ACTIVANT RESEARCH

Long Live Software

How incumbency creates the right to challenge

May 2026

AUTHORS:

Jono Vickery

Vice President, Research

About Activant

Activant is a research-led global investment firm that partners with high-growth companies. Since 2015, we have invested in category-defining businesses during their most critical phases of growth, partnering with founders who have won their initial battles and are ready for the next challenge.

Our approach pairs deep, proprietary research with patient, flexible capital and hands-on operational partnership. We work alongside founders and leadership teams to refine strategy, strengthen operations, and accelerate sustainable growth.

Activant Research is dedicated to uncovering the most exciting emerging technologies, sectors, and companies we believe will shape the future. Our research-driven perspective informs everything we do, helping us invest at meaningful inflection points and support founders in building enduring, category-leading businesses.

You can find out more about Activant and our research at https://activantcapital.com/.

Long Live Software!?

Notion can be vibe-coded in a weekend. The workflows it supports are well suited to be off-boarded to Claude/ChatGPT and its unit of value does not allow the company to move beyond a seat-based model.

Palo Alto Networks, on the other hand, runs a high-stakes workflow mediated by regulations and compliance, and ingests ~15 petabytes of proprietary telemetry data daily: data that improves the product for all users. They run firewalls that sit in production systems, not a Claude interface, and their AI systems like Cortex XSIAM are perfectly suited to usage-based monetization, already surpassing $500 million in ARR.[1]

These companies do not face the same risks.

Software is being disrupted, but the market has priced disruption relatively equally across companies like Notion and companies like Palo Alto Networks. Telling them apart is the job of this piece.

Software survival in the AI era runs through five moats: pricing tied to observable ROI, proprietary data that compounds across customers, workflow depth that resolves business logic AI cannot fake, ecosystem centrality that makes the platform the substrate for AI-native point solutions, and vertical specialization horizontal models cannot cross. Together, they determine which incumbents earn the right to fight for the labor budget AI is now unlocking.

Pricing Tied to Observable ROI

In our last article, we argued that seat-based software businesses need to innovate their pricing models. The challenge is that not everyone can.

Firstly, CFOs don't ***want*** variable pricing models. They want to ensure they can predict their bill for the year and avoid exorbitant cost spikes. Atlassian reported in their Q2 2026 earnings call that **"customer's preferred method of payment is still an understandable, predictable pricing pattern, which tends to be a seat-based model".**[2]

Four in five IT leaders report unexpected charges from consumption pricing, and 61% have cut projects due to unplanned cost increases.[3] When CFOs can't predict a line item, they cap it and adoption stalls.[4]

Software businesses need to earn the right to price on a variable metric, like tokens. They need to ensure that the value delivered to the customer aligns with that metric increasing. Atlassian can't pivot to usage-based pricing because the metric that captures value in a collaboration platform is **the number of users collaborating**.

The moat that arises is being embedded in a workflow where your software can drive measurable ROI, increasing observable metrics tied to your customer's success.

ServiceNow's platform, for example, is tied to the **Mean Time to Resolution (MTTR)** of IT issues: a metric with a direct, dollar-denominated impact on the enterprise. When ticket resolution flips from a multi-day, human-mediated task to an automated agent action, the impact is observable and measurable, MTTR declines. There is no debate over whether the customer is generating ROI or not, and ServiceNow earns the right to charge more for the tokens consumed or the resolution. That dynamic has helped Now Assist grow the number of customers spending $1 million+ by 130% YoY, with management setting a Net New Annual Contract Value (NNACV) target of $1.5 billion for full-year 2026.[5]

That said, a wave of start-ups as well as large AI Labs are competing to offer agentic tools like Now Assist. Even if companies are embedded in a monetizable workflow, they need deeper defenses to ensure that they can seize the opportunity.

Proprietary Data

One of the most widely misunderstood moats in software has always been "data moats," and the market has long debated whether they exist. The AI era makes this debate much simpler: any data owned by the customer is not a moat. Modern LLMs are adept at parsing XML, YAML, JSON, and SQL and translating seamlessly between them, making it easy to extract, refactor and migrate data to a new platform. A global insurer recently migrated a legacy SAS database to Databricks 50% faster than pre-AI, using no human intervention.[6]

Switching costs that were really a euphemism for migration complexity have shrunk by an order of magnitude, and the customer's own data has become substantially more portable than it was a year ago.

In the AI era, the only real data moat is proprietary: data the vendor owns, curates, or aggregates across its customer base. Proprietary data compounds across customers and improves the product for the entire base, creating genuine network effects. Every additional customer makes the product better for every

existing one, and AI-native challengers cannot bootstrap the same dataset without spending years winning the same enterprise relationships and the capital to do so.

