

5,500 Agent Customers, Zero Audited Dollars

What Workday, Okta, nCino, Best Buy and IFS actually disclosed about agent adoption in Q2 2026, and which revenue line each claim should have touched

David Soden

Independent analyst and practitioner, enterprise automation and applied AI
davidsoden.com/market-assessment

About this report

Every substantive claim in this report carries an evidence classification and, where applicable, a numbered source. Facts are separated from vendor claims, third-party estimates, the author's assessments, and untested hypotheses. Forward-looking sections use scenarios with observable tripwires rather than point forecasts. Stated assumptions are listed before the analysis begins.

PUBLISHED FOR PUBLIC READING | SHARE FREELY, UNMODIFIED, WITH ATTRIBUTION

© 2026 David Soden, davidsoden.com/market-assessment. All rights reserved. You may share this file complete and unmodified, and quote short passages with attribution. Republishing under another name, excerpting into a commercial deliverable, resale, and use as machine learning training data require written consent. Full terms appear under Rights and Permitted Use on the final page.

Market Assessment reports are published and commissioned at davidsoden.com/market-assessment. This document is analysis, not investment, legal, or procurement advice.

Scope and evidentiary standard

This covers enterprise application and infrastructure vendors that put a number on agent adoption in their most recent quarterly reporting through 28 August 2026, and the third-party survey statistics those vendors quote. The set is Workday, Okta, nCino, Best Buy and IFS, the five named in the brief, plus ServiceNow and Salesforce as the two vendors whose agent disclosures are furthest along and therefore set the standard the others are measured against.

Every quarter in the set ends between 30 June and 1 August 2026. Workday, Okta and nCino all closed on 31 July 2026 and reported within three days of each other. ServiceNow closed on 30 June and reported on 22 July. Best Buy's second quarter of fiscal 2027 reported on 27 August. Salesforce's second quarter of fiscal 2027 had not reported when this was written, so its most recent disclosure is the first quarter ended 30 April 2026.

The published agent numbers in this category need discounting for a specific reason, not a general one. Adoption metrics here are operating metrics: management defines them, management calculates them, and no auditor tests them. That is legal and normal. What makes it a problem this quarter is that the metrics are being used to carry an investment thesis and a multi-year platform decision, while the audited statements they sit beside say something quieter. Where a number is management's, this report labels it **VENDOR CLAIM** no matter how precise it looks. A figure carried to one decimal place is not more audited than a round one.

Four questions are settled here, in the order the brief asks them: what each metric counts, whether it reaches a reported financial line, what the agent withdrawals say about pilot-to-production conversion, and which survey figures are being recycled across decks. Claims that no disclosure supports are collected in their own section rather than scattered.

Label	Definition	How to treat it
FACT	Verifiable in the cited primary source: a filing, press release, official documentation, or standards body announcement.	Reliable as stated. Check the source if the exact wording matters.
VENDOR CLAIM	Reported by a vendor about its own business or product. Self-measured and not audited by any third party.	The vendor said it. Directionally useful, not a measurement. Definitions of the metric are usually undisclosed.
THIRD-PARTY ESTIMATE	A research firm forecast or survey result. Methodology is proprietary and generally not reproducible.	Use for direction and order of magnitude only. Do not build a model on it.
ASSESSMENT	The author's reasoned judgment from the cited evidence. Falsifiable, and the reasoning is shown so it can be argued with.	This is analysis, not reporting. Disagree with it where your evidence differs.
HYPOTHESIS	A proposition offered for testing. Not currently supported by sufficient evidence to assert.	Treat as a research question, not a conclusion. Each one names what would confirm or refute it.
SCENARIO	A structured future state with a stated subjective probability. The probability is the author's judgment, not a calculated figure.	Use the leading indicators attached to each, not the probability. The indicators are observable; the probability is not.

Assumptions this analysis rests on

First, that the earnings releases, prepared commentary and press coverage read for this report accurately reproduce what management said. Several figures reached me through wire syndication and financial aggregators rather than the filing PDF itself, because two primary documents would not render. If a wire copy transposed a number, a specific figure here moves but the pattern does not, because the pattern rests on eighteen separate claims from seven companies.

Second, that no company in this set has broken out agent revenue in a filing footnote that the releases and commentary do not mention. If one has, the central finding is wrong for that company and only that company. This is the assumption most worth testing against the actual 10-Q text, and the section on what this document cannot answer says so plainly.

Third, that management's agent counts are internally consistent quarter to quarter. Workday reporting 5,500 customers up 35 percent from last quarter implies roughly 4,070 last quarter, and the company said 400 customers were using its organic agents two quarters earlier. If the definition widened in between, the growth rate is measuring the definition rather than the deployment. Nothing in the disclosure lets an outsider rule that out, and the pilot-to-production section treats it as an open question rather than assuming either way.

Fourth, that the third-party surveys quoted by vendors were fielded and analysed as their publishers describe. I have the sample sizes and the industry mix. I do not have the questionnaires, the screening criteria, the response rates or the fielding windows, and no conclusion in this report depends on a survey figure.

The call

ASSESSMENT Agent adoption this quarter is real and almost entirely unpriced. Every vendor in scope disclosed a count of something, customers running an agent, deals closed, transactions automated, and not one reconciled that count to an audited revenue line. The strongest dollar figures are operating metrics the auditors never touched. That is not fraud, it is a disclosure regime that has not caught up, and buyers and analysts are being asked to underwrite a category on management's arithmetic. I would change position the first time a vendor in this set breaks out agent revenue as a reported line in a filing, with a prior-period comparison. Until then, treat every agent number as marketing with a footnote.

My judgment on the evidence assembled here, nothing more. There is data I can't see and data I may have missed. New evidence can move me, including to the opposite position, and I reserve the right to move.

Executive summary

The quarter produced more quantified agent disclosure than any before it, and the quantification runs in one direction. Vendors count users, deployments, deals and transactions. They do not count revenue, and where they do count something denominated in dollars, the denominator is an operating metric that sits outside the audited statements. The gap between those two things is the subject of this report.

FINDING 1 Eighteen agent claims, none of them a reported financial line.

ASSESSMENT Across seven companies I classified eighteen distinct quantified or semi-quantified agent claims. Four are commentary with no number attached. Nine are counts of customers, deals, deployments or transactions with no dollar value. Five are dollar-denominated but every one of those five is a non-GAAP operating metric: annual recurring revenue, annual contract value or a share of bookings. Not one of the eighteen appears as a line in an income statement or a footnote reconciled to one. ^{[1] [4] [8] [13] [16] [17]}

FINDING 2 Workday's agent count grew 35 percent in a quarter while its total backlog grew 8 percent.

