OSF DIGITAL

AGENTIC ENTERPRISE



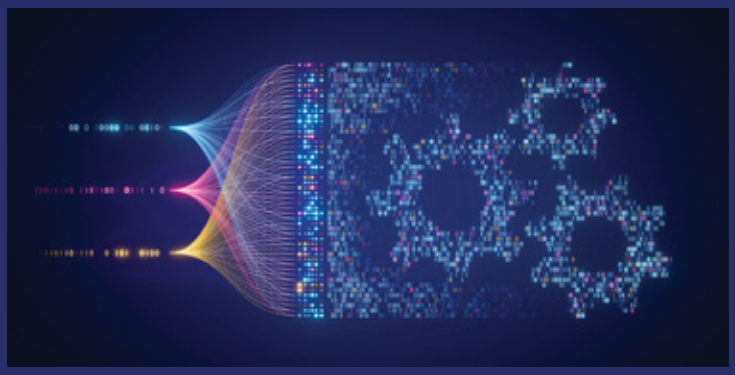
THE END OF THE WAIT

From months of delivery to workflows built at the speed of conversation

THE EXECUTION GAP

Why AI pilots succeed –
but enterprises fail to scale them

CONTENTS



3

The Agentic Enterprise

Building intelligent applications, agent swarms, and autonomous business at execution speed.



6

The AI Execution Gap

Why most AI investments stall — and the operating model fix that unlocks real ROI.



10

From Org Chaos to Operating Rhythm

How cross functional governance turns AI and data into predictable business outcomes.



15

The Framework for the Agentic Enterprise

A blueprint for moving from manual processes to deterministic AI workflows that enterprises can trust at scale.



19

Avoiding the Cost of Getting Transformation Wrong

Delivery experience is the foundational architecture of the agentic enterprise.



The Agentic Enterprise: Speed, Swarms, and the End of the Wait

The "Wait" is Finally Over

We all know the drill. You have a brilliant idea to improve a business process—maybe it's a better way to handle returns or a smarter way to qualify leads. You map it out, you get excited, and then you hit the wall: IT.

You're told it'll take six months to scope, another six to build, and by the time it launches, the market has moved on. This is the **"Execution Gap"**. It's the black hole where strategy goes to die because we've been trapped in the "classical order of software," where 80% of our technical resources are stuck just keeping the lights on, leaving almost nothing for innovation.

But what if I told you that the wait is over? We are shifting to a new paradigm called the **Agentic Enterprise**. In this world, we don't build software in months; we generate workflows in **seconds**. We are moving from a world where you manage software to a world where you mobilize an "unlimited army" of intelligent agents to do the work for you.

This isn't just an upgrade. It's a gamechanger. Here is how it works.



1 The Shift: From "Clicking" to "Conversing"

For the last 30 years, enterprise software has been passive. It sits there and waits for you to click a button. If you stop working, the software stops working.
The Agentic Enterprise flips this on its head. We are moving to **Intelligent Applications**. These aren't just tools; they are teammates. They have **agency**, meaning they can act on your behalf without you holding their hand every step of the way.

Imagine telling your system, ***"Help me fix the supply chain delay in the southeast region."***

- **The Old Way:** You'd open ten tabs, check email, download a spreadsheet, and manually update the ERP.

- **The Agentic Way:** The application understands your **intent**. It uses embedded intelligence to look at the data, figure out what "southeast delay" means, and proposes a fix. It's an **Adaptive Experience** that molds itself to you, not the other way around.

2 New Concepts You Need to Know

To lead in this new era, you need to add three new concepts to your vocabulary: **Swarms, Intelligent Applications, and Autonomous Business.**

The Swarm (The "Army of Agents")

One AI agent is helpful, like a smart personal assistant. But the real power comes when they work together. We call this a Swarm (or in some circles, a Multi-Agent System).
Think of it like a digital office. You have a "Sales Agent," a "Legal Agent," and a "Finance Agent." When a complex problem hits—like a massive custom order—one agent doesn't try to do it all. It decomposes the task and delegates it.

- The Sales Agent qualifies the deal.
- The Legal Agent checks the contract against the **Semantic Data Model** to ensure compliance.
- The Finance Agent approves the credit terms.

Ils se transmettent le travail, collaborant pour résoudre le problème instantanément. C'est ainsi que nous mobilisons cette « armée d'agents intelligents » pour augmenter la capacité à l'infini lorsque la demande l'exige.

