

SAP IDEA EXCHANGE

Volume 2

Innovate. Collaborate. Inspire

In association with
Fayetteville State University





Contents

Foreword	4
Executive summary	6
Chapter 1 The intelligent core: Unlock SAP S/4HANA potential with generative AI	9
Chapter 2 The future of supply chains: Where S/4HANA, AI, and humans work together	15
Chapter 3 Unlock the power of data in SAP S/4HANA transformations: Why starting early makes all the difference	22
Chapter 4 Build the right tech foundation for SAP S/4HANA transformation	27
Chapter 5 The human blueprint for AI-driven SAP S/4HANA success	33
Chapter 6 Govern the AI: Responsible intelligence for SAP S/4HANA enterprises	39
Chapter 7 Scalable AI in the enterprise: The role of small language models	45

Foreword



Ulysses Taylor

*Dean, Broadwell College of Business and Economics,
Fayetteville State University*



Tarang Puranik

*Executive vice president, global service offering head
(SAP) and head Europe, Infosys*

“ Technology sets the pace of change. People determine how far an enterprise can go with it.”

A moment of transformation

Enterprises are at an important turning point. AI is no longer operating on the sidelines: it is shaping decisions, redesigning workflows, and redefining operational excellence. A big part of this shift is SAP S/4HANA, now enhanced by generative AI. Together, these capabilities connect supply chains, finance, analytics, and customer engagement to turn data into actionable intelligence.

The human factor

As this magazine illustrates, intelligent enterprises are not built by technology alone. They are created through the integration of different systems, design, and human capability. AI plays a critical role in accelerating automation, but delivering meaningful outcomes requires human judgment and responsible governance.

A partnership with purpose

This partnership between Infosys and Fayetteville State University (FSU) is to create an academic record on how to prepare emerging talent for a world where intelligent systems are core to business performance. This collaboration looks to equip the next generation of professionals with the skills, insights, and applied experience needed to lead enterprise transformation.

This publication brings together perspectives from SAP practitioners, university faculty, and writers from the Infosys Knowledge Institute. Their insights converge on how responsibility lies with enterprise leadership to make decisions that affect outcomes, risk, and sustainability.

What you'll discover

From unlocking SAP S/4HANA's potential with generative AI to creating the right building blocks for efficient management and governance in the context of SAP S/4HANA, the chapters in this publication explore how technology and talent come together to shape the future-ready enterprise.

Across supply chain planning, process mining, data orchestration, and analytics-driven design, a consistent theme emerges that technology delivers its greatest value when it amplifies human potential and is supported by continuous learning.

Looking ahead

As you explore these chapters, they can be viewed as a blueprint for future collaboration that emphasizes responsible design, human-AI synergy, and the overall future of SAP solutions.

The future of enterprise transformation is collaborative, intentional, and deeply human. This publication is a step toward shaping that future.

Executive summary



Enterprise transformation is entering a decisive phase as organizations accelerate change through SAP S/4HANA and AI. These technologies promise faster decisions, predictive operations, and more adaptive systems. SAP S/4HANA has become the digital core for modern enterprises. Its impact grows when paired with AI-driven automation and analytics that reduce manual effort, enable real-time decisions, and improve efficiency. However, system modernization by itself is not enough. Successful transformation depends on the broader ecosystem that surrounds technology.

This SAP magazine brings together insights from Infosys leaders and practitioners, along with professors from Fayetteville State University. It explores how organizations can unlock greater value from SAP S/4HANA by embedding AI responsibly across processes, platforms, and people. Building an intelligent enterprise is a collective effort that depends on collaboration across institutions.

Partnerships are essential when change is the only constant. AI evolves quickly, business models shift constantly, and regulations continue to grow more complex.

Universities and industry now need to work together rather than separately. While academia can bring its share of critical thinking, ethical perspective, and talent development, helping organizations question assumptions and design technology

responsibly, industry can contribute to real-world context, scale, and execution, showing how ideas perform in live operations and turning them into measurable results. These partnerships create a continuous cycle where talent, technology, and governance develop in step, making transformation sustainable rather than short-lived.

Rather than treating AI as an add-on to enterprise resource planning (ERP) tools, we emphasize integration: between S/4HANA and generative AI, between human decision-makers and AI agents, and between technology platforms and governance. The message is clear: intelligent operations depend on capabilities built through education, development, and partnership, not simply through software deployment.

The seven chapters outline a shared agenda for building intelligent enterprises:

- **Chapter 1: The intelligent core: Unlock SAP S/4HANA potential with generative AI**
This opening chapter explores how SAP S/4HANA evolves into an intelligent core when combined with generative AI, bringing together transactional integrity and AI-driven reasoning to enhance decision-making and responsiveness.
- **Chapter 2: The future of supply chains: Where S/4HANA, AI, and humans work together**

Here, the authors examine future operations where AI agents handle repetitive and analytical tasks while humans focus on judgment, exception handling, and strategic decisions.

- **Chapter 3: Unlock the power of data in SAP S/4HANA**
transformations: Why starting early makes all the difference The authors emphasize data readiness as the cornerstone of AI-enabled S/4HANA success, ensuring accuracy, trust, and relevance for reliable outcomes.
- **Chapter 4: Build the right tech foundation for SAP S/4HANA transformation**
It focuses on designing robust, scalable, and future-ready architectures that support evolving AI capabilities and business needs.
- **Chapter 5: The human blueprint for AI-driven SAP S/4HANA success**
The authors address organizational and cultural change, helping employees trust AI agents and integrate them into daily workflows.
- **Chapter 6: Govern the AI: Responsible intelligence for SAP S/4HANA enterprises**
This focuses on governance, security, ethics, transparency, and compliance as essential to responsible AI-enabled S/4HANA operations.



- **Chapter 7: Scalable AI in the enterprise: The role of small language models in SAP and hybrid landscapes**

Finally, this chapter showcases how organizations can apply fit-for-purpose small language models to optimize SAP-centric architectures while improving efficiency and control.

Together, these chapters set out a central message: that SAP S/4HANA and AI can modernize processes, but only people can modernize decision-making. The organizations that will lead the next decade will be those that invest not just in systems, but in relationships, institutions, and shared learning that give those systems life. Intelligent transformation is a collaborative journey, and it succeeds when we build it together.

Chapter 1

The intelligent core: Unlock SAP S/4HANA potential with generative AI

Dr. Murat Adivar

*Professor, Fayetteville State University,
North Carolina*

Organizations today are not short on data. The area where they struggle most is using that data promptly enough to keep pace with day-to-day requirements. Across industries like manufacturing, utilities, retail, and the public sector, decision-makers are less concerned with pursuing new transformation programs; instead, they are more concerned with delays in reaching the final go-ahead, followed by approvals and execution. When conditions change, many organizations still find it hard to react in time.

SAP S/4HANA has become central to addressing this challenge. By bringing finance, operations, supply chain, and customer information into a single platform, it provides a complete view of how the business is running. For many organizations, this has already improved consistency and control. The remaining gap lies in how quickly that information is translated into action.

A strong digital core

SAP S/4HANA gives organizations a solid operational foundation. It enables data to move across functions in real time, reducing delays caused by disconnected systems. Leaders can see what is happening across different departments like procurement,

production, finance, and customer operations more clearly.

Yet visibility alone does not resolve execution challenges. Many core processes still rely on manual steps, and as transaction volumes grow, the effort required to manage those steps increases. Approvals, validations, and exception handling often sit outside the system or require additional coordination. Over time, these steps slow response times and make it harder to scale.

Manual processes create friction

If anyone walks through a large manufacturing operation, it becomes clear where manual work reaches its limits. Maintenance teams rely largely on experience to spot issues. Although SAP is in place, planners often feel more confident tracking changes in spreadsheets. This increases the likelihood of small delays that can compound quickly. A single equipment issue can disrupt production schedules, supplier coordination, and customer commitments.

[An Indian cement manufacturer](#) encountered this challenge as its operations expanded. While SAP provided a steady backbone, manual steps across procurement and operational workflows reduced visibility and slowed response times. Teams spent significant effort managing approvals and

checks instead of focusing on improving outcomes. Similar issues can happen in other sectors with overreliance on manual processes to complement SAP systems.

How AI helps close the execution gap

Generative artificial intelligence (AI) addresses this gap in SAP S/4HANA environments by removing the manual effort required to turn system information into action. While S/4HANA provides visibility and structure, generative AI provides an extra layer of validation that works at the point where data needs to be interpreted, validated, summarized, or prepared, operating directly within SAP workflows to support these steps.

