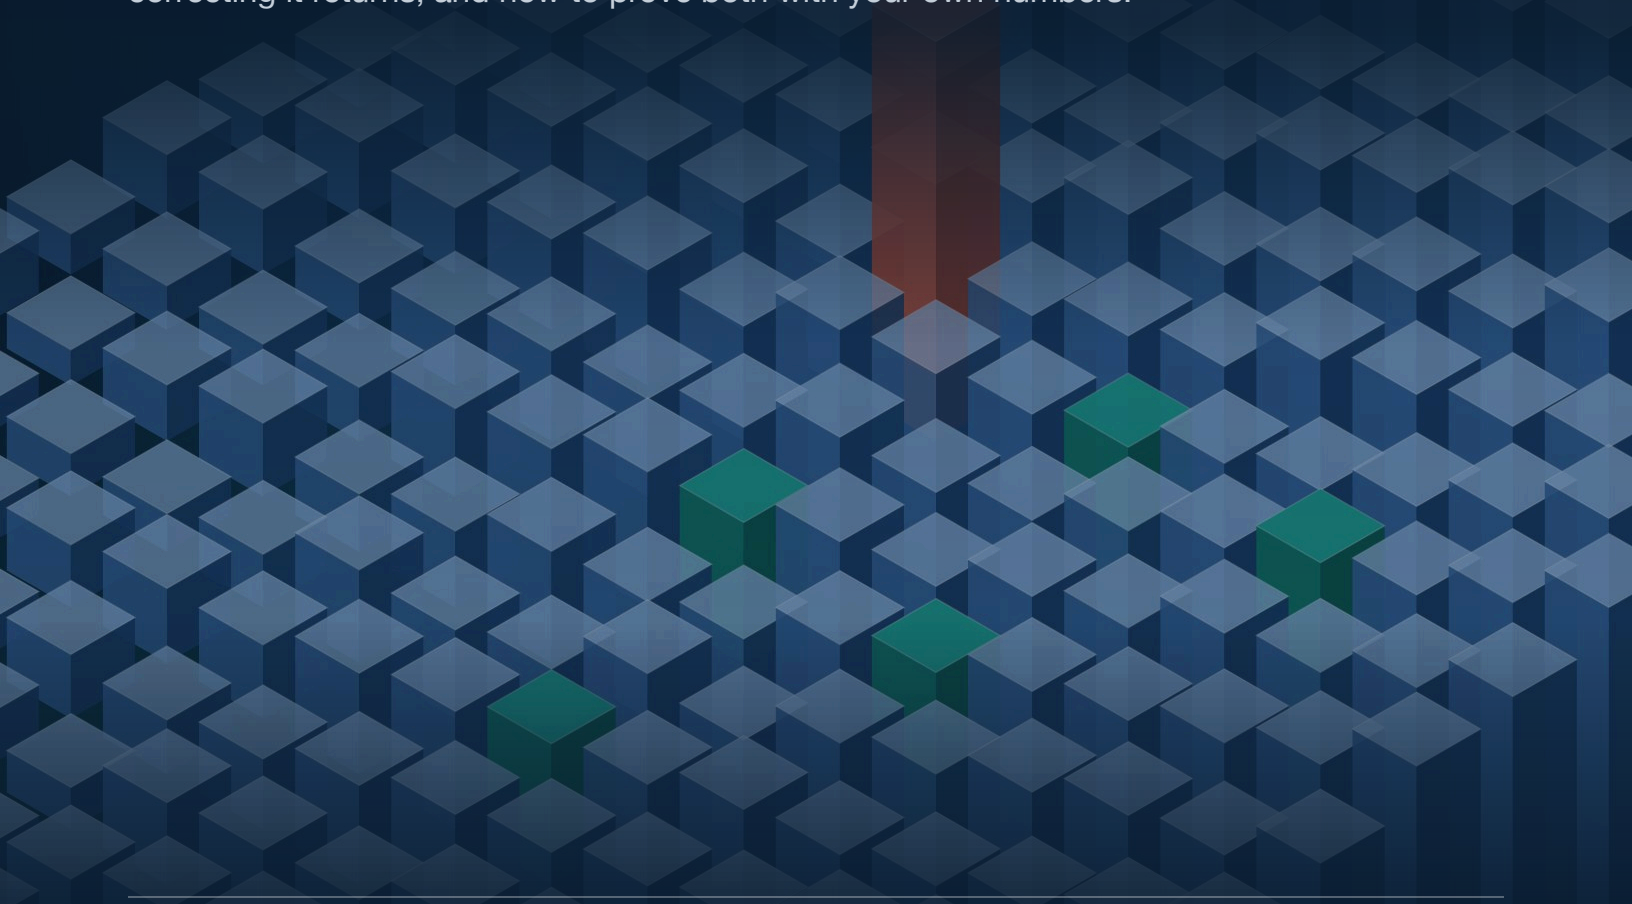
REPORT TEMPLATE WITH WORKED EXAMPLE

Content Drift ROI Report

What inaccurate and outdated content costs a large enterprise, what correcting it returns, and how to prove both with your own numbers.



$11.15M
modeled annual cost of drift, illustrative large enterprise

$6.69M
returned in year one on a $2.0M program, expected case

3.6 months
to pay back the program

© 2026 CONTINUITY INTELLIGENCE, INC.

EXECUTIVE SUMMARY

In the worked example, content drift costs a large enterprise $11.15M a year. Correcting it returns $6.69M in year one.

On a $2.0M program that is $4.69M net in year one, a 235% ROI, with payback in 3.6 months.