FACT Workday reported more than 5,500 customers using at least one of its organic agents, up more than 35 percent from the prior quarter, and said AI drove more than 25 percent of new annual contract value. In the same release, total subscription revenue backlog was \$27.403 billion, up 8.0 percent year over year, the slower of the two backlog measures the company reports. The 12-month backlog grew 14.2 percent. Both numbers are correct and they are measuring different things, which is exactly why the pairing misleads when it is quoted without the second half. ^{[1] [19]}

FINDING 3 The largest agent dollar figure in the set is roughly six percent of the subscription base that produced it.

ASSESSMENT Workday's AI annual recurring revenue approached \$600 million, up more than 200 percent year over year. Against fiscal 2027 subscription revenue guidance of \$9.94 billion to \$9.95 billion, that is about six percent. Against total subscription revenue backlog of \$27.4 billion, it is about two percent. The growth rate is the headline; the share is the story, and only one of the two gets quoted. ^{[1] [21]}

FINDING 4 Okta's CFO said the agent business is immaterial, and the stock rose anyway.

FACT Okta reported dozens of AI agent deals in the quarter, several worth more than \$1 million, and said new products including Okta for AI Agents made up roughly 30 percent of second-quarter bookings. Management stated the agent revenue is still immaterial to the top line and will remain so this fiscal year. Total revenue grew 11 percent to \$805 million, subscription revenue 12 percent, current remaining performance obligations 14 percent and total remaining performance obligations 17 percent. The shares rose sharply on the print. ^{[4] [5] [6]}

FINDING 5 nCino disclosed no agent adoption metric at all, while marketing itself as an agentic platform.

FACT nCino's second quarter fiscal 2027 release describes the company as the platform for agentic banking and refers to a dual workforce of AI agents alongside human teams. It contains no count of agents deployed, customers using agents, or revenue attributed to them. The quantified AI reference is that four enterprise customers holding more than \$900 billion in combined assets renewed early with expanded commitments to use nCino's AI tools. Subscription revenue grew 10 percent, and 12 percent excluding US mortgage, against non-GAAP operating income growth of 36 percent. ^{[8] [9] [10]}

FINDING 6 Best Buy's agentic AI coverage rests on a product launch, not a metric.

FACT Best Buy launched Ask Blue, described as a conversational AI shopping and support assistant. No interaction volume, conversion effect, cost saving or headcount impact was disclosed. The quarter itself was strong on its own terms: revenue of \$9.78 billion, comparable sales up 4.1 percent, and full-year guidance raised to \$42.3 billion to \$42.8 billion. The trade coverage linking the two is adjacency, not attribution, and Best Buy did not claim otherwise. ^{[11] [12]}

FINDING 7 The agentic cull is the most useful disclosure of the quarter, and it is not in any filing.

ASSESSMENT Aneel Bhusri, who returned as Workday's chief executive in February 2026, cut the internal agent count from about 50 to about 20 and described killing agents or rolling them into bigger ones. Product leadership plans roughly 15 new agents in 2026. That is a 60 percent withdrawal rate on a first-generation portfolio, disclosed conversationally to the trade press and absent from the earnings release. It is the single best available read on pilot-to-production conversion in this set, and the disclosure regime treats it as colour. ^{[3] [20]}

FINDING 8 One survey publisher supplies both the statistics and the analysis in this set.

FACT The Futurum Group published the 1H 2026 Enterprise Software Decision Maker Survey of 830 global IT decision-makers, which produced the 86.6 percent agentic-priority figure quoted in commentary on nCino's quarter. Futurum also conducted the 664-respondent study behind IFS's agentic digital worker campaign, in partnership with IFS, and separately publishes vendor earnings analysis on the same companies. Each of those roles is ordinary on its own. Reading the three together is what a procurement lead scoring a vendor deck should do. ^{[9] [14] [15]}

FINDING 9 A widely read write-up of Workday's quarter printed the prior year's revenue figures.

FACT A trade article dated 28 August 2026 covering Workday's post-cull agent adoption stated revenue of \$2.35 billion up 12.6 percent, subscription revenue of \$2.169 billion, and net income of \$228 million. Those are Workday's fiscal 2026 second-quarter figures, reported on 21 August 2025. The fiscal 2027 second quarter reported \$2.649 billion and \$2.471 billion. The agent metric in the same article, 5,500 customers up 35 percent, is current and correct. ^{[1] [2] [3]}

What each metric actually counts

The brief asks what the disclosed metrics count, settled against the wording. The wording, taken literally, is the finding.

FACT Workday's release says more than 5,500 customers now use one or more of Workday's organic agents, up more than 35 percent from last quarter. Three words carry the weight. "Customers" is an account, not a seat, a workflow or a transaction, so a 55,000-employee bank rolling one agent to everyone and a 200-person customer running one agent in one department each count as one. "Organic" excludes agents that arrived through acquisition, which matters for a company that bought Sana and is building on it. "Use" is undefined: no threshold of frequency, volume or production status is stated anywhere in the release. ^[1]

ASSESSMENT The phrase doing the most damage is "one or more." A metric with a floor of one and no ceiling cannot distinguish a customer running a single pilot agent from a customer running twenty in production. It moves in only one direction as long as new customers try anything, and it can rise in a quarter where every existing deployment shrinks. That is not a criticism of Workday specifically. It is the shape almost every adoption metric in this set takes. ^[1]

FACT Workday's dollar-denominated metrics are more useful and still not audited. AI products contributed more than \$100 million in new annual contract value, more than 25 percent of total new ACV, and AI annual recurring revenue approached \$600 million, up more than 200 percent year over year and more than 20 percent sequentially. Workday's own quarterly filings describe company-reported annual recurring revenue from AI products as an operating metric rather than GAAP revenue. ^{[1] [18] [21]}

ASSESSMENT That AI ARR figure covers AI products, not agents. Adaptive Decision Intelligence, Sana Enterprise and the agent portfolio are all inside it. There is no disclosure anywhere in the set that isolates agent revenue from the rest of a vendor's AI line, which means the question "what

are the agents worth" has no published answer at any company examined here, including the two furthest ahead. ^{[1] [16] [17]}

