


Setting Firm Foundations for Migration: insightsoftware's Axel Streichardt on How to Successfully Manage an SAP Transition



Axel Streichardt

*General Manager & VP, SAP Products
insightsoftware*

As the general manager and VP of SAP products at insightsoftware, Axel Streichardt interfaces directly with many of the key challenges businesses face while preparing to migrate to SAP S/4HANA. As the end of ECC maintenance approaches next year, companies continue to find themselves embarking on 18-month projects. To complement this transition, insightsoftware offers services focused on gaps in cross-system integration, self-service analytics, advanced analytics and trend analysis, BI tool integration, AI-enhanced intelligence, and process automation.

A [recent ASUG survey](#) found that while 39% of SAP landscapes are building foundational knowledge for AI implementation and 41% are amid active experimentation or piloting, only 10% of those surveyed have been able to achieve enterprise-scale rollout.

Following SAP's announcement of its vision for the Autonomous Enterprise at this year's SAP Sapphire & ASUG Annual Conference, Streichardt sat down to discuss the data remediation that has to precede every migration, AI's role in identifying previously uncaught errors, and the future of integration with SAP products.

This interview has been edited and condensed for length and clarity.

Q: You've spent more than 20 years working with SAP, at SAP itself, and at AWS, EMC, and Pure Storage before insightsoftware. Across all those seats, what do SAP customers consistently get wrong about a transition like this?

Customers structure migration around capability readiness: Do we have the right infrastructure? Do we have to be in the cloud? Is the application configured? What they don't ask is, "What data do we actually have, and is it fit for purpose on the other side?" That's the discrepancy—they have no idea how broken their data is until they're halfway through. If 30% is broken, that's a six-week remediation effort, and someone owns this data going forward or it decays again. Whatever data quality sins you let slip will come back to haunt you when you try to activate AI 18 months later.

Q: At SAP Sapphire, SAP put its vision for the autonomous enterprise at the center of everything: agents orchestrating end-to-end processes, applications that execute work rather than record it. What's your honest read on that vision?

At SAP Sapphire, I was talking about the \$15 million gap that every CFO has. That's real, not abstract—it's losing money because you can't see the gaps in your data. In one conversation, a CFO said, "I have visibility for 60% of our spend. The rest is dark data. I don't know if we're overpaying suppliers, missing volume discounts, or carrying redundant inventory, because I can't connect the dots between procurement, finance, and supply chain."

Margin leakage happens because your data sits in silos, and supply chain issues are the worst. A customer discovered they were systematically paying

premium freight costs for expedited shipments that were going to the wrong warehouse because their order system wasn't talking to their inventory system. That cost them \$2.3 million annually. The data gap was invisible, but then we pivoted to SAP autonomous enterprise agents orchestrating workflows. Connect your master data so finance, supply chain, and operations are talking, then layer AI on top to automate finding problems. We've been doing that foundation work for over 20 years. Adding AI is the logical next step, which is why we launched AI automation capabilities at the end of June to surface cost and supply-chain challenges automatically that usually go undetected.

Q: Only about 1 in 10 customers have scaled AI past a pilot, and most are after something more modest than full autonomy—AI that supports people rather than replaces them. How do you square where SAP is pointing with where customers actually are?

CFOs and supply chain leaders want AI and have funding for AI projects, but first they need to trust it. Trust only comes one way: The AI has to be right. So, the conversation usually starts differently—not “Here’s an AI agent that will automate your decisions,” but “Here’s AI that will give you the right answers to your hardest questions.” Once the customer sees that the analysis is accurate, you can move to proactive and have the system automatically flag financial anomalies and supply chain challenges.

Q: Master data is one of the biggest [unsolved problems our members report](#). This is your home turf. When you get into a customer’s environment, what’s actually broken in how their data sits?

Master data lives in silos, so there’s no single source of truth for identity. A global manufacturer has GL codes in SAP, cost center codes in their supply chain system, account codes in a legacy order system—same cost, three different identifiers. When finance

asks, “What did we spend on this cost center?”, the answer depends on which system you’ve asked. Nobody owns reconciliation, and there’s no data validation process across systems, so bad data propagates from system to system.



Q: A lot of your customers are still mid-migration, and moving to S/4HANA is [the single biggest challenge our members name](#). Meanwhile, RISE now commits them to switching on Joule assistants in year one. How does the AI conversation land for a company that hasn't finished the migration?

It doesn't, really. One CIO made it very clear to me: “We are still loading master data and fixing cutover issues. Joule is year two at the earliest.” Migrations and Joule are competing for the same resources, and migration wins. So, we're helping customers today with AI to accelerate the migration itself during that brutal cleanup and master data remediation phase. That gets them to clean data faster when Joule arrives. They can start using AI today to win their migration.

Q: ASUG research found that [how a company approaches AI](#) predicts success better than its size or budget does. What separates the customers pulling ahead from those stuck in pilots?

The ones that win start with a process, not a model: “Here’s the process that’s broken. If we fix this, here’s the business value. Can AI help?” Often the answer is “no,” because you need governance first, but when the answer is “yes,” they have clear success criteria from day one. Pilots that fail usually start with, “Hey, let’s build an ML model to predict X.”

Second, successful customers measure business outcomes, not model accuracy. Pilots that fail measure “We are reducing month end close by x days.” Model accuracy is table stakes. Outcomes are what really matter.

Third, successful customers have a point person with actual authority. Pilots that scale have a CFO or CEO, or even a VP of finance, who sponsors the work and takes accountability for results.