$11.15M
Annual cost of drift
Contacts, review time and revenue at risk

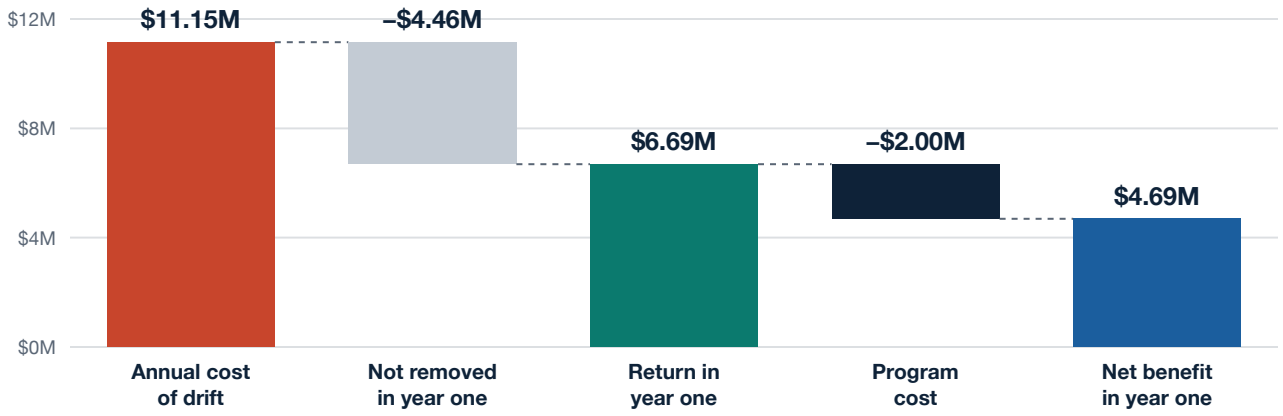
$6.69M
Return in year one
60% of the cost removed

$4.69M
Net benefit in year one
After a $2.0M program

3.6 mo
Payback
235% ROI in year one

EXHIBIT 1

From the annual cost of drift to net benefit in year one, expected case



Worked example: an illustrative large enterprise with 10 million assisted contacts a year, 40,000 published assets and $2 billion in recurring revenue. Not a customer result. The conservative and upside cases return a net benefit of $0.63M and $11.80M (section 4).

What the expected case means for an executive team

Cost of delay
Drift costs about $930,000 a month, or $30,500 a day, until it is corrected.

Scale
The cost equals 0.56% of recurring revenue. Avoidable contacts alone are 5% of a $135M assisted-support budget.

Sensitivity
Each percentage point of contacts caused by content is worth $1.35M a year.

Downside
In the conservative case, with no revenue effect counted, the program still returns more than it costs.

© 2026 Continuity Intelligence, Inc. | continuityintelligence.com

01

PART ONE / SIZING

How big is a large enterprise?

Public figures run from under two million to 170 million customer contacts a year. The worked example sits near the low end.

170M customer calls a year reported by one carrier

SECTION 1 / SIZING THE EXAMPLE

The worked example sits below every consumer-scale figure we found

It assumes 10 million assisted contacts a year: about five times the enterprise-software marker and a quarter of the smallest consumer marker.

EXHIBIT 2

Publicly reported customer contact volumes against the worked example

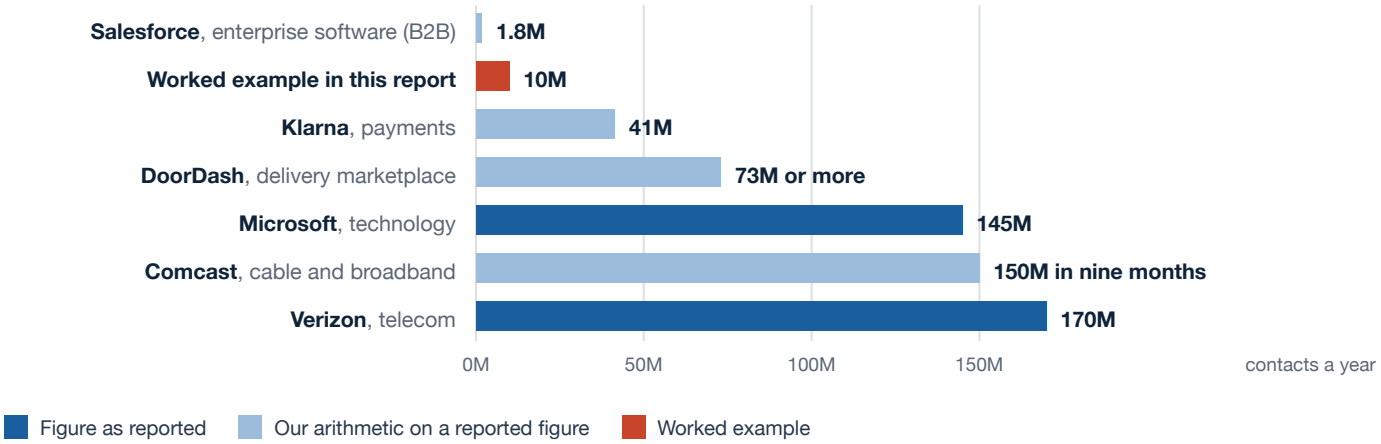


EXHIBIT 3

What each company reported

COMPANY	WHAT IT REPORTED	PER-YEAR READING	SOURCE AND CHECK
Verizon Telecom	About 170 million customer calls a year; 60,000 call agents	170M a year	Reuters, June 18, 2024 Opened, via The Star
Comcast Cable and broadband	30 million+ fewer service interactions year to date, a 20% fall in contact rate	About 150M in nine months, implied	CX Dive, Oct 22, 2026 Opened
Microsoft Technology	About 145 million customer interactions a year	145M a year	UC Today, July 10, 2025 Opened. Microsoft's own page no longer loads
DoorDash Delivery marketplace	Hundreds of thousands of requests for assistance a day; 37 million+ monthly consumers	73M or more a year	AWS case study Opened
Klarna Payments	AI assistant: 2.3 million conversations in its first month, two-thirds of all chats; 150 million consumers	About 41M chats a year	Klarna press release, Feb 27, 2024 Opened
Salesforce Enterprise software (B2B)	AI agent on its help site: 2 million+ conversations, October 2024 to November 2025; 4% went to a human	About 1.8M a year	Salesforce blog, Nov 6, 2025 Opened

Companies are named as public scale markers only. Nothing in this report describes their content, costs or results. Per-year readings marked About or Implied are our arithmetic.