Organizations are embedding generative AI into SAP S/4HANA using [SAP Business AI](#) on the [SAP Business Technology Platform](#) (BTP). SAP Business AI is SAP's portfolio of AI capabilities embedded into business applications and processes, combining generative AI, predictive models, and rule-based logic. These capabilities operate on live enterprise data and are applied within existing workflows rather than being introduced as separate tools.

In practice, the need is to deploy generative AI at points in the process where execution slows, typically before approvals, validations, or exception handling, or, in short, at the

point where decisions are made. It supports tasks such as interpreting requests, validating input, summarizing exceptions, and preparing context for review. These tasks still require human oversight, but the effort needed to complete them is eased.

From theory to practice

For the manufacturer, integrating generative AI simplified work inside SAP rather than adding new process layers. It supported approval preparation, data checks, and process validation within procurement workflows. As transaction volumes increased, teams kept pace without creating additional workarounds, and the system remained manageable as operations scaled.

For easing [finance and shared services](#), AI assistants are embedded in SAP to support high-volume inquiries by interpreting requests expressed in natural language, retrieving information from SAP S/4HANA, and preparing responses using current transactional data. These assistants typically run on SAP BTP, where generative AI models handle language understanding and summarization, while SAP business logic and authorization controls determine what data can be accessed and how responses are framed. Humans retain control over exceptions, but routine interpretation and information retrieval are handled by the system.

“In practice, the need is to deploy generative AI at points in the process where execution slows, typically before approvals, validations, or exception handling, or, in short, at the point where decisions are made.”

In [retail](#), generative AI uses operational and customer data from SAP S/4HANA to support product recommendations and in-store decision-making. Large language models (LLMs) are used to interpret customer context and interaction signals, while domain-specific logic grounded in SAP master data, pricing, and inventory determines which recommendations can be generated. These capabilities are embedded into SAP-supported retail workflows, allowing store associates and sales teams to access contextual insights directly, without relying on separate analytics tools.

In [professional services](#), AI copilots called [SAP Joule](#) are being embedded into SAP workflows to assist with analysis and preparation tasks that support service delivery. Joule operates on SAP BTP, connecting securely to SAP S/4HANA

through interfaces and services. It combines generative AI capabilities for summarization and reasoning with SAP-specific business data and processes to help consultants retrieve relevant data, prepare analyses, and draft materials. For professional services companies, the value lies in preparing information and surfacing context so that it can complement human judgment and retain accountability for decisions.

Across these use cases, generative AI functions as an interpretation and preparation layer within SAP S/4HANA. It does not make decisions or replace business rules. Its role is to reduce the manual effort required to move work to a decision point, allowing organizations to execute more consistently as volumes and complexity increase.

Design for reliable execution

Seeing benefits from generative AI starts with making deliberate choices about where and how it is used.

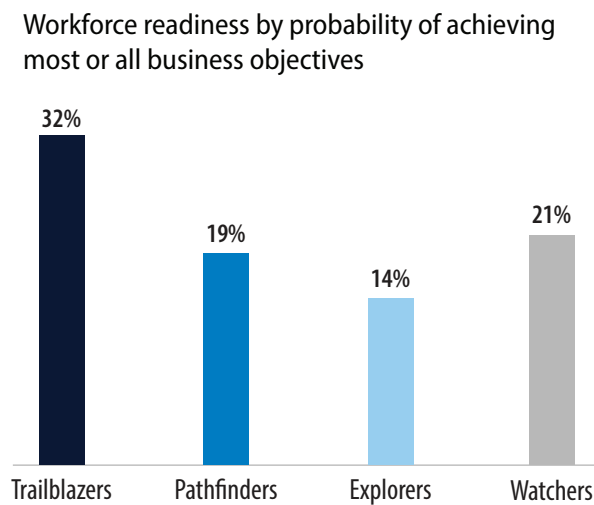
The most effective starting points are specific steps in core S/4HANA processes where manual effort consistently slows things down, such as approvals, validations, exception handling, or large volumes of routine inquiries. These steps tend to be repeatable and measurable, which makes them practical candidates for AI support.

Generative AI works best when it is built directly into existing SAP workflows. As discussed through the case studies above, AI helps interpret information, assemble context, or prepare work inside the system which teams can use naturally as part of their day-to-day work. Introducing separate tools or parallel solutions often creates extra friction and limits adoption, even if the technology itself is sound.

However, technology alone is not enough. Research from the Infosys AI Business Value Radar 2025 shows that organizations see the strongest results when AI adoption goes hand in hand with changes to processes, operating models, and workforce readiness. Teams need clarity on how AI-assisted outputs should be used and when human judgment is required. Simplifying processes before adding AI and helping people adjust to new ways of working can increase the success rate of AI initiatives. Figure 1 shows that organizations that are more advanced in their AI journey report a higher likelihood of achieving most or all business objectives, with Trailblazers at 32% compared to 19% for Pathfinders, 14% for Explorers, and 21% for Watchers.

Governance also needs to be addressed early. AI outputs should be based on trusted SAP data, follow existing business rules, and remain traceable. Human oversight should be clear, especially for approvals

Figure 1: Readiness boosts outcomes



Source: Infosys AI Business Value Radar

“Research from the Infosys AI Business Value Radar 2025 shows that organizations see the strongest results when AI adoption goes hand in hand with changes to processes, operating models, and workforce readiness.”

and exceptions, so there is no ambiguity about accountability. Ethical considerations, including transparency, explainability, and bias, are best handled as part of the design, rather than being treated as an afterthought.

Finally, leadership alignment also needs to be prioritized. Business and technical teams

need to be clear about what they want to improve, whether that is faster turnaround times, more consistent outcomes, or the ability to scale without adding complexity.

When those goals are shared and measured, generative AI becomes a practical extension of SAP S/4HANA that supports execution while maintaining control and trust.

Conclusion

Embedding generative AI into SAP S/4HANA changes how organizations manage everyday work. S/4HANA provides structure and consistency. Generative AI reduces the manual effort that slows decisions and execution. Together, they help organizations respond more steadily as conditions change.

Adding generative AI does introduce new technology, but the value comes from how it is integrated. When generative AI is embedded into existing S/4HANA workflows, applied to well-defined tasks, and governed by clear guardrails, it improves execution without adding unnecessary complexity. In this form, generative AI supports interpretation, validation, and routing work while keeping decision authority and accountability within established processes. The result is a more practical intelligent core that scales with the business and supports how teams operate.



Chapter 2

The future of supply chains: Where S/4HANA, AI, and humans work together

Radhakrishnan Jagadeesan

*Industry principal, SAP practice,
Infosys, Texas*

Supply chains are under pressure. Between unpredictable demand swings, shipping delays, complex global regulations, and rising sustainability expectations, businesses are feeling the heat. Many have already invested in digital tools — but traditional systems often operate in silos, struggle to synchronize planning and execution, and rely heavily on manual intervention to keep things moving. This is especially true in sensitive areas like cold chains or regulated logistics, where timing, traceability, and compliance are nonnegotiable.

SAP S/4HANA replaces fragmented, human-dependent processes with an intelligent, unified digital core. By integrating real-time data, end-to-end process automation, and embedded analytics, it eliminates many of the bottlenecks created by outdated architectures and manual workflows. The result is a supply chain that can operate with speed, precision, and resiliency as per modern business demands.

And when you layer in AI agents, powered by technologies such as small language models (SLMs), machine learning (ML), and cognitive automation, you get something even more powerful: a supply chain that doesn't just react, but learns, adapts, and makes intelligent decisions in real time.

Where do traditional systems fall short?

Many enterprises still struggle to fully modernize their supply chain operations. The challenges are now more about how to achieve true coordination, end-to-end visibility, and real-time adaptability across complex networks. From disconnected planning and execution to the complexities of cold chain logistics and evolving demands of global compliance, supply chains face persistent friction that slows decision-making and impacts performance.

Here are some of the most pressing pain points holding businesses back:

- **Planning-execution disconnect:** Forecasts and execution decisions often operate in isolation, causing inventory misalignment, stockouts, or excess.
- **Cold chain and regulated logistics:** Temperature-sensitive goods need real-time monitoring, predictive alerts, and compliant documentation — all of which are hard to manage manually.
- **Global visibility and traceability gaps:** Legacy systems struggle to track batch lineage, recall impacts, or multitier supplier compliance.
- **Delivery date complexity:** Traditional systems can't dynamically

analyze factors like production time, transportation routes, and batch availability. This leads to unreliable delivery dates and missed commitments.

- **Human oversight needed:** Despite automation, exception handling — like supplier delays or customs clearance — still relies on human judgment. But without the right insights, responses are slow and error prone.