Connect with Axel Streichardt
on LinkedIn.



Q: You frame AI as increasing trust in the data. At the same time, governance and security are the top concerns customers cite when it comes to adopting AI. How do you reconcile those, and why should a CFO trust an AI-generated insight more than the spreadsheets they’ve used for 15 years?

Fifteen years of spreadsheets have not solved the \$15 million gap I talked about. A spreadsheet is passive; it shows you what you already know to look for. A CFO then has to hunt through rows of data to find anomalies, variances, and hidden problems. AI is

different. Governance doesn’t constrain AI; it enables trust. If I can see how the AI model was trained, which rule it’s overriding, or who approved the exception, I can trust it.

Here’s a concrete example: We implemented automated GL reconciliation for a finance service firm. The AI system flags reconciling items outside normal variances and bands and surfaces root cause anomalies automatically. Every flagged item gets a reason code and validation rules, and approvals are locked with business justification. Finance went from 40 hours a week on GL-reconciliation to under an hour.

Q: Your positioning has been to complement SAP rather than replace it. SAP has acquired Reltio and Dremio, and the company is building SAP Business Data Cloud to combine SAP and non-SAP data, an area close to what you’ve described as your own. How do you see the relationship evolving from here?

SAP owns the transaction system and process layer, but what SAP doesn’t own is the work of integrating SAP data with supply chain platforms and non-SAP ERP legacy systems like Oracle or Salesforce.com. We’ve been doing data integrations for 20-plus years across thousands of customer engagements. That experience can’t be acquired—it has to be built. That’s why SAP partners with us as a SolEx partner. They don’t want to replace us; they want to work with us.

Q: Microsoft’s Copilot is in active use at most organizations today, while adoption of SAP’s own embedded AI is still early. Your products work across both. How should a customer think about where to place their bets while the AI layer of their landscape is still taking shape?

You don’t want to tell someone that the investments they did, the millions of dollars they spent, should be replaced by something else. That’s why the

answer isn't choosing one over the other; it's letting both coexist. Copilot is a general intelligence layer: summarize this variance report, draft an email explaining this. SAP is embedded AI trained on SAP data and processes. We help customers integrate both into their standards, processes, and governance. A real example: a customer is using SAP AI to identify the top variance to suggest the GL codes. Finance analysts review the recommendation, then use Copilot to draft an explanation for the CFO. So, SAP AI isn't being replaced by Copilot; they're complementary. One finds the problem, the other helps communicate the solution, and we are the platform that makes both work together seamlessly.

Q: Picture the practitioner who came home from Sapphire energized and now faces their actual landscape on Monday. What are the unglamorous first steps that never make it into a keynote slide before any of this is real?

Audit your data. This probably takes four to six weeks, and it's boring. You're running SQL queries, sampling records, checking for orphan entries. Start with your most critical data sets: GL accounts, cost centers, customer master, supplier master. Pick one sample. Once you understand the scope of the problem, you have a baseline.

Second, document one critical process end-to-end. Pick a process that's manual, with lots of human stats involved, error-prone—for instance, close operation—then document it, not as a process map but as an actual walkthrough. Who does what, when, with what systems? Where are the handoffs? This is one week of work.

Third, fix your master data, starting with the highest-risk items. Don't try to fix everything. Pick the data that creates the most business impact when it's wrong for finance, GL accounts, and cost centers. Audit those specific data sets, identify the junk, and create a remediation plan. This is six to eight weeks

of work. Baseline: expect 20–30% of master data to need remediation; companies that do this upfront save 60% of migration rework downstream.

Q: A lot of vendors talk about these AI capabilities, and when you ask whether they're really there, the answer is often no. Where do you see the gap between ambition and reality right now, and what should customers be most discerning about?

What frustrates me the most is when vendors put a little icon on the slides for data cleanup, data context, data enhancement, then spend the rest of their presentation on flashy dashboards. Everything is focused on the results, not on how to get to these results. You can't automate your way around the foundational work of data cleanup context and enhancements. So, customers should be discerning about vendors who gloss over the 80% data cleanup context and enhancement, because it's invisible and hard to demo. Ask them, "How many customers have you actually done this with? What experience do you have in data cleanup and enhancement at scale?" That's where the real value lives, even if it's the hardest to show.

Q: If every customer reading this took away one thing before chasing that vision of SAP's autonomous enterprise, what would it be?

Get the fundamentals right first. The autonomous enterprise vision is real, no question. It's also three to five years away for most organizations. You can't get there from here by buying more software; you get there by doing hard work first. Don't chase autonomous enterprise—chase the next level of operational maturity.

Visit the [insightsoftware](https://www.insightsoftware.com) website.





insightsoftware is a global provider of comprehensive solutions for the Office of the CFO. We believe an actionable business strategy begins and ends with accessible financial and operational data. With solutions for SAP customers across financial planning and analysis (FP&A), accounting, IT, and supply chain operations, we transform how teams operate, empowering leaders to make timely and informed decisions. With data at the heart of everything we do, insightsoftware enables SAP customers to automate processes, deliver trusted insights, boost predictability, and increase productivity.



ASUG is the world's largest SAP user group. Originally founded by a group of visionary SAP customers in 1991, its mission is to help people and organizations get the most value from their investment in SAP technology. ASUG currently serves thousands of businesses via companywide memberships, connecting more than 130,000 professionals with networking and educational resources to help them master new challenges. Through in-person and virtual events, on-demand digital resources, and ongoing advocacy for its membership, ASUG helps SAP customers make more possible.