© 2026 Continuity Intelligence, Inc. | continuityintelligence.com

SECTION 1 / SIZING THE EXAMPLE

What the avoidable-contact line is worth at each scale

Contact volume is the input that differs most between enterprises, by nearly two orders of magnitude in the public figures.

EXHIBIT 4

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 cost line 1 is scaled. Review time, revenue at risk and program cost do not rise in step with contact volume.

How to read your own scale

Enterprise software vendors

Contact volume is low, so line 1 alone ($1.35M at 5%) does not carry a $2.0M program. The case rests on review time, retention and risk, with a program sized to match.

Large consumer platforms

At 40 million contacts a year, line 1 alone is $27.0M at 5%, and each percentage point is worth $5.4M.

Fortune 50 service providers

At carrier scale, 170 million calls a year, each percentage point of calls caused by content is worth about $23M a year.

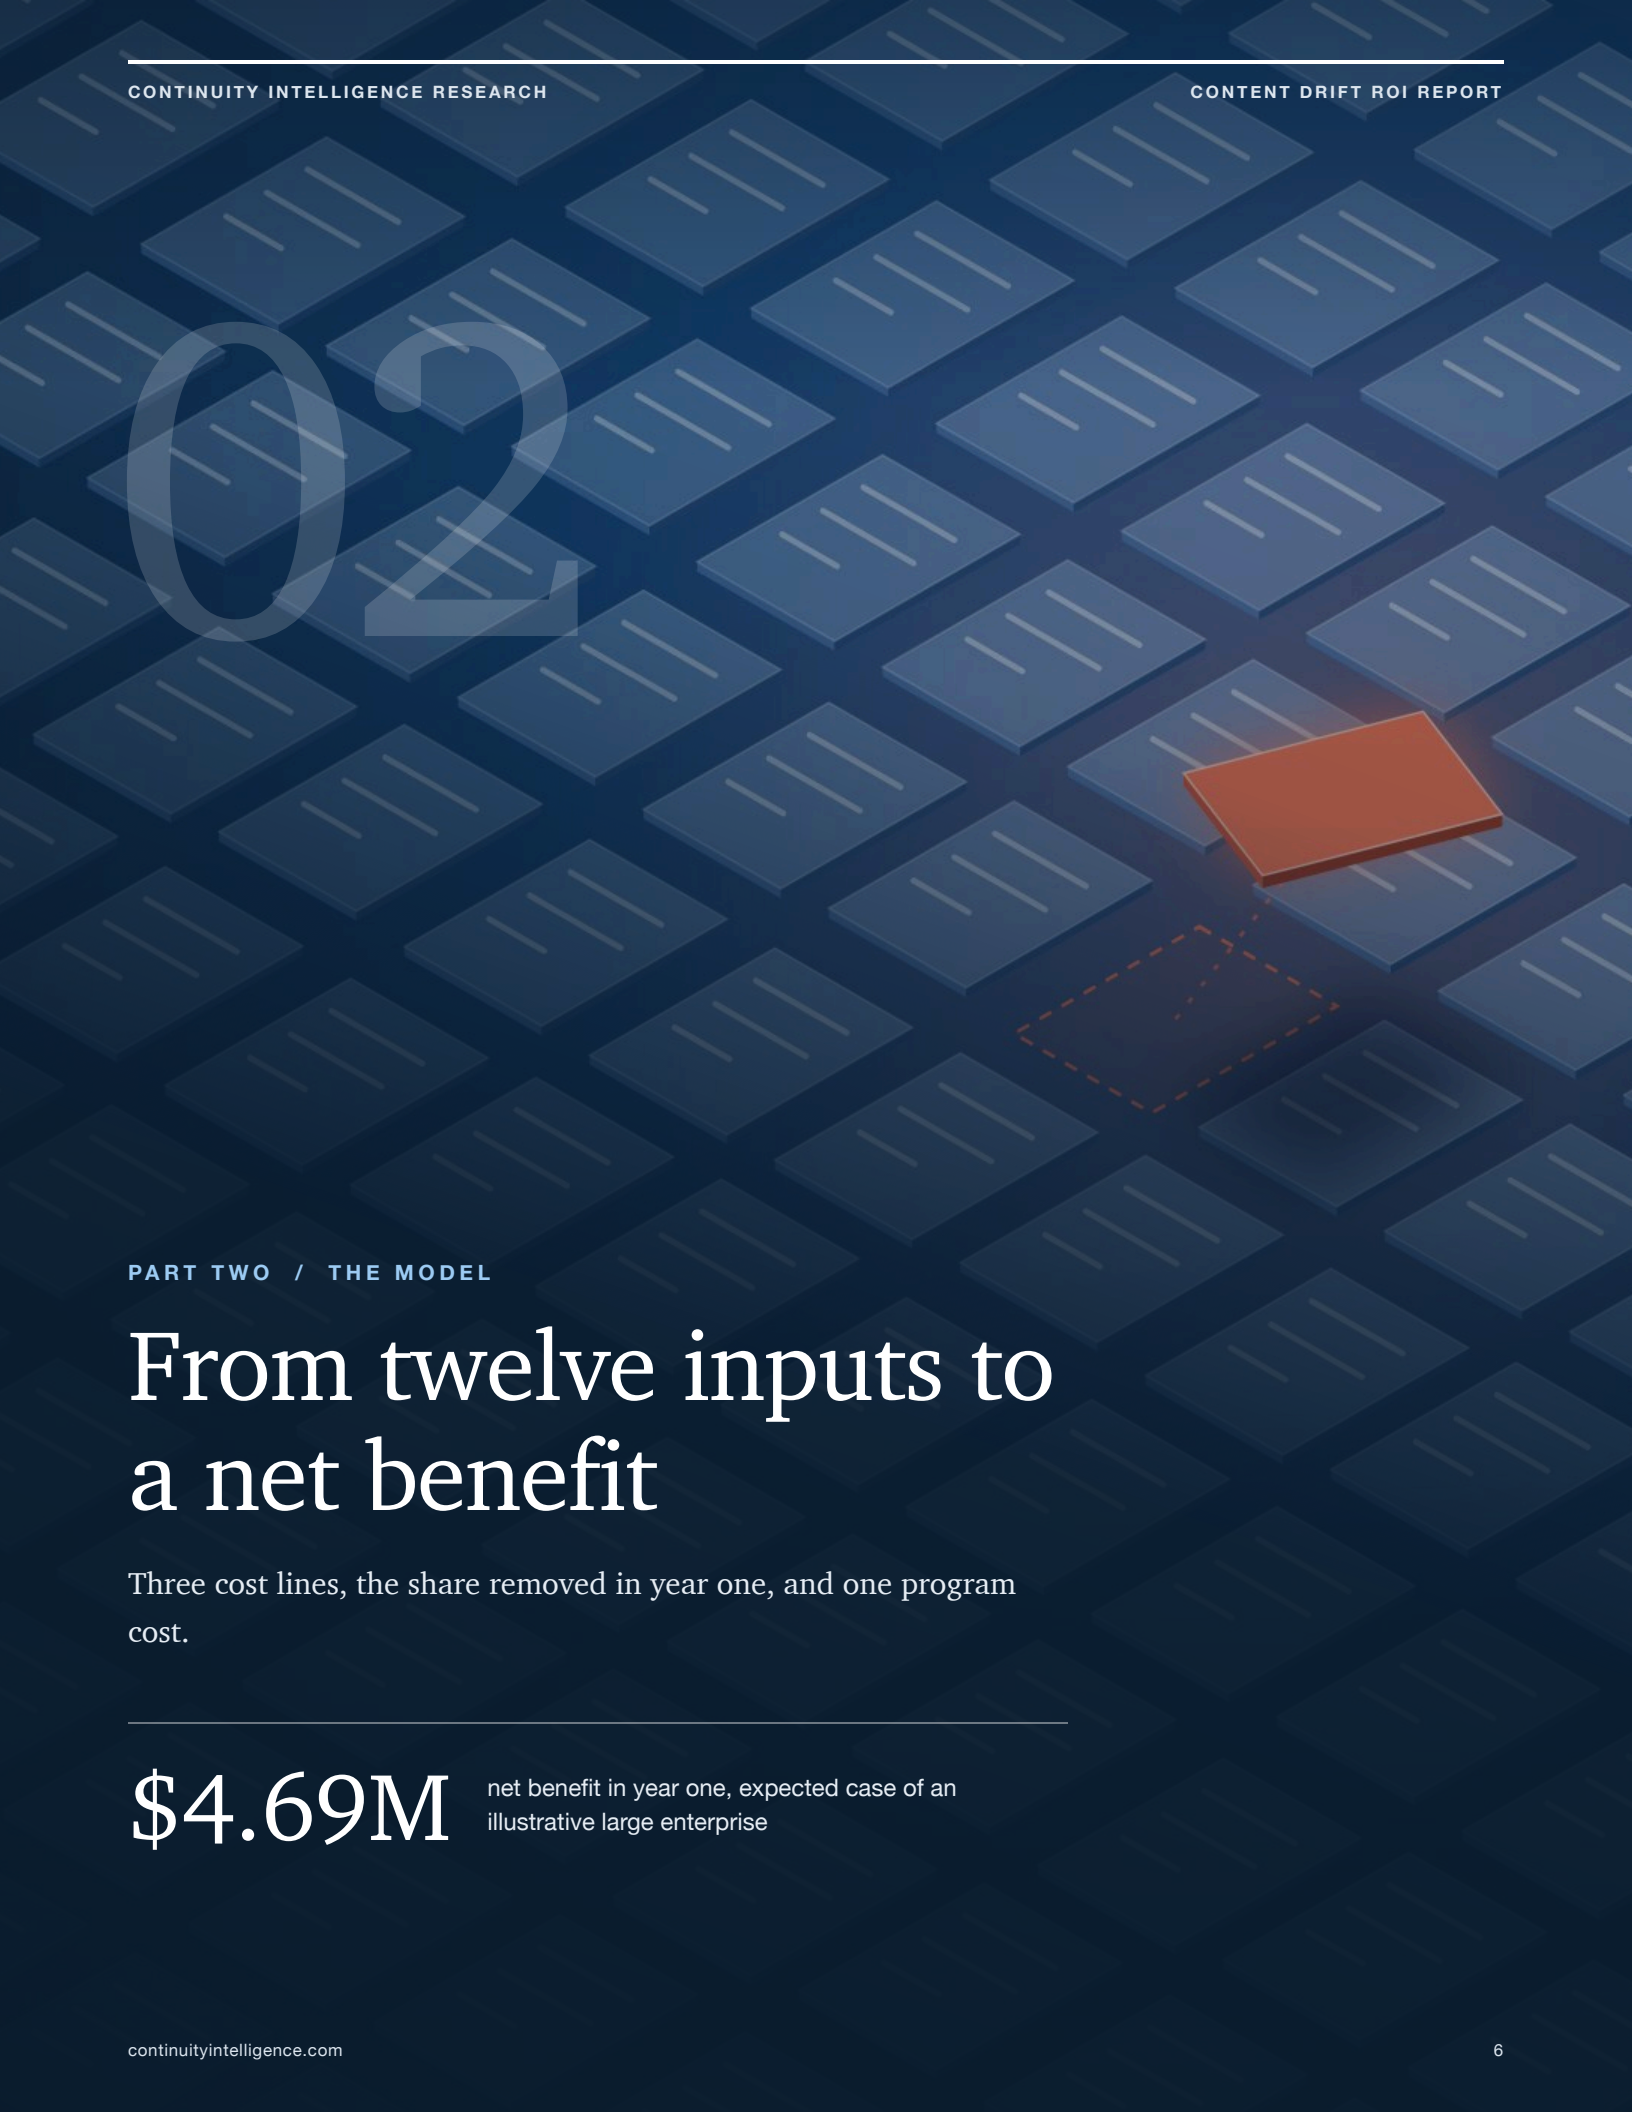
Contacts per customer

Klarna's figures work out to about one chat for every four consumers a year. Multiply your own customer base by your own rate before using any row above.

Arithmetic on public figures. These rows multiply a reported volume by an assumed share and the Gartner median cost. They are not findings about any named company, and the share of contacts caused by content still has to be measured.

© 2026 Continuity Intelligence, Inc. | continuityintelligence.com

02



PART TWO / THE MODEL

From twelve inputs to a net benefit

Three cost lines, the share removed in year one, and one program cost.

$4.69M net benefit in year one, expected case of an illustrative large enterprise