SAP and AI power resilient supply chains

Bridging these gaps in supply chain management calls for creating systems that support the people running them. SAP S/4HANA and AI agents enable this by turning fragmented processes into an integrated, adaptive network, where human judgment and machine intelligence operate side by side.

These technologies don't replace expertise, they amplify it. Planners and operators gain real-time insights, predictive models, and simulation tools that help them anticipate change and act decisively. Instead of reacting to disruptions, businesses can sense, explain, and adapt in sync with market realities.

By combining SAP's digital core with AI-driven technologies such as SLMs, ML, and cognitive automation, organizations can connect planning and execution, simplify

compliance, and create systems that learn continuously while keeping humans in control.

Smarter planning


Traditional planning cycles rely heavily on historical data, making them slow to react when market dynamics change. With SAP [Integrated Business Planning \(SAP IBP\)](#)

“ **These technologies don’t replace expertise, they amplify it. Planners and operators gain real-time insights, predictive models, and simulation tools that help them anticipate change and act decisively. Instead of reacting to disruptions, businesses can sense, explain, and adapt in sync with market realities.**

and embedded AI capabilities, forecasting has become a living process between planners and intelligent systems. AI agents continuously scan external signals — weather patterns, geopolitical shifts, social sentiment, and local events — to detect anomalies early.

When the system spots something unusual, SLMs assist planners by explaining the anomaly in plain language and suggesting intelligent overrides. Instead of static spreadsheets, planners now have real-time simulation tools that test alternative sourcing





With SAP Integrated Business Planning (SAP IBP) and embedded AI capabilities, forecasting has become a living process between planners and intelligent systems.

or inventory scenarios instantly. AI copilots guide users through what-if analyses, helping them evaluate trade-offs across cost, lead time, and sustainability. This kind of system provides data-driven recommendations, but humans remain in control, applying their experience and judgment to make the final call.

A great example of this in action comes from global food producer [BRF](#), which transformed its planning process with SAP IBP and AI-driven demand sensing. By moving away from static, historical forecasts to a responsive system, BRF was able to adjust according to shifting consumer trends and external factors like weather or supply fluctuations in near real time. This helped human planners to collaborate with algorithms, resulting in faster planning cycles, improved inventory balance, and more accurate forecasts across its global network.

Connected execution

Execution is where strategy meets reality and where AI and humans collaborate most closely. With SAP S/4HANA as the digital core, AI agents gain direct access to real-time operational data, embedded processes, and transactional signals allowing them to bridge the gap between planning and physical execution far more effectively than in traditional systems.

Here's how AI and SAP S/4HANA work together across the supply chain:

- In warehouses, AI agents tap into [SAP S/4HANA's extended warehouse management \(EWM\)](#) data to optimize picking routes, slotting, and bin allocation. Workers receive dynamic, voice-assisted guidance that reflects real-time inventory positions and task priorities coming directly from S/4HANA.
- In transportation, AI analyzes freight orders, delivery schedules, and carrier data maintained in [SAP S/4HANA transportation management](#) to predict route disruptions and propose optimized alternatives. Logistics teams can then choose the most practical response with full visibility into cost, time, and compliance implications.
- In cold chain logistics, AI continuously

processes internet of things (IoT) sensor and batch-level data integrated with SAP S/4HANA's batch management and quality modules. If temperature anomalies appear, AI triggers alerts before a breach occurs, enabling proactive human action while ensuring every step remains traceable in S/4HANA.

- In compliance-heavy operations, AI enhances global visibility and traceability by automatically capturing events, generating documentation, and preparing audit-ready reports based on S/4HANA's embedded compliance frameworks, while humans validate and oversee the process.

One example of this approach in practice is [Fusion Consulting's](#) cold chain solution, built on SAP BTP. AI monitors temperature data from sensors and tags, spotting anomalies far earlier than manual checks could. When something looks off, it sends an immediate alert so the operations team can intervene, preventing product risk before it materializes. The system handles continuous monitoring while people apply judgment and expertise to nullify that.

Intelligent promise management

Customer expectations for accuracy and transparency demand precise coordination between systems and people.

“Execution is where strategy meets reality and where AI and humans collaborate most closely.”

With AI-enhanced [Available-to-Promise \(ATP\)](#), SAP S/4HANA dynamically considers production capacity, logistics constraints, batch availability, and even carbon footprint before confirming a delivery date.

If disruptions occur due to a late supplier, a weather event, or a port closure, the system doesn't just flag the issue. The system recalculates and reallocates in real time, splitting or reprioritizing orders based on service levels and customer importance, while operators take care of informing and assuring customers. This illustrates how AI handles complexity while humans provide the judgment and the empathy.

[ArcelorMittal](#), one of the world's largest steel manufacturers, modernizes its fulfillment process with ATP. When disruptions occur, planners work with AI to reassess routes and reallocate stock, ensuring realistic, transparent delivery commitments while maintaining customer trust.

How to get started with AI agents

The vision for an intelligent, adaptive supply chain is clear, however, the next question is: how do you get started?

Successfully operating AI in supply chain environments requires a thoughtful, phased approach, where people and intelligent systems work side by side.

Here's how enterprises can begin their journey with SAP S/4HANA, AI agents, and humans, all together:

- **Unify planning and execution where it matters most:**
Begin with the areas that feel the most pressure, like inventory planning, cold chain logistics, or batch traceability. These are often where data is scattered, decisions are delayed, and compliance is critical. SAP S/4HANA brings clarity by extracting actionable insights from this complexity. This ensures that systems surface the right insights, while human planners validate and act on them, turning fragmented decision-making into coordinated, real-time collaboration.
- **Deploy AI agents across execution layers:**
AI agents thrive in fast-moving environments like transportation, warehouse management, and traceability

“Customer expectations for accuracy and transparency demand precise coordination between systems and people.”

where small decisions have big impacts. These agents act as digital teammates, forecasting delivery delays, suggesting optimal routes or layouts, and identifying potential exceptions before they happen. Meanwhile, human operators remain in command, reviewing AI-driven recommendations and applying situational context that machines alone can't provide.

- **Enable cold chain resilience with sensor-AI integration:**
Temperature-sensitive logistics depend on real-time awareness and human oversight. By integrating IoT sensors with AI agents, enterprises can automate monitoring, alerts, and corrective actions when temperature thresholds are breached. If a sensor flags a risk, AI recommends solutions like rerouting shipments or adjusting storage while human supervisors confirm or fine-tune the decision.

- **Use SLMs for human-AI collaboration:**
SLMs are ideal for supporting frontline teams with contextual, domain-specific insights. They can assist humans in exception handling by interpreting anomalies, generate documentation for audits, and help identify root causes in operational issues. Because SLMs are lightweight and focused, they integrate smoothly into SAP workflows without disrupting existing processes.
- **Embed human-in-the-loop checkpoints in critical flows:**
Build workflows where AI tools provide smart recommendations, but the final call rests with experienced experts. This keeps accountability clear, builds trust, and ensures complex situations (like ATP decisions) get the nuanced thinking they deserve.
- **Leverage SAP BTP and integration suite for orchestration:**
A collaborative supply chain is supported by a strong backbone for integration and

governance. BTP and Integration Suite connect AI services, business processes, and external data sources into one cohesive ecosystem. They allow humans and intelligent systems to work effectively across SAP and non-SAP environments, ensuring interoperability, scalability, and security as AI capabilities expand.

Turn vision into value

The goal of an adaptive supply chain isn't to replace people, it's to make their work more effective. The best way forward is to start with focused steps: introduce AI where it can deliver clear value and keep humans involved in every critical decision.

With SAP S/4HANA and AI agents working alongside human expertise, organizations can move beyond the limits of traditional systems. The result is a supply chain that learns and improves continuously, where technology and people share insights, make decisions together, and respond to change as one integrated team.

Chapter 3

Unlock the power of data in SAP S/4HANA

Manish Srivastava

*Industry principal, SAP practice,
Infosys, Bridgewater*

Enterprises across industries are accelerating efforts to modernize their core systems, with SAP S/4HANA emerging as a key enabler. Organizations increasingly recognize that upgrading their core unlocks greater agility in decision-making, operational efficiency across functions, and scalability in business processes. The share of companies actively migrating to S/4HANA has surged from **21% in 2024 to 40% in 2025**, signaling that modernization is no longer optional but a strategic imperative.

At the same time, AI is reshaping how businesses work, promising smarter insights, streamlined operations, and better customer experiences. But none of this works without quality data.

