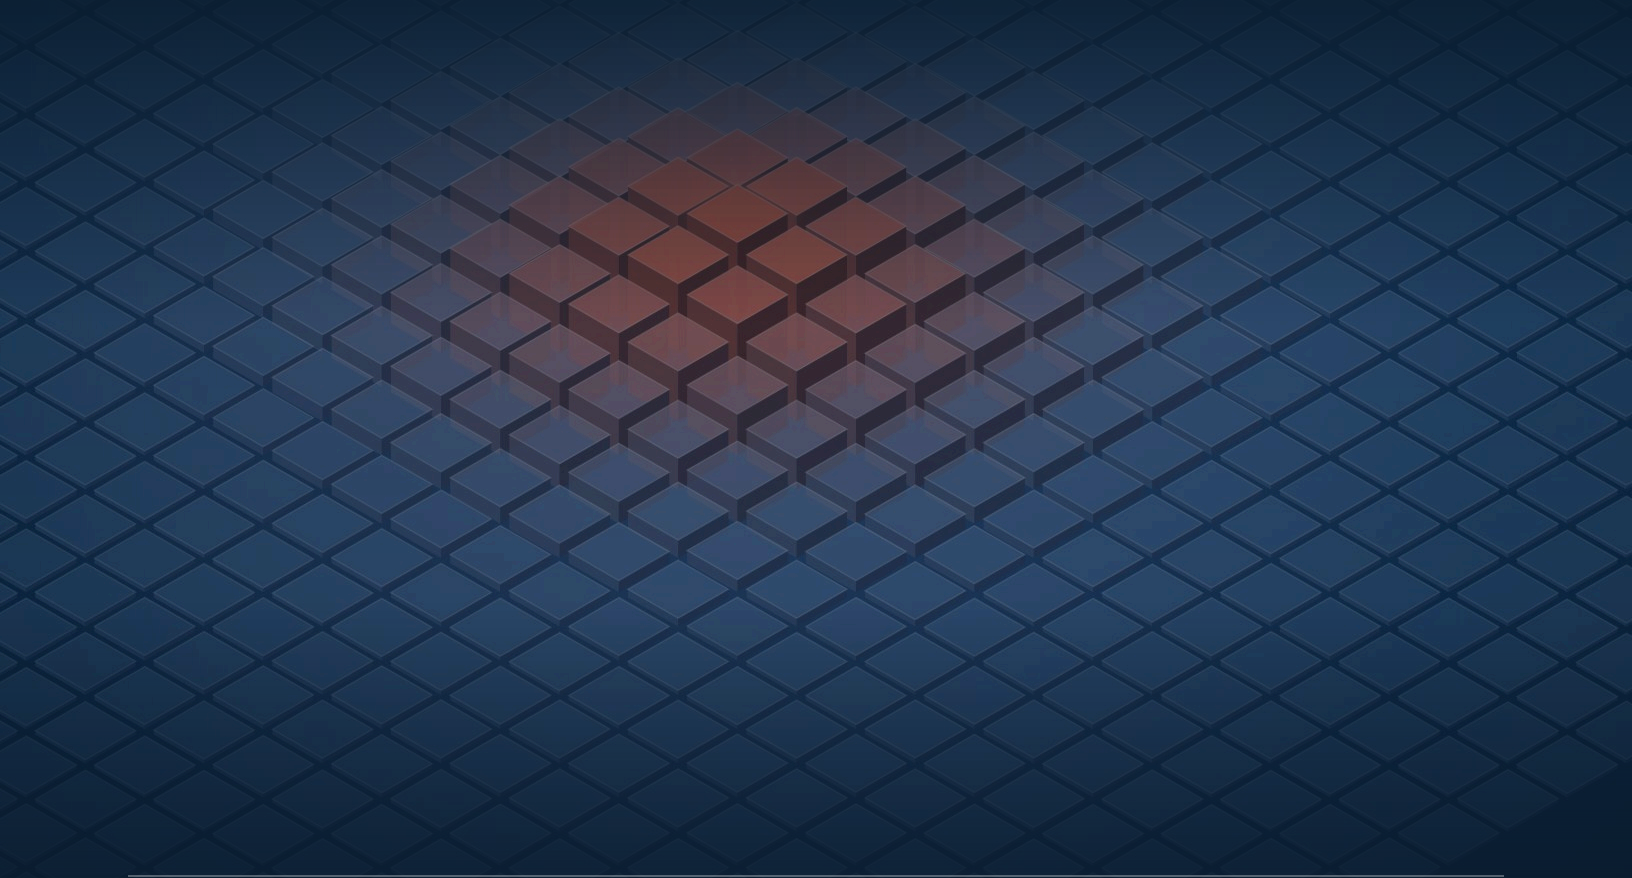
CASE STUDY

When the Help Center Drifts, the AI Assistant Follows

A B2B software company's customer content gained 2.3 outdated facts a week. Given seven of the ones open today, Claude passed every one on to the customer.