These compounding advantages are strongest in security. CrowdStrike collects cross-customer telemetry amounting to trillions of security events per week.[7] This gets distilled into threat intelligence and graph artifacts, where an attack on one customer updates the logic to protect every other. At trillion-event scale, no individual customer or new entrant can replicate the threat graph.

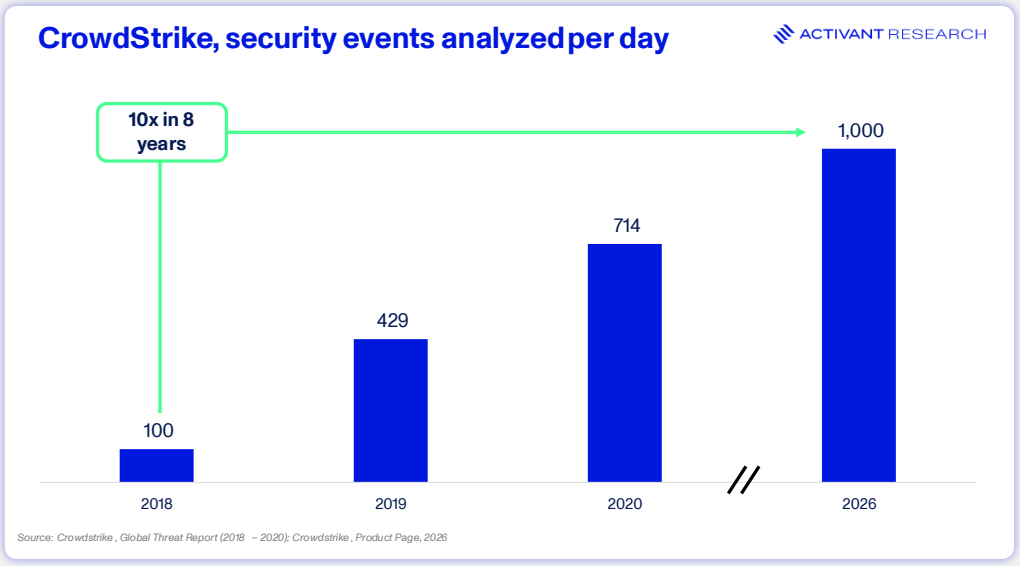

In vertical SaaS, ServiceTitan has aggregated $68bn+ of transaction volume across ~9,000 customers, generating cross-customer benchmarking on labor, pricing, and conversion that power personalized quarterly reports for each customer.[8] Klaviyo's predictive models are trained across more than 8 billion consumer profiles and over half a trillion customer interactions, allowing customers to immediately spot patterns in their own data.[9]

Critically, this framework rules out the data that most enterprise software vendors actually hold. A CRM record in Salesforce, a ledger entry in NetSuite, an HR profile in Workday, an issue in Jira, these all belong to the customer, can be exported, and (with AI) can now be reconstituted in a competing platform faster than ever before. The real moat in those systems lies in the configurations, automations, integrations, and audit trails wrapped around the data inside the platform.

Workflow Depth

When an AI SDR reaches into Salesforce to find prospects to reach out to, it isn't just grabbing a row of data, it's fetching resolved, stateful representation of business logic. The absence of that resolution is exactly where AI agents fail.

Over a year ago we wrote that data was the undervalued element of AI, and that prediction is still holding up. On average, organizations report that only 20% of their unstructured data and 32% of their structured data is "AI-ready" — correctly tagged, governed, and contextually resolved for use by an autonomous system.[10]

Most enterprise data is missing the required provenance, and that's a real risk for deploying agents.

Consider the AI SDR again. It surfaces a prospect with a "strong fit" score and no recent activity, which would be a perfect outbound target on the surface. The agent sends the email. What it didn't see: the entity is a wholly owned subsidiary of a corporation currently in the final legal review of a $4M contract with the same vendor; the parent's procurement office has explicit no-contact rules during live negotiations; and the named buyer moved out of decision authority three weeks ago. The deal slips, the relationship is bruised, and someone in RevOps spends a week reverse-engineering what went wrong.

The AI SDR didn't just need a row of data. It needed the relationship hierarchy, the vendor's account sharing rules, and the business logic that prevents outreach during live negotiations, all of which live inside Salesforce. AI needs deep, up-to-date context governed by business logic, and that's precisely where the configuration, automation, integration, audit trail, and operational logic inside software platforms shines.

In the case of Salesforce, the average organization has 24,000 metadata items like custom objects, triggers, validation rules, and permission sets.[11] Custom objects extend the generic CRM's data model and make the schema itself non-portable. Triggers run automatically when a field changes, executing deterministic business logic that should not be outsourced to a probabilistic system: generating invoices on close, syncing renewal dates, locking edits during approval windows. Validation rules enforce data integrity. Sharing rules govern who, and what, can see each record, ensuring compliance.