Data as the backbone of transformation

As organizations migrate to SAP S/4HANA, data is the key foundation. A strong data strategy is what separates successful transformations from costly missteps. Clean, well-governed data enables AI to deliver meaningful insights, helps operations run smoothly, and ensures compliance with ever-evolving regulations. Still, many enterprises focus on systems and processes first and assume the data will sort itself out. However, it doesn't.

Poor data quality can derail even the most well-funded digital initiatives. From inaccurate reporting to failed automation, the consequences are real and expensive. According to [Gartner](#), poor data quality costs organizations an average of \$12.9 million per year.

A global manufacturing company learned this the hard way. During its transformation planning, the team found that supplier master data was scattered across systems, filled with duplicates, outdated entries, and inconsistencies. The result: procurement delays, higher compliance risks, and no clear view of supplier relationships.

The reality of enterprise data

And this example isn't unique. Across industries, data has the potential to be a strategic asset — but in reality, it's often fragmented, inconsistent, and poorly governed.

Here's what many organizations struggle with:

- **Unclear ownership:** Data responsibilities are scattered or undefined, making it hard to enforce standards or drive improvements.
- **Low-quality master data:** Core business data like customer,

“According to Gartner, poor data quality costs organizations an average of \$12.9 million per year.”

product, and supplier information is riddled with duplicates, errors, and outdated entries.

- **Complex migration paths:** Moving data from legacy systems to S/4HANA is both a technical challenge and a strategic one.
- **Compliance and security gaps:** With regulations tightening, poor data governance can lead to serious legal and financial consequences.

Without a solid data foundation, digital transformations risk falling short. AI models trained on bad data produce unreliable results. Dashboards show misleading trends. And inefficiencies creep in, undermining the goals of transformation.

Build a proactive data strategy

How should organizations prepare to implement SAP S/4HANA? The answer lies in a proactive, structured data strategy that begins well before blueprinting kicks off.

S/4HANA isn't a simple technical upgrade. It changes how data flows, how processes work, and how decisions are made. Unlike legacy enterprise resource planning (ERP) systems, it uses a simplified data model, integrated analytics, and real-time processing. That means any issues with data — be it quality, ownership, or duplication — get exposed quickly and can have a direct impact on process integrity and business performance.

Take [USG Boral](#), for example. It had multiple legacy ERPs and manual processes across regions. Infosys helped the company move to SAP S/4HANA using preconfigured tools. By aligning data early, it reduced risk and built a solid foundation for future digital growth.

Or consider [a global manufacturing company](#) that partnered with Infosys BPM. This company made data readiness a core part of its migration strategy. With early cleansing, expert involvement, and thorough testing, the company achieved zero downtime, cut disk space by 75%, and unlocked \$5.5 million in business value.

To prepare your data early in the process, an effective S/4HANA data strategy looks like:

- **Define clear data governance roles:**
S/4HANA's integrated modules, including finance, logistics, and procurement

modules, require consistent master data across functions. Assigning data owners and stewards ensures accountability across business units and IT, so key entities like customers, vendors, materials, and general ledger accounts are maintained with consistent standards.

- **Cleanse and monitor data quality:**
In S/4HANA, bad data can impact process automation, reporting, and compliance. For example, duplicate vendor records can delay payments, and incorrect material master data can disrupt supply chain execution. Early profiling and cleansing of data, before system design takes place, prevents legacy issues from being carried forward.

“S/4HANA isn't a simple technical upgrade. It changes how data flows, how processes work, and how decisions are made. Unlike legacy enterprise resource planning (ERP) systems, it uses a simplified data model, integrated analytics, and real-time processing.”

- Establish strong master data management:**

S/4HANA's efficiency depends on high-quality master data. Implementing a robust master data management framework ensures that core entities are synchronized across modules and external systems such as customer relations management, providing a single source of truth for accuracy and reporting.
- Support AI and embedded analytics readiness:**

S/4HANA comes with embedded analytics and supports AI through its tech foundation, SAP BTP. But these tools are only as good as the data the companies use. Ensuring semantic consistency, clear data models, and high-quality inputs makes it easier to leverage predictive analytics, process automation, and intelligent recommendations.
- Secure handling of data:**

S/4HANA's in-memory database allows rapid access to real-time data, but that also means managing volume becomes critical. Implementing smart data retention and archival strategies helps reduce system footprint, improve performance, and ensures that compliance with regional frameworks such as Europe's General Data Protection Regulation is met from the start.

Start early to stay ahead

One clear lesson from successful S/4HANA transformations is: start your data work early, even before blueprinting begins. S/4HANA introduces new data models, such as business partners replacing separate customer/vendor records, real-time processing, and simplified reporting. The earlier your data is aligned with these models, the smoother your project will be.

Here are four specific actions you can take early to prepare:

- Define the data organization structure and assign data stewards:**

In S/4HANA, changes in one module can ripple across others. Having designated data stewards for each domain such as customer, material, or financial data, ensures there's a go-to person responsible for maintaining consistency and coordinating across teams. Their involvement is especially crucial during design and testing phases.
- Begin data quality assessment and cleansing:**

Tools like [SAP Data Services](#) and [Information Steward](#) can help you identify common issues such as duplicate records, missing fields, outdated entries, and inconsistent naming conventions. By resolving these problems early, you

ensure smoother data migration and reduce the risk of errors in S/4HANA's high-integrity environment.

- **Create a simplification strategy for reporting and analytics:**

S/4HANA provides embedded analytics and supports intuitive SAP Fiori dashboards. However, legacy systems often contain outdated or redundant reports. To streamline your reporting landscape, review existing reports, eliminate unnecessary ones, and align your analytics strategy with S/4HANA's virtual data models (known CDS views). This not only simplifies future reporting but also lowers maintenance costs and improves decision-making efficiency.

- **Identify data for archiving:**

S/4HANA runs on an in-memory platform, storing and processing data directly in RAM instead of relying on slower disk storage. This makes everything fast, but it also means memory is a premium resource and can get expensive. To keep things running smoothly without overloading the system, it's a smart move

to archive data that's no longer actively used. Think closed transactions, old purchase orders, or inactive master data. [SAP's Information Lifecycle Management](#) tool can help you figure out what can be safely achieved while still staying compliant with data policies. This makes your system lean, fast, and ready for action.

Let data lead the way

A successful SAP S/4HANA transformation isn't just about migrating functionality — it's about rethinking the way your organization uses data. Clean, governed, and accessible data powers real-time insights, embedded analytics, and automation that S/4HANA promises.

And the key to getting there? Start early, plan strategically, and treat your data as a first-class citizen in the project.

By laying this groundwork before blueprinting, you're giving your S/4HANA implementation a head start and ensuring it delivers long-term value from day one.

Chapter 4

Build the right tech foundation for SAP S/4HANA transformation

Giri Gopaldaswamy

*Industry principal, SAP practice
Infosys, Chennai*

From digitizing data to enabling real-time, intelligent decision-making, enterprise IT is proving transformative. And now, with AI taking the center stage, organizations are rethinking how they build and manage their technology frameworks — especially when SAP S/4HANA is part of the equation.

About [32% of organizations](#) have already made the move to SAP S/4HANA, and another 27% are in the middle of implementation. That's solid momentum. But 21% are still exploring their options.

So, while the shift is clearly underway, many businesses are still figuring out the basics: What's the business case? What's the right strategy? How much will it cost? And how does it all fit into our long-term tech and AI goals?

Those businesses taking their time over implementation risk things getting messy without a clear architecture and strategy in place. This in turn can mean costly, frustrating detours.

The intention is right — modernize, automate, innovate — but without a clear technology framework and enterprise architecture, the landscape quickly becomes a patchwork of disconnected systems,

overlapping tools, and siloed data. This kind of setup makes it hard to scale AI initiatives, maintain system stability, or ensure security. Worse, it can lead to misalignment between IT and business goals.

The [Infosys AI Business Value Radar 2025](#) highlights this challenge vividly. While more than 50% of AI use cases are delivering some business impact, only 19% fully meet their objectives. Another 32% are on the cusp — which means nearly half of AI initiatives are either underperforming or failing to scale or are yet to mature. These figures would include use cases that are still in the planning, proof of concept, and pilot phases.

The report shows that success requires transforming the underlying business architecture, data infrastructure, and operating models. AI use cases that require deeper transformation — like upgrading data architecture or rethinking operating models — tend to be more successful. IT-related applications such as operations, cybersecurity, and software development are leading the way (in terms of the likelihood of success), while human-centric areas like marketing and customer service lag due to fragmented systems and lack of integration.