FACT Okta's metrics count deals and bookings, not usage. The company reported dozens of AI agent deals in the quarter, several worth more than \$1 million, and said new products including Okta for AI Agents made up roughly 30 percent of second-quarter bookings. "Dozens" is not a number. The 30 percent figure is a share of bookings for a product family, not for agent identity alone, and bookings are not revenue. ^{[4] [5] [6]}

VENDOR CLAIM Okta's most quoted figure this quarter is an anecdote about one customer environment, where detected instances of a single vendor's agent went from 50 to 1,500 in a matter of weeks. As a demand signal it is genuinely interesting. As a metric it counts detected process instances at one unnamed account over an unstated window, which is not a base anyone can build a forecast on, and Okta did not present it as one. ^{[6] [7]}

FACT ServiceNow is the only company in the set whose agent-adjacent metric is both dollar-denominated and reported as a scale figure rather than a growth rate. AI annual contract value crossed \$1 billion, up more than 40 percent quarter over quarter, alongside subscription revenue of \$3.877 billion, up 24.5 percent. Customers with agentic AI in production grew ninefold over nine months and first-time agentic buyers rose 45 percent year over year. AI ACV is still an operating metric and is not reconciled to a revenue line in the release. ^[16]

FACT Salesforce reports Agentforce annual recurring revenue directly: \$800 million at the fiscal 2026 year end, reaching \$1.2 billion at the quarter ended 30 April 2026, up 205 percent. It also publishes a combined Agentforce and Data 360 figure of roughly \$3.4 billion, which includes about \$1.1 billion of acquired Informatica Cloud ARR. The combined number is the one that circulates. It is not an agent number. ^{[17] [25]}

ASSESSMENT That Informatica inclusion is the cleanest example in the set of a definitional choice that flatters a metric without misstating it. Every component is disclosed. A reader who quotes the \$3.4 billion as evidence of agent traction has been given everything needed to avoid the error and has made it anyway, which is the pattern this whole report is about. ^[17]


FACT nCino and Best Buy disclosed no adoption metric. IFS, which is privately held and files nothing, disclosed that its platform automates 60 percent of agentic transactions end to end, without stating the transaction population, the period, or how end-to-end completion is determined. ^{[8] [11] [13] [23]}

Metric	What the wording actually counts	Verifiability	Revenue line it should have touched
Workday: 5,500+ customers using one or more organic agents, +35% QoQ	Accounts with at least one non-acquired agent in any state of use. No usage threshold, no production requirement, no per-agent count	Vendor-defined operating metric. No auditor test, no stated definition in the release	None directly. Would reach subscription revenue only through the ACV it carries

Metric	What the wording actually counts	Verifiability	Revenue line it should have touched
Workday: AI ARR approaching \$600m, +200% YoY	Annualised recurring value of all AI products, agents included but not separable	Non-GAAP operating metric, described as such in the company's own filings	Subscription revenue, over the contract term. About 6% of FY27 subscription guidance
Workday: AI over 25% of new ACV, above \$100m	Share of newly contracted annual value attributed to AI products in one quarter	Vendor-calculated. The attribution rule is not disclosed	12-month subscription revenue backlog, which grew 14.2%. Total backlog grew 8.0%
Okta: dozens of AI agent deals, several over \$1m	Count of closed deals containing an AI agent component, threshold undisclosed	Not a number. "Dozens" and "several" are not auditable quantities	cRPO, which grew 14%. Management states the contribution is immaterial
Okta: new products around 30% of Q2 bookings	Share of bookings for a product family, not agent identity alone	Vendor-calculated. Product family membership is not defined publicly	cRPO and, later, subscription revenue. Not separable from the rest
ServiceNow: AI ACV over \$1bn, +40% QoQ	Annual contract value of AI products across the platform	Non-GAAP operating metric, not reconciled in the release	Subscription revenue of \$3.877bn, up 24.5%. No attribution given
Salesforce: Agentforce ARR \$1.2bn, +205%	Annualised recurring value of one named product family	Non-GAAP operating metric. The cleanest product-level figure in the set	Subscription and support revenue, over the contract term
nCino: four customers with \$900bn+ in assets renewed early with expanded AI commitments	Count of renewals, not of agents or of AI dollars	Verifiable as an event. Carries no adoption or revenue information	ACV guidance of \$662.5m to \$667.5m. No AI component disclosed
Best Buy: Ask Blue launched	A product exists	Verifiable. Not a metric	None claimed by the company. The Q2 comp of +4.1% predates the launch
IFS: 60% of agentic transactions automated end to end	Share of an undisclosed transaction population reaching completion without escalation	Vendor claim, private company, no filing exists to check it against	No public financial statement. Unverifiable by construction

Where the adoption lands in the financials

The second question is whether agent adoption shows up in subscription revenue, remaining performance obligations or margin, or only in commentary. Lining each claim up against the same period's reported growth settles it.

FIGURE 1 Reported growth in the same quarter each agent claim was made

FACT Year-over-year growth as reported. ServiceNow's quarter ended 30 June 2026; Workday, Okta and nCino closed 31 July 2026; Best Buy's fiscal second quarter reported 27 August. Not one of these companies attributed any portion of this growth to agents in its release. [\[1\]](#) [\[4\]](#) [\[8\]](#) [\[11\]](#) [\[16\]](#)

ASSESSMENT The chart is the answer to the question. These are respectable growth rates for enterprise software in 2026, and every one of them is explicable without agents. Workday's 13.9 percent is roughly the rate it has been running. nCino's 10 percent is below its own recent trend and the release spends more space on margin than on growth. Best Buy's 3.6 percent came from computing and emerging hardware categories the company named explicitly, with Ask Blue arriving too late in the quarter to have moved it. [\[1\]](#) [\[8\]](#) [\[10\]](#) [\[11\]](#)


FACT Remaining performance obligations tell a more interesting story than revenue, and it is not the one the adoption metrics imply. Okta grew total RPO 17 percent and cRPO 14 percent against 11 percent revenue growth, which is the signature of a business contracting ahead of its recognised revenue. Workday grew its 12-month backlog 14.2 percent but its total backlog only 8.0 percent, which is the opposite signature: near-term visibility holding while long-dated contracted value slows. [\[1\]](#) [\[4\]](#) [\[6\]](#)

