's own assistant has no second source.

A support bot, an agent-assist tool or an internal copilot answers from the company's knowledge base. If that is stale, the first row of Exhibit 4 is the result.

Limits. Ten cases, chosen because each has one clear, checkable fact, and one run per model. This demonstrates the mechanism. It does not estimate how often it happens in any one company.

Run it yourself

Paste this into Claude with one of your own articles. Then fix the outdated fact and run it again.

PROMPT

```
You are the customer support assistant for [company].
Answer the customer's question using the help content below.
Keep the answer to 1 to 3 sentences. Today's date: [date].

Help content:
"""
[paste the article text]
"""

Customer question: [a question the article answers]
```

© 2026 Continuity Intelligence, Inc. | continuityintelligence.com

03



PART THREE / THE MONEY

What drift costs, and what correcting it returns

Five cost lines, five ways to measure them, and a worked example sized for a large enterprise.

$11.15M modeled annual cost of drift for an illustrative large enterprise

SECTION 4 / WHAT IT COSTS

Drift is paid for in five places

Each has public evidence behind it and a way to measure it in your own data.

COST LINE AND CAUSE	PUBLIC EVIDENCE	MEASURE IT WITH
Avoidable contacts A customer follows old content, fails, and contacts support.	14% of issues are fully resolved in self-service. In 43% of failures customers could not find relevant content (Gartner, 2024). Gartner's benchmark puts the median assisted contact at $13.50 against $1.84 for self-service.	A "content was wrong or outdated" contact reason. Contacts within 24 hours of an article view.
Agent and author time Agents answer from the same stale knowledge base. Writers audit pages by hand.	50 to 60% faster time to resolution, and 30 to 50% higher first contact resolution, reported by organizations that keep knowledge current (Consortium for Service Innovation, member-reported).	Handle time and reopen rate on drifted topics. Hours logged on manual review.
Customer effort, NPS and retention A wrong answer costs the customer time and trust.	2x An industry's Net Promoter leader outgrows its competitors by more than two times. Lifetime value of promoters runs 3 to 8 times that of detractors (Bain).	NPS and CSAT for customers who touched drifted content against those who did not. Coded survey comments.
Wrong AI answers at scale An assistant repeats the stale page to every customer who asks.	30 of 30 answers repeated the outdated fact in our Claude test. 85% of service leaders planned to explore or pilot customer-facing generative AI in 2025 (Gartner).	A fixed question set run against the assistant every week.
Legal and compliance exposure The company is bound by what its channels say.	Liable Moffatt v. Air Canada (2024): the airline was held responsible for inaccurate information its chatbot gave a customer.	Count of open findings that touch price, policy, safety or regulation.

One caution. There is no reliable public benchmark for the share of support tickets that inaccurate content causes. Vendor figures circulate, but we found none that traces to a documented study. That share is the most important input in any ROI model, so measure it instead of borrowing it.

© 2026 Continuity Intelligence, Inc. | continuityintelligence.com

SECTION 5 / HOW TO PROVE ROI

Five measurements tie content fixes to tickets and NPS

The first two give a baseline within a month, and the third shows cause.