Intelligent Applications

These are the platforms where the agents live. Unlike the static apps of the past, Intelligent Applications are defined by **Autonomous Orchestration**. They don't just store data; they go out and do things. They connect to APIs, retrieve records, and update systems across the enterprise without you needing to write a single line of integration code.

Autonomous Business

This is the destination. An **Autonomous Business** isn't about removing humans; it's about shifting humans from "operators" to "orchestrators." In this model, the business processes—ordering, shipping, invoicing—run themselves. Your job shifts to designing the rules and guardrails, while the agents handle the execution speed and scale that humans simply can't match.

3 Speed: Building in Seconds, Not Days

This is the metric that matters most. In the Agentic Enterprise, the "time to build" collapses. Because these agents understand human language, you don't need to write a requirements document for a developer. You just describe what you want: ***"Build a workflow that alerts me when inventory drops below 10% and automatically emails the vendor."***

Boom. The agent understands your intent, looks at the available API blocks (thanks to **Composable Architecture**), and assembles the workflow in **seconds**.
This means you can:

- Launch a new customer onboarding flow in a morning.
- Respond to a complex RFP in seconds.
- Pivot your entire operation to react to a competitor in real-time.

4 Why Wasn't this Possible Before?

You might be asking, ***"We've had automation for years. Why is this different?"***

It comes down to three things we didn't have until now:

1. **The Brain (LLMs):** Before Generative AI, software couldn't "reason." If a process changed slightly, the bot broke. Now, LLMs allow agents to handle ambiguity and make judgment calls.
2. **The Context (Semantic Data Models):** Old software was dumb about data. It saw rows and columns but didn't understand that "Client A" and "Customer A" were the same person. The Semantic Data Model gives agents the context they need to understand the meaning of business data, turning messy spreadsheets and emails into "queryable intelligence."
3. **The Hands (APIs):** We finally have systems that are API-First. Agents can't click mice; they need code-based handles to grab onto. Modern composable architecture gives them those handles.

The Gamechanger

The modern AI-powered enterprise platform is a gamechanger because it finally closes the Execution Gap. It democratizes development, allowing business experts—not just coders—to build powerful applications at the speed of conversation.

We are moving from a world of scarcity (limited IT hours) to a world of abundance (unlimited agentic labor). The companies that win the next decade won't be the ones with the best five-year plan. They will be the ones who can build, launch, and iterate their ideas in seconds.

Welcome to the Agentic Enterprise.



The AI Execution Gap: Why Enterprises Are Investing in AI but Failing to Scale — And the Operating Model Fix

Corporate investment in artificial intelligence is reaching record highs, yet most of these projects remain stalled in the pilot phase. For enterprise leaders, the challenge has shifted from acquiring the technology to integrating it into the business's operational fabric—a disconnect known as the AI Execution Gap.

While approximately 78% of organizations report implementing AI in some capacity, only a small fraction—estimated at roughly 5%—are realizing significant financial or operational value at scale.

While companies often encounter technical hurdles, the failure to scale is primarily driven by the organizational operating model. To bridge this gap, leadership must transition from treating AI as an experiment to integrating it into the core business strategy and prioritizing AI adoption across the organization.

Defining the AI Execution Gap

The AI Execution Gap describes the space between a successful "proof of concept" (PoC) and a fully integrated, value-generating production environment. In a controlled pilot, AI often performs well because the data is clean, the scope is narrow, and the stakeholders are highly motivated. However, scaling often requires AI to be integrated with legacy systems, disparate data sources, and workflows involving manual tasks.

Recent industry analysis suggests that 60% of companies see minimal returns on their AI investments. This stagnation is typically driven by three operational failures:

1. **Dead End Pilots:** Many projects are designed to prove technical feasibility rather than business scalability. Without a strategy for integrating AI into business processes, these projects may stagnate and fail to reach production.
2. **The Precision-Usability Gap:** While model accuracy is essential for establishing user trust, technical precision alone does not guarantee business impact. Accurate models will not deliver value until teams learn how to properly leverage the new technology to improve business processes.
3. **Organizational Silos:** When data scientists, IT departments, and business units work in isolation, AI tools often fall short of meeting business stakeholder expectations. When stakeholders are not aligned, technical solutions often fail to address the specific business challenges that drive enterprise ROI.

The Solution: Establishing an AI Center of Excellence (CoE)

To close the execution gap, enterprises require a centralized structure to manage the complexity of AI deployment. The **AI Center of Excellence (CoE)** is the primary mechanism for this "Operating Model Fix."