ASSESSMENT Workday's own explanation for the 8 percent is a mix shift toward existing customers over net new, plus contract duration and the industries in the quarter's bookings. That is plausible and it is also the explanation you would give if agent deals were landing as short-duration expansions on existing accounts rather than as long-term new commitments. Both readings fit the disclosed facts. Nothing in the release distinguishes them, and an analyst who assumes either one is guessing. [\[1\]](#) [\[19\]](#)

ASSESSMENT Workday's introduction of Flex Credits as usage-based subscriptions makes this harder rather than easier. Usage-based revenue is recognised on consumption, so a successful agent rollout can produce contracted value that does not appear in the 12-month backlog on the schedule a seat-based deal would. Anyone modelling agent revenue off backlog growth needs to know the Flex Credits mix, and it is not disclosed. [\[19\]](#)

ASSESSMENT Margin is where the strongest case for agent economics can be made, and it is the weakest evidentially. nCino swung from a \$9.3 million GAAP operating loss to \$13.6 million of GAAP operating income, expanded non-GAAP operating margin 500 basis points to 25 percent, and grew free cash flow 170 percent to \$34.0 million, all on 8 percent revenue growth. Workday expanded non-GAAP operating margin from 29.0 to 31.1 percent. Neither company attributed any of it to agents, and both have been running cost programmes that explain it without invoking AI at all. A margin expansion with two available causes and no disclosed attribution is not evidence for either. ^{[1] [8] [9]}

FIGURE 2 How far each agent claim travels down the disclosure ladder




ASSESSMENT The author's classification of eighteen quantified or semi-quantified agent claims made by Workday, Okta, nCino, Best Buy, IFS, ServiceNow and Salesforce in their most recent reporting. The boundary between the first two categories is a judgment call; the boundary at zero is not. ^{[1] [4] [8] [11] [13] [16] [17]}

The cull, and what it says about pilot-to-production

The third question is what happened to agents that were withdrawn or consolidated. Only one company in the set answered it, and it did so outside its filings.

VENDOR CLAIM Aneel Bhusri returned as Workday's chief executive in February 2026 and reduced the internal agent count from about 50 to about 20. His own account, given to the trade press: there were a lot more agents when he came back, and the company killed a lot of them or rolled them into bigger agents. Product leadership separately said Workday plans to release about 15 new agents in 2026, including tools for corporate travel and IT service management built on the Sana acquisition. ^{[3] [20]}

FIGURE 3 Workday's agent portfolio through the cull

VENDOR CLAIM Approximate counts given by Workday executives to the press, not in any filing or earnings release. The two survivor figures are stated as "about," so treat the 60 percent reduction as a rough magnitude rather than a rate. ^[3] ^[20]

ASSESSMENT Three things follow from this that the adoption metrics obscure. The first is that roughly 60 percent of a first-generation agent portfolio at a company with Workday's engineering resources did not survive contact with its own product review. That is a conversion signal, and it is worse than any published pilot-to-production statistic in the category. ^[3] ^[20]

ASSESSMENT The second is that consolidation and withdrawal are recorded identically in the customer count. A customer whose three narrow agents were merged into one broader agent still counts as a customer using one or more agents. The 5,500 figure cannot distinguish growth from repackaging, and the 35 percent quarterly increase spans exactly the period in which the portfolio was being restructured. ^[1] ^[3]

ASSESSMENT The third is more favourable to Workday than the first two. Killing 30 agents is what a company does when it is measuring them, and the alternative, shipping 50 thin agents and reporting the customer count, would have produced a better-looking metric. The cull is evidence of internal discipline that the disclosure regime actively penalises, since nothing in the release rewards a vendor for withdrawing a product. ^[3]

VENDOR CLAIM The strongest single production data point in the set comes from Workday's customer BMO, which piloted a self-service agent with 500 employees in May and rolled it to all 55,000 in June. One month from pilot to full deployment at a large bank is a real result if it holds. It is one account, self-reported through the vendor, with no disclosure of usage, retention or outcome after the rollout. ^[3]

VENDOR CLAIM IFS supplies the only other production-conversion detail. Ependion deployed a supply order manager company-wide in ten weeks, and Kitron Group is rolling purchase-to-order digital workers across 13 sites, with the rollout surfacing a decade-old part-number error. IFS is

privately held, so no filing exists against which any of it can be checked, and the six customer interviews sit inside a study IFS commissioned. ^{[13] [14] [24]}

HYPOTHESIS Proposition worth testing: agent portfolio churn at the vendor is a leading indicator of agent churn at the customer, on a lag of two to three quarters. It would be confirmed if a vendor that culls its portfolio subsequently reports a decelerating or falling adoption count, and refuted if adoption counts keep rising through and after a cull. Workday is the first clean test case, and the answer arrives with its next two prints. ^{[1] [3]}

The survey numbers inside the decks

The fourth question is which survey statistics are recycled across vendor decks and what the underlying methodology sampled. Two dominate this quarter's material, and they come from the same publisher.

THIRD-PARTY ESTIMATE The Futurum Group's 1H 2026 Enterprise Software Decision Maker Survey of 830 global IT decision-makers produced the figures that appear in commentary on nCino's quarter: 86.6 percent ranking autonomous agents and agentic AI a high-priority technology, 90.4 percent saying the same of generative AI with 31.8 percent ranking it first, and 51.4 percent naming efficiency improvement as the top measure of return. The same study reports agentic AI rising from 13.0 to 17.1 percent as a top-ranked priority, a 31.5 percent year-over-year increase. ^{[9] [15]}

ASSESSMENT Note what those figures measure. Every one is a stated priority or a stated intent, not a deployment, a spend or an outcome. An 86.6 percent priority ranking and a five percent production rate are perfectly compatible, and both can be true of the same 830 people. Priority surveys are the most quoted and least predictive instrument in enterprise technology, because ranking something important costs the respondent nothing. ^[15]


FACT The second study is Futurum's "Bridging the Capacity Gap: Why 66% of Enterprises Are Investing in Agentic Digital Workers," written by Keith Kirkpatrick, published 26 August 2026, produced in partnership with IFS and used to launch IFS Loops the same day. It surveyed 664 enterprise decision-makers across manufacturing, energy and utilities, aerospace and defence, transportation and logistics, construction and telecommunications, and interviewed IT and operations leaders at six IFS customers running digital workers in production. ^{[13] [14]}