2.3

new outdated facts a week in steady state, across 32 weekly scans

0 of 21

Claude answers were correct when grounded in the drifted articles

21 of 21

were correct once the one outdated fact in each article was corrected

© 2026 CONTINUITY INTELLIGENCE, INC.

AT A GLANCE

The help center drifted every week, and the AI assistant followed

A B2B software company's customer content gained 2.3 outdated facts a week. Given seven of the ones open today, Claude passed every one on to the customer: 0 of 21 answers were correct. With the articles corrected, 21 of 21 were.

137
sources in scope
Help center, marketing site, product videos and partner academy

32
weekly scans
By Continuity Intelligence, April 8 to November 16, 2026

47
verified findings open
On November 16, 2026

20 days
median time to clear a finding
Under weekly monitoring

THE CHALLENGE

The product changed faster than the content that describes it. In eight months the company moved customers to a new version of its product, changed what each plan includes, changed how usage is billed, renamed controls and retired options.

Each change had to reach every place the old fact lived: help articles, marketing pages, tutorial videos and partner training. Some did not. A support article still promised a one-hour chat reply after the pricing page set four hours for most customers. Tutorial videos still taught a billing rule the company had replaced. A setup guide still pointed to an icon that now does something else.

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

ORGANIZATION
B2B software company, name withheld

NEW DRIFT
2.3 a week
findings in steady state

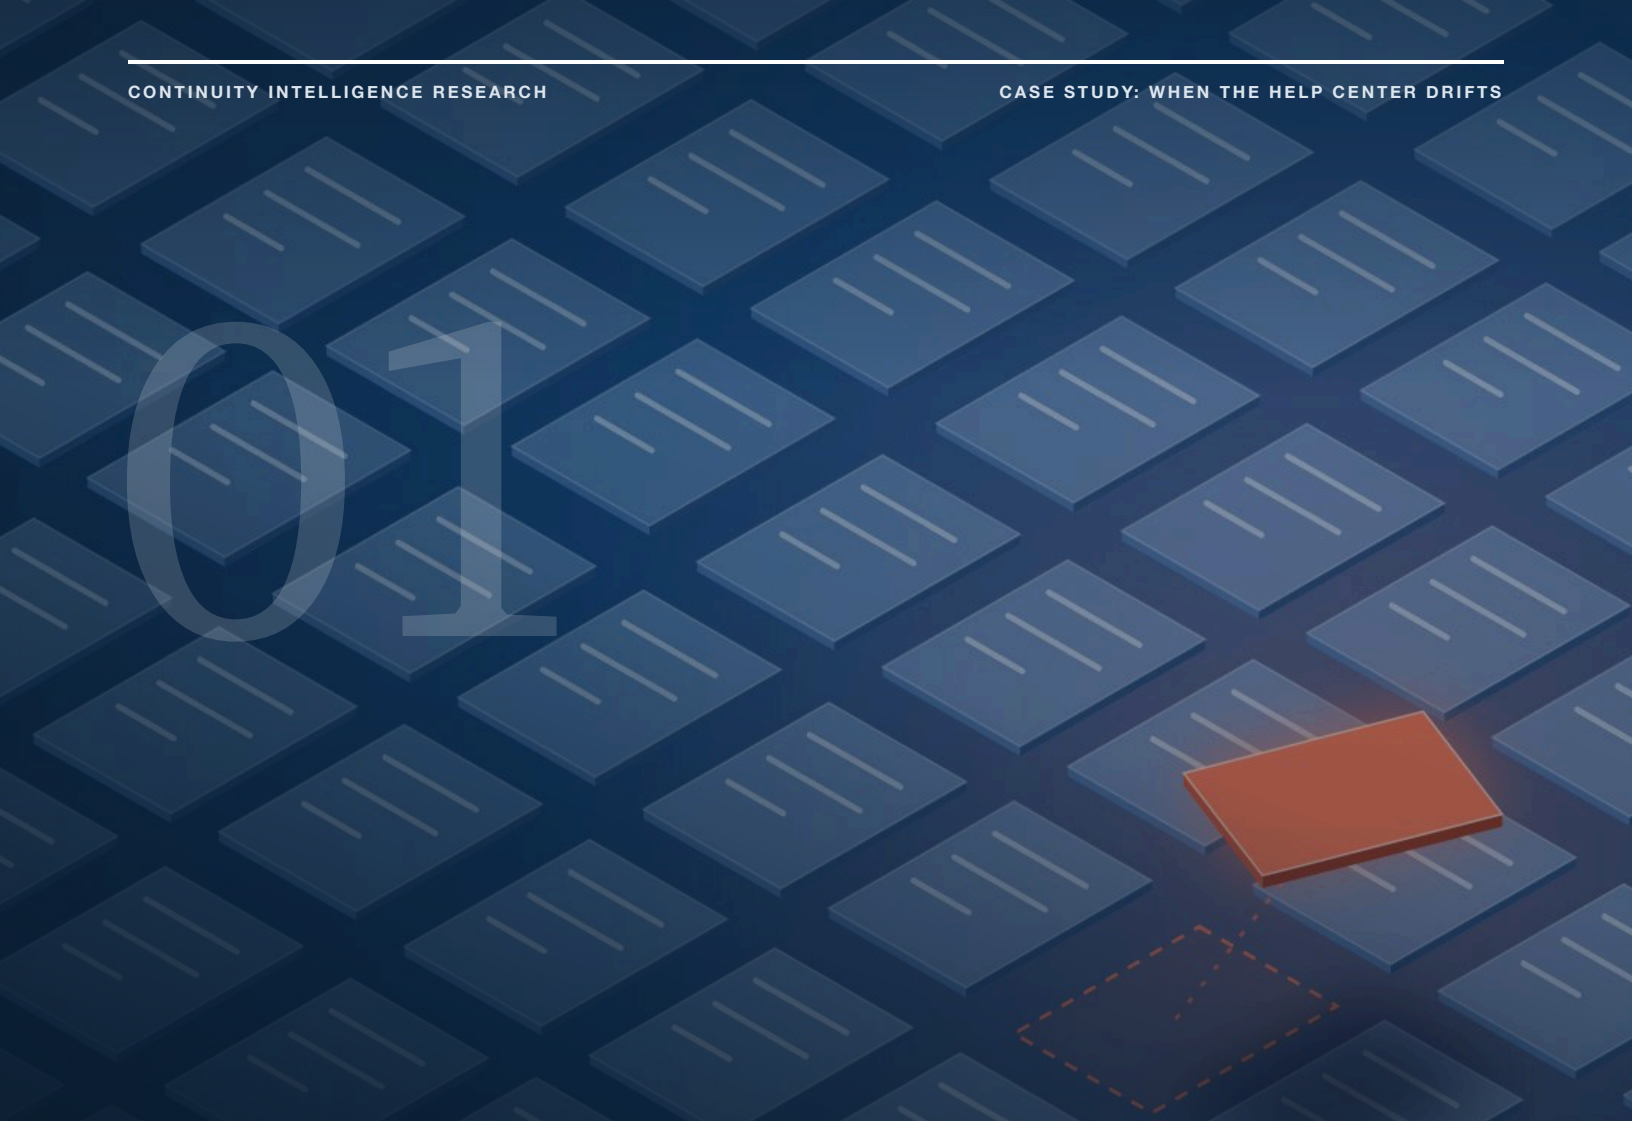
CLAUDE ANSWERS, DRIFTED ARTICLES
0 of 21
correct

CLAUDE ANSWERS, CORRECTED ARTICLES
21 of 21
correct

TICKET AND NPS IMPACT
Not yet measured. This case study does not estimate it. The last two sections show how to measure it.

© 2026 Continuity Intelligence, Inc. | continuityintelligence.com

01



PART ONE / THE EVIDENCE

Thirty-two scans and one test

What weekly monitoring found in one company's content, and what Claude did with seven of the open findings.