High-functioning AI CoEs are typically cross-functional teams composed of technical leaders, business strategists, and legal/compliance stakeholders. The CoE mandate is to bridge the gap between technical potential and **business results** by providing a unified framework for the entire organization.

Strategic Functions of the CoE

- **Standardization:** The CoE defines the enterprise "AI stack," ensuring that different departments aren't using incompatible tools or redundant data sources.
- **Governance and Risk Management:** Centralized oversight allows the company to manage data privacy, security, guardrails, and bias at scale, rather than leaving these critical risks to individual project teams.
- **ROI Prioritization:** The CoE acts as a filter, directing resources toward use cases with the highest potential for impact, such as high-volume business



Scaling ROI Through Intelligent Process Automation

While AI has broad applications, Business Process Automation (BPA) remains one of the strongest use cases for driving immediate and measurable ROI. Traditional automation was limited to rigid, rule-based tasks. AI evolves this into "Cognitive Automation," where systems can interpret unstructured data—such as emails, legal documents, and images—to handle complex enterprise workflows.

A mature operating model utilizes AI-driven BPA not just for efficiency, but as a catalyst for end-to-end process transformation. By integrating AI with core systems and data sources, organizations position themselves to enable high-value, autonomous decision-making. To achieve this level of execution, AI tools leverage not only structured enterprise data but also unstructured data that is siloed in spreadsheets, documents, and even emails.



Operating Model Dependencies

In addition to the creation of a CoE, leaders must establish core principles to ensure AI can scale.

1. Unified Data Definitions

AI performance is dependent on the **quality and consistency of data**. Scaling becomes complex when different departments use conflicting definitions for the same metrics. The CoE should shepherd the creation of a "common data language"—a unified set of definitions that ensures the AI interprets information consistently across the organization.

2. Active Governance

Scaling AI introduces certain regulatory and reputational risks. Traditional governance models, which often review projects after development is complete, are too slow for the AI lifecycle. Instead, governance must be "baked in" to the development process. The CoE ensures that AI solutions include checks for bias, data lineage, guardrails, and compliance from the start.

3. End-to-End Workflow Redesign

Enterprises often make the mistake of "paving the cow path"—using AI to automate a process that is already broken or inefficient. To scale successfully, business leaders must be willing to reimagine workflows from the ground up. This often means transforming a process with many manual steps to one with automated tasks, utilizing a human-in-the-loop approach for orchestration and decision-making.

Case Studies: Success vs. Operational Failure

Success: Air India's Integration Strategy

Air India transitioned from simple chatbots to a generative AI agent integrated into core operational systems. By centralizing data access and securing an executive mandate for the agent to execute transactional tasks—such as flight bookings and seat reassignments—the airline automated 97% of nearly 4 million interactions. This deep integration allowed the carrier to maintain service levels even as passenger traffic doubled.

Failure: Zillow's Valuation Model

Zillow's "Offers" project failed because it scaled advanced pricing models without a robust operational governance framework. While the algorithms were technically sophisticated, they lacked the human-in-the-loop overrides necessary to adjust for sudden market volatility. The lack of centralized oversight led to systematic overpayment for properties, resulting in losses exceeding $880 million before the program was terminated.

5 Actionable Tips for Business Leaders

1. **Charter a Balanced CoE:** Your Center of Excellence must have equal representation from the business and technical sides. If it is purely an IT function, it may fail to deliver significant business value.
2. **Prioritize High-Impact, Low-Complexity Quick Wins:** Instead of solving the most difficult problems immediately, start with projects that offer measurable impact with less effort and shorter timelines. These successful pilots build organizational confidence and provide the internal buy-in necessary to fund and tackle larger initiatives.
3. **Establish Unified Data Standards:** The CoE should be a catalyst for creating a common data language across all departments. Standardization is important for ensuring that AI-driven results are both trustworthy and operationally sound.
4. **Automation and Transformation:** View AI projects as opportunities to move beyond the automation of existing tasks. Use these initiatives to redesign business processes—improving efficiency, quality, and the measurability of outcomes.
5. **Cultivate AI Fluency in Management:** Scaling AI requires leaders who understand the value of leveraging AI agents across the organization. Invest in training for middle management to ensure they can effectively oversee AI-augmented teams.

Adapting the Organization to the Technology

The AI Execution Gap is not a technical problem; it is a management challenge. Enterprises that successfully scale AI do so when they stop viewing it as a standalone tool and start treating it as a core operational discipline. By establishing a Center of Excellence and improving the operating model, leaders can finally turn AI's potential into a scalable, high-ROI reality.