SECTION 2 / INPUTS

Twelve numbers drive the report

Input C matters most and has no public benchmark, so it must be measured.

EXHIBIT 5

How the report is built

STEP 1
Twelve inputs
A to L, from help desk, finance and a content scan.

→

STEP 2
Cost of drift
Contacts, review time and revenue at risk, lines 1 to 3.

→

STEP 3
Return
Cost of drift times the share removed in year one, K.

→

STEP 4
ROI and payback
Return less program cost L, over program cost.

ID	INPUT	EXAMPLE	WHERE THE CUSTOMER'S NUMBER COMES FROM	BASIS FOR THE EXAMPLE
A	**Assisted contacts per year**	10,000,000	Help desk reporting	ILLUSTRATIVE
B	**Cost per assisted contact**	$13.50	Support operations or finance	BENCHMARK Gartner median, February 2024
C	**Share of contacts caused by wrong or outdated content**	5% (range 3% to 8%)	Contact-reason tag, or ticket text search	ASSUMPTION
D	**Published assets in scope**	40,000	Continuity Intelligence scan	ILLUSTRATIVE
E	**Drift rate**	10%	Continuity Intelligence scan	CI DATA Median across 41 libraries
F	**Hours a year spent on manual content review**	20,000	Content and enablement time logs	ASSUMPTION
G	**Loaded hourly cost**	$60	Finance	ASSUMPTION
H	**Annual recurring revenue**	$2,000,000,000	Finance	ILLUSTRATIVE
I	**Annual revenue churn**	8%	Finance	ILLUSTRATIVE
J	**Share of churn where a content failure contributed**	2% (range 0% to 4%)	NPS comment coding and exposure analysis	ASSUMPTION
K	**Share of drift cost removed in year one**	60% (range 50% to 75%)	Before-and-after measurement	ASSUMPTION
L	**Annual program cost, software plus internal time**	$2,000,000	Quote plus staff time	ILLUSTRATIVE

