AI foundations: A guide to business transformation through AI-powered automation

HYLAND™

Contents

Executive summary

Everywhere you turn, the conversation is dominated by artificial intelligence. The promises are grand — imagine complete autonomous operations, transformative generative capabilities and a total rewrite of the global economy. Yet, for the average business leader responsible for unstructured data and content, the reality feels different. It feels … unachievable.

While many organizations are picking up AI tools and running forward at full throttle, an **MIT report** found that just 2 of 8 major sectors showed meaningful structural change, and while big firms lead in piloting AI, they lag in scale-up.

Unfortunately, the ripple effect around this perceived failure can be pervasive. The irony is that AI initiatives rarely fail because the technology itself is ineffective; too often, it's due to a lack of proper preparation and inability to feed the model adequate data.

No wonder many organizations feel their first AI use case must be the corporate equivalent of sending a rocket to the moon. We hear it in nearly every "get to know you" call — they're waiting for the perfect, massive AI application that will revolutionize their entire enterprise overnight.

This hesitation, while understandable, is a mistake.

The true power of AI in the enterprise today is not about moonshots. It's about eliminating wasted time and identifying the processes that have been broken for years — the mundane, the repetitive, the part of your operations that everyone knows is problematic even as-is — and applying new intelligence to fix them. It's about showing up for your AI tool as prepared as possible.

This guide is designed to cut through the AI hype. It is not a technical manual for coding large language models (LLMs), nor is it a futurist's prediction of the next decade. It is a practical roadmap for business leaders who need to understand how to start, how to scale and how to govern AI right now.

Over the following chapters, we will explore:

1. **Where to start:** Moving from "freeze" to action by identifying the right use cases.
2. **The art of prompting:** Why this is your organization's new shared language.
3. **Security reality:** Navigating compliance in EMEA and the U.S. without fear.
4. **Data integration:** The nuanced relationship between structured and unstructured data.
5. **The human element:** Why technology is only 30% of the equation.
6. **Real-world examples:** Using AI in content management.

The goal is simple: To help you move from asking "What is everyone else doing?" to asking "If we could solve this problem with our data, what would we do?"

Let's go.

Where do I start with AI?

The most common question business leaders ask today is also the most difficult to answer in a vacuum:
“What is my first use case?”

The paralysis stems from a misunderstanding of what AI is meant to do in a business context. We often conflate automation with agentic AI.

Don’t conflate automation with agentic AI

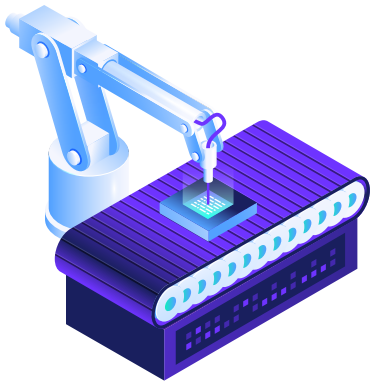

Automation is the digitization of a standard operating procedure. It follows a rigid set of rules.

Agentic AI, however, mimics the human decision-making that happens while reading that standard operating procedure. The ability to reason, adapt and act based on context is embedded in the tool.

To find your starting point, do not look for a process that requires a genius-level agent. Look for the hot spots in your organization — the mundane tasks that consume disproportionate amounts of human energy.

Identifying the opportunity

When identifying potential first use cases, focus on scenarios where the impact of improvement is tangible, yet the complexity of implementation remains manageable. For example, start by looking for processes involving unstructured data where you have historically accepted a high failure rate or high manual intervention cost.

Tap into more data

We know that 80% of enterprise content is unstructured, which means your organization is running on just a portion of the data you have. By tapping into that unused data with AI, you activate it and fill in gaps in existing processes that aren't working for you. Once that unstructured data is structured, your downstream AI tools can make the most of it.

These use cases often revolve around tasks where unstructured data, such as scanned documents, emails or legacy records, plays a central role. Processes within these tasks typically involve repetitive actions, manual data handling or frequent decision-making bottlenecks.

Target workflows with noticeable inefficiencies

Workflows that are well known as problematic are prime starting points. Most teams can point a finger at offending workflows pretty quickly — think of workstreams that:

- Cause delays in processing critical documents
- Introduce data entry errors
- Consistently underperform on compliance standards due to inconsistent information
- Involve high volumes of paper or fragmented systems

Often, these environments are ripe for automation and integration improvements.

The “if you could” framework

If you are struggling to identify these spots, gather your business process owners and ask them a simple question:

> If you could do anything with your unstructured data, what would you do?”

Ignore the technology aspect for a moment. Don’t think about ***how***.

Would you classify 2.5 million archived documents to find risk exposure?

Would you let customers self-serve their dispute resolutions?

Would you extract data from contracts to stop revenue leakage?

Once you identify the desire, the technology conversation becomes easier. These are not moonshots; they are quick wins that build organizational confidence and clear the path for true transformation.

Small adjustments, big wins

Explore areas where small gains can produce outsized benefits. For example, automating invoice processing — a well-established and high-success endeavor — can lead to significant time savings, faster approvals and reduced human error. Similarly, streamlining patient record management in healthcare can drive both operational efficiency and improved care delivery.

A word of caution: AI is still an innovation-led early technology, and not all good ideas around its use will generate an acceptable ROI. Use case development also comes with a warning to validate at all points that there is good ROI potential in final delivery. Most customers who are early adopters make use of the new technology to overcome issues that have existed for an extended time — a good place to look first, rather than moonshot ideas.

Look for teams who will champion AI

Crucially, partner with stakeholders across departments to validate these opportunities. MIT found that a team’s unwillingness to adopt new tools was the top barrier to scaling AI in the enterprise. The input and emotional buy-in of your team helps uncover pain points and ensures alignment with broader organizational goals.

By focusing initially on well-defined, high-impact problems, your organization can establish a strong foundation for broader adoption, demonstrating clear value while minimizing risk.