In short, the rush to adopt AI without doing the groundwork — aligning architecture, upgrading infrastructure, and preparing people — leads to fragile implementations

that don't scale, don't deliver value, and don't last.

How to build the right foundation

Take a [global toy manufacturer](#), for instance. Its ERP system had been around for nearly 20 years — customized, patched, and stretched to its limits. When the company decided to modernize, it didn't chase AI straight away. Instead, with Infosys as its partner, it streamlined the entire landscape through a structured migration to SAP S/4HANA. Thousands of redundant custom objects were cleaned up, data volumes dropped by more than 70%, and real-time reporting replaced slow, overnight batch runs. The result was a leaner, faster, and far more resilient core — one capable of powering future AI-driven insights and automation without breaking a sweat.

A [large US utility company](#) took a similar approach, although with different goals. Its systems were functional but rigid, making it hard to keep up with the changing demands of energy distribution, sustainability, and customer expectations. The company chose a bluefield transformation to SAP S/4HANA and BW/4HANA on AWS, keeping the useful parts of its existing system while replacing outdated elements. This sits between a greenfield approach, where you rebuild everything from scratch, and a brownfield approach, where you upgrade the current

system step by step. The company could build a digital core that unified finance, operations, and analytics. It also automated code remediation, simplified integrations, and harmonized data models, laying the groundwork for an intelligent and adaptive enterprise.

These examples show that real transformation starts beneath the surface. It's about aligning systems, data, and processes before layering AI on top. The payoff is more than cleaner code or faster reports; it's an architecture that supports intelligence across the enterprise.

A solid foundation for a smart future

So how do you make sure your transformation doesn't fall into that "tech chaos" — where systems multiply, integrations break, and AI runs on unstable foundations? The answer is the right groundwork: a clear, flexible architecture that connects technology with business strategy.

That's what [SAP's Enterprise Architecture Framework \(EAF\)](#) is built for. It acts as a blueprint for sustainable, intelligent growth and provides a clear approach to designing, implementing, and evolving your technology landscape so that it's stable today and ready for future needs.

SAP's EAF brings together structure and

“That's what SAP's Enterprise Architecture Framework is built for.”

flexibility. It combines methodologies such as TOGAF for organizing enterprise structure, BPMN for mapping and improving workflows, and UML for designing and visualizing software systems. It also supports key tools including SAP LeanIX EAM for managing architecture, Signavio Process Explorer for workflow refinement, and the SAP Business Accelerator Hub and Discovery Center for quick access to prebuilt content and innovation services. Add to that a wealth of reference architecture content, strong governance practices, and expert services to guide implementation — and you get a framework that helps you design your business and technology.

With this foundation, organizations can finally start layering in AI as a natural extension of their enterprise strategy. Because when architecture, data, and business goals move in sync, AI scales, sustains, and transforms.

How to integrate AI into SAP

With a strong enterprise architecture, the next step is to bring intelligence, with clear

purpose. AI should be built into the SAP landscape in a way that fits naturally with how the organization already works.

The Infosys AI Business Value Radar 2025 shows that successful AI adoption depends on having the right foundation in place. Organizations that take the time to rethink their operating models, upgrade their data architecture, and prepare their workforce tend to see much stronger results from their AI initiatives.

Here's a practical approach to doing it right:

- **Revisit strategic vision and governance:**

Start by aligning your AI strategy with your business objectives. This means setting clear goals, putting governance frameworks to manage data ethics and compliance, and involving stakeholders from across business, IT, and data teams to ensure alignment and ownership.

In the case of the global toy manufacturer, it focused first on cleaning up and modernizing all of its systems. This helped ensure strong alignment and accountability across teams, which in turn created a solid foundation for AI to be added later, without complications.

- **Extend your enterprise architecture:**
Your architecture needs to be ready for

“The Infosys AI Business Value Radar 2025 shows that successful AI adoption depends on having the right foundation in place.”

AI. SAP's EAF can be expanded to include AI-specific components and reference models tailored to your industry. Key upgrades to consider include:

- SAP BTP AI core and foundation: Use these tools to build, train, and deploy AI models efficiently.
- SAP BTP integration suite and APIs: Ensure seamless connectivity across systems and data sources.
- Robust data strategy: Establish strong foundations for data quality, governance, and accessibility, critical for successful AI adoption.
- **Embed intelligence into business processes:**
SAP's platform is evolving rapidly. The shift from Fiori to Joule is bringing more intuitive, conversational interfaces

powered by generative AI. This opens up opportunities to build intelligence directly into business processes and improve decision-making and efficiency.

Examples include autonomous finance operations, predictive maintenance, and AI-powered supply chain optimization — all running on the SAP S/4HANA backbone and integrated across enterprise functions.

- **Operationalize and monitor AI:** Once AI is embedded, it's important to make it a part of day-to-day operations. This involves deploying

models into production, monitoring their performance, and continuously refining them. SAP offers tools to track key performance indicators, user behavior, and system health to help keep everything on track.

Equally important is preparing your people. The Infosys report shows that organizations investing in workforce readiness through training, AI literacy, and cultural change can improve AI success rates by up to 18 percentage points. Change management acts as a strategic enabler of AI adoption, not a support task.



Conclusion

Success in AI-led transformation depends on the right foundation: a clear strategy, a scalable architecture, and a workforce prepared for change.

SAP S/4HANA, supported by SAP's EAF, offers the structure needed to bring AI meaningfully into business operations. Organizations investing in groundwork, from data readiness to change management, are seeing the most positive results.

By aligning technology with business goals and embedding intelligence thoughtfully, enterprises can move beyond experimentation and achieve lasting value from AI.



Chapter 5

The human blueprint for AI-driven SAP S/4HANA success

Dr. Murat Adivar

*Professor, Fayetteville State University,
North Carolina*

Support for SAP ECC will end in 2027, driving organizations to transition to SAP S/4HANA. While business leaders are focusing on timelines, migrations, data conversion, and business continuity, they also need to understand how this shift will change the way decisions are made.

S/4HANA embeds intelligence directly into core processes. Artificial intelligence (AI) capabilities are designed for faster decisions, fewer handoffs, and tighter coupling between insight and action. As organizations move to SAP S/4HANA, they inherit not just a new platform, but a new operating expectation in which humans and AI must work together. When that expectation is not met, the value of the transformation is diminished rather than realized.

Why technology is not the constraint

With the 2027 deadline approaching, progress is already visible. [Around a third of organizations](#) have completed their transition to S/4HANA, and many more are actively planning their move. This suggests that the technical aspects of migration are well established, even as many organizations are still working through the preparation required to execute them. While many are on track with their migration, the challenge

often begins after go-live, when AI-driven insights start appearing inside everyday workflows. Organizations should not expect immediate benefits unless they are prepared to adapt to how decisions are made.

Infosys AI Business Value Radar 2025 shows that AI success depends on organizational transformation. The report finds that companies that change their operating model, update their data architecture, and prepare their workforce deliver significantly better outcomes.

Many organizations still operate with structures shaped by slower systems and delayed reporting. When insight was harder to produce and riskier to interpret, data preparation and validation were concentrated in specialist teams, and decisions moved through multiple approval layers. Those roles were designed to manage risk and accuracy, but they now slow action in environments where insight is available in real time. S/4HANA changes this by eliminating the need to move data between operational systems and analytical environments. Because transactions and analysis work from the same live dataset, insight becomes available within everyday processes, rather than arriving later through reports or dashboards.

SAP S/4HANA can surface insight within everyday work, but it cannot resolve

“ Infosys AI Business Value Radar 2025 shows that AI success depends on organizational transformation.”

organizational ambiguity. According to research published in Harvard Business Review, most businesses struggle to realize the value of AI because people, processes, and internal incentives do not change alongside it. When existing structures are unchanged after AI implementation, people hesitate to act on system-generated insight. Faced with unclear expectations, they default to established approval paths and familiar processes. The result is that AI is technically available but rarely decisive. This is not because the insight is poor, but because the organization has not adjusted how decisions are made.

How people should engage with AI

The experience at [Team Liquid](#), a Netherlands-based esports company, offers a practical view of using the Joule copilot and Joule agents to solve organizational ambiguity and achieve better business outcomes.

“ According to research published in Harvard Business Review , most businesses struggle to realize the value of AI because people, processes, and internal incentives do not change alongside it.

The first step was integrating SAP’s Joule copilot into existing workflows used by analysts, coaches, and operations staff. This changed the nature of interaction. Analysts no longer spent most of their time assembling and cleaning data. Instead, analysts worked with AI-generated insights that were already prepared, allowing them to focus on checking results, exploring anomalies, and explaining patterns to coaches. Their role shifted from producing information to interpreting it.