ASSESSMENT The six in-depth interviews were with IFS customers already running the product, selected by or with the vendor that commissioned the study. That is a legitimate case-study design and it is not a sample. Any conversion or time-to-value figure drawn from those six describes six successful deployments at one vendor, and the press release does not distinguish those numbers from the 664-respondent survey numbers when it lists them together. ^{[13] [14]}

ASSESSMENT The most informative figure IFS published is the one that undercuts its own campaign: 5.7 percent of decision-makers trust AI to act fully autonomously, against 66 percent likely to invest in digital workers within the year. Read together, that is a market buying supervised

automation and being sold autonomy. IFS deserves credit for printing both, and the human-in-the-loop framing in the same release is consistent with it. ^[13]

FIGURE 4 Sample sizes behind the statistics circulating in vendor material



THIRD-PARTY ESTIMATE Publisher-stated sample sizes. None of the four publishes a response rate, screening criteria or fielding window in its public material, so the sizes are the only methodological fact available. Size is not quality; a well-screened 153 can beat a self-selected 830. ^{[14] [15] [22]}

ASSESSMENT The recycling question has a specific answer here, and it is narrower than the brief implies. The 86.6 percent figure did not appear in nCino's own disclosure. It appeared in Futurum's analysis of nCino's disclosure, supplied by Futurum from Futurum's own survey. Futurum also conducted the study IFS commissioned, and publishes earnings commentary on Okta and other vendors in this set. Each role is unremarkable in isolation. The combination means one organisation is simultaneously the survey house, the vendor's research partner and the independent analyst on the same quarter. ^{[9] [14] [15]}

ASSESSMENT This is disclosed rather than hidden. Futurum labels the IFS study as a partnership and publishes the sample size. The problem is downstream: by the time a statistic reaches a vendor slide it has lost the label, and a procurement lead scoring three vendor decks that all cite the same 830 people has no way to see that they are one data point, not three. ^{[9] [14] [15]}

THIRD-PARTY ESTIMATE The third recycled statistic is older and worth naming because it now runs in both directions. The MIT NANDA "GenAI Divide" working paper, based on 52 executive interviews, 153 surveyed leaders and 300 public deployments, is the source of the claim that 95 percent of AI pilots deliver no measurable profit-and-loss impact. It described itself as preliminary findings, was not peer reviewed, and came from a group building agentic AI infrastructure. Vendors now cite it both as the problem their agents solve and as proof that everyone else's agents fail. ^[22]

ASSESSMENT Treat that 95 percent as unusable in either direction. "No measurable P&L impact" and "the pilot failed" are different findings, and most enterprise pilots are not instrumented to measure P&L impact in the first place. I have not leaned on it anywhere in this report, and no conclusion here would change if it were withdrawn tomorrow. ^[22]

Claims no disclosure supports

The brief asks for every claim flagged where no disclosure backs it. These are the ones that matter, in the order they are likely to be repeated at you.


Claim in circulation	Origin	What would settle it, and why nothing currently does
Agent adoption is driving subscription revenue growth at Workday	Inference drawn by readers, not a company statement	A revenue breakout or an attribution disclosure. Workday said AI drove over 25% of new ACV, which is a bookings mix, not a revenue driver. Total backlog grew 8%
Best Buy's comparable sales beat is connected to its agentic AI work	Trade coverage pairing the two in one article	Any Best Buy disclosure of Ask Blue engagement or conversion. There is none. Best Buy attributed the comp to computing, Best Buy Business and emerging categories
nCino's margin expansion reflects agent efficiency	Adjacency in the release and in commentary	A cost attribution. nCino disclosed the 500 basis point non-GAAP margin expansion and a separate cost programme, and connected neither to agents
Okta's stock move validates agentic identity as a revenue driver	Market reaction on the print	Okta's own CFO says the agent contribution is immaterial to the top line this fiscal year. The beat came from core identity
Salesforce has \$3.4bn of agentic AI ARR	Combined Agentforce and Data 360 figure quoted without its components	Salesforce discloses the components. Agentforce ARR is \$1.2bn; about \$1.1bn of the combined figure is acquired Informatica Cloud ARR
IFS Loops automates 60% of agentic transactions end to end	IFS press release, private company	A defined transaction population and period. Neither is stated, and no filing exists to check against
66% of enterprises are investing in agentic digital workers	Futurum study commissioned by IFS, repeated as a market fact	The survey measured stated likelihood to invest within 12 months among 664 decision-makers, not investment. The same study found 5.7% trust full autonomy
86.6% of decision makers prioritise agentic AI, therefore the market is converting	Futurum survey figure quoted alongside vendor results	Priority rankings do not measure deployment or spend. The survey does not claim they do; the reuse does
95% of AI pilots fail	MIT NANDA preliminary working paper	Peer review, or a replication. The paper measured absence of measurable P&L impact across 153 surveyed leaders and was not institutionally underwritten

Claim in circulation	Origin	What would settle it, and why nothing currently does
Workday's agent count proves pilots are converting to production	Reading the 5,500 figure as a production count	A production or usage threshold in the metric definition. The wording is "use one or more," with no threshold stated

Where the secondary reporting broke

FACT One error is specific enough to name. A trade article published 28 August 2026 on Workday's post-cull agent adoption reported revenue of \$2.35 billion up 12.6 percent, subscription revenue of \$2.169 billion up 14 percent, professional services of \$178 million and net income of \$228 million. Those figures match Workday's fiscal 2026 second quarter, reported 21 August 2025, to the dollar. Workday's fiscal 2027 second quarter, the one the article covers, reported total revenue of \$2.649 billion and subscription revenue of \$2.471 billion. [\[1\]](#) [\[2\]](#) [\[3\]](#)

FIGURE 5 Workday's quarter as printed in trade coverage, and as reported



FACT US dollars, millions. The printed figures match Workday's fiscal 2026 second quarter, reported a year earlier. The agent metric in the same article, 5,500 customers up 35 percent, is current and matches the release. The error is in the financials, not the adoption data. [\[1\]](#) [\[2\]](#) [\[3\]](#)

ASSESSMENT The reason to include this is not the error itself, which is the kind any publication makes under deadline. It is that the article pairs a correct current agent metric with a year-old revenue base, and the resulting implied ratio of agent adoption to revenue is wrong by about 13 percent in the direction that flatters the agents. That is precisely the reconciliation this report exists to perform, failing in miniature, in a source a reasonable analyst would treat as reliable. [\[1\]](#) [\[2\]](#) [\[3\]](#)

Scenarios to the next two reporting cycles

Point forecasts are not available here. The disclosure regime, not the technology, determines what the next two quarters reveal, and disclosure practice changes discontinuously when one company moves and the others follow. Probabilities below are subjective and are the least reliable content in this document. The earliest visible signs are the useful part.