Spotlight: Quick AI wins

The mailroom dilemma:

Consider a digital mailroom. Traditional automation might achieve a 45% effectiveness rate because it relies on rigid templates. If an invoice arrives in an unexpected format, the system breaks. An AI-enhanced approach does not need a template. It reads the document like a human would, potentially driving effectiveness up to 80% or 90%.

The “1%” problem:

In retail and logistics, 99% of invoices might match the purchase order. The 1% that do not match often cost as much to resolve as the other 99% combined. This is a prime target for AI. The goal is to turn that 1% exception-handling into a standard, low-cost transaction.

Regulatory Q&A:

Imagine a global facilities management team. Instead of manually searching through thousands of PDF manuals for a compliance answer, they could query a knowledge base in natural language. “What is the safety protocol for this equipment in Germany?” The AI retrieves the specific answer from the unstructured data instantly.

Key takeaways

- Start with the mundane.
- Start with the tasks people hate doing.
- Start with the processes that are prone to human error due to fatigue or boredom.

Prompting: The new corporate language

We should all start treating prompting as a core business skill. Too often, we use enterprise AI tools the same way we use Google: typing in a few keywords and hoping for the best. This results in generic, unhelpful outputs.

Prompting is the mechanism by which we transfer implied business knowledge — the underlying language of your organization — into the AI system. It is how we translate the unwritten rules of your company culture and operational preferences into a digital format that an agent can act upon.

From keywords to context

A prompt is simply an extraction of what a subject matter expert does daily. However, different industries (and even different departments or team members) use the same vocabulary to mean vastly different things.

Take the word “hedge.” To a gardener, it is a boundary. To a currency trader, it is a risk mitigation strategy. To a fund manager, it is an investment vehicle. If your prompt does not establish the context, the AI may hallucinate or provide irrelevant information.

Developing strong prompting practices drives consistent agent performance. If you want a repeatable, auditable result, you need a repeatable, well-structured prompt.

Perhaps the biggest challenge in effectively wielding prompts is in fully understanding the business processes that will benefit and ensuring the correct processes and choices are understood. The desired outcomes should be mapped out at the same time, so that elements are correctly captured and tested within the prompt. This helps you — as much as possible — create a predictable, repeatable outcome.