From Org Chaos to Operating Rhythm: Building the Cross-Functional Governance Needed for AI, Data Cloud, and Multi-Cloud Success

Organizational chaos is not a spectator sport. Think of it as a participatory event, the corporate equivalent of a brawl breaking out after a rough tackle on the soccer field. The burst of opposing opinions and attitudes often leads to a messy pile of players with their heads mashed into the turf.

Replace the soccer match with broad-based initiatives to deploy AI, data platforms and multi-cloud architectures, and you have the framework for organizational chaos. Because there is rarely a single or a “right” way to approach multi-cloud success, companies often fall into dysfunctional traps that prevent AI and data projects from kicking off with enough support and ensuring the best outcomes after launch.

The traps and pitfalls that organizations face include disconnected initiatives, inconsistent data definitions, unclear accountability, stalled innovation, and rising technical debt. This white paper presents proven recommendations to help you create and scale:

- Governance structures that are key to decision making
- Cross-functional alignment programs
- Data hygiene
- KPIs and optimization
- Change management and adoption acceleration

The Organizational Patterns That Undermine AI Investments

Does your company suffer from Organizational Chaos when it comes to AI and data projects? Here's a self-diagnosis.

When governance is weak or absent in tech projects, specific patterns often appear. Here are some of the most common symptoms seen in organizations struggling with AI and cloud adoption. Which of these issues is bogging down your company?

1. Conflicting priorities

Sales wants personalization. Marketing wants new segments. Service wants automation. IT wants stability. Without governance, teams compete rather than collaborate, slowing delivery and creating frustration.

2. Fragmented data ownership

Data definitions, permissions, pipelines, and quality expectations vary across teams. Data projects stall when no one agrees on what a "customer" or "account" really means.

3. Shadow IT and uncontrolled experimentation

Teams adopt tools independently, create duplicate solutions, or deploy AI models without proper validation. Risk grows while consistency declines.

4. Inefficient allocation of resources, also known as the "Pilot Trap"

Without unified prioritization, teams pursue low-value pilot projects while strategic initiatives pile up. Engineering teams become overloaded, and business leaders lose trust in delivery timelines.

5. Platform under-utilization

Companies invest in robust technology, but use only a fraction of the capability because governance around adoption, enablement, and scaling is unclear. Too often, there is no single governance body working across stakeholder groups to help optimize the use of AI and data.

6. Slow decision-making

Work requires multiple layers of approvals without clarity of who has authority. Opportunities pass by while decisions sit idle.



The Way to Cure Organizational Chaos

These symptoms of org chaos are not technical problems—they are organizational issues. Solving them requires cross-functional governance that creates alignment, transparency, predictable operating rhythms and, especially, optimism across stakeholder groups.

Cross-functional governance should focus on three primary areas:

- **Data Quality & Silos:** Data is often scattered across departments and legacy systems, making consolidation difficult and leading to poor data quality that undermines AI model accuracy.
- **Cultural Resistance & Skill Gaps:** Employees may resist new AI tools, and organizations often lack the in-house expertise to manage complex AI systems and their associated risks.
- **Governance Gaps & Compliance Risks:** Without clear policies, AI systems can present risks such as bias amplification, privacy breaches, and regulatory non-compliance.



The Details of Cross-Functional Governance

Cross-functional governance is a structured, repeatable model for how strategy, data, platforms, and workloads are managed across the company. It ensures the right people make the right decisions at the right time. Governance should not feel like a bureaucracy but instead like a partner driving innovation and efficiency.

A strong governance model connects four layers of decision-making and execution:

1. Strategic Governance

Aligns AI and cloud investments to business goals, value realization, and enterprise roadmaps.

2. Platform Governance

Oversees data and AI platforms. Ensures interoperability, security, and sustainability.

3. Data and Model Governance

Establishes ownership, quality standards, lifecycle management, and monitoring for data and AI models.

4. Delivery Governance

Creates consistent intake, prioritization, sprint cadences, and success measurement across teams.

When these layers operate together, organizations gain a recognizable operating rhythm—one where teams determine priorities together, know how to get work approved, have visibility into dependencies, and track KPIs to prove that decisions align with business needs.