This report is a template. Replace each input with the customer's own number and recompute the result with the formulas shown. The worked example is an illustrative large enterprise with 10 million assisted contacts a year, 40,000 published assets and $2 billion in recurring revenue. It is not a customer result.

© 2026 Continuity Intelligence, Inc. | continuityintelligence.com

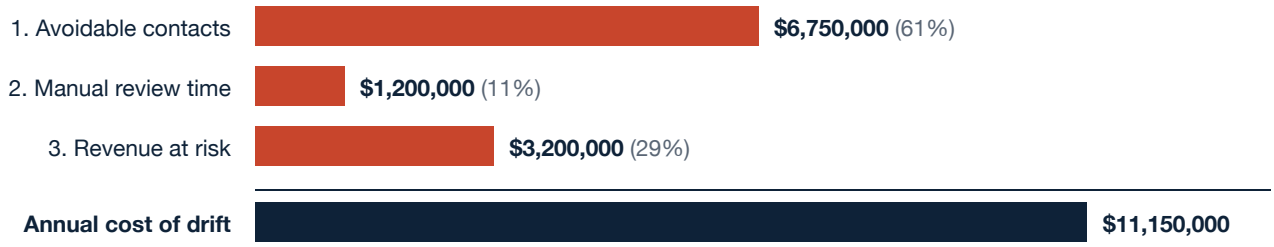
SECTION 3 / COST OF DRIFT

Three lines carry a dollar figure. Two are reported as counts.

Only the dollar lines enter the ROI.

EXHIBIT 6

Annual cost of drift by line, expected case



Worked example, illustrative large enterprise. Shares are of the annual cost of drift.

COST LINE	FORMULA	WORKED EXAMPLE	EVIDENCE
1. Avoidable contacts	A x C x B	10,000,000 x 5% x $13.50 = $6,750,000	Measured once the contact tag is live
2. Manual review time	F x G	20,000 x $60 = $1,200,000	Measured from time logs
3. Revenue at risk	H x I x J	$2,000,000,000 x 8% x 2% = $3,200,000	Modeled until the NPS linkage is run
Annual cost of drift	**Lines 1 to 3**	**$11,150,000**	
4. Wrong AI answers (count)	AI conversations a year x share that draw on a drifted asset	4,000,000 x 10% = 400,000 answers at risk	Estimate. Assumes the assistant draws on assets in proportion to the library. Replace with the weekly answer audit
5. Legal and compliance exposure (count)	Open findings that touch price, policy, safety or regulation	From the scan	Listed by finding, not priced

Line 4 rests on a tested mechanism. In Continuity Intelligence's November 2026 test, Claude repeated the outdated fact in 30 of 30 answers when the page it was given had drifted, and in 0 of 30 once the page was corrected.