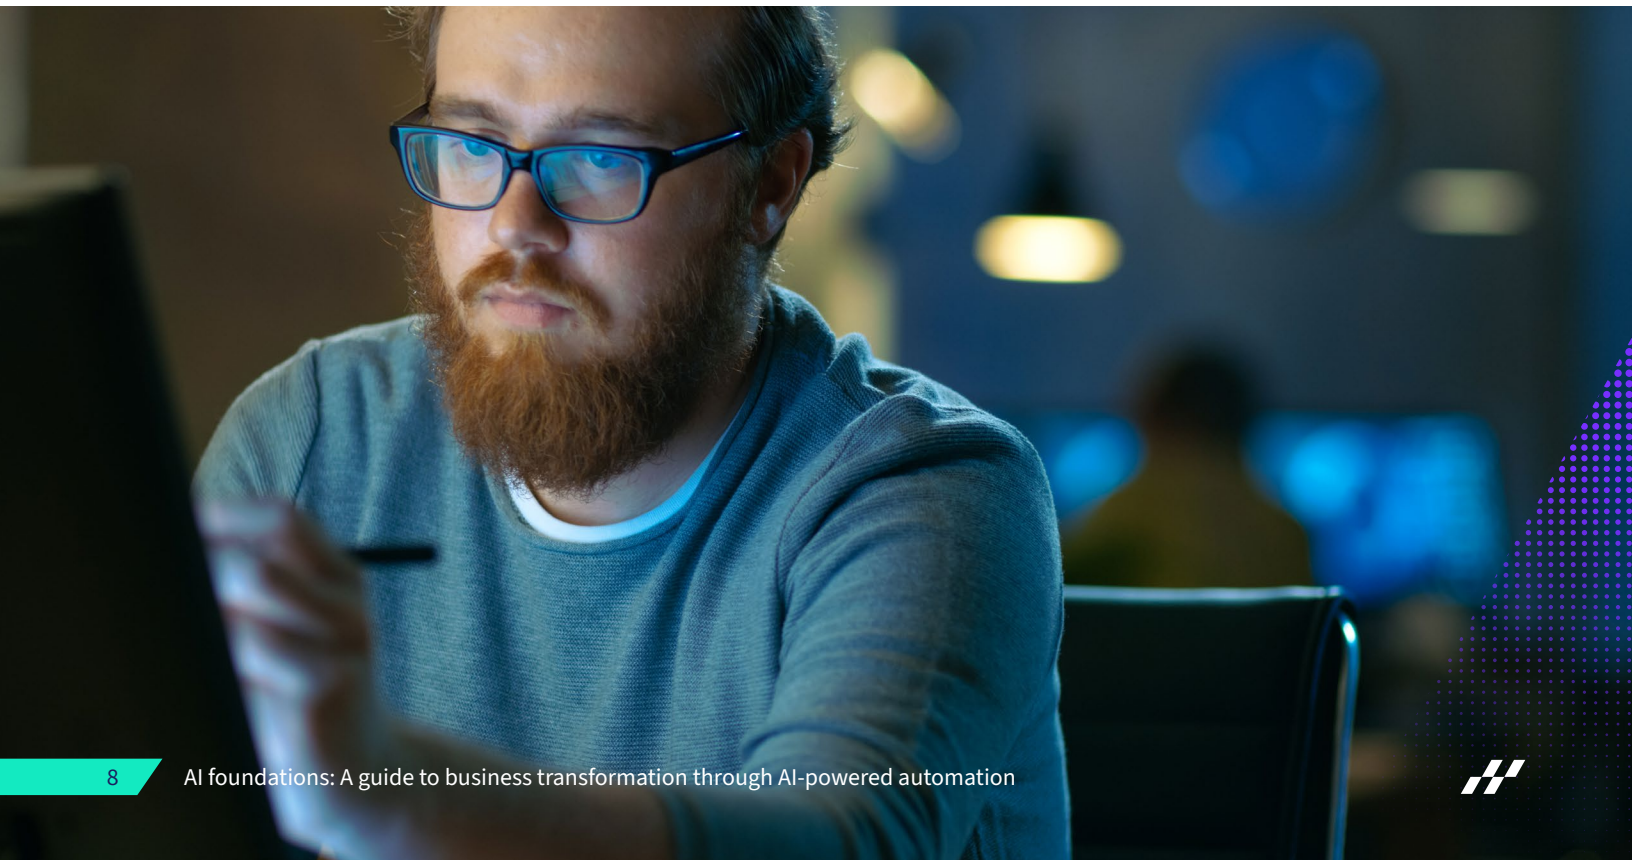

Establishing a shared language

As organizations mature in their AI use, prompting will likely become standardized. We will move away from individual prompting toward a structured library of approved prompts that reflect the company's risk appetite and tone.

For example, organizations might enforce "restricted use" words or mandatory frameworks to ensure AI outputs align with corporate policy.

With the use of agentic AI, you can lean into a network of agents that are tuned to specific verticals or personas. Hyland Agent Mesh accomplishes this, so that your finance agent understands that a "hedge" is financial; while a facilities agent understands a "hedge" is agricultural. Your prompting strategy must feed these specific personas.

Understanding your organization's internal culture and its appetite or ability to change (at all levels) is critical. This internal reflection is yet another factor of AI that is more of a people focus than a technology focus. If missed or ignored, these people-oriented changes likely will be costly when it comes to adoption and ultimate performance.

Hackathons as knowledge capture: One of the best ways to build strong prompting capabilities is to run internal prompt hackathons dedicated to specific working practices. Challenge your teams to solve a specific business problem using AI. Do not just record the winner; record the process. The failed prompts are just as valuable as the successful ones because they reveal the nuances of the business logic that needs to be codified.

Key takeaways

- Prompting is not just typing; it is programming with natural language.
- It requires a deep understanding of the business process.
- Successful organizations will train their people to think like prompt engineers.

Security and data management in Europe

Security conversations regarding AI often devolve into fear. In Europe, the EU Artificial Intelligence Act is well on its way to being fully adopted. In the U.S., data privacy concerns dominate boardrooms. However, compliance is not about avoiding technology; it is about implementing the correct controls and picking suppliers that match your needs and expectations.

AI is predominantly a cloud-based technology due to the shared-cost nature of LLM development. To enable the use of AI, this has pushed some organizations to reconsider their approach to utilizing cloud technologies.

New considerations need to be addressed with suppliers to ensure all aspects of security are covered when linked to AI, such as:

- Ensuring data is encrypted in transit and at rest to the highest standard practically available
- Using key management for protection
- Ensuring applications are tested for vulnerabilities
- Confirming all requirements are validated and standards are met

In comparison, organizations choosing to build their own applications must face these requirements on their own, demanding vigilance to ensure no corners are cut to achieve results beyond proof of concept. In the world of AI, infrastructure choices and the deployment of relevant security technologies that don't degrade performance are increasingly important criteria for vendor evaluation and selection.

The Hyland Trust Center clearly shows which standards are in place with guidance on how technologies are used.