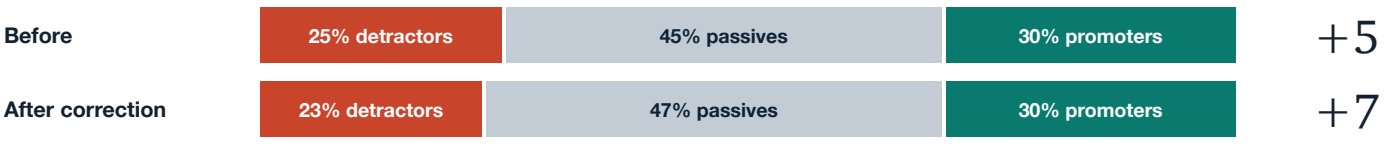
	MEASUREMENT	WHAT YOU DO	WHAT IT PROVES	FIRST RESULT
1	**Content-caused contact tag**	Add one contact reason: “content was wrong or outdated”. Agents apply it and paste the article link.	The count and cost of contacts that drift creates.	30 DAYS
2	**Ticket text search**	Search 12 months of tickets and chats for the outdated terms a scan finds: old plan names, old prices, retired features.	A baseline, without waiting for new tags.	1 WEEK
3	**Corrected against uncorrected articles**	For each corrected article, compare contacts per 1,000 views and helpfulness votes for 30 days before and after. Uncorrected articles are the control.	That the fix moved the number, not the season.	60 DAYS
4	**AI answer accuracy**	Run a fixed question set against the assistant every week, before and after correction.	Wrong answers removed from the AI channel.	1 WEEK
5	**NPS and CSAT linkage**	Code survey comments for wrong or outdated information. Compare scores for customers who viewed drifted articles.	The NPS points and retention at stake.	ONE SURVEY CYCLE

The return is then a single calculation

$$\text{ROI} = \frac{(\textbf{contacts avoided} \times \text{cost per contact}) + (\textbf{hours recovered} \times \text{hourly cost}) + \textbf{revenue retained} - \text{program cost}}{\textbf{program cost}}$$

EXHIBIT 6

On NPS, the arithmetic is direct: every 1% of respondents moved from detractor to passive adds one point



Illustrative example. NPS is the share of promoters minus the share of detractors. Here coded comments show that 2% of all respondents are detractors because of wrong or outdated information, so two points are recoverable by fixing it.

© 2026 Continuity Intelligence, Inc. | continuityintelligence.com

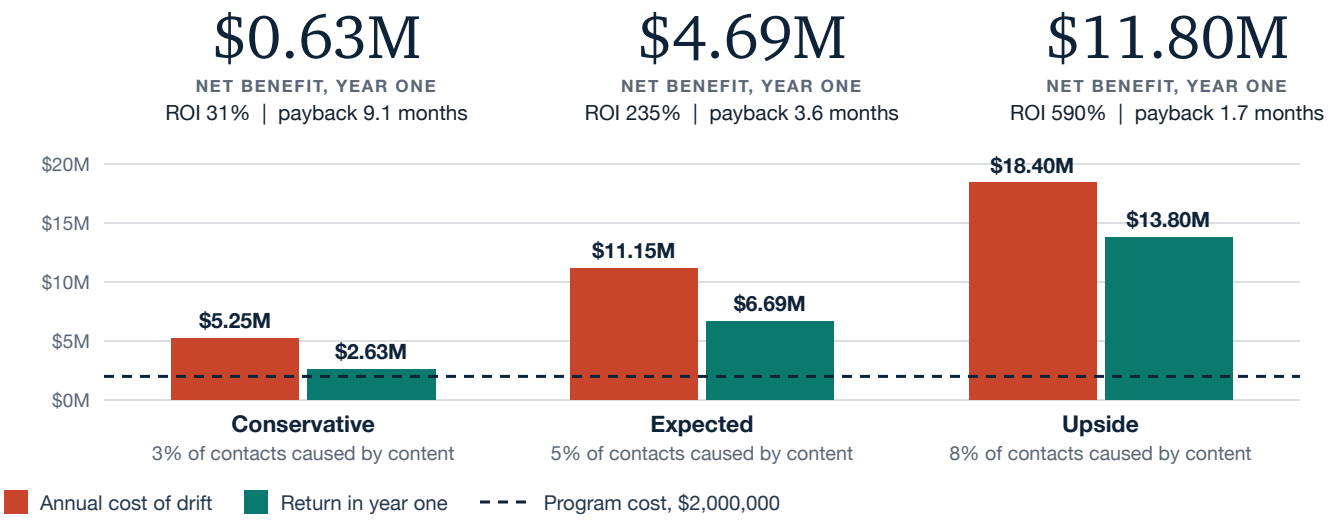
SECTION 6 / THE RETURN

What correcting it returns

For an illustrative large enterprise with 10 million assisted contacts a year, correcting drift returns between $2.6M and $13.8M in year one, against a $2.0M program. The spread comes almost entirely from one input: the share of contacts that content causes.