Coaches, in turn, gained direct access to insights without waiting for reports or intermediaries. Coaches could explore scenarios, ask follow-up questions, and make decisions while remaining accountable for outcomes. The system did not replace judgment or automate decisions. It shortened the distance between insights and action.

Trust developed through repeated use rather than instruction. Because outputs were consistent and understandable, teams learned when the system was helpful and when context required a different call. Over time, AI became part of the decision process rather than being used occasionally or set aside.

The same pattern appeared in operational areas such as expense management, where automation reduced manual effort without removing responsibility. Staff focused on exceptions and decisions rather than routine processing. Across these use cases, people engaged with AI as a practical aid to their work, not as an abstract capability or enforced mandate.

What stands out is that these outcomes were enabled by clear expectations around roles, authority, and responsibility.

How work must change

This experience shows that overcoming organizational barriers requires a human-centered playbook that enables people and AI agents to work together effectively on the S/4HANA foundation. The components of the playbook include:

- Clarity matters first. People need to understand how AI supports their judgment, where human discretion

remains essential, and how responsibility is shared. When expectations are clear, hesitation to work on AI recommendations is reduced, and engagement increases.

- Continuous upskilling, delivered in the flow of work, builds confidence in using AI copilots for better decisions. Seeing how recommendations are formed, and understanding their limits further builds confidence to work with AI in real situations.
- Workflow redesign ensures that intelligent agents and humans operate within streamlined, modernized processes rather than legacy approval chains. Clarifying ownership and removing unnecessary handoffs allows people to act on their own while keeping accountability intact.
- Finally, accountability must be reinforced rather than softened. People need to feel able to question AI outputs, refine how they are applied, and take responsibility for decisions informed by them. Trust grows when outcomes are visible and ownership is clear.

How to implement the S/4HANA playbook

Once organizations accept that implementing AI-enabled S/4HANA requires

changes in how people engage with decisions, the role of leadership becomes more specific. The task here is to shape the conditions under which it can be used confidently and consistently. To unlock the full value of AI-enabled S/4HANA transformations, leaders should:

- Begin effective transformations by anchoring change in business outcomes rather than system features. S/4HANA places analytics and AI-driven recommendations directly into operational workflows, which changes how decisions are expected to be made. Leaders should be clear about which decisions should improve as a result, whether that means responding to issues sooner, making more consistent choices, or reducing unnecessary handoffs. When people understand why their work is changing, they are more likely to engage in new ways of working.
- Engage teams early and continue that engagement after go-live to help ensure that changes in roles and expectations are understood and addressed. Involving employees during design and preparation helps surface concerns about trust, authority, and accountability while there is still time to address them. These conversations after go-live allow organizations to adjust based on how the system is being used.



- Build learning directly into the flow of work rather than separating it from daily activities. As people use AI copilots and embedded insights in real situations, they begin to understand how recommendations are formed and how their own judgment fits alongside them. This kind of learning develops confidence gradually and is far more effective than one-time training sessions that cannot anticipate every scenario.
- Finally, track behavioral adoption over time. Usage reports and dashboards can show that capabilities are available, but system logs do not reveal whether insight is shaping decisions. The strongest signal of success is seeing teams consistently rely on embedded insights and AI recommendations as part of everyday decision-making, because sustained behavioral change is what turns new capability into lasting value.

Conclusion



The move to SAP S/4HANA is already underway for many organizations, driven by the need to modernize core systems. For meaningful results, the speed of migration needs to be complemented by how organizations redesign work to accommodate these changes.

Team Liquid's experience shows that AI delivers value when expectations are clear, roles evolve, and accountability aligns with decision-making. As insight becomes embedded in everyday workflows, organizations must be explicit about who acts on it, how judgment is applied, and where responsibility sits. Without this clarity, even capable systems fail to change behavior.

S/4HANA provides the foundation for more informed and timely decisions, but it does not determine how those decisions are taken. That responsibility sits with leaders who shape how people work with the system, learn through use, and adapt processes to match new ways of operating.

Chapter 6

Govern the AI: Responsible intelligence for SAP S/4HANA enterprises

Dr. Murat Adivar

*Professor, Fayetteville State University,
North Carolina*

Rather than replacing human decision-making, AI increasingly acts as a companion that augments human judgment. Yet trust remains a concern: **95% of businesses report** having experienced at least one AI-related incident, raising concerns about control, accountability, and risk.

These concerns become more obvious in core enterprise systems. When AI capabilities are applied within SAP S/4HANA to generate forecasts, risk scores, and planning recommendations, decisions propagate across different functions of SAP S/4HANA, such as finance, supply chain, and operations. However, without clear governance, accountability for those decisions can become ambiguous.

The challenge now is to establish effective AI governance in SAP S/4HANA to realize value while preserving trust, human judgment, and accountability.

From insight to influence

Organizations moving to SAP S/4HANA expect more than faster transactions and simplified architecture. They expect AI to improve how decisions are made, with predictive forecasts reducing uncertainty, intelligent recommendations ensuring

faster response, and embedded copilots supporting daily work across finance, procurement, supply chain, and operations.

SAP S/4HANA makes this possible by combining real-time data with analytics and ML models. As a result, the system not only reflects what has already happened; it can influence what happens next in terms of guiding cash flow decisions, shaping production plans, and affecting credit limits and supplier selection.

These capabilities deliver value only when users trust and act on the outputs. Trust depends on more than accuracy. People need to understand how recommendations are produced, when they should be followed, and when judgment is required.

But trust cannot be assumed. With so many enterprises having already experienced at least one AI-related incident, such as incorrect recommendations, biased outcomes, and failures in controls or compliance, organizations must make sure guardrails are in place to stop these incidents from happening — and those guardrails must be in place from the start.

Governance is difficult to operationalize

A majority of organizations have begun addressing AI governance, with a recent survey showing [that 77% are currently](#)

“ A majority of organizations have begun addressing AI governance, with a recent survey showing that 77% are currently working on it.

[working on it](#). However, the challenge lies in making governance work consistently in day-to-day operations.

Infosys research shows that only 2% of enterprises meet responsible AI maturity standards like [OWASP](#), or the [NIST RMF](#). This gap exists because governance is not a one-time framework or policy document. It requires consistent execution across data, models, and workflows.

Governance failures typically arise from operational weaknesses. Inconsistent or poorly maintained master data can propagate errors across AI-driven processes. Models trained on partial or biased data can reproduce those patterns in recommendations. Unclear decision ownership makes accountability difficult once AI outputs are embedded into everyday work.

In SAP S/4HANA environments, these challenges are more pronounced because AI-driven recommendations are inserted directly within core workflows through SAP's copilot, Joule. When recommendation logic, input data, or confidence indicators are not transparent to end users, trust in system outputs degrades. In response, some teams introduce manual workarounds or bypass automated recommendations to mitigate perceived risk, while others over rely on the system and treat outputs as deterministically correct. In both cases, decision quality and process efficiency suffer, and the transformation fails to realize its intended business value.

Without transparency and clear ownership, AI can create doubt in users. While the intent was to speed up decisions, these challenges slow them down. In practice, missing or weak governance indirectly becomes the reason AI introduces risks in business processes.

A continuous governance model

The solution is not to slow AI adoption and limit its use. Instead, organizations must govern AI as an ongoing practice.

SAP offers a practical example of how AI governance can be operationalized. The company has established a [formal Responsible AI framework](#) that includes defined decision rights, ethical review

processes, human accountability requirements, and continuous oversight. Every AI use case undergoes structured risk assessment and governance review before deployment and throughout its life cycle. This approach illustrates how the responsible-by-design approach builds on the cybersecurity principle of secure by design. It applied the same idea of addressing risk early and systematically rather than adding it as an extra control layer after the implementation is completed.

Below is a reference model for how enterprises can structure governance when deploying AI in SAP S/4HANA environments.

Data integrity

AI recommendations depend directly on the quality of enterprise data. Master data such as suppliers, materials, customers, pricing conditions, and financial hierarchies must be accurate, consistent, and governed. This requires clear ownership, standardized definitions, validation rules, and traceability across systems and regions.

Tools such as [SAP Master Data Governance](#) support this effort by enforcing consistency and approval controls. Data integrity ensures AI outputs are grounded in reliable information. Otherwise, it can create errors that spread quickly and become difficult to trace.

Model transparency

Users need to understand why the system recommends a particular action. Transparency requires clarity around key drivers, assumptions, and limitations. People should be able to see what data influenced a recommendation and how confident the model is in its output.