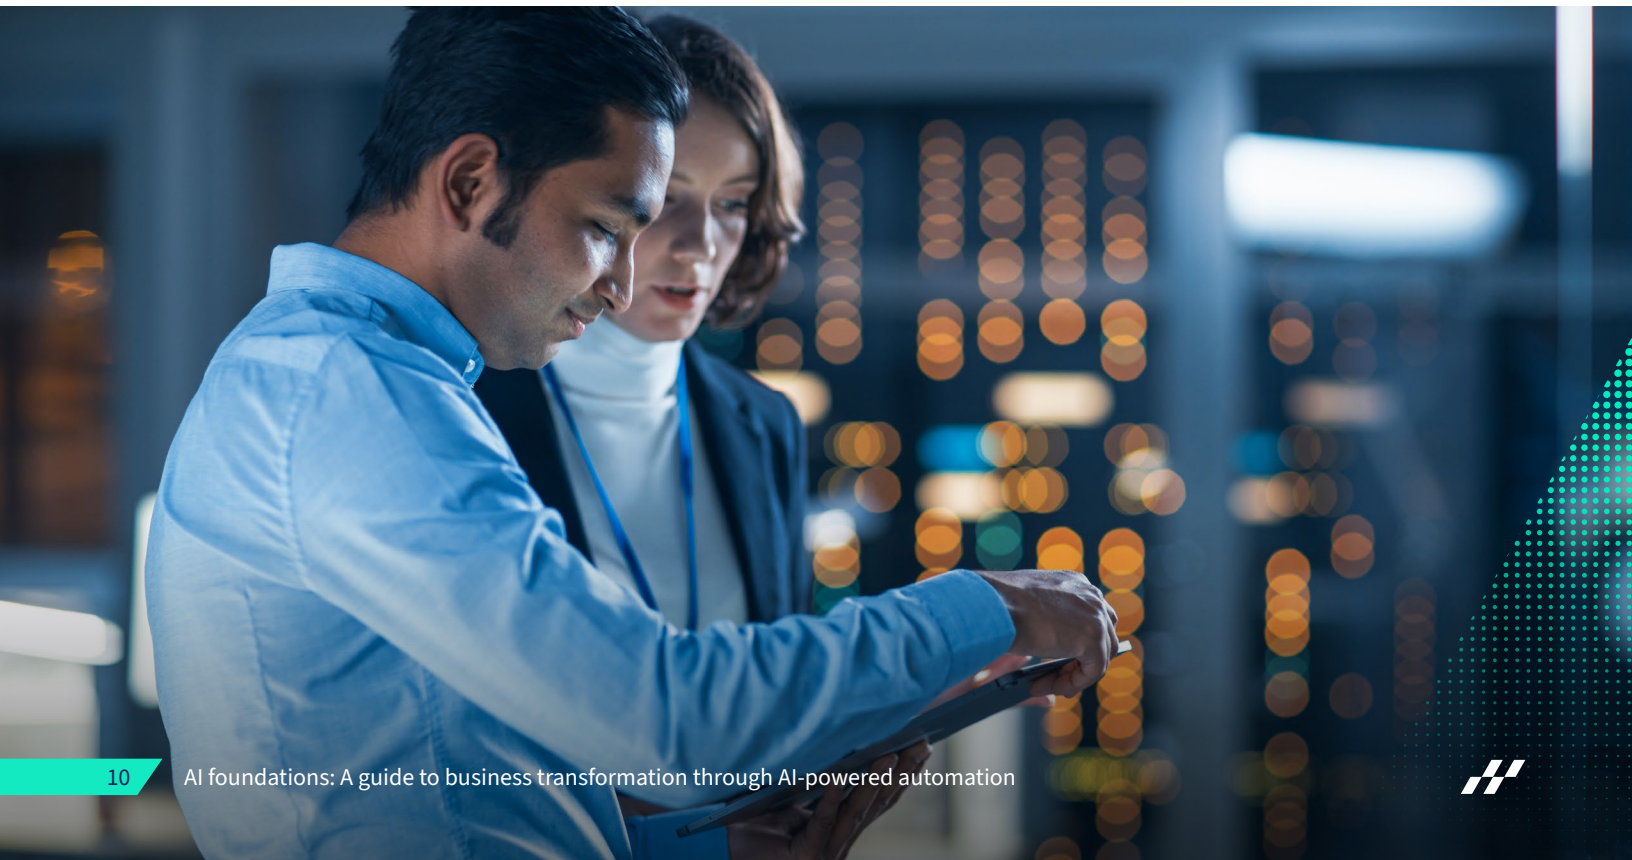

Moving beyond “trust us”

When evaluating AI vendors, especially for operations in Germany (requiring C5 compliance) or broader EMEA, “trust us” is not an acceptable answer. You need transparent descriptions of controls.

- **Isolation of data:** Ensure your vendor guarantees that your data is not being used to train the public foundation model. Your data must remain your data.
- **Zero-retention policies:** Look for architectures where the AI processes the data for the transaction and then forgets it immediately.
- **Vulnerability testing:** AI software is still software. It requires the same rigorous pen-testing, vulnerability scanning and access control audits as any other enterprise application.

Data sovereignty in a European context

For organizations operating in Europe, data sovereignty is closely tied to regulatory accountability, including requirements set out by GDPR and, increasingly, the **Digital Operational Resilience Act (DORA)**. These frameworks focus on how data is stored, processed, protected and governed, rather than assuming data is confined to a single geographic location. In practice, AI and cloud based architectures often involve cross border data processing, which must be managed through appropriate legal, technical, and organizational controls.

While the EU’s **General Data Privacy Regulations (GDPR)** doesn’t explicitly address artificial intelligence, its data protection principles can be compatible with AI and big data. By integrating privacy by design and conducting thorough risk assessments, European businesses can confidently drive innovation while maintaining robust data protection standards.

Ultimately, security in AI is about asking the right questions. Do not ask “Is AI safe?” Ask “How do you manage encryption keys?” “How do you prevent data leakage between tenants?” “How do you ensure our PII is isolated from the LLM?” The answers to these questions will separate enterprise-ready partners from experimental tools.

Hyland and EU requirements: Hyland supports these requirements through EU-based data centers, adherence to recognized security and governance standards, and a privacy-by-design approach that is embedded in our architecture. This ensures any international data transfers remain aligned with applicable European regulatory obligations. In our AI offerings, Hyland uses only AWS components that are C5-certified.

Key takeaways

- The EU AI Act doesn’t prohibit AI systems; it provides parameters on AI use cases.
- Your AI partner should offer robust mechanisms to manage PII securely and prevent any unintended data exposure.
- Focus on asking detailed, technical questions to assess an AI partner’s readiness for enterprise security and data management requirements.

Combining structured and unstructured data

For decades, ERP systems have been the source of truth that organizations depend on to keep things running. The **structured data** stored within them (think rows, columns, numbers and dates) is what mission-critical operations are based on, from HR records and accounts payable to case management and customer-facing decision-making.

Unstructured data, which represents approximately 80% of enterprise content, is the data found in less conforming formats like emails, PDFs, contracts and images. It can be vast, messy and difficult to harness.

Now, AI gives us the ability to easily extract data from the unstructured world and pour it into the structured world. The question is: Should we?

Organizations must evaluate if they actually need this combination. Questions for consideration include:

- Will combining your structured and unstructured data benefit your business?
- Does combining both sets of data outweigh the cost of investment?
- Have you identified the outcome you desire before deciding to combine these data sources?

Hyland recommends you identify what outcomes you're seeking from AI. Once you know the outcomes, the question becomes: How can we best arrive there? Oftentimes, it's by combining your structured and unstructured data. The AI models need that input to arrive at your identified goal.