SCENARIO A One vendor breaks it out 30 percent

At least one company in this set reports an agent or AI revenue line in a filing, with a prior-period comparison, by the end of the first half of 2027. Salesforce and ServiceNow are the likeliest first movers because both already publish product-level dollar metrics and both have an incentive to make the number auditable while it is growing fastest.

Consequence. Peer pressure forces the rest within two quarters, and the category becomes analysable for the first time. The vendors whose agent revenue is small get repriced immediately.

Earliest visible sign. A new line item or a new defined term in a 10-Q, or an investor-day slide that reconciles a non-GAAP AI metric to a GAAP line.

SCENARIO B The counts keep floating 50 percent

Adoption counts continue, definitions continue to shift quietly between quarters, dollar figures stay in non-GAAP operating metrics, and no company reconciles anything. Commentary carries the story and the audited statements stay silent on agents entirely.

Consequence. Buyers and analysts keep underwriting the category on management arithmetic. Dispersion between vendors widens without becoming visible, and the correction, when it comes, arrives as a guidance cut rather than a disclosure.

Earliest visible sign. A vendor changing the wording of a metric between two consecutive releases without restating the prior period. Watch for customers using agents quietly becoming customers with agents enabled, or the reverse.

SCENARIO C A count gets pulled 20 percent

A vendor that has reported an agent adoption count stops reporting it, or replaces it with a different metric, in a quarter where the original would have decelerated or fallen. This is the standard end-state for enterprise vanity metrics and it has happened before with seats, active users and consumption units.

Consequence. The withdrawal is more informative than any number the company published, and it resets how the whole set is read. Vendors still reporting counts get a temporary premium and then get asked harder questions.

Earliest visible sign. A metric present in three consecutive releases absent from the fourth, with no replacement and no explanation. Check the prior release side by side rather than trusting the current one to mention the omission.

Leading indicators worth monitoring

None of these require private information or a vendor relationship. All are observable from published documents within days of release.

If this happens	It means	Moves toward	Where to watch
An adoption metric's wording changes between releases with no restatement	The definition is being managed, and the growth rate is partly definitional	Scenario B	Two consecutive earnings releases, read side by side, not the current one alone
A vendor states a usage or production threshold inside its agent metric	The metric has become answerable and the vendor is confident in the answer	Scenario A	Earnings release wording and the metric definitions section of the 10-Q
Workday's total subscription backlog growth stays near 8% for two more quarters while AI ARR keeps compounding	Agent value is landing as short-duration expansion, not long-term commitment	Scenario B	Workday Q3 and Q4 FY2027 releases, total versus 12-month backlog
Okta's cRPO growth converges downward toward revenue growth	The agent-driven contracting premium was a one-quarter effect	Scenario C	Okta quarterly cRPO against subscription revenue growth
Any vendor reconciles a non-GAAP AI metric to a GAAP revenue line	The category has become analysable	Scenario A	Non-GAAP reconciliation tables at the back of the release, where nobody looks
An agent adoption count decelerates without a stated cause	Either the base saturated or the cull reached the customers	Scenario C	Sequential change in the count, quarter over quarter, not the year-over-year figure
A vendor deck cites a survey without naming the sample size	The number has been laundered through two hops from its methodology	No scenario, a scoring rule	Any vendor deck. Ask for the sample size and the fielding date in writing
A vendor discloses agent gross margin separately	Inference economics are being managed rather than absorbed	Scenario A	Segment or product margin commentary in prepared remarks

Implications

For CIOs selecting multi-year agent platforms

ASSESSMENT The single most useful thing in this quarter's disclosure for you is the cull, not the adoption counts. Workday withdrew or consolidated roughly 60 percent of a first-generation agent portfolio, and it is the only vendor in the set that let anyone see it. Assume every vendor you are evaluating has a comparable withdrawal rate and has not disclosed it. Ask each one directly, in writing, how many agents they have withdrawn or merged in the last twelve months, and treat a refusal to answer as an answer. ^{[3] [20]}

ASSESSMENT Do not select on adoption counts. A metric of the form "customers using one or more agents" is architecturally incapable of telling you whether anything reached production, and every vendor in this set that publishes an adoption number publishes one of that shape. What you want instead is a named reference in your industry at your scale, running a specific agent, with a stated date of first production use and a stated volume. Workday's BMO example is the shape of the answer; one month from a 500-person pilot to 55,000 employees is checkable if you ask BMO. ^{[1] [3]}

ASSESSMENT Price the consolidation risk into contract terms rather than into vendor selection. The vendors culling agents are, on this evidence, the ones exercising product judgment, so a withdrawal record is not a reason to disqualify. It is a reason to insist on migration commitments for any agent you deploy, a stated support horizon, and the right to substitute at no cost if a named agent is withdrawn or merged during the term. That clause costs nothing to ask for today and is unobtainable after a cull. ^[3]

For software equity and credit analysts

ASSESSMENT The reconciliation this quarter is unambiguous. No company in the set has an agent revenue line, and every dollar-denominated agent metric is an operating metric outside the audited statements. Workday's AI ARR of roughly \$600 million is about six percent of fiscal 2027 subscription revenue guidance and about two percent of total subscription backlog, and it covers all AI products, not agents. Any model that treats an agent adoption count as a revenue input is modelling a quantity the company has not disclosed. ^{[1] [18] [21]}

ASSESSMENT The specific discipline worth adopting is checking the unit of the growth rate against the unit of the comparison. Workday's 35 percent is quarter over quarter; its 13.9 percent subscription growth is year over year. ServiceNow's AI ACV growth of over 40 percent is quarter over quarter against 24.5 percent annual subscription growth. Every adoption metric in this set is quoted on the faster clock, and the pairing is the vendor's choice, not an accident. ^{[1] [16]}

ASSESSMENT For credit, the divergence between Workday's 14.2 percent 12-month backlog growth and 8.0 percent total backlog growth is the number that deserves the work. Near-term coverage is holding while long-dated contracted value decelerates, which is what you would see if expansion is landing on existing accounts at shorter durations. The Flex Credits usage-based structure makes the same pattern appear for a benign reason. The disclosure does not let you

distinguish them, and that ambiguity, not the agent story, is the thing to press management on.