Designing the Right Operating Rhythm for AI and Data Projects

An operating rhythm is a structured set of processes that keep teams aligned. Although the details will differ by organization, the following components are required:

1. Executive Steering Committee

Purpose: Set enterprise objectives for AI, cloud, and data

Cadence: Monthly or quarterly

Participants: Executive sponsors and leaders of other key functions

Focus/Outcomes: Funding decisions, KPIs, risk posture, and strategic direction

2. Cross-Functional Governance Council

Purpose: Define policies, data standards, platform usage, and prioritization

Cadence: Bi-weekly

Participants: IT, Data, AI, Security, Architecture, Marketing, Sales, Service, Product, Finance

Focus/Outcomes: One source of truth for strategy, standards, and prioritization. The members of the Governance Council will also act as evangelists to sell in change to the rest of the organization.

3. Platform Working Groups

Purpose: Address the day-to-day operational needs of tech platforms and business operations, and ensure consistency of integrations

Cadence: Weekly

Focus/Outcomes: Clear change control, shared best practices, and reusable patterns

4. Delivery Rhythm

Purpose: Drive consistent planning, execution, and demonstration

Cadence: Bi-weekly sprints, monthly reviews, quarterly planning

Focus/Outcomes: Predictable throughput and aligned dependencies

With these components in place, technology investments stop being siloed and start becoming coordinated, measured, and scalable.

Governance Specifics by Tech Segment: AI and Data Cloud, Multi-cloud Architecture

AI and Data Cloud

AI and Data Cloud introduce unique governance needs that extend beyond traditional cloud management.

Key focus areas include:

- **Data Ownership** – Every critical domain (customer, product, order, service, etc.) requires clear business owners and stewards
- **Model Lifecycle Management** – Versioning, testing, monitoring, bias checks, and retirement policies
- **AI Safety and Guardrails** – Prompt governance, role-based access, model outputs, and human-in-the-loop where needed
- **Data Quality and Lineage** – Continuous monitoring ensures reliable AI outcomes and trustworthy dashboards
- **Reusable Feature Stores** – Prevent duplicate feature engineering and ensure consistent model inputs
- **Governed Data Products** – Enterprise data that is packaged, documented, and accessible with defined SLAs

Without these practices, AI outputs become unpredictable and Data Cloud becomes difficult to scale.

Multi-Cloud Architecture

As organizations adopt multi-cloud strategies, governance challenges multiply.

Critical considerations:

- Consistent security controls across clouds
- Unified cost management and optimization
- Cross-cloud integration standards and patterns
- Disaster recovery and business continuity requirements
- Centralized observability and performance monitoring
- Application lifecycle management across multiple infrastructures

Multi-cloud succeeds only when delivery teams follow shared standards and architectures. Governance ensures that each cloud is used for the right purpose and managed predictably.

How Do You Know When You Are Closer to Operating Rhythm Than Chaos?

Most companies progress through four maturity stages:

1. Ad-Hoc
Teams work independently, priorities conflict, and platforms are inconsistent.

2. Emergent Governance
Some standards exist, but adoption is uneven and not universally enforced.

3. Coordinated Governance
Cross-functional bodies, unified intake, and shared data definitions start creating alignment.

4. Operationalized Rhythm
AI, data, and cloud capabilities are embedded in business operations. Value delivery becomes predictable and scalable.

Reaching the final stage requires executive sponsorship, cultural change, and continuous iteration, but the payoff is transformative.

