



FROM THE RESEARCH BENCH

The metric didn't die. The cohort did.

Software companies are said to be quietly retiring net retention disclosure. Across 104 companies and 438 annual filings, the rate did not move.

Sinan Isoglu, MBA (Quantic)

Commercial growth leader, lecturer and doctoral researcher

PUBLISHED

18 August 2026

UPDATED

27 August 2026

READING TIME

7 min

WORDS

1,534

<https://isoglu.com/insights/journal/the-metric-didnt-die-the-cohort-did/>

Management summary	2
What customer cohort dynamics did the longitudinal disclosure audit measure?	3
Why did public software companies suddenly stop reporting dollar retention rates?	4
How should analysts counter the standard argument that metric omission is benign?	5
What specific causal claims does this empirical disclosure study not make?	6
How can commercial finance build resilient cohort retention metrics?	7
Where are the empirical boundaries of cohort metric decay research?	7
References	8

MANAGEMENT SUMMARY

A widely repeated claim holds that public software companies are quietly retiring net dollar retention disclosure as the number gets worse. The claim is specific, credible and supported by real cases, but it has never been counted. This piece counts it: 104 companies from a published benchmark cohort, 438 annual filings, four fiscal years, one coding rule. The share disclosing a retention metric did not fall. It went from 58.3% to 61.5%, and the share publishing both a net and a gross measure rose from 6 companies to 11. What did happen is that 30 of the original 134 companies stopped filing altogether, and 29 of those 30 were acquired rather than failed. Disclosures do vanish from a tracked list. Mostly because the companies do.

Keywords: Net revenue retention · Voluntary metric disclosure · Benchmark peer sets · SEC annual filings · Commercial decision-making

There is a claim going around that public software companies are quietly killing off net dollar retention disclosure. As the number got worse, the argument runs, the metric got harder to find: dropped outright, described in words instead of digits, or moved to a slower reporting cadence where the decline is less legible.

It is a good claim, and the instinct behind it is the right one: a number that gets worse is a number somebody would rather not print. That instinct is why evidence beats anecdote even when the anecdote is well chosen. It names mechanisms rather than gesturing, and it names companies. Workday told the SEC it “no longer intend[s] to include disclosure related to net revenue retention rates”. Fastly said its metrics were “no longer used by management”. Those are real, and I found both of them independently in the filings before I ever read the claim.

But the claim has never been counted. It circulates as a pattern illustrated by cases, not as a rate. So I counted it. What follows is a companion to an earlier piece on how retention is measured from the vendor’s side of the table: that one was about what the number means, this one is about whether it is there at all.


What customer cohort dynamics did the longitudinal disclosure audit measure?

I took a published benchmark cohort: the 135 companies in the appendix of a 2024 study of how public SaaS companies report retention, chosen because it is one of the few such lists that names its constituents and can therefore be re-read by someone else. Then I pulled every annual report those companies have filed with the SEC from fiscal 2022 onward and read all of them with one script and one rule.

That is 438 annual filings across 104 companies, each one checked for a named retention or expansion metric with a stated definition. One rule applied to every filing, which matters more than it sounds: comparing my count to somebody else’s count would measure the two of us as much as it measures the companies.

The result is a flat line.

Figure 1 The disclosure rate over four fiscal years



Author's own replication. Frame: the 135 companies of Appendix A in Ordway Labs, 'Dollar-Based Net Retention Rate: How Public SaaS Companies Report' (published July 2024, analysis performed December 2023), of which 134 rows parsed from the report's PDF. 438 annual reports (10-K, 20-F, 40-F) read from SEC EDGAR at 18 August 2026 for the 104 baseline companies still filing. The series is the balanced panel of the 96 companies with a filing in each of fiscal 2022–2025; entrants and leavers are excluded rather than allowed to move the denominator. A company counts as disclosing where a named net or gross retention metric appears in its annual report. Frame and codebook were fixed and dated before any filing was read. Every row carries its accession number and the sha256 of the source document, so each value can be recomputed. Single coder; a second pass is owed. Corpus: 80 Reference/Source library/files/retention-replication-2026-08.

Fifty-eight per cent to sixty-two per cent, wobbling by a point or two a year. Over the exact window the euthanasia claim covers, fiscal 2022 to fiscal 2024, the line is 58.3, 60.4, 59.4. Nothing is being retired at cohort level.

And the composition moved the wrong way for the thesis. The number of companies publishing both a net and a gross retention measure went from 6 to 11. A retreat from measurement does not typically involve adding a second measure.


Why did public software companies suddenly stop reporting dollar retention rates?

They are real. I can see them. They are just not what they look like.

Of the 134 companies in that original cohort, 30 no longer file an annual report at all. Splunk, VMware, Zendesk, Qualtrics, Avalara, New Relic, Smartsheet, HashiCorp, Squarespace, Zuora, Informatica, Jamf, Vimeo and seventeen others.

If you track a list of software companies and their retention disclosures across these four years, disclosures genuinely do vanish from your dataset. Roughly a fifth of them. Not because anybody went quiet.

Figure 2 How the thirty that stopped filing left



Each company's own SEC filings, read at 18 August 2026: the merger proxy (DEFM14A), the completion recorded in a final 8-K or 6-K, the Schedule 13E-3 for a going-private transaction, and the Chapter 11 petition, with the Form 25-NSE and Form 15 deregistrations where filed. One source URL per row. The 30 are the baseline companies with no annual report for a period ending on or after 1 January 2025; the count is a floor, because a company that filed and was acquired afterwards counts as present. Each route is assigned from the filing itself, never from press coverage. The single Chapter 11 is 2U, July 2024. Corpus: exits-2026-08-18.json.