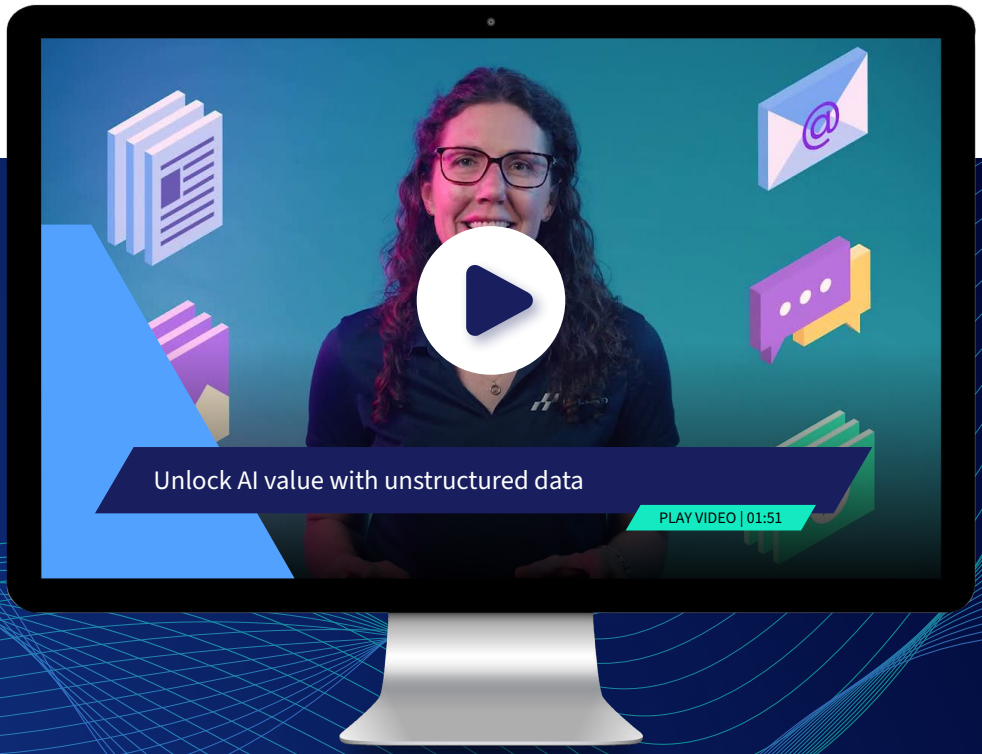

The case for structuring the unstructured

Structured data systems can't handle ambiguity. If you introduce data extracted from a document that is 99% accurate, you are still allowing a 1% error rate into your system. Traditional methods of data processing can't navigate around the discrepancies. For example, consider:

- **Rounding errors:** A human entering data might round a figure to two decimal places. An AI scraping a spreadsheet might capture eight decimal places. When that data hits the ERP, the conflict can cause reconciliation errors that cost millions.
- **Conflict resolution:** If the ERP says a customer address is X, and the new contract says Y, who wins? Without a "human in the loop" or strict validation rules, automating this merge can degrade the integrity of your core systems and slow everything down.

When done correctly, combining your structured and unstructured data drives immense value. The gap is bridged with intelligent document processing (IDP) through:

- **Capture and extract:** Using AI, your unstructured is "read" and intelligently processed. The useful bits are extracted and unified with your structured data.
- **Enrichment:** Your structured data is often incomplete. AI can mine historical documents to fill in the gaps — classifying customers, identifying contract renewal dates or extracting terms that were never captured in the original manual entry.
- **Validation:** Instead of blind overwriting, use AI to flag discrepancies. "The invoice says $10,000, but the PO says $9,500." The AI acts as a sophisticated auditor, presenting the anomaly to a human agent for a decision.

The case for separation

Sometimes, the best strategy is to keep the data separate but accessible. Use AI to query the unstructured data for insights — without necessarily forcing that data into the ERP tables.

Key takeaways

- 80% of your enterprise data is unstructured, making it unusable in its current form.
- The goal is not just to combine data for the sake of it. It is to define the business outcome.
- If merging the data improves "cost to serve" or "revenue realization," do it — but wrap it in validation protocols that protect your source of truth.

Implementing AI requires technology and change management

The Boston Consulting Group reported game-changing findings about AI implementation that every leader should memorize:

- 10% of the value from data analytics comes from advanced algorithms (the LLM).
- 20% from the data and tech backbone (the application).
- 70% comes from change management, new processes and people upskilling.

We spend billions on the 10% and the 20%, yet we often ignore the 70%.

The psychology of change

Successful AI adoption is not a software install; it is a psychological transition. Change management frameworks often describe a curve: Denial, defense, discarding, adaptation, internalization.

Many organizations try to force their employees straight to internalization. They announce a new tool at a kickoff meeting and expect everyone to be on board by Monday. This approach fails.

Employees must be allowed to move through the stages.

They will defend their old ways of working and try to discard the new tool.

Leadership must guide them to adaptation — taking what they know about their job and applying it to the new environment.

Augmentation, not replacement

The fear that kills AI projects is the fear of replacement. Leaders must be transparent about the goal: moving from automation to augmentation.

We are moving toward an era of agentic operations, where humans and AI agents collaborate. The AI handles the mundane data processing (the 99%), and the human handles the nuance (the 1%).

To succeed, you must invest in the 70% — that significant portion of implementation that encapsulates the human element of change. Think of the 70% as the foundation of the content we share, the ideas we generate and the reasons we provide for adopting AI. Investing in this area means:

- Equipping your teams with the skills to adapt
- Fostering a culture of experimentation
- Addressing the psychological journey of change

By focusing on the 70%, you're not just installing AI, you're empowering your people to embrace it, collaborate with it and drive meaningful results. Train your teams on prompting. Run hackathons to let them "break" the tool and learn its limits. Acknowledge the "defense" stage. When employees feel they are the pilots of the AI, rather than its victims, you move from simple installation to genuine business transformation.

This can feel like an uphill battle. There are many stories about AI redundancies so your teams may feel at risk. However, when evaluating you likely will find there are many ways to augment your workforce rather than replace them. This augmentation allows you to gain additional revenue and profitability while your team adjusts to the changes.

A key statement that resonates with many leaders is that properly embracing AI puts them in the driver's seat (i.e., AI isn't happening ***to*** them, it's happening ***because*** of them.) Employees should feel the same and feel empowered to play and ask questions of the new technology. They may just invent a new career for themselves.

Key takeaways

- The key to AI implementation is in centering humans.
- Change management can be the determining factor in AI success.
- Human and AI agent collaboration is the future.