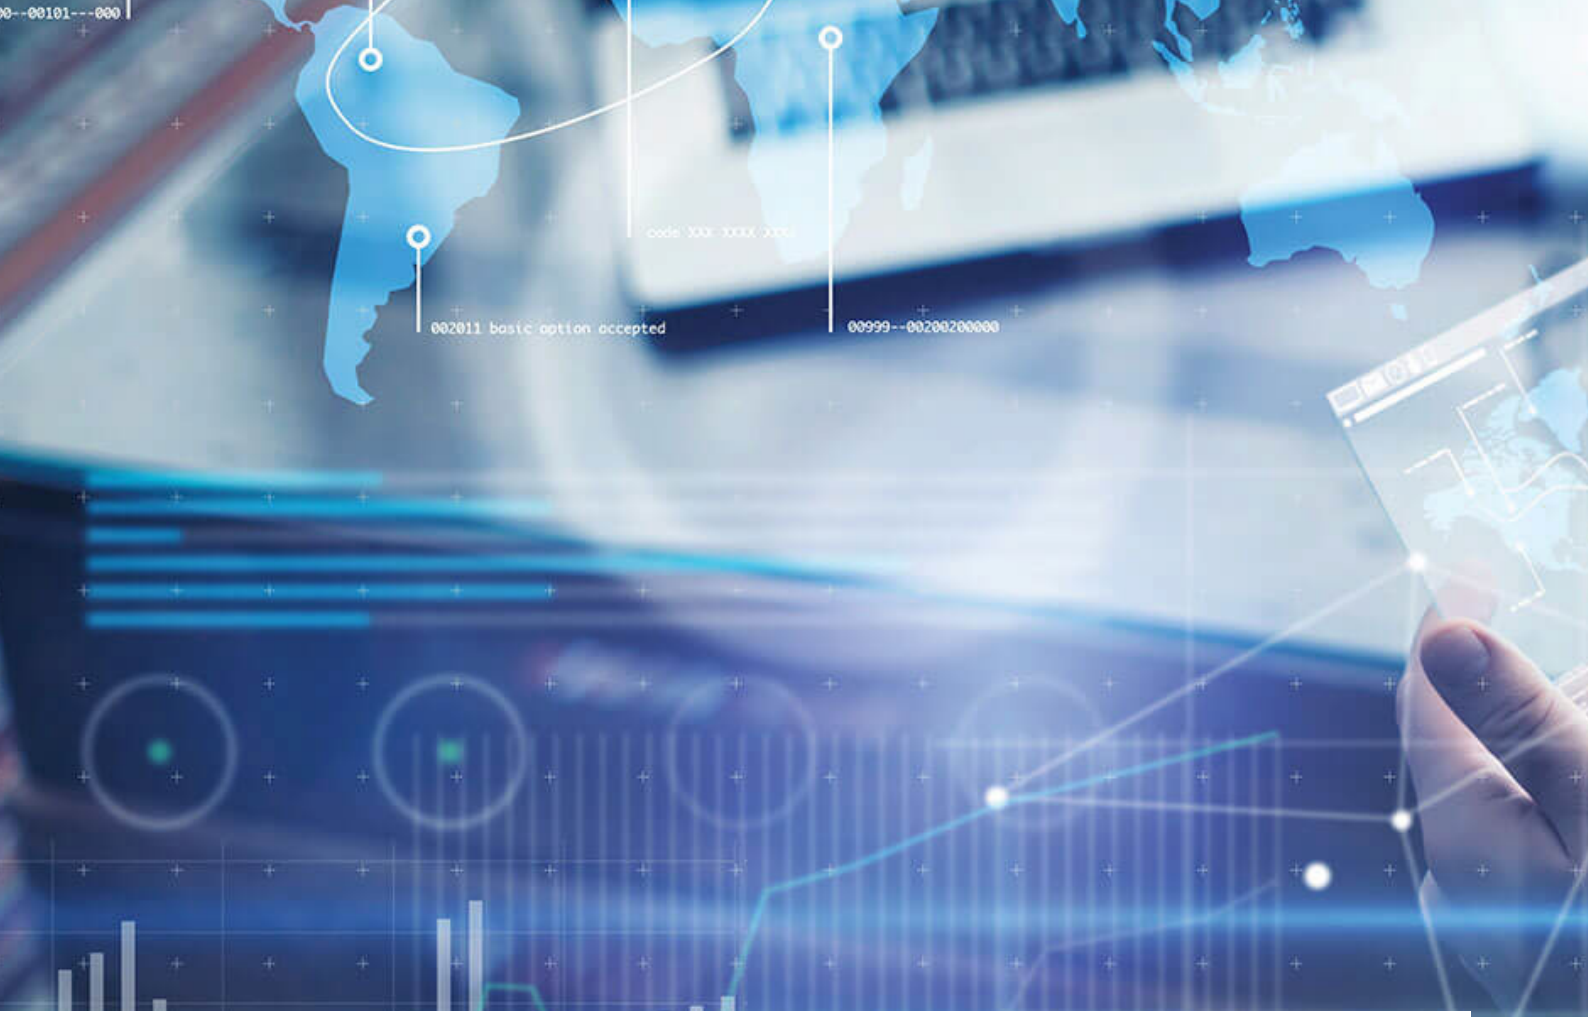
The Business Impact of Strong Governance

OSF has found that organizations that establish a unified operating rhythm experience measurable benefits:

- • 20–40% faster delivery cycles
- Improved trust in data and AI results
- Reduced duplicated effort across teams
- Lower operational and cloud costs
- Higher platform adoption and ROI
- Clear accountability and reduced risk
- Better alignment between business and technology strategies

Data Cloud, AI, and multi-cloud architectures offer enormous potential, but they cannot be fully realized through technology alone. What separates successful organizations from struggling ones is not the size of their tech stack; it is their commitment to developing and nurturing an operating rhythm. Cross-functional governance provides the structure, alignment, and decision-making clarity needed to deliver value repeatedly and predictably.