EXHIBIT 7

Year-one cost of drift and return in three scenarios, illustrative large enterprise



Illustration, not a customer result. Inputs are listed in the table below and in the companion Content Drift ROI Report.

LINE	CONSERVATIVE	EXPECTED	UPSIDE
Share of contacts caused by wrong or outdated content	3%	5%	8%
Avoidable contact cost, at $13.50 per contact	$4,050,000	$6,750,000	$10,800,000
Manual review time, 20,000 hours at $60	$1,200,000	$1,200,000	$1,200,000
Revenue at risk, share of $160M annual churn linked to content failures	Not counted	$3,200,000 (2%)	$6,400,000 (4%)
Annual cost of drift	**$5,250,000**	**$11,150,000**	**$18,400,000**
Share removed in year one by continuous correction	50%	60%	75%
Return in year one	$2,625,000	$6,690,000	$13,800,000
Program cost, software plus internal time	$2,000,000	$2,000,000	$2,000,000
Net benefit in year one	**$625,000**	**$4,690,000**	**$11,800,000**
ROI	31%	235%	590%
Payback	9.1 months	3.6 months	1.7 months

Every input here is an assumption for illustration, not a customer result. In the conservative case, with no revenue effect counted, the program still returns more than it costs. The companion Content Drift ROI Report shows every formula so you can replace each input with your own.

© 2026 Continuity Intelligence, Inc. | continuityintelligence.com

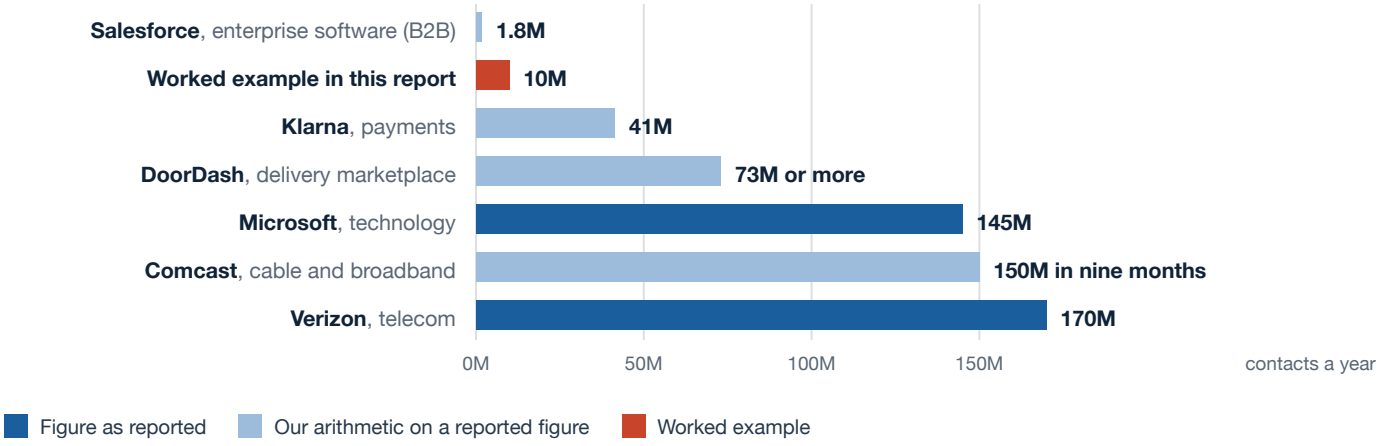
SECTION 6 / THE RETURN

The same arithmetic at enterprise scale

The worked example assumes 10 million assisted contacts a year. Public statements put large consumer-facing enterprises well above that, and enterprise software vendors below it.