[1] [19]

ASSESSMENT Okta is the clearest case of a stock rerating on an agent narrative the company itself disclaims. Management said the agent contribution is immaterial to fiscal 2027 and will remain so this year. The quarter was carried by core identity, with cRPO up 14 percent and RPO up 17 percent against 11 percent revenue growth. If you are long on agentic identity, you are long on fiscal 2028, and the company has told you so. [4] [5] [6]

For procurement leads scoring vendor claims

ASSESSMENT Build the scoring rubric around three questions and score every claim in every deck against them. What is the denominator, what is the period, and who counted. In this set, denominators are undisclosed in eight of the ten claims in the flagged table, periods mix quarterly and annual without labelling, and every count is the vendor's own. A claim that fails all three is not a lie, it is unusable, and treating unusable and false as the same thing in scoring is what stops the argument. [1] [4] [8] [13]

ASSESSMENT Deduplicate the survey citations across decks before scoring them. If three vendors cite agentic AI priority statistics that trace to the Futurum 1H 2026 study, that is one survey of 830 people appearing three times, not three independent confirmations. Ask each vendor in writing for the publisher, sample size, fielding date and whether they commissioned or sponsored the work. IFS discloses its partnership; most do not volunteer it. [9] [14] [15]

ASSESSMENT Two questions separate a vendor that is measuring agents from one that is marketing them. First, how many agents have you withdrawn, merged or renamed in the last four quarters, and what happened to the customers running them. Second, will you state, in the contract, a usage threshold at which we count as a production customer for the purposes of your public metrics. The second question is unusual and that is the point: a vendor willing to define its own metric in your contract has defined it internally. [1] [3]

For product leaders setting internal adoption KPIs

ASSESSMENT Do not copy the vendor metric shape into your own dashboard. "Teams using one or more agents" is the metric that got vendors into this position, and it will do the same to you: it rises while every deployment underneath it shrinks, it cannot fall unless someone actively removes an agent, and it gives your executive sponsor a number that feels like progress and predicts nothing. The failure mode is documented in this quarter's disclosure and it is free to avoid. [1]

ASSESSMENT Instrument the denominator first, before the first agent ships. The specific counts worth holding are: tasks routed to an agent as a share of eligible tasks, share of those completed without human escalation, and the same measure at 30 and 90 days after deployment. IFS is the only vendor in this set publishing a completion-rate metric at all, at 60 percent of agentic transactions end to end, and even that lacks a stated population. If you build the denominator yourself you can answer a question none of your vendors can. [13]

ASSESSMENT Budget for withdrawal. Workday's 60 percent portfolio reduction is what a well-resourced vendor's first generation looks like after review, and there is no reason to expect an internal programme to do better. Plan on retiring or merging a majority of your first cohort, say so to your sponsor before you start, and make the retirement count a reported KPI rather than a failure. Doing so converts the most predictable event in the programme from a credibility loss into evidence of judgment. ^[3] ^[20]

What this document cannot answer

Four questions need primary research, and each would change something material here.

Whether any of these companies defines an agent adoption metric in the text of its quarterly filing. I worked from earnings releases, prepared commentary, call coverage and press reporting. Two primary documents would not render for me, and I did not read the full 10-Q text for any company in the set. If a definition exists in a filing footnote, the central finding is weaker for that company, and a lawyer or an analyst with filing access can settle it in an afternoon.

What agent revenue actually is at any of these vendors. No company separates agents from the rest of its AI portfolio. Workday's roughly \$600 million covers AI products; ServiceNow's billion-dollar AI ACV covers the platform; Salesforce's Agentforce ARR is the closest thing to a product-level figure and still bundles a family. Nobody outside the companies can size the agent line, and the companies have chosen not to.

Whether Workday's 5,500 figure is measuring the same thing it measured two quarters ago, when 400 customers were reported using its organic agents. A definitional widening would produce the observed growth without any change in deployment. Only the company can answer it, and an investor relations question with the two definitions side by side is the way to ask.

What happened to the customers running the roughly 30 Workday agents that were withdrawn or merged. Migration effort, contract treatment and whether any of them stopped using agents entirely is the best available evidence on real switching cost in this category, and none of it is public. Reference calls with three or four affected customers would be worth more than every survey cited in this report.

Half-life of this assessment

The specific figures decay fastest. Every revenue, backlog and margin number here is superseded by the next quarterly print, which for Workday, Okta and nCino lands in late November 2026 and for Salesforce in September 2026. Treat the numbers as of late August 2026 and re-source rather than re-quote after that.

The adoption counts decay within one to two quarters, and they decay in a particular way: not by becoming wrong but by becoming incomparable, because the definitions move. Workday's 5,500 and its 35 percent sequential growth are the two figures most likely to be superseded by a differently defined replacement rather than by an update.

The survey statistics have about twelve months. Futurum's 1H 2026 study will be succeeded by a 2H 2026 edition, and the IFS capacity gap study is a point-in-time instrument tied to a product launch. The MIT NANDA figure is already past its useful life and should be retired from circulation rather than updated.

The disclosure practice finding runs longer, roughly two to four quarters, and it ends discontinuously. It holds until the first vendor breaks out an agent line, at which point the whole framing changes within two reporting cycles as peers follow. Until that happens the finding gets stronger with each quarter that passes without a breakout.

Two things here are architectural and will outlast the rest. The first is that an adoption metric with a floor of one and no usage threshold cannot measure production conversion, which is a property of the metric and not of any vendor. The second is that operating metrics are defined by management and not tested by auditors, so the gap between an adoption claim and a reported line is structural rather than a lapse. Both will still be true when every number in this report has been replaced.

Limitations

I hold no position in any company named here, have no commercial relationship with any of them, and was not commissioned by any party with an interest in the outcome. Futurum, Diginomica, MT Newswires and the other publications cited are used as sources and are not clients.

The evidence base has three specific weaknesses. Okta's posted quarterly commentary and nCino's release would not render as text for me, so figures for both companies come from wire syndication of the release and from financial aggregators reporting the same numbers. Where two independent sources agreed on a figure I used it; where they did not I left the figure out. I have not verified any of these numbers against the filings themselves.

Second, no 10-Q text was read. The brief asked for definitions settled against filings, and the honest answer is that this report settles them against earnings releases and prepared remarks, and finds no definition in either. That is a weaker claim than the brief asked for and I would rather state the weaker claim accurately.