47 verified findings open on November 16, 2026

SECTION 1 / WHAT THE SCANS FOUND

Drift arrived every week and never stopped

Across 32 weekly scans the count of open findings ranged from 16 to 47 and never reached zero.

EXHIBIT 1

Open drift findings at each of 32 weekly scans



Source: Continuity Intelligence scan data, April 8 to November 16, 2026. Raw open-finding counts per scan.

51
new findings in the 22 steady-state weeks from May 3 to September 28, or 2.3 a week.

84
findings cleared under weekly monitoring, in a median of 20 days.

about 30
findings would build up between reviews under a quarterly audit at the same arrival rate.

Open findings fell from 38 to 16 between mid-June and early July, when 40 pages were revised in a single week. They rose again as the product kept changing and as more content came into scope.

© 2026 Continuity Intelligence, Inc. | continuityintelligence.com

SECTION 2 / THE CLAUDE TEST

Every answer from a drifted article was wrong

Every answer from the corrected article was right. The test used seven findings that were open on November 16, 2026.

EXHIBIT 2

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

WHAT CLAUDE WAS GIVEN	HAIKU 5.5	SONNET 5.5	OPUS 5.5	ALL ANSWERS
Article as published	●●●●●●●	●●●●●●●	●●●●●●●	0 of 21
Article with the fact corrected	●●●●●●●	●●●●●●●	●●●●●●●	21 of 21

● Answer gave the outdated fact ● Answer gave the current fact

For each finding we gave Claude the article text as its knowledge base, the way a support assistant is grounded in help content, and asked one customer question. We ran Claude Haiku 5.5, Sonnet 5.5 and Opus 5.5 on November 17, 2026. Then we corrected the one outdated fact and asked again.

EXHIBIT 3

One question, one changed fact: chat response time

GROUNDED IN THE ARTICLE AS PUBLISHED

HELP CONTENT GIVEN TO CLAUDE

Chat: Within 1 hour during open hours.

CUSTOMER

I'm on the Business plan. How quickly will I get a reply on chat?

CLAUDE SONNET 5.5

Chat replies come within 1 hour during open hours. The help content I have doesn't list a different response time for the Business plan, so I can't confirm a faster or slower time for your plan.

✕ OUTDATED The pricing page sets four hours for this plan.

GROUNDED IN THE CORRECTED ARTICLE

HELP CONTENT GIVEN TO CLAUDE

Business plans get Live Chat with a reply within 4 hours during business hours. Enterprise plans get a reply within 1 hour.

CUSTOMER

I'm on the Business plan. How quickly will I get a reply on chat?

CLAUDE SONNET 5.5

On the Business plan you get Live Chat, with a reply within 4 hours during business hours.

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

Claude Sonnet 5.5's answers in full. Article excerpts are shortened.

© 2026 Continuity Intelligence, Inc. | continuityintelligence.com

SECTION 2 / THE CLAUDE TEST

The seven questions

Claude Sonnet 5.5's answers. Product names are masked.

	TOPIC	THE ARTICLE SAYS	THE CURRENT FACT	ANSWER FROM THE ARTICLE	ANSWER AFTER CORRECTION
1	**Chat response time**	Within 1 hour	4 hours on the Business plan, 1 hour on Enterprise	✕ "Chat replies come within 1 hour during open hours"	✓ "a reply within 4 hours during business hours"
2	**Plan needed for an integration**	Starter or higher	Pro or higher	✕ "your Starter plan qualifies"	✓ "requires a Pro plan or higher, so it isn't available on the Starter plan"
3	**A renamed control**	The heart icon adds an item to the team list	The bookmark icon does. The heart saves to personal favorites	✕ "Click the heart icon"	✓ "Click the bookmark icon"
4	**How usage is billed**	Every clip counts unless deleted before the cycle ends	Only downloaded minutes count	✕ "an undownloaded clip counts unless you delete it before the cycle ends"	✓ "a clip you generate but never download doesn't count against your plan"
5	**Character limit**	1,000 characters	5,000 characters per take	✕ "the character limit increases to 1000"	✓ "A single take in [product] can be up to 5,000 characters"
6	**File format by plan**	MP3, WAV or OGG for everyone	WAV and OGG need Business or Enterprise	✕ "Yes, when you download in [product] you can choose .wav"	✓ "WAV downloads require a Business or Enterprise plan"
7	**A retired option**	Still listed as available	Retired in May 2026	✕ "Yes, [name] is listed"	✓ "I don't see ... [name] in our current list"

> *No answer warned that the article might be out of date. Six of the 21 added a note that the article gave no further detail, and still stated the old fact.*

Each wrong answer has a price

RESPONSE TIME

A Business customer who is told one hour waits four and complains.