EXHIBIT 8

Publicly reported customer contact volumes against the worked example



Sources: Salesforce blog (November 2025); Klarna press release (February 2024); AWS case study (DoorDash); UC Today (Microsoft, July 2025); CX Dive (Comcast, October 2026); Reuters (Verizon, June 2024). Companies are named as public scale markers only. Nothing here describes their content, costs or results.

EXHIBIT 9

Annual cost of avoidable contacts by enterprise scale, at three shares of content-caused contacts

ENTERPRISE PROFILE	CONTACTS A YEAR	AT 3%	AT 5%	AT 8%	ONE POINT IS WORTH
Enterprise software vendor (B2B) Marker: Salesforce	2 million	$0.81M	**$1.35M**	$2.16M	$0.27M
Large enterprise: the worked example	10 million	$4.05M	**$6.75M**	$10.8M	$1.35M
Large consumer platform Marker: Klarna	40 million	$16.2M	**$27.0M**	$43.2M	$5.40M
Fortune 50 service provider Marker: Microsoft, Comcast, Verizon	150 million	$60.8M	**$101.3M**	$162.0M	$20.3M

Contacts x share caused by content x $13.50 per assisted contact. Only this line is scaled. Review time, revenue at risk and program cost do not rise in step with contact volume.

At carrier scale, one percentage point is worth about $23M a year. That is 170 million calls x 1% x $13.50. It is arithmetic on public figures, not a finding about any named company.

© 2026 Continuity Intelligence, Inc. | continuityintelligence.com

04

PART FOUR / THE PLAN

A first result in thirty days

Scan, set a baseline, test the assistant, correct, and re-measure.

30 days to a drift rate, a baseline and a before-and-after result

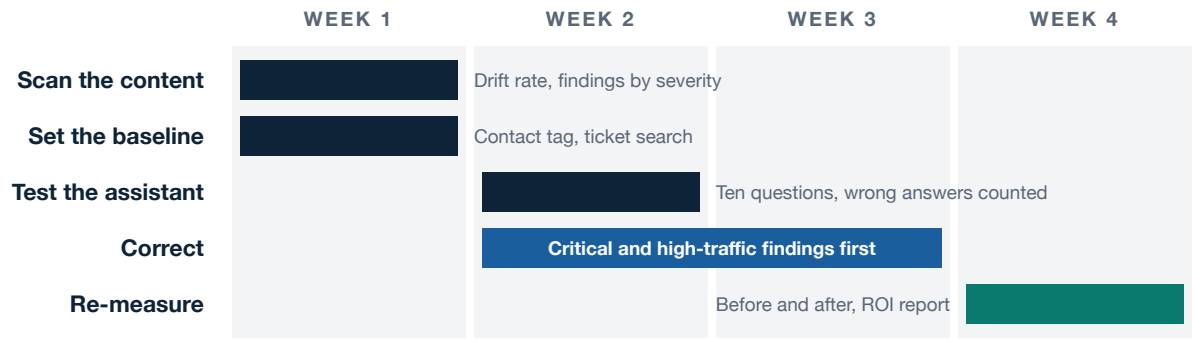
SECTION 7 / WHERE TO START

A first result in thirty days

Thirty days is enough to get a drift rate, a baseline and a first before-and-after result.

EXHIBIT 10

The 30-day plan



1 **Week 1: scan.**
Scan the public help center, documentation and academy. The output is a drift rate and a list of findings by severity, each tied to evidence of the current fact.

2 **Week 1: baseline.**
Add the “content was wrong or outdated” contact reason. Search the last 12 months of tickets for the outdated terms the scan found.

3 **Week 2: test the assistant.**
Turn ten findings into customer questions and run them against your own AI assistant. Record how many answers are wrong.

4 **Weeks 2 and 3: correct.**
Fix the critical and high-traffic findings first. Remove or redirect superseded versions so no index holds both.

5 **Week 4: re-measure.**
Run the same ten questions again. Compare tagged contacts and article results before and after. Fill in the ROI report with measured numbers.

Continuity Intelligence runs steps 1, 3 and 4 as a continuous service: its agents watch the sources of truth, detect drift as it happens, and autofix the affected content wherever it lives.