By moving from organizational chaos to an intentional operating rhythm, companies unlock the true power of their platforms and create a foundation for continuous innovation. The organizations that thrive in the AI era will be those that treat governance not as a constraint, but as a strategic enabler of speed, safety, and scale.



The Framework for the Agentic Enterprise: Moving from Manual Tasks to Deterministic Workflows

Most organizations have sophisticated back-end systems yet still depend on spreadsheets and email chains to bridge the gaps between them. These shadow processes are the most significant drag on operational efficiency and the most valuable target for agentic AI.

The path forward is not more automation. It is deterministic AI execution, workflows that produce predictable, auditable outcomes without requiring humans to verify every step. Achieving this requires three foundational disciplines: grounding agents in live systems of record, breaking complex processes into tightly scoped tasks, and building in continuous validation from the start.

The window for deliberate, strategic action is narrowing. Organizations that spend the next twelve months running more pilots will find themselves competing against peers who used that time to build operational infrastructure for AI execution.



The Gap Between Systems and Execution

Many organizations are finding that the initial excitement surrounding AI experimentation has hit a plateau. While the technology is impressive, the primary challenge for the modern enterprise is moving beyond simple chat interfaces and into the realm of reliable, autonomous execution.

Most business leaders face a significant disconnect. IT departments have successfully deployed sophisticated ERP systems and robust databases. Yet, the day-to-day operations of the global supply chain are still largely powered by "shadow processes"—the many spreadsheets and email chains used to bridge functional gaps. To build a true **agentic enterprise**, leaders must shift their focus from basic business process automation (BPA) toward **AI execution**. This requires building workflows that are not just automated, but deterministic.

Moving from Assistance to Agency

For years, automation was defined by rigid "if-this-then-that" logic. This worked for simple data entry but failed whenever it encountered the nuance of human communication or unstructured data. When generative AI arrived, it was initially treated as a creative partner—useful for drafting emails or marketing content, but often too unpredictable to be used for operational tasks where accuracy is essential.

The Agentic Enterprise occupies the space between these two extremes. By deploying AI agents, organizations can finally address tasks that require context without sacrificing accuracy. The objective is a **deterministic outcome**: a result that is as predictable and auditable as a traditional software program, but with the reasoning capabilities of a human expert.

The Requirement for Determinism

In the world of supply chain and procurement, accuracy is the only metric that matters. If an agent is responsible for verifying line items on a commercial invoice or checking a supplier's insurance compliance, the process must be foolproof. This is the primary hurdle in **Scaling AI**. If a human must double-check every output, the efficiency gains are lost. Determinism provides the trust required to let AI operate at scale without constant oversight.

Building the Infrastructure for AI Execution

Scaling AI workflows requires an architectural approach built on three core pillars.

1. Data Grounding

An agent is only as effective as the tools it can access. Reliable execution requires connecting agents directly to your systems of record—ERP, CRM, and PLM—via secure APIs. Rather than allowing an AI agent to "reason" its way to a part number, it should be programmed to use a lookup tool to retrieve the exact SKU from the database. This ensures the agent is operating on facts, not probabilities.

2. Task Atomization

The most common mistake in AI implementation is over-scoping. Asking one agent to "manage procurement" is a recipe for failure. A more effective blueprint involves breaking complex processes into "atomic" tasks. By isolating specific actions—such as data extraction, policy comparison, or supplier communication—you can optimize the prompt and the toolset for that specific step, increasing the overall success rate.

3. Testing and Iteration

Reliability is built through rigorous validation. Organizations should move away from hidden system prompts toward a transparent environment where logic is evaluated against diverse datasets and can be verified. Using a "Chain of Thought" framework ensures that the agent follows a standardized, logical path to its conclusion, making its decisions auditable and repeatable.

5 Practical Steps for Moving to AI Execution

1. **Map Your "Shadow" Workflows:** Identify tasks that rely on manipulating data in spreadsheets, completing forms or documents, and using email for execution. These manual hand-offs between systems are the most valuable candidates for agentic intervention.
2. **Start with High-Volume Verification:** Focus on repetitive validation tasks—like checking safety certifications or tracking shipments—where the rules are clear, but the data is unstructured.
3. **Move to "Human-on-the-Loop":** Transition your team from being "doers" to "reviewers." The agent prepares the work to 90% completion, with humans doing the final approval.
4. **Standardize the Hand-off:** Require your agents to produce structured outputs, such as JSON or standardized tables. This ensures that the data can flow back into your ERP without manual cleanup.
5. **Treat Failures as Fine-Tuning Data:** Every time an agent fails to parse a document, it provides a specific data point to refine the prompt logic or tool permissions in the next version.