© 2026 Continuity Intelligence, Inc. | continuityintelligence.com

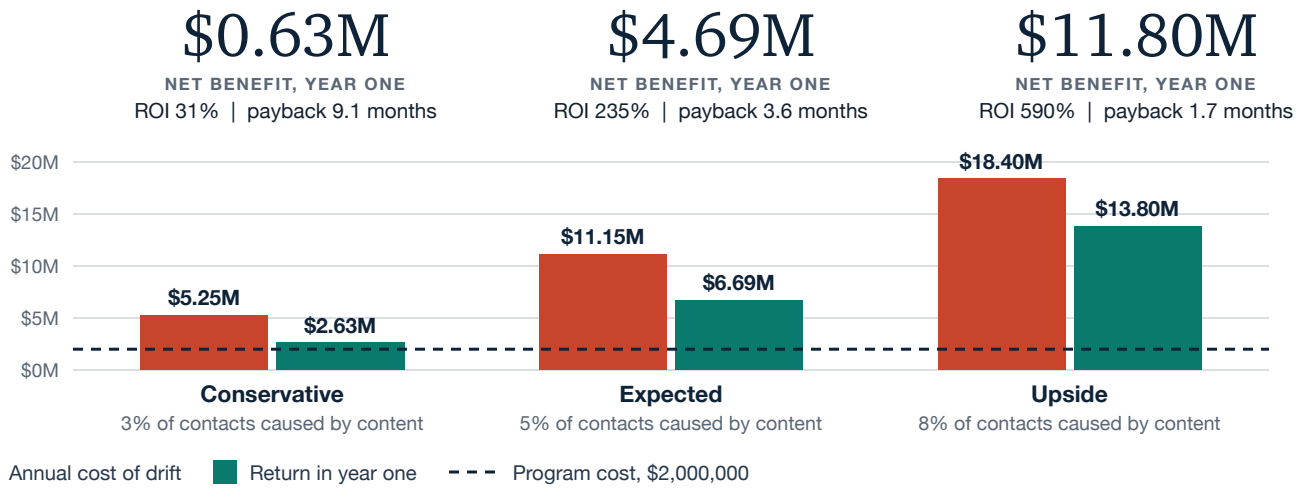
SECTION 4 / RETURN FROM CORRECTION

The measured share of content-caused contacts decides the return

The year-one return ranges from $2.6M to $13.8M across the three scenarios, against a $2.0M program.

EXHIBIT 7

Year-one cost of drift and return in three scenarios



LINE	CONSERVATIVE	EXPECTED	UPSIDE
C. Share of contacts caused by content	3%	5%	8%
1. Avoidable contacts	$4,050,000	$6,750,000	$10,800,000
2. Manual review time	$1,200,000	$1,200,000	$1,200,000
3. Revenue at risk (J)	Not counted	$3,200,000 (2%)	$6,400,000 (4%)
Annual cost of drift	**$5,250,000**	**$11,150,000**	**$18,400,000**
K. Share removed in year one	50%	60%	75%
Return = cost of drift x K	$2,625,000	$6,690,000	$13,800,000
L. Program cost	$2,000,000	$2,000,000	$2,000,000
Net benefit = return less L	**$625,000**	**$4,690,000**	**$11,800,000**
ROI = net benefit / L	31%	235%	590%
Payback = L / return x 12	9.1 months	3.6 months	1.7 months
Net benefit over three years, year-one figures held flat	$1,875,000	$14,070,000	$35,400,000

Break-even. On hard costs alone, at the conservative settings, break-even is about 207,000 content-caused contacts a year. That is 2.1% of the example's volume. Below that, the case rests on retention and risk, and the report should say so.

© 2026 Continuity Intelligence, Inc. | continuityintelligence.com

SECTION 5 / THE NPS LINK

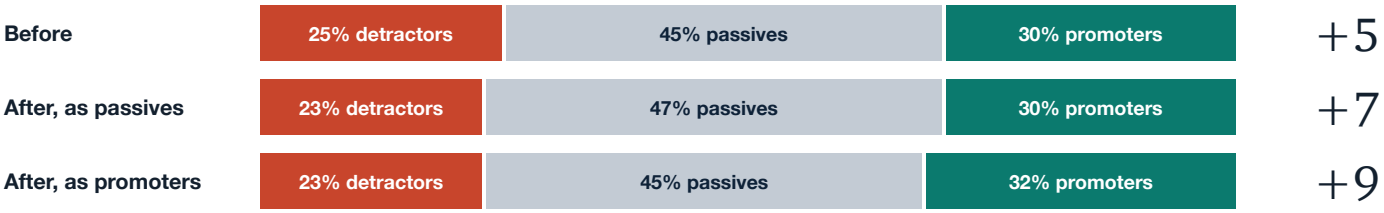
Every 1% of respondents moved from detractor to passive adds one NPS point

Two points if they become promoters. The work is to find how many detractors content created.

1. Code detractor and passive comments for content causes: “wrong”, “outdated”, “didn’t match what I saw”, “the article said”.
2. Count content-caused detractors as a share of all respondents. Call it d.
3. Recoverable NPS is between d points (they become passives) and 2d points (they become promoters).
4. Revenue retained = content-caused detractor accounts x (detractor churn rate less passive churn rate) x revenue per account.
5. After correction, track NPS monthly for customers who viewed a corrected article against those who did not.