[SAP S/4HANA Embedded Analytics](#) supports transparency by linking insights back to underlying transactions and data sources. When users can interpret and question recommendations, trust develops through use rather than instruction.

Human-in-the-loop controls

AI is effective at identifying patterns and proposing options. It is not accountable for outcomes. Human-in-the-loop controls ensure that responsibility remains with people, particularly in high-impact areas such as financial approvals, compliance reporting, and supply risk decisions.

Clear thresholds, escalation paths, and override mechanisms define where automation is appropriate and where review is required. Well-designed workflows make these controls explicit and consistent. They also capture feedback when users disagree with AI outputs, creating opportunities to improve models over time.

Continuous monitoring

AI behavior changes over time, even when models are not actively modified. While data patterns shift, business conditions also evolve. These external factors introduce new risks. Continuous monitoring tracks accuracy, bias, performance, and drift so that issues are detected early.

This form of governance is not static. It adapts as conditions change and ensures that AI remains aligned with business intent rather than previous assumptions.

Embed governance into daily operations

Defining governance principles is only the first step. To make governance effective in SAP S/4HANA programs, leaders must focus on how it operates in practice.



Governance should be treated as an ongoing effort, with regular checks on AI models, risks, and policies as business needs change.



Governance should be treated as an ongoing effort, with regular checks on AI models, risks, and policies as business needs change. Tools that explain how AI works should be built into daily work, so teams can easily

understand AI results and question them as part of their routine.

Work should be designed so people make the final decisions, and AI only helps by organizing information or offering suggestions. This way, it's always clear who is responsible for the outcome.

Employees should feel safe to speak up if something looks wrong or doesn't make sense, without worrying about consequences. When people are allowed to question results and share feedback, problems are caught earlier, and the system keeps getting better.

Finally, governance should be centralized through a dedicated responsible AI office. As highlighted in [Infosys's Responsible AI Radar](#), establishing a central function with clear ownership enables consistent oversight across the AI life cycle. Supported by a unified platform for managing models, enforcing standards, monitoring performance, and tracking risk, this approach helps organizations apply governance consistently across the enterprise.

Conclusion

SAP S/4HANA provides a solid base for intelligent operations. AI builds on that base, but governance determines how safely and confidently those capabilities are used. When transparency and accountability are designed in from the outset, organizations avoid rework, manual overrides, and loss of trust that often slow AI-led transformation.

When AI influences real decisions, trust does not emerge by default. It is built deliberately, embedded in processes, and maintained carefully. This is a shared responsibility of leadership and the teams using AI.



Chapter 7

Scalable AI in the enterprise: The role of small language models

Manoj Rajan

*Industry principal, SAP practice,
Infosys, Texas*

Bringing AI into core enterprise systems is no longer a futuristic idea: companies are doing this right now to stay competitive. One of the smartest ways to do this is through small language models (SLMs) that offer a practical, efficient way to embed intelligence into everyday operations without the headache of heavy infrastructure or high resource demands.

By adopting SLMs strategically, organizations can innovate faster while keeping a firm grip on compliance, security, and operations.

Why SLMs are gaining ground

SLMs are a lighter, faster form of AI designed to do specific jobs really well. They focus on clearly defined tasks and deliver quick, accurate results with much lower computing power.

Because they are small and efficient, SLMs can run close to where business data already resides — within company systems, on-premises environments, or private clouds. This makes them easier to manage, more secure, and better suited for organizations with strict data and compliance requirements.

These strengths make SLMs a natural fit

for SAP use cases like processing invoices, sorting IT tickets, or improving master data. They respond quickly, integrate smoothly with SAP platforms, and don't slow down core business processes.

When used with SAP BTP and [SAP Joule](#), SLMs can automate routine tasks and provide helpful insights directly inside familiar SAP applications such as SAP Fiori, without changing how users work.

Another key advantage is accessibility. Since [SLMs need less infrastructure and lower investment](#), teams can adopt AI faster and at lower cost. This allows more business units—and even mid-sized organizations to use AI without long setup times or large budgets.

In 2025, [SAP researchers](#) showed that smaller AI models can sometimes outperform larger ones when used in retrieval-augmented generation (RAG) setups. RAG is an approach where an AI system first searches for relevant information from documents or databases and then passes that information to a language model to produce a more accurate response.

In this setup, smaller models such as MiniLM-v6 are used to find and rank the most relevant information, while a larger language model focuses on generating the final answer. The study found that MiniLM-v6 performed better than the much larger BGE-

Large model in these retrieval and re-ranking steps.

Because smaller models are faster, more efficient, and require less computing power, they are well suited for SAP's large-scale enterprise environments, where performance, cost, and scalability matter.

The takeaway is that small can be smart. Well-tuned, lightweight models can play a strategic role in hybrid AI architectures, handling specific tasks efficiently while supporting larger systems in complex SAP cloud landscapes.

LLMs versus SLMs in SAP: A strategic fit

The choice between LLMs and SLMs isn't about picking sides — it's about choosing what fits best for the task at hand. LLMs are great when you need to tackle complex, cross-functional challenges like summarizing multi-document reports or forecasting across business units. But they come with a heavier footprint — more infrastructure, higher latency, and potential data privacy due to external hosting.

SLMs, on the other hand, are tailor-made for real-time, domain-specific tasks that are common in SAP environments such as classifying invoices in SAP S/4HANA, triaging IT tickets in [SAP Solution Manager](#), or enriching material master data.

And because SAP systems often deal with sensitive data and require low-latency responses, SLMs are a natural fit. That's why many enterprises are embracing a hybrid AI strategy, combining LLMs for strategic insights and SLMs for operational excellence.

The promise and the pitfalls of SLMs

Despite their potential, the path to enterprise-wide SLM adoption is not without challenges. Organizations often encounter several roadblocks:

- **Limited model capacity:**
SLMs by design have fewer parameters than LLMs. While this makes them faster and more efficient, it also limits their ability to handle highly complex or ambiguous tasks without fine-tuning.
- **Integration complexity:**
Enterprises typically operate in heterogeneous environments, combining SAP and non-SAP systems. Integrating SLMs across these platforms while maintaining data consistency and process integrity can be daunting.
- **Data privacy and compliance:**
With increasing regulatory scrutiny, enterprises must ensure that AI deployments adhere to strict data governance, privacy, and compliance standards.

- **Lack of operational frameworks:**
Many organizations lack the infrastructure to fine-tune, monitor, and govern SLMs. They need [machine learning operations \(MLOps\)](#) practices to bring all technology components together, reducing operational complexity and increasing the speed at which you can take AI products into production. Without these practices, it becomes difficult to align model performance with business objectives.

A strategic, modular approach to SLMs

To overcome these implementation hurdles, tech-specific enterprises such as Microsoft are planning to adopt a modular, strategic approach to SLM deployment — an approach that balances innovation with operational rigor.

Different approaches could include:

- **SAP-native integration with BTP and Joule:**
For organizations that have already embedded SAP in the business, platforms like SAP BTP and Joule offer in-built capabilities to integrate SLMs. These platforms provide pre-built connectors, security frameworks, and orchestration tools that simplify deployment and reduce time-to-value.

- **Federated models for hybrid landscapes:**

In hybrid environments, businesses can leverage open-source SLMs and tools such as SAP Integration Suite to unify their landscape. Open-source models are well suited to this approach because they offer the flexibility to run workloads on-premises, in the cloud, or at the edge, without vendor lock-in.

They also support data security and compliance by keeping sensitive information within enterprise systems, while remaining cost-effective and compatible with existing setups.

This flexible, federated model allows workloads to run closer to where data resides, resulting in faster responses and smoother coordination across complex environments.

- **Efficient fine-tuning with LoRA and adapter modules:**

To enhance model performance without incurring high computational costs, enterprises are turning to parameter-efficient, fine-tuning techniques such as low-rank adaptation (LoRA) and adapter modules.

This allows enterprises to rapidly customize intelligent agents for specific workflows without retraining entire models, making AI

deployment faster and more cost-effective.

- **Robust MLOps for governance and compliance:**

To keep SLMs running reliably in SAP environments, a solid MLOps setup is key. This includes tracking model versions, monitoring performance in real time, and enforcing compliance with data privacy rules. With these practices in place, enterprises can ensure their SLMs stay aligned with business goals and regulatory standards.

What organizations should focus on?

If you are thinking about bringing SLMs into your enterprise architecture, the key is to start smart and stay strategic. Here are some areas to focus on:

- **Start with the right use cases:**

Begin with tasks that are high-impact but low-risk. Think invoice classification, HR query resolution, or IT ticket triage — areas where SLMs can deliver quick wins and prove their value early on. This helps build confidence in what AI can achieve.