PLAN ENTITLEMENT

A Starter customer tries an integration that is not on the plan and opens a ticket.

BILLING RULE

A customer deletes clips to protect an allowance that deletion no longer affects.

© 2026 Continuity Intelligence, Inc. | continuityintelligence.com

02



PART TWO / THE STAKES

What it costs, and what the fix returns

What has been measured, what has not, and how to measure the rest in your own data.

21 of 21 answers correct once the articles were corrected

SECTION 3 / THE COST

What has been measured is exposure

The dollar cost for this company has not been measured, and this case study does not invent one. What has been measured is how many outdated facts are live, and for how long.

QUANTITY	VALUE	STATUS
Outdated facts live on November 16, 2026	**47**	MEASURED
New outdated facts per week	**2.3**	MEASURED
Time a finding stays live under weekly monitoring	**Median 20 days**	MEASURED
Time a finding waits before a quarterly audit sees it	**45 days on average**	CALCULATED Half of a 90-day cycle, if findings arrive evenly
AI answers wrong when grounded in an open finding	**21 of 21**	MEASURED IN THE TEST
Support contacts caused by these findings		TO BE MEASURED Ticket search and contact tag
NPS detractors who cite wrong or outdated information		TO BE MEASURED Survey comment coding

EXHIBIT 4

Time an outdated fact stays live



At 2.3 new findings a week, that is about 120 findings a year, each live 25 fewer days: roughly 3,000 fewer finding-days of exposure a year.

Against a quarterly audit, weekly monitoring shortens the time an outdated fact stays live from at least 45 days to 20.

Once contacts are counted, the conversion is simple. At Gartner's median of $13.50 per assisted contact, every 1,000 content-caused contacts a year cost $13,500.

© 2026 Continuity Intelligence, Inc. | continuityintelligence.com

SECTION 4 / THE FIX AND THE RETURN

The content was the only variable

Correcting the content changed the AI's answers from 0 of 21 right to 21 of 21, with no change to the model or the prompt.

0 of 21
answers correct from the articles as published

→ **21 of 21**
answers correct once one fact in each article was corrected

Each finding carries its evidence and the corrected wording, checked against the live product. Continuity Intelligence's agents autofix the content where they are connected to the system that holds it.

RETURN	RESULT	STATUS
Findings cleared under weekly monitoring	**84, in a median of 20 days**	MEASURED April 8 to October 25, 2026
AI answers correct after correction	**21 of 21, up from 0 of 21**	MEASURED IN THE TEST
Support contacts avoided		TO BE MEASURED Compare corrected and uncorrected articles for 30 days
NPS movement		TO BE MEASURED Next survey cycle, by exposure to corrected articles

The companion *Content Drift ROI Report* turns the last two rows into a dollar return once the counts exist. Its worked example, an illustrative large enterprise, is not this company.

© 2026 Continuity Intelligence, Inc. | continuityintelligence.com

SECTION 5 / MEASURE IT IN YOUR OWN DATA

A first count of content-caused contacts takes days, not months

The seven findings map directly to ticket searches.

FINDING	SEARCH TICKETS AND CHATS FOR	COUNT
Chat response time	“1 hour”, “still waiting”, “response time”, from Business plan accounts	Complaints about slower replies than promised
Plan needed for an integration	The integration’s name with “Starter”, “not available” or “upgrade”	Contacts from customers on the lower plan
A renamed control	“heart”, “bookmark”, “can’t add”, “team list”	How-to contacts on the step
How usage is billed	“deleted”, “clip count”, “charged”, “minutes”	Billing questions and disputes
Character limit	“1000”, “character limit”, “too long”	How-to contacts on script length
File format by plan	“WAV”, “OGG”, “download format”, from Pro plan accounts	Contacts and upgrade friction
A retired option	The option’s name with “missing” or “can’t find”	Contacts since the retirement date

Five steps

1. Run the seven searches over the last 12 months. That gives a first value for the share of contacts that content causes.
2. Add the contact reason “content was wrong or outdated” so the count continues without searching.
3. Code the last NPS and CSAT comments for the same seven topics, and for the words “wrong”, “outdated” and “didn’t match”.
4. Correct the seven articles. Compare contacts on those topics for 30 days before and after.
5. Enter the counts in the ROI report.