EXHIBIT 8

Worked example: 2% of respondents are detractors because of wrong or outdated information



Survey base	20,000 responses: 30% promoters, 45% passives, 25% detractors. NPS = +5
Content-caused detractors	2% of respondents (400 people) cite wrong or outdated information
NPS after correction	Detractors fall to 23%. NPS = +7 if they become passives, +9 if promoters
Customers affected	2% of 2,000,000 customers = 40,000
Churn gap	Detractors churn at 15%, passives at 7%: an 8-point gap
Revenue retained	40,000 x 8% x $1,000 a year per customer = $3,200,000 a year

The example figures are assumptions. This is a second route to cost line 3, and in this example both routes give $3,200,000. Use this route or input J, not both.

Bain’s research supports the direction of the link, not a fixed conversion rate. It reports that an industry’s Net Promoter leader outgrows competitors by more than two times, and that promoter lifetime value runs 3 to 8 times that of detractors. Bain also advises each company to measure its own figures.

© 2026 Continuity Intelligence, Inc. | continuityintelligence.com

03

PART THREE / THE PROOF

Turning assumptions into measured numbers

Twelve ways to prove the return, and a 90-day plan that issues two reports.

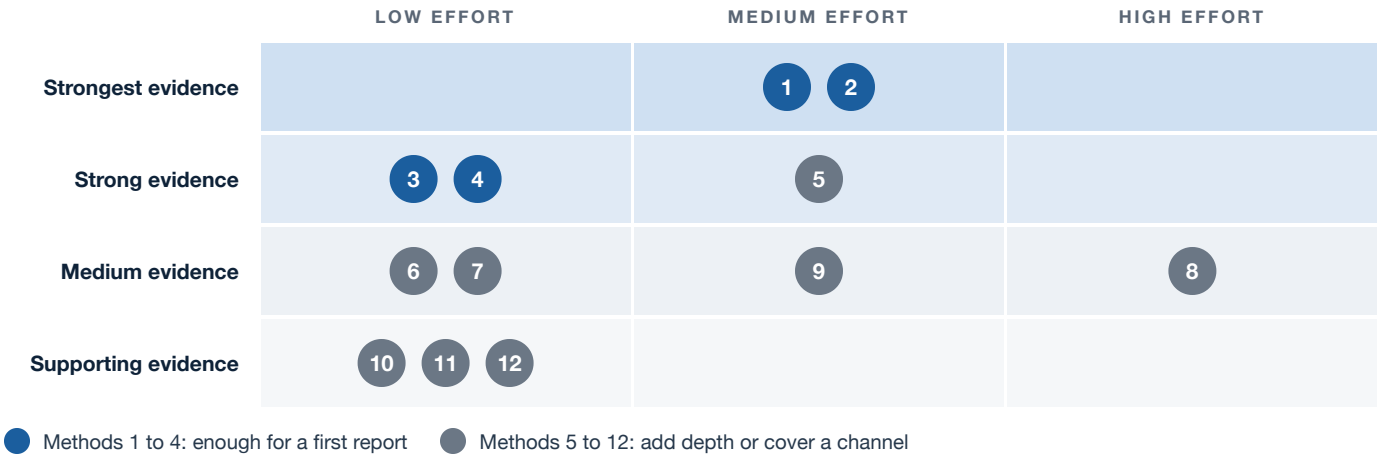
90 days from assumptions to a measured return

SECTION 6 / HOW WE PROVE IT

Twelve ways to prove the return, strongest first

EXHIBIT 9

Proof methods by effort and strength of evidence



Method 2 delays only low-severity fixes. Critical findings are corrected at once and measured with method 1. Method 8 rates medium because other factors differ between the two groups.

#	METHOD	WHAT IT PROVES
1	**Corrected against uncorrected articles.** Contacts per 1,000 views, 30 days before and after, with a control group	The fix caused the drop
2	**Staged correction.** Fix half of the low-severity findings first, the rest 30 days later	Same, with a cleaner control
3	**Contact-reason tag.** “content was wrong or outdated”, with the article link	Count and cost of content-caused contacts
4	**AI answer audit.** A fixed question set run weekly against the assistant	Wrong answers removed from the AI channel
5	**Article-linked contacts.** Tickets opened within 24 hours of viewing a drifted article	Contacts that the tag misses
6	**Ticket text search.** 12 months of tickets and chats searched for old plan names, prices and retired features	An immediate, backward-looking baseline
7	**Survey comment coding.** NPS and CSAT comments classified for wrong or outdated information	NPS points at stake
8	**NPS by exposure.** Scores for customers who viewed drifted articles against those who did not	Size of the loyalty gap
9	**Agent time study.** Handle time and reopen rate on drifted topics against other topics	Cost inside the support team
10	**Time to correct.** Days from a source change to the content fix	Exposure window closing
11	**Audit hours avoided.** Time logged on manual review before and after	Labor saving
12	**Risk register.** Open findings that touch price, policy, safety or regulation	Exposure that is not priced

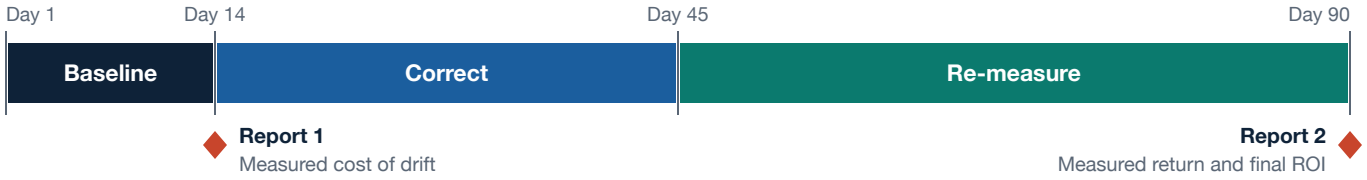
© 2026 Continuity Intelligence, Inc. | continuityintelligence.com