Operational Success: Real-World Applications

This shift is already reshaping how global firms manage complex operations.

- **Compliance Automation:** A large-scale manufacturer was manually auditing thousands of supplier safety documents, a process that took weeks and was prone to oversight. By implementing a Document Agent, they automated the extraction of coverage limits and expiry dates. This reduced the qualification lead time by 70% while ensuring 100% compliance with corporate standards.
- **Intelligent Reconciliation:** A logistics provider used an AI agent to cross-reference shipping manifests with physical warehouse receipts. What used to take a dedicated team hours of manual comparison is now handled in seconds, with the AI only flagging discrepancies requiring human intervention.

The infrastructure for the Agentic Enterprise is ready. The next step is to stop treating AI as a chatbot and start treating it as a core component of your execution strategy.





Choosing the Right Transformation Partner for a Multi-Cloud, AI-Enabled Future

Most enterprise transformation programs do not stall because the technology fails. They stall because the operating model around the technology was never designed to hold it. Data sits in places where it cannot be activated. Integration layers are fragile. Governance is unclear. And when AI capabilities are layered on top, the cracks widen rather than close.Data sits in places where it cannot be activated. Integration layers are fragile. Governance is unclear. And when AI capabilities are layered on top, the cracks widen rather than close.

Choosing a partner used to be a question of which firm could implement a specific cloud. That question is no longer the right one. The decision now is which partner can help you make multiple clouds work as a single operating model, with data, workflows, and decisions that move cleanly across the business.

What to look for: Operational readiness, not feature checklists

Building an Agentic Enterprise requires more than enabling AI features. It requires a partner who can integrate, govern, orchestrate, and scale without introducing instability. Enterprise-grade automation does not operate inside a single cloud. It moves across customer journeys, supply chains, revenue processes, service operations, and marketing ecosystems. If one layer is weak, performance degrades. If integration is fragmented, processes stall. If governance is missing, risk increases.

The right partner is the one who can tell you, before the work starts, where your data will hesitate, where your processes will stall, and where your governance will be tested. Automation delivers value only when it operates inside structured environments, and structured environments are the result of deliberate design, not optimism.

Proven execution across the clouds that run your business

Cross-cloud delivery is where most partner conversations get vague. A partner can claim breadth on a website. The harder question is whether the same teams have stood up each of those clouds in production environments in your industry.

This is the standard to which OSF Digital holds itself and where our credentials are relevant. We hold Expert-level Salesforce Competencies across Commerce, Marketing, Sales, Service, Platform, and Data 360, alongside Expert recognition in Consumer Goods, Retail, Manufacturing, and High Tech, with additional Accredited Competencies in Agentforce, Analytics, Financial Services, and Media. That coverage matters because the teams designing your commerce experience understand how it connects to service. The people configuring your service operation understand the data your marketing and sales teams rely on. The architects integrating your systems understand the industry workflows underneath them.

Products



Industries



Experience that prevents failure before it compounds

Enterprise transformation involves interconnected decisions spanning systems, teams, data, and processes. The cost of getting those decisions wrong is rarely visible the moment they are made. It shows up later as the AI initiative that cannot find clean data, the customer journey that breaks at a handoff, or the new capability that cannot be governed at scale.

A partner with real depth across the architecture helps you avoid those costs before they compound. The right partner is not the one with the most awards on a website. It is the one whose people have done this work before, in your industry, across the clouds you actually use, and can tell you what will break before it breaks.

If you are evaluating partners for a multi-cloud or AI-enabled transformation, connect with OSF Digital to assess your current readiness and define a practical path forward.

About OSF Digital

OSF Digital is a Salesforce-first consulting partner helping organizations turn AI, data, and trust into measurable business impact. With 20 years of hands-on Salesforce experience and deep multi-cloud expertise, OSF guides clients from strategy through implementation to managed services across financial services, retail, manufacturing, and more. Known for award-winning delivery and a distinctive partnership mindset, OSF acts as an expert extension of its clients' teams. With regional hubs across North America, LATAM, EMEA, and APAC, OSF combines global reach with personal, high-touch delivery.

Stay current.
Stay current on Agentforce and AI topics for 2026. Scan the QR code to subscribe to the OSF Digital newsletter or visit **osf.digital**.


OSF DIGITAL