

Software Misunderstood

By Rajat Sharma, Managing Director, Software Coverage, Technology, Media and Telecom, Investment & Corporate Banking, Mizuho Americas

Recent earnings, and reactions to those earnings, across the application software and data infrastructure ecosystems have sharpened a pattern that, just a few months ago, was hard to identify when looking at any single software sub-sector purely in isolation.

Five Points That Stand Out:

1. The data and observability infrastructure layer in software is not being disrupted by AI—it is rather being amplified. Consumption of content/usage compounds with every new workload. The architecture that looked fragile in what was in the past a slower-growth environment, turns out to be the most natural fit for the world being built. This amplification is not yet at a steady state—enterprises today are using unbudgeted spend on tokens and accepting trade-offs on accuracy and governance that over the long run will be hard to hold. As reference architecture normalizes, the same forces which are driving consumption will demand cost discipline, accuracy and governance. All of these factors compound back into the data infrastructure layer.

2. Gravitational “lock-ins,” either in data and deployment or in workflows and transactions, appear to be the most durable moat in the new agentic era. The companies which have achieved one of these two anchors appear to be growing faster, at higher margin profiles and with more committed forward revenues than at any point in their history. The workflow anchor itself is being tested in two waves: the first wave involves automation of manual processes and the displacement of “single-use-case departmental SaaS;” and the second wave, involves the orchestration of agents by individuals where productivity improvement with respect to spend on AI is inquired by Boards and C-Suite.

3. The application layer without gravitational anchoring does not hold at scale. This is no longer a thesis—it is realized in reported results—and should matter to every software business that is honest about where its own anchor sits. Compounding this pressure: the SaaS data anchor itself seems to be eroding as customers move their operational data into the warehouse, leaving integrations and user interface as the residual moats—both real, neither durable. The era of best-of-breed point solutions is giving way to a new era of platforms where software works seamlessly across departments rather than inside a single-use-case function.

4. Nobody has, to the best of my knowledge, built the system of record for the blended workforce—humans and agents—and the window to claim that position will close faster than the companies best positioned to own it appear to recognize. When an enterprise runs 500 humans alongside 5,000 AI agents in parallel, questions arise—who onboards these agents, who governs their compliance, who tracks their performance, who offboards them? The answer today is: nobody, systematically. In my opinion, that is not a gap that stays open for long. Snowflake's May 2026 announcement to acquire Natoma appears to be a step in that direction.

5. The death of the seat model is not primarily a pricing problem, but rather a potential budget category migration, that is quietly transferring the enterprise software buying decision from the CIO to the CFO. The vendors who figure this out first will restructure their entire go-to-market around it. The ones who don't will keep selling to the wrong person. These lessons are being played out through corporate earnings.

What Is the Data Layer Telling Us?

Snowflake's Q1 FY2027 (ended April 2026) was its strongest quarter under its CEO Sridhar Ramaswamy. Snowflake Intelligence, its agentic AI product, was described by the company as the fastest adoption ramp in company history. The consumption model, long criticized for unpredictability, is proving to be exactly the right architecture for AI workloads—because AI consumption compounds as models improve and workflows deepen. Similarly, company performance suggests that Databricks is well positioned with its Databricks Unit (DBU) framing and educating the marketplace which in turn is driving growth and customer adoption. And finally, Datadog's Q1 CY2026 growth with \$1bn of quarterly revenue has reinforced the value of its observability layer, with an acceleration in the non-AI-native core business growth. Datadog's Bits AI SRE agent has been described by the company in several recent earnings calls as reaching a large number of paying customers within weeks of its availability. With each new AI agent deployed in the enterprise, it generates observable, billable infrastructure workloads. Datadog sits at the exact layer where those workloads surface.

What Is the Orchestration and Deployment Layer Telling Us?

Palantir's Q1 CY2026 results illustrate the boot camp model with forward deployed engineers has crossed from clever tactic to structural advantage. The enterprises that now embrace it have adopted a deployment-speed benchmark that its other software vendors are implicitly being measured against. ServiceNow's Q1 CY2026 business trends, along with the recent acquisitions, also point to a progression and amplification of this trend: customers expand AI spend in addition to its license renewals. Management has said explicitly: the platform is structurally complete. From here it is execution, not acquisition.

What Counterpoint Is the Market Underweighting?

In the strongest AI demand environment on record, execution misses get punished as structural issues—read as proof of no data moat, no workflow gravity. That fear compounds with the prevailing view that application-layer AI is a wrapper—compressed from below and displaced from above, simultaneously. The current tape is the clearest signal yet of my premise: the application layer without gravitational anchoring does not hold at scale.

The Two Truths the Earnings Data Cannot Tell, but the Organizational Reality Can:

On the blended workforce: the conversation in venture and in the market broadly has been about agents replacing humans, or agents assisting humans. Both framings miss the more precise and more immediately operational question. When agents proliferate inside an enterprise at scale, in production, created by business users, they need to be onboarded into systems, credentialed, governed for compliance, audited for performance and eventually decommissioned. That is a workforce management problem, not a technology problem. The company that builds the system of record for the blended workforce doesn't just capture a new product category—it becomes the compliance and governance layer for the next era of enterprise operations. The infrastructure for this does not exist yet in any coherent form today, but as the CEOs and Boards demand increased adoption of AI and as the business users vs. “coders” create their own army of AI agents, it will need to get solved. The eventual winner will need to satisfy two requirements that today's contenders largely do not: governed access to data across multi-platform infrastructure rather than a single warehouse, and easy interoperability with the unstructured sources where most agent context actually lives.

On the budget migration: per-seat software is a headcount-linked expenditure. It lives in the same budget as people, a category managed by HR, sanctioned by the CIO, rationalized by headcount planning. Consumption-based AI is an operations expenditure—variable, tied to usage, rationalized by output and ROI. That is a fundamentally different budget line, owned by a fundamentally different executive, namely the CFO. The vendors who are still presenting AI ROI cases to CIOs are already behind. The go-to-market motion that wins the next decade of enterprise AI runs through the office of the CFO—because that is where the budget is migrating and where the decision authority is consolidating. One unresolved question sits underneath the migration: today's consumption variables are largely cost-plus passthroughs of vendors' own infrastructure economics, not measures aligned to customer outcomes. Whether the next generation of pricing converges on outcome-linked variables—or stays anchored to vendor cost structures dressed up as utility—will determine which vendors get to sustainably price as workloads scale, and which get squeezed.

This communication is for information purposes only and should not be construed as investment, legal, audit or tax advice nor as a solicitation, offer, recommendation, or commitment to buy or sell any investment instruments, to effect any transactions, or to conclude any legal act of any kind whatsoever. We make no representation or assurance regarding the accuracy, completeness or suitability of the information in this communication and have no obligation to provide updates or changes. This information should not be used to make any investment decisions.

Please see <https://www.mizuhogroup.com/americas/disclosures> and <https://mizuhoamericas.com/email-disclaimer> for important additional disclosures and terms and conditions relating to this communication and your reliance upon the contents herein.

See how we look
Beyond The Obvious

Find out more >