3 real-world applications of AI in content management

Discover how leading organizations use AI-native solutions to unlock the full potential of their enterprise content. These practical case studies of anonymous Hyland customers highlight the transformative impact of AI, from automating complex workflows to improving outcomes and governance. Explore these success stories to see how you can apply intelligent innovation to your own operations and drive meaningful outcomes.

The German online fashion brand that accelerated invoice approvals

A rapidly growing German online fashion brand faced a common retail bottleneck: a massive influx of vendor invoices that threatened to overwhelm its accounts payable team. To keep pace with demand, the retailer needed to replace slow, error-prone manual processing with an intelligent, automated approach.

The retailer received documents in unpredictable formats and multiple languages, making manual data entry tedious. Relying on human validation inevitably led to costly errors, delayed vendor payments and a frustrating lack of scalability. Furthermore, the company struggled to connect this disjointed process with its existing enterprise resource planning and electronic data interchange workflows.

To conquer this complexity, the brand implemented **Hyland IDP**, built on the **Content Innovation Cloud™**.

Embracing AI and agentic automation: The solution effortlessly navigates multiple languages and unpredictable invoice structures. Hyland IDP instantly automated data extraction, validation and exception handling, drastically reducing the need for manual intervention. The AI-native technology integrated seamlessly into the company's existing ERP workflows, providing a scalable foundation to support ongoing expansion. With real-time dashboards offering complete visibility, the retailer successfully accelerated approval cycle times, improved overall data accuracy and empowered its team to work smarter.

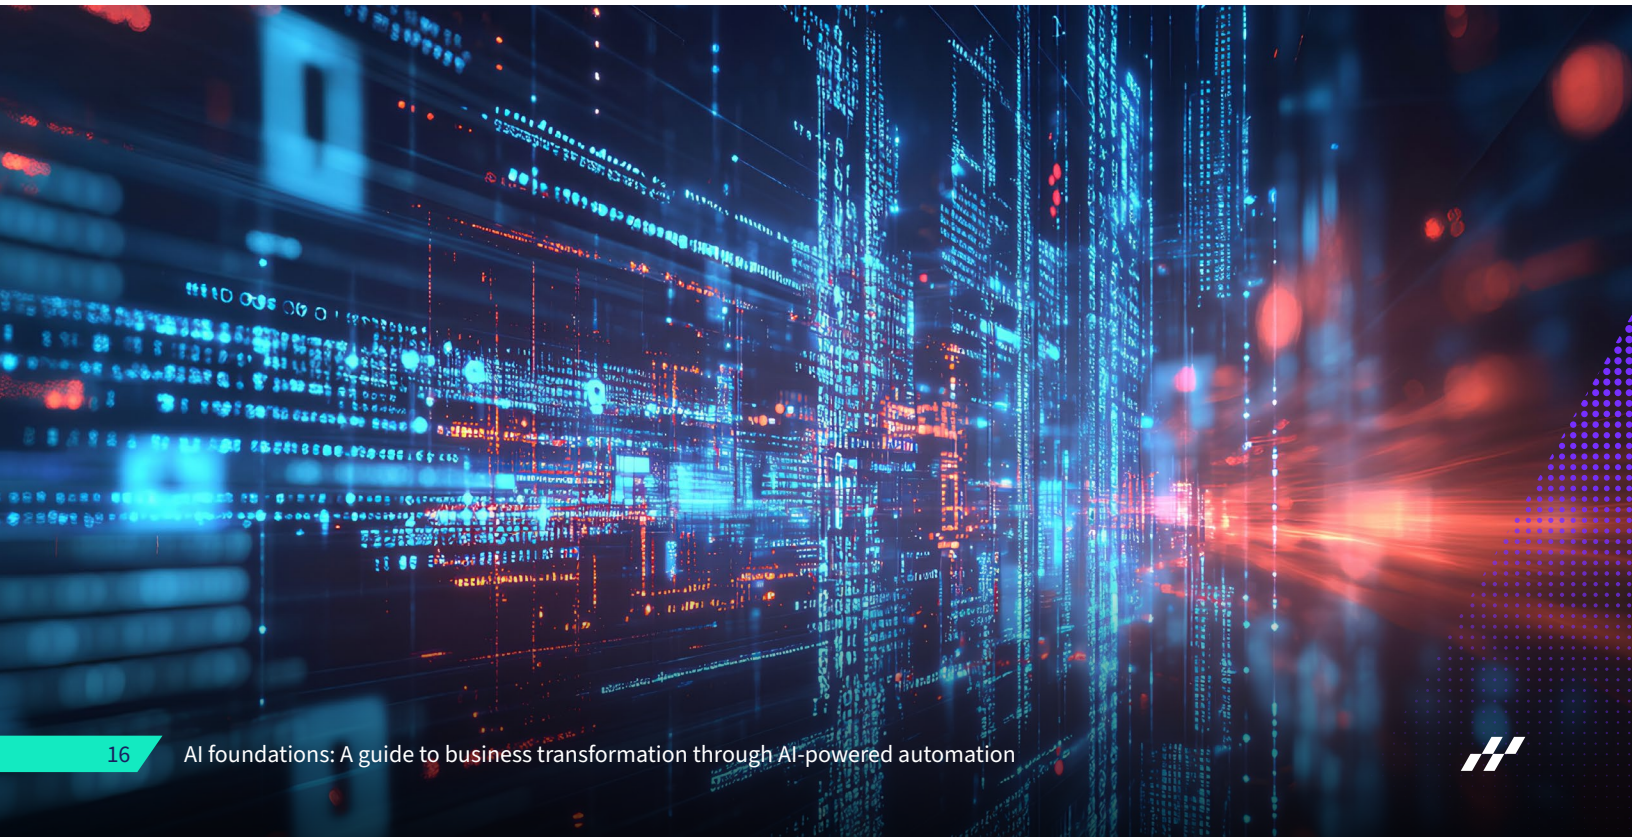

The South African healthcare payer that increased automation to over 80%

For a leading South African healthcare payer, a heavy reliance on manual claims verification created unsustainable processing delays.