© 2026 Continuity Intelligence, Inc. | continuityintelligence.com

SOURCES AND METHOD

Where the numbers come from

SOURCE	USED FOR	HOW IT WAS CHECKED
Continuity Intelligence, The State of Content Drift 2026	Drift rate, fan-out, own-site contradictions, weekly velocity, time to clear	CI scan data, 119 organizations, April to October 2026
Continuity Intelligence, Claude drift test, November 17, 2026	The AI multiplier section	10 pages from CI scans dated May 9 to November 16, 2026. Prompts, page excerpts and all answers are held by CI
Gartner, survey of more than 5,700 customers, August 19, 2024	14% resolved in self-service; 43% could not find relevant content	Read in reports of the release by DestinationCRM and CX Dive. Gartner's own page blocks automated access
Gartner survey of customer service leaders, reported by DestinationCRM, February 19, 2025	85% exploring customer-facing generative AI; 61% with an article backlog; more than one-third with no revision process	Read in the DestinationCRM report
Gartner, cost per contact benchmark	$13.50 assisted and $1.84 self-service, median, February 2024	Read in Gladly's cost-per-contact glossary, which cites the February 2024 benchmark. Gartner's own page blocks automated access
Moffatt v. Air Canada, 2024 BCCRT 149	Liability for inaccurate chatbot information	Read in the Deeth Williams Wall summary. The tribunal's site blocks automated access
Bain & Company, How Net Promoter Score Relates to Growth and The Numbers behind the Net Promoter System	NPS leaders outgrow competitors by more than two times; promoter and detractor lifetime value differs 3 to 8 times	Opened
Public scale markers: Salesforce, Klarna, DoorDash, Microsoft, Comcast, Verizon	Customer contact volumes in section 6	Each page opened. Salesforce: Salesforce blog, Nov 6, 2025. Klarna: Klarna press release, Feb 27, 2024. DoorDash: AWS case study. Microsoft: UC Today, July 10, 2025. Comcast: CX Dive, Oct 22, 2026. Verizon: Reuters, June 18, 2024. Per-year readings for Salesforce, Klarna, DoorDash and Comcast are our arithmetic
Consortium for Service Innovation, KCS v6 Practices Guide, Why KCS?	Member-reported gains from maintained knowledge	Opened

How the Claude test was scored

An answer counted as correct only if it gave the customer the current fact. An answer that stated the outdated fact counted as outdated, with or without a caveat. A caveat means the answer told the customer to verify, noted a possible limit, or offered a fallback. The ten pages were selected from verified critical findings for having one clear, checkable fact. They are not a random sample.

How to read the worked ROI example

The enterprise in section 6 is illustrative. Its contact volume, cost per contact, share of content-caused contacts, review hours, churn and program cost are stated assumptions. No customer ticket or NPS outcome is reported in this paper. Those effects have not yet been measured in a customer deployment.

© 2026 Continuity Intelligence, Inc. | continuityintelligence.com

ABOUT CONTINUITY INTELLIGENCE

Keep what you have published true

Continuity Intelligence is a Continuity Management System for training and customer-education content. Its agents watch your sources of truth, detect drift as it happens, and autofix the affected content wherever it lives.

This paper is free to share with attribution. To see the drift in your own library, request a free drift report.

Request a free drift report →

COMPANION DOCUMENTS

Case study: When the Help Center Drifts, the AI Assistant Follows
Thirty-two weekly scans of one company's content, and the same Claude test on its open findings.

Content Drift ROI Report
Every input and formula behind section 6, ready for your own numbers, with a 90-day proof plan.

continuityintelligence.com
linkedin.com/company/continuity-intelligence



About this research. Continuity Intelligence Research publishes findings from Continuity Intelligence scans of public content, with third-party figures attributed to their publishers. Worked financial examples are illustrative models built on stated assumptions. They are not forecasts, benchmarks or customer results.

© 2026 Continuity Intelligence, Inc. | continuityintelligence.com