Twenty-nine of the thirty were bought. One failed.


That matters, because it kills the obvious explanation. The intuitive story about a shrinking peer set is natural selection: the weak ones die, the survivors look good, the average flatters. That is not what happened here. The companies that left were, overwhelmingly, the ones somebody paid a premium to own. The single failure in the group is 2U, which filed for Chapter 11 in July 2024.

How should analysts counter the standard argument that metric omission is benign?

If a fifth of the population left, and I am reporting a disclosure rate among the ones who stayed, then my flat line has exactly the disease I am describing. If the leavers were disproportionately the companies that never disclosed, their departure lifts the survivors' average all on its own, and I would be showing you a composition effect and calling it stability.

The honest thing is to test it rather than to argue about it, so I pulled the final annual report of all 30 departed companies and coded them on the same rule.

Figure 3 Leavers and stayers, coded on one rule



The final annual report of all 30 departed companies, pulled from SEC EDGAR and coded on the same rule as the panel: 18 of 30 disclose, 60.0% (17 net, 1 both, 12 none). The survivor bar is the balanced panel at fiscal 2025, 59 of 96, 61.5%. This is a descriptive selection check for the four-year series. Single coder; a second pass is owed. Corpus: departed-coding-2026-08-18.json.

They disclosed at 60.0%; survivors disclosed at 61.5%. The point estimates are close in this sample, but without an uncertainty calculation I cannot call them statistically indistinguishable or conclude that departures do not bias the series. The comparison is a descriptive check, not a formal no-bias test. The flat line survives its own strongest objection only in that limited sense.

What specific causal claims does this empirical disclosure study not make?

The claim I started from is not wrong, and I want to be precise about that rather than generous.

It was advanced through named mechanisms and named cases, not as an aggregate statistic. Workday and Fastly did what it says they did. Treating a case-based thesis as though it had asserted a rate, and then refuting the rate, would be an argument with something nobody said.

Three limits, and each one bites:

This reads annual filings only. One of the mechanisms in the original claim is a shift from quarterly to annual reporting, and a study of annual reports is blind to that by construction. If companies are retreating from retention disclosure, the 10-Q is where I would expect to see it first, and I cannot see it here.

The frames are different. The original looked at 82 companies; this looks at 104 from a different list, and the overlap between them is not established, because the 82-company list is not published. A cohort difference remains a live explanation for the disagreement, and I could not rule it out.

The comparison cannot be run in reverse. The original analysis does not publish its 82-company list, so while I can check my cohort against its argument, nobody can check its cohort against mine. That is a limitation of the exchange rather than of either side, and it is the reason the frame for this piece was taken from a study that does name its constituents.

How can commercial finance build resilient cohort retention metrics?

The finding underneath all of this is not really about retention metrics. It is about what happens when you track anything across a population that is quietly losing members, which is the same failure mode as a corpus that quietly goes out of date.

A peer set has a half-life. The one here lost 22% of its constituents in about three years, and it lost them to acquisition, which means the loss is invisible in exactly the way that matters: no bankruptcy headline, no obvious gap, just a company that stops appearing in the data. Every benchmark built on a cohort from three years ago is quoting a group that has materially changed composition, and almost none of them say so.

So when a metric appears to be vanishing from your peer set, check whether the peer set is vanishing first. In practice that is one extra column in the comparison table on the next board deck: how many of last year's comparators still file? If the answer has moved, the trend in the other columns is partly a fact about who is left, and you will not find that out from the trend line.

Where are the empirical boundaries of cohort metric decay research?

Boundary. The replication describes one published cohort and annual filings, not the whole market. Preserve the cohort frame and add independent coding or a second panel before generalising the trend.

Evidence base. The analytical frame also draws on these additional sources: Gustafson 2024; Ordway Labs 2024; 2U 2024. The links identify the exact works; they support the mechanisms and boundary conditions discussed here, not every claim in isolation.

REFERENCES

- Gustafson, C. J. (2024). NDR is slowly being euthanized. Mostly Metrics. <https://www.mostlymetrics.com/p/ndr-is-slowly-being-euthanized-3cf4>
- Ordway Labs. (2024). Dollar-based net retention rate: how public SaaS companies report. <https://ordwaylabs.com/resources/research/how-public-companies-calculate-net-revenue-retention/>
- 2U, Inc. (2024). Current report on Form 8-K, July 25, 2024. U.S. Securities and Exchange Commission. <https://www.sec.gov/Archives/edgar/data/1459417/000119312524219689/d820461d8k.htm>

END OF ESSAY

HOW TO CITE THIS ESSAY

Isoglu, S. (2026, August 18). The metric didn't die. The cohort did. Sinan Isoglu.
<https://isoglu.com/insights/journal/the-metric-didnt-die-the-cohort-did/>

KEEP READING

From the research bench

The definition has a shorter memory than the number

<https://isoglu.com/insights/journal/the-definition-has-a-shorter-memory/>

From the research bench

Evidence over anecdote: what a number has to survive.

<https://isoglu.com/insights/journal/evidence-over-anecdote/>

From the research bench

What peer review deleted, the web kept.

<https://isoglu.com/insights/journal/what-peer-review-deleted/>



Sinan Isoglu, MBA (Quantic)

Commercial growth leader, lecturer and doctoral researcher