The audit layer compounds the moat further. Salesforce Event Monitoring captures every field change, login, report export, and API call, with user attribution and timestamps. An AI Agent manipulating JSON or markdown can't be relied upon in an audit of a financial restatement, a clinical decision, or a redress action

under data protection law. **In regulated contexts, the audit trail is not a feature of the software, it is the reason the software exists.**

Workday's payroll engine resolves jurisdictional tax rules, union agreements, benefits accruals, and statutory reporting across hundreds of countries and tens of thousands of locality combinations, and offers 850 prebuilt process automations.[12] SAP's transaction logic runs so deeply through large enterprises that implementation failures can hold up production facilities and drive millions in losses.

Workflow depth is what makes a customer's own data more valuable inside the vendor's platform than anywhere else, despite the triviality of moving the data itself. And with strength in this area, SoRs ensure that they become the center of large ecosystems, rather than being disintermediated by those building on top.

Ecosystem Centrality

An SoR with one dominant agentic interface risks disintermediation. That interface controls the user, observes the schema, and over time has the leverage to migrate the underlying data into an architecture it prefers. An SoR with two hundred intelligence layers plugging into it becomes the critical substrate that makes the entire constellation of agents possible and the place where governance, identity, audit, permissioning, and business context actually live.

Shopify's app store has ~18,000 apps that run on top of it, roughly 1,400 of which are built specifically for Shopify.[13,14] These apps derive their value entirely from how they interact with the merchant's store data inside Shopify's ecosystem. Collectively, they cover features including marketing, site experience, logistics, tax, compliance, and niche features like subscription management. For a merchant to consider moving their core order, inventory, and payments functionality out of Shopify, they would need to consider whether all of these additional ecosystem vendors can move too. When ~8% are literally built for Shopify, that's unlikely. Their ecosystem is the moat.

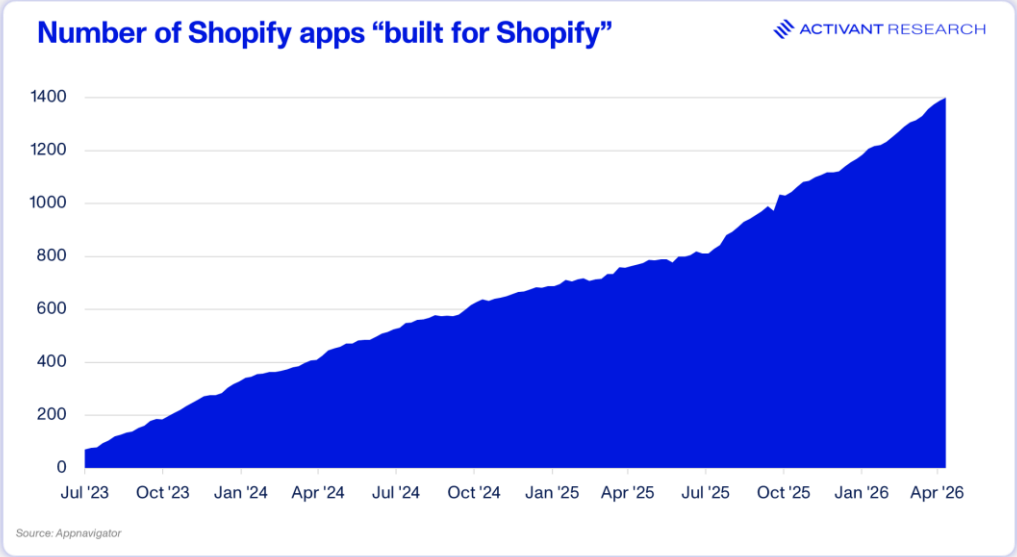

The defensive moves that enable this are platform primitives like data fabrics, MCP support, and agent-system-of-record concepts that make the SoR the natural orchestration layer for AI-native point solutions. This pillar addresses both platform transition and indirect competition simultaneously and is arguably the most strategically important for incumbent SoRs.

Vertical Specialization

If horizontal AI is the disruption threat, verticalized AI is the framework that defends against it. Verticalized software platforms can defend against AI disruption through product depth and niche and fragmented distribution.

Deep specificity: Veeva's product is a relational database wrapped in a decade of FDA-validated workflows, MedDRA dictionaries, multi-channel HCP engagement logic, and 21 CFR Part 11 controls. When Salesforce launched Life Sciences Cloud and acquired IQVIA's pharma CRM in 2024, Veeva continued to grow. The workflow specificity, not the underlying database, was the product. AI-native entrants that win in vertical software will look the same: hundreds of integrations, edge cases, and compliance rules that take years to encode and do not transfer across domains. We noted workflow depth as a moat earlier, and vertical companies take that one step further.