- **Check your infrastructure readiness:**

Take stock of your current systems. SLMs work especially well in setups with siloed, high-value data and limited compute resources, as their smaller size and less GPU power and energy requirement

make them equipped to fit into all kinds of environments. Whether you're deploying on-premises, in a virtual private cloud, or across hybrid environments, make sure your infrastructure can support secure and scalable AI.

- **Measure ROI the smart way:**

Before you deploy, set baseline metrics like error rates, cycle times, and how often humans have to step in. After deployment, track how those numbers improve. Return on investment (ROI) isn't just about cost savings; it's also about boosting productivity, reducing risk, and sparking innovation. Tools like MLOps dashboards can help you monitor performance and validate impact over time.

- **Fine-tune efficiently:**

Use techniques like LoRA, prefix tuning, or adapters to customize your models without burning through compute. These methods let you adapt SLMs to specific tasks while keeping things lightweight and manageable.

- **Keep security and compliance front and center:**

If you are in a regulated industry, it is important to build in security features like role-based access, encryption, and audit logging right from the start. These help control who can see what,

keep data safe, and track any changes. Running SLMs on your own servers also helps meet strict data rules like keeping sensitive information within your country or organization, and keeps your data protected and compliant.

- **Think about life cycle management:**

SLMs aren't something you just set up once and forget about. You will need a plan to manage things like different versions, keeping the model size efficient, and making sure it works well for your specific use case. Tools that support MLOps can make this easier by helping you track performance, retrain the model when needed, and keep everything running smoothly over time.

- **Plan for scale and flexibility:**

Your SLM strategy should be able to grow and adapt as your business evolves. A [modular setup](#) lets you build your system from scratch, so you can start with what you need now and add more as things change. With this kind of design, it's easier to connect with both SAP and non-SAP systems, support new business needs, and stay flexible.

- You don't have to go all-in from day one — you can begin with a monolithic architecture and gradually move to a more modular, distributed system that is easier to manage and more resilient.

Conclusion



SLMs represent a powerful, scalable pathway to enterprise AI adoption. Their agility, efficiency, and adaptability make them particularly well suited for complex, data-rich environments. By taking a strategic, modular approach grounded in robust integration, fine-tuning, and governance, enterprises can unlock the full potential of SLMs to drive innovation, efficiency, and competitive advantage.

As AI continues to evolve, the organizations that succeed will be those that not only adopt the right technologies but also embed them thoughtfully into their operational fabric. With the right strategy, SLMs can become a cornerstone of intelligent enterprise transformation — especially when deployed with SAP's ecosystem.

About Fayetteville State University



Fayetteville State University (FSU), founded in 1867 and a constituent institution of the University of North Carolina System, has a more than 155-year legacy of advancing education, workforce readiness, and regional economic development. Located in Fayetteville, North Carolina, home to Fort Bragg, the largest US Army installation in the world, FSU serves a broad population of learners, including traditional students, adult learners, and military-affiliated professionals transitioning to civilian careers.

The Broadwell College of Business and Economics at FSU is accredited by AACSB International, the premier global accrediting body for business education. The university's MBA program has been recognized by US News & World Report, ranking among the top programs in North Carolina, including a recent ranking of number 5 in the state, reflecting its strong focus on applied learning and career outcomes.

FSU has established itself as a national leader in enterprise systems education through its SAP-enabled curriculum and industry-focused programs. The university hosts one of the 14 SAP Next-Gen Labs in North America and integrates SAP S/4HANA across undergraduate and graduate business curricula. A cornerstone of this ecosystem is the Center for ERP and Advanced Analytics (CERPAA), which serves as a hub for experiential learning, SAP certification, industry partnerships, and workforce reskilling initiatives aligned with employer needs.

Through CERPAA, FSU delivers hands-on training in enterprise systems, business process integration, analytics, and digital transformation technologies. The university is also the only higher education institution in North Carolina offering the SAP IEE2E Business Process Integration certification pathway and has developed innovative workforce initiatives, including a US Army Career Skills Program (CSP) approved training pipeline that prepares transitioning service members for ERP consulting and technology careers. These programs reflect FSU's commitment to preparing skilled professionals who contribute to organizational performance and economic growth.

FSU and Infosys

The relationship between Fayetteville State University and Infosys reflects a strong alignment in preparing industry-ready talent and supporting workforce development in enterprise technology and consulting. Infosys has been a key employer partner for FSU graduates, with multiple students and alumni joining the company and building successful careers in consulting and digital transformation roles.

Infosys representatives have engaged with the university community through campus visits, recruitment activities, and participation in academic-industry events, strengthening connections between students and industry opportunities. This engagement has helped create clear career pathways for students pursuing enterprise systems and technology careers through FSU's programs and CERPAA initiatives.

The collaboration also includes thought leadership engagement. Dr. Murat Adivar, director of CERPAA, has been invited to speak at Infosys headquarters, contributing academic perspectives on workforce development, enterprise education, and university-industry collaboration. These interactions reflect a growing strategic relationship centered on talent development, innovation, and industry-academia alignment.

As FSU continues to expand its SAP-enabled education ecosystem through CERPAA, professional certification programs, and workforce initiatives, the relationship with Infosys represents an important connection between education and industry practice. Both organizations share a commitment to preparing future professionals with the business and technology capabilities required in a rapidly evolving global economy.

Authors



Dr. Murat Adivar is a Professor at Fayetteville State University and the founding director of the Center for ERP and Advanced Analytics, where he leads ERP, SAP S/4HANA, and analytics-driven education and applied research. He specializes in building SAP-enabled curricula, workforce pipelines, and university–industry partnerships. He is chair of the SAP North America Academic Community (SAPNAAC) Partnership Committee and is an ASUG champion, working closely with SAP, consulting partners, and enterprises to advance SAP education. Email: madivar@uncfsu.edu



Radhakrishnan Jagadeesan is a senior principal and industry principal at Infosys Consulting with extensive experience leading enterprise-wide digital transformation initiatives. Based in Dallas, Texas, he has held senior leadership roles across business consulting, digital transformation, and supply chain execution. Radhakrishnan’s global experience includes leadership roles in North America and Asia, including several years in Singapore, giving him a strong international perspective on transformation programs. Email: Radhakrishnan_J@infosysconsulting.com



Manish Srivastava is a practice director at Infosys, responsible for leading client services for SAP data, analytics, and AI solutions on the SAP platform across North America. He brings deep expertise in helping organizations modernize in alignment with enterprise transformation goals. Manish works with senior client leaders to define value-driven strategies and roadmap. His experience spans international consulting assignments across the US, Europe, and India, and delivery models that combine onshore and offshore teams across multiple industries. Email: Manish_Srivastava01@infosys.com



Giri Gopaldaswamy is a seasoned SAP S/4HANA architect and industry principal with more than 20 years of deep finance domain experience supporting global enterprises. He has led the implementation, optimization, and maintenance of finance systems across multiple ERP platforms in complex, multi-country organizations. Giri has helped enterprises modernize financial operations while balancing compliance, scalability, and business continuity. His work spans global delivery leadership, customer engagement, and execution. Email: Giri_Gopaldaswamy@infosys.com



Manoj Rajan is an industry principal and principal technology Architect with more than 25 years of experience delivering SAP-led transformation and ERP modernization initiatives. He specializes in the oil and gas value chain, leading complex programs. Manoj is a thought leader in SAP-centric AI, advocating the practical use of SLMs to enable enterprise automation integrated with SAP BTP and Joule-aligned patterns. His work emphasizes moving to scalable, measurable outcomes while maintaining strong governance and architectural discipline. Email: Manoj_Rajan@infosys.com

Editorial and production

Kate Bevan | Infosys Knowledge Institute, London

Pramath Kant | Infosys Knowledge Institute, Bengaluru

Sandeep | Infosys Knowledge Institute, Bengaluru

Pragya Rai | Infosys Knowledge Institute, Bengaluru

Taniya Chandra | Infosys Knowledge Institute, Bengaluru



For more information, contact askus@infosys.com



© 2026 Infosys Limited, Bengaluru, India. All Rights Reserved. Infosys believes the information in this document is accurate as of its publication date; such information is subject to change without notice. Infosys acknowledges the proprietary rights of other companies to the trademarks, product names and such other intellectual property rights mentioned in this document. Except as expressly permitted, neither this documentation nor any part of it may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, printing, photocopying, recording or otherwise, without the prior permission of Infosys Limited and/ or any named intellectual property rights holders under this document.