SECTION 7 / THE 90-DAY PROOF PLAN

Ninety days turns the assumptions in this report into measured numbers

EXHIBIT 10

Three phases, two reports



PHASE	DAYS	WHAT HAPPENS	OUTPUT
Baseline	1 to 14	Scan the content. Add the contact-reason tag. Search 12 months of tickets for outdated terms. Run the AI answer audit. Code the last NPS cycle's comments	Drift rate, measured value of C, AI accuracy score, content-caused detractor share
Correct	15 to 45	Fix critical and high-traffic findings first. Retire superseded versions. Stage low-severity fixes in two groups	Findings cleared, time to correct, list of corrected articles with dates
Re-measure	46 to 90	Compare corrected and uncorrected articles. Re-run the AI audit weekly. Read tagged contacts and the next NPS cycle	Measured K, contacts avoided, NPS movement, final ROI

The report is issued twice: at day 14 with a measured cost of drift, and at day 90 with a measured return.

What the executive sponsor decides at each report

At day 14

Whether the measured cost of drift justifies the program, using the customer's own value of input C in place of the 5% assumed here.

At day 90

Whether to continue, on the measured share of cost removed (K), contacts avoided and NPS movement.

© 2026 Continuity Intelligence, Inc. | continuityintelligence.com

ASSUMPTIONS AND SOURCES

What is assumed and what is sourced

Every figure in the worked example is either an input the customer replaces or arithmetic on those inputs.

Assumptions to replace with measured values

Inputs C, F, G, J, K and L are assumptions made for the example. So are the AI conversation volume in line 4 and the survey mix and churn rates in the NPS example. Inputs A, D, H and I set the size of the illustrative enterprise. Input A is sized against the public figures in section 1. No public benchmark gives a reliable figure for C. Vendor figures circulate, but we found none that traces to a documented study.

The three-year figure multiplies the year-one net benefit by three. It assumes no growth in contact volume, no further gain in the share of cost removed and a flat program cost.

Sources

SOURCE	USED FOR	HOW IT WAS CHECKED
Continuity Intelligence, The State of Content Drift 2026	Input E: median drift rate of 10% across 41 libraries	CI scan data, April to October 2026
Continuity Intelligence, Claude drift test, November 17, 2026	Line 4: 30 of 30 answers outdated from drifted pages, 30 of 30 correct after correction	10 pages from CI scans, three Claude models. Prompts and answers held by CI
Gartner, cost per contact benchmark, February 2024	Input B: $13.50 median per assisted contact	Read in Gladly's cost-per-contact glossary, which cites the benchmark. Gartner's own page blocks automated access. Use the customer's own figure where it exists
Public scale markers: Salesforce, Klarna, DoorDash, Microsoft, Comcast, Verizon	Section 1: customer contact volumes used to size input A	Each page opened. Details and dates in Exhibit 3
Bain & Company, How Net Promoter Score Relates to Growth and The Numbers behind the Net Promoter System	The NPS link	Opened

The companion white paper, *The Cost of Content Drift*, gives the full evidence and the test method.

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

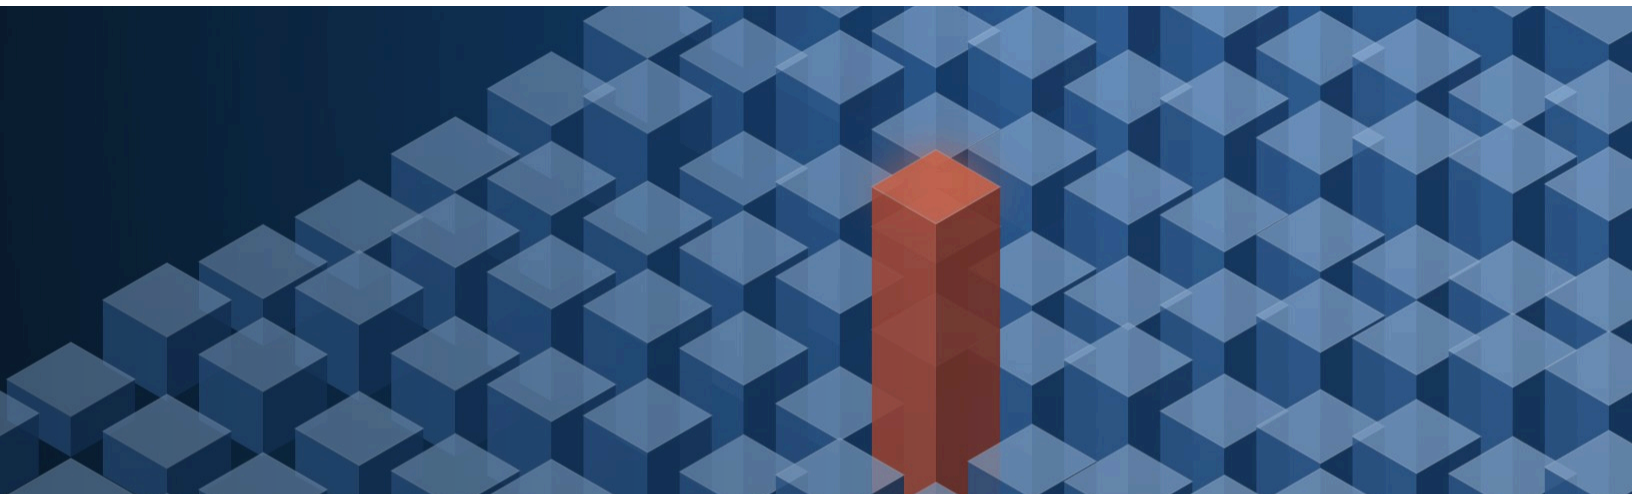
Request a free drift report →

COMPANION DOCUMENTS

White paper: The Cost of Content Drift
The evidence behind each cost line, the Claude drift test and the five measurements that prove ROI.

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.

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