Niche, fragmented distribution: Horizontal AI distributes through APIs and self-serve, which works in tech-forward verticals only. Owner.com, an Activant portfolio company, leverages educational online content to build trust with small independent restaurant owners, where it now serves thousands of locations on the platform. ServiceTitan has 5 million small HVAC, plumbing, and electrical

contractors, won through industry trade shows, dealer networks, and field reps who speak the customer's language. Owning the customer in a fragmented, high-touch vertical creates a distribution moat horizontal players cannot cross without standing up the same field GTM.

Move into the Intelligence Layer

When software companies defend their incumbency with observable ROI, proprietary data, workflow depth, ecosystem centrality, and niche distribution, they win the right to move into the intelligence layer.

Datadog is the prime example: in Q1 FY26 customers representing 80% of ARR had adopted AI observability tooling, with usage of tools like their AI SRE agent and MCP server tools growing 2x-4x QoQ. This momentum, inserting themselves into the intelligence layer rather than being disrupted by it, enabled the company to accelerate revenue growth by 700bps from the year-ago quarter.[15]

By integrating copilots across every surface where enterprises already work, Microsoft has grown annualized AI revenues to $37 billion, making its AI business larger than both OpenAI and Anthropic.[16,17] ServiceNow's Now Assist, as noted, is targeting $1.5 billion ARR for FY26, approaching 10% of total ARR.

It's widely acknowledged that 365 Copilot is not the *best AI tool in the world*, and ServiceNow's users don't *love* the platform. Their success comes in spite of that, because they are integrating AI into the surfaces where real people actually get work done, through tools that procurement has already approved and on data architectures already cleared by compliance.

This is the incumbency moat. If existing software companies can leverage the moats discussed in this paper, they will not only avoid being disrupted by AI, but significantly benefit from it, reaccelerating revenue and finding new TAM to pursue.

The counterfactual is Salesforce, which has done everything to organize the business around agents, but the success of Agentforce has been sub-par, with ARR reaching just ~2% of total and the company ceding significant CRM market share over the past two years.[18,19]

That's why the final moat in this framework is really about execution. The competitive advantages of incumbency win you the right to challenge for AI revenues, but they don't automatically mean that you win. Incumbents need to not only rest on what they've already built, but also pull up their sleeves and build great products for the AI era.

A Framework for Success

When we began this series, software was trading 40% below its November 2025 peak, with the market applying a uniform discount to a sector that is, in reality, about to undergo its most discriminating reshuffling in two decades. That distinction is the point.

The five moats we have laid out, pricing tied to observable ROI, proprietary data that compounds across customers, workflow depth that resolves business logic AI cannot fake, ecosystem centrality that makes the platform the substrate for hundreds of AI-native point solutions, and vertical specialization that horizontal models cannot cross into are not just a checklist for survival. They are the conditions under which a software company earns the right to fight for the labor budget AI is now unlocking.

The companies that fail this test will look closer to the market's bearish pricing and, in some cases, deserve to. A platform whose value lives entirely in a UI that an LLM can re-render, whose data belongs to the customer and travels with them, whose workflows are shallow enough to be reconstructed from a wrapper, will not be saved by incumbency alone. The 40% drawdown is not wrong as an aggregate signal; it is wrong as a uniform one.

Software is dead. Long live software.

[1] Jefferies, Key Takes from our PANW meeting at RSA, 2026

[2] Atlassian, Q2 2026 Earnings Conference Call, 2026

[3] Mitchell Mackey, How AI is reshaping SaaS stratgy, 2026

[4] Monetizely, The Doomed Evolution of Salesforce's Agentforce Pricing, 2026

[5] ServiceNow, Q1 2026 Earnings Conference Call, 2026

[6] TD Cowen, Services GenAI360, The Tide Is Choppy & Won't Rise Evenly, 2026

[7] CrowdStrike, Form 10k, 2026

[8] ServiceTitan, FY ended January 31, 2025 10-K

[9] Klaviyo, Q3 FY2025 Earnings Conference Call

[10] Snowflake, The ROI of GenAI and Agents, 2026

[11] Hubbl, Understanding metadata: The key to unlocking a healthy Salesforce org, 2025

[12] Workday, The Workday business process framework, 2026

[13] Craftberry, Shopify App Store Statistics, 2026

[14] Appnavigator, Statistics, 2026

[15] Datadog, Q1 FY26 Earnings Conference Call, 2026

[16] Microsoft, Q3 FY26 Earnings Conference Call, 2026

[17] Anthropic's revenue is rumored to have reached $44bn, but the latest public disclosures from the company stand at $30bn

[18] Salesforce, Q4 FY26 Earnings Presentation, Activant Analysis

[19] See analysis in article 1