It struggled with a low straight-through processing rate of approximately 45%, forcing it to rely on tedious manual verification. Diverse document templates and handwritten claims slowed turnaround times to two to three days per claim. Without a smarter approach to unstructured data, the payer faced ongoing inefficiencies and an inability to meet rising customer demands. It needed a reliable way to break through these bottlenecks and transform its unstructured data into a streamlined operational advantage.

An AI-driven solution for scalability: The advanced AI capabilities of the Content Innovation Cloud with Hyland IDP helped the payer minimize rigid template configuration and drastically reduced manual intervention. It managed unpredictable formats and handwritten notes, extracting data accurately to accelerate claims approvals. Overall, Hyland helped drive efficiency and establish a highly scalable foundation, successfully driving the organization's automation rate to over 80%.

The European infrastructure operator that automated manual transmittal validations

For a major European infrastructure operator, manual, document-by-document reviews created severe bottlenecks that limited their ability to scale operations efficiently.

Because the team relied on human validation, it struggled to maintain consistency and auditability across various projects. These bottlenecks severely limited their scalability and left the company exposed to compliance risks and mounting operational costs as project demands grew.

An AI-driven approach to governance: Using the Content Innovation Cloud with Hyland IDP and process automation, it automated metadata extraction, intricate content checks and digital signature verification. The operator transitioned to an exception-driven model that only requires human intervention when necessary. This transformation reduced manual effort and supports future growth while also improving information governance and consistency across projects.

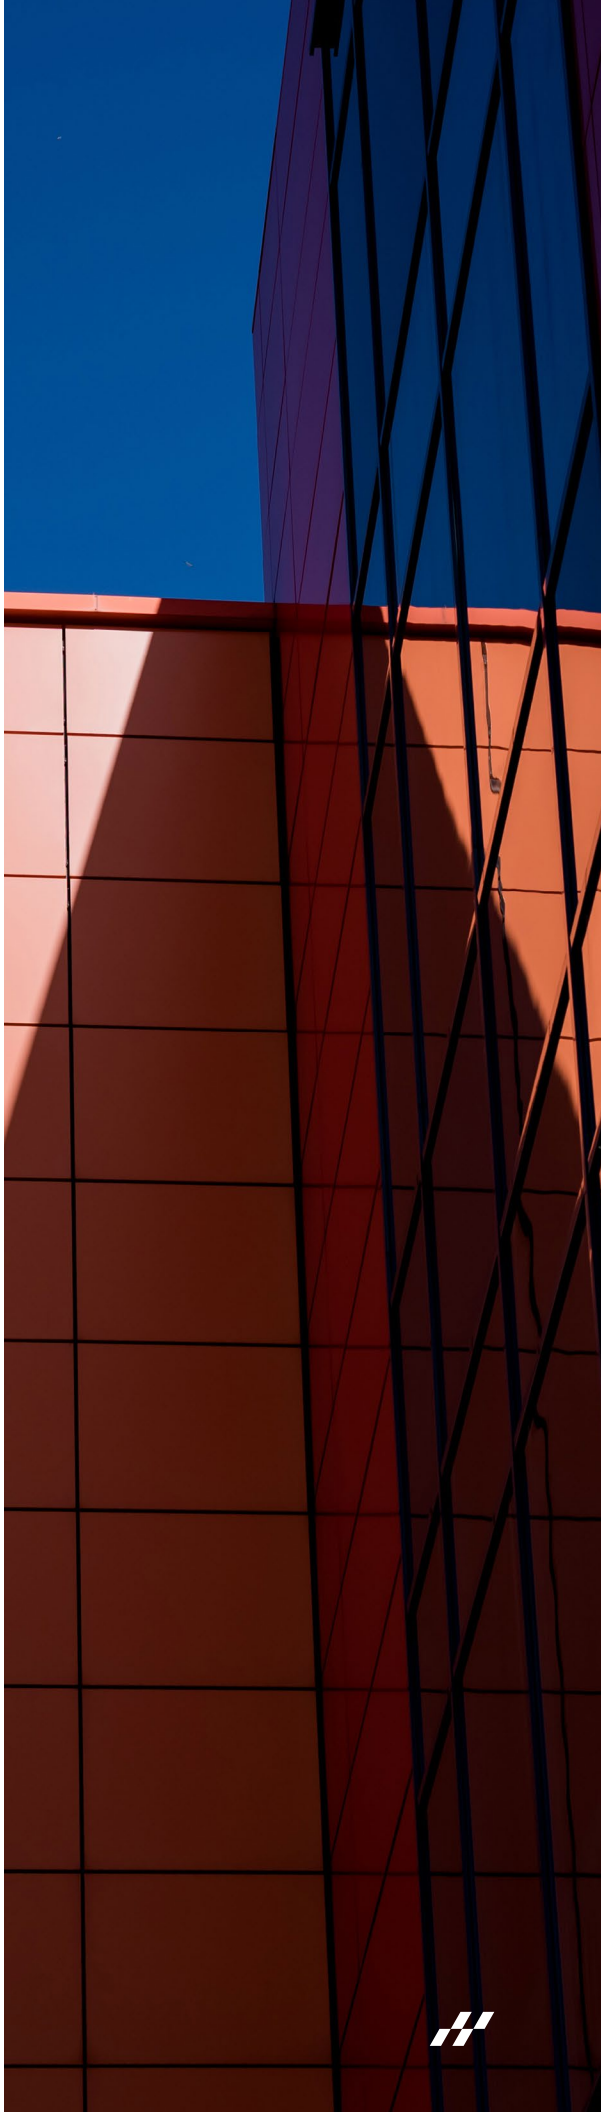

Unlock your potential with Hyland

The journey to AI maturity is not about buying the flashiest model; it is about solving real business problems with your content.