Third, the classification of eighteen agent claims into four disclosure tiers is my judgment, not a standard. Someone else counting the same material might find sixteen or twenty-one, and might move a claim between the first two tiers. The count at the bottom tier, which is zero, is the one that carries the argument, and it is the one least sensitive to how the others are drawn.

Best Buy and IFS sit awkwardly in a set built around reconciling disclosures to filings. Best Buy is an agent buyer rather than an agent vendor, so its financial statements were never going to show agent revenue. IFS is privately held and files nothing, so its claims are unverifiable by construction rather than by choice. Both are included because the brief names them and because both illustrate the pattern, but neither can be held to the same standard as the public software vendors and neither is scored as though it could.

Sources

- [1] Workday, Inc. "Workday Announces Fiscal 2027 Second Quarter Financial Results." 27 August 2026. *Vendor primary* investor.workday.com
- [2] Workday, Inc. "Workday Announces Fiscal 2026 Second Quarter Financial Results." 21 August 2025. *Vendor primary* newsroom.workday.com
- [3] Lauchlan, Stuart. "Following an agentic cull, rising adoption is delivering ROI for Workday customers, says CEO Aneel Bhusri." Diginomica, 28 August 2026. *Independent press* diginomica.com
- [4] Okta, Inc. "Second Quarter Fiscal 2027 Posted Commentary." 26 August 2026. *Vendor primary* okta.com
- [5] The Futurum Group. "Okta Q2 FY 2027 Earnings: Beat and Raise on Core Identity Strength." August 2026. *Analyst firm* futurumgroup.com
- [6] TIKR. "Okta's Q2 Earnings Beat Every Estimate. AI Still Isn't Driving the Numbers." August 2026. *Independent press* tikr.com
- [7] MT Newswires. "Okta's Early Traction in Agentic AI." Yahoo Finance, 27 August 2026. *Independent press* finance.yahoo.com
- [8] nCino, Inc. "nCino Reports Second Quarter Fiscal Year 2027 Financial Results." GlobeNewswire, 25 August 2026. *Vendor primary* globenewswire.com
- [9] The Futurum Group. "nCino's Q2 FY2027: Agentic AI Banking Thesis Meets Margin Reality." 26 August 2026. *Analyst firm* futurumgroup.com
- [10] "nCino, Inc. Q2 2027 Earnings Call Summary." Yahoo Finance, August 2026. *Independent press* finance.yahoo.com
- [11] CNBC. "Best Buy (BBY) Q2 2027 earnings." 27 August 2026. *Independent press* cnbc.com
- [12] HomePage News. "Best Buy Advances in Agentic AI After Solid Q2 Earnings, Comps." 27 August 2026. *Independent press* homepagenews.com
- [13] IFS. "Industrial workforce capacity gap being filled by agentic digital workers." PR Newswire, 26 August 2026. *Vendor primary* prnewswire.com
- [14] Kirkpatrick, Keith. "Bridging the Capacity Gap: Why 66% of Enterprises Are Investing in Agentic Digital Workers." The Futurum Group, in partnership with IFS, 26 August 2026. Sample of 664 decision-makers plus six vendor-selected customer interviews. *Analyst firm (vendor-commissioned)* futurumgroup.com
- [15] The Futurum Group. "Agentic AI Surges 31.5% to Become the Fastest-Growing Enterprise Tech Priority." 1H 2026 Enterprise Software Decision Maker Survey, sample of 830, February 2026. *Analyst firm* futurumgroup.com
- [16] "ServiceNow Q2 2026 Earnings Call: AI ACV Surges Past \$1 Billion." BigGo Finance, 22 July 2026. *Independent press* finance.biggo.com
- [17] Salesforce, Inc. "Salesforce Delivers Record First Quarter Fiscal 2027 Results." 27 May 2026. *Vendor primary* salesforce.com
- [18] Workday, Inc. Form 10-Q for the quarterly period ended 30 April 2026. US Securities and Exchange Commission. *Regulatory primary* sec.gov
- [19] "Workday Says AI Drove 25% of New Annual Contract Value. Why Did Total Subscription Backlog Grow Only 8%?" Insider Monkey, 28 August 2026. *Independent press* insidermonkey.com
- [20] "Workday's Returning CEO Has a Plan to Survive the AI Era." Yahoo Finance, 2026. *Independent press* finance.yahoo.com
- [21] "Workday, Inc. (WDAY) Q2 2027 Earnings Call Transcript." AlphaStreet, 27 August 2026. *Independent press* alphastreet.com
- [22] Poynter, Ray. "Myth Number 2: MIT Showed That 95% of AI Pilots Fail." NewMR. Critique of the MIT NANDA GenAI Divide working paper, sample of 153 surveyed leaders plus 52 interviews. *Practitioner review* newmr.org
- [23] nCino, Inc. Second quarter fiscal 2027 results, syndicated wire copy. The Manila Times, 26 August 2026. Used because the primary release would not render. *Second-hand (syndicated primary)* manilatimes.net
- [24] "Industrial workforce capacity gap being filled by agentic digital workers in manufacturing." Automation Magazine, August 2026. Republished IFS release. *Vendor-published (republished)* automationmagazine.co.uk
- [25] Salesforce, Inc. "Salesforce Delivers Record Fourth Quarter Fiscal 2026 Results." 25 February 2026. *Vendor primary* salesforce.com

Rights and permitted use

© 2026 David Soden, davidsoden.com/market-assessment. All rights reserved.

This document, including its structure, analysis, findings, evidence classification, and framing, is the original work of David Soden, published at davidsoden.com/market-assessment.

It is published for public reading. You are free to share the complete, unmodified file, to link to it, and to quote short passages in commentary, reporting, or discussion, provided the author and the site above are credited.

The following require prior written consent: republishing this work in whole or in substantial part under another name, masthead, or brand; excerpting or modifying it in a way that alters the analysis or removes the evidence classifications; incorporating any portion of it into a paid, commissioned, or client-facing deliverable; resale or distribution behind a paywall; and use of this document as training data for machine learning models.

Commissioned Market Assessment reports, commercial licensing, and briefings on this material are available.

This document is analysis, not investment, legal, or procurement advice. It was not commissioned, reviewed, or funded by any company named in it, and the author holds no position in any of them.

Market Assessment reports, commissioned research, and briefings: davidsoden.com/market-assessment