© 2026 Continuity Intelligence, Inc. | continuityintelligence.com

METHOD AND SOURCES

How this case was built

Scans

Continuity Intelligence scanned the company's public content weekly from April 8 to November 16, 2026. Each finding is a claim on a page that no longer matches the live product or a newer page from the same company, with the evidence recorded. New findings per week and time to clear are as published in The State of Content Drift 2026, which covers the scans through October 25.

Claude test

Seven findings open on November 16, 2026 were chosen for having one clear, checkable fact. The prompt asked Claude to act as the company's support assistant and answer from the article text supplied, in one to three sentences, with the date given as November 17, 2026. An answer counted as correct only if it gave the current fact. Each model was run once per condition. Prompts, article excerpts and all 42 answers are held by Continuity Intelligence.

Limits

The test shows what an assistant does when it is grounded in a drifted article. It does not measure how often customers ask these seven questions. Ticket and NPS effects are not measured in this case study.

Cost per contact

The $13.50 median for an assisted contact is from Gartner's cost per contact benchmark, February 2024, as cited in Gladly's cost-per-contact glossary. Gartner's own page blocks automated access.

Anonymity

The company's name and product names are withheld. Quotations are Claude's own words, with product names replaced by bracketed labels.

The companion white paper, *The Cost of Content Drift*, sets this case beside nine others and the wider evidence.

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

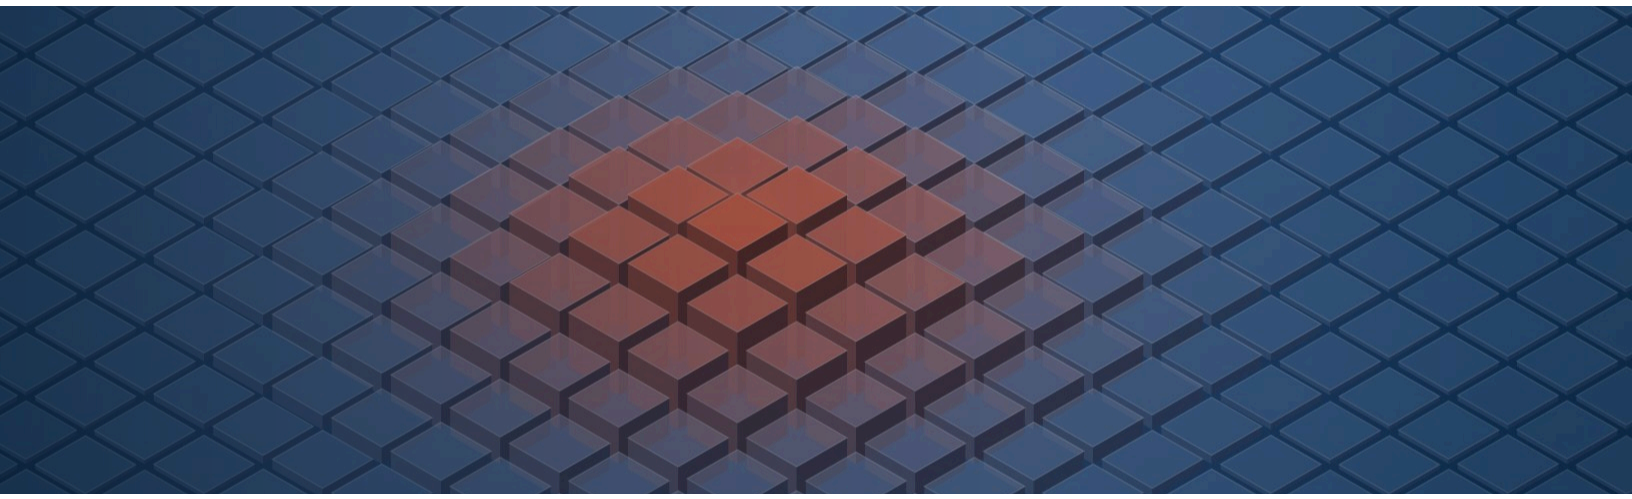
Request a free drift report →

COMPANION DOCUMENTS

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

Content Drift ROI Report
A standard report with every input and formula, a worked example for a large enterprise and 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