Hyland Content Innovation Cloud is built to bridge the gap between unstructured content and intelligent action. By leveraging our deep experience in content services with cutting-edge agentic automation, we help you secure your data, streamline your problematic processes and empower your people.

Do not let the noise paralyze you. Start with the mundane. Start with your data.

Explore Hyland's AI capabilities.

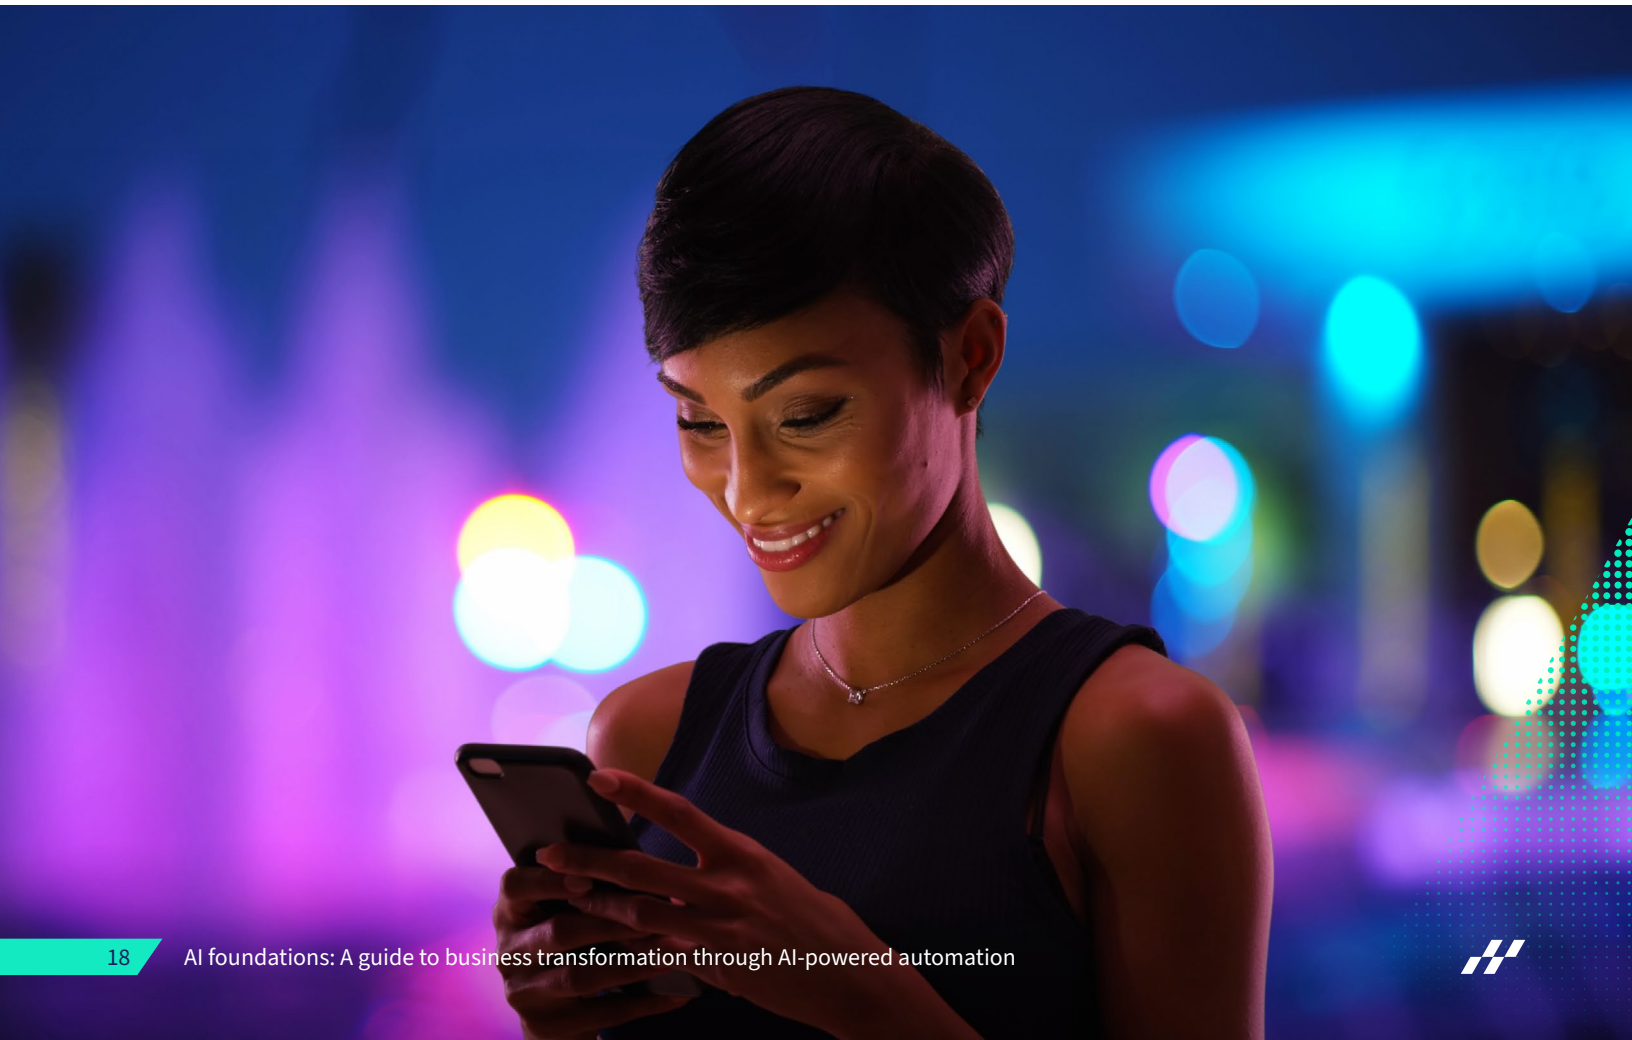

Looking for more? You may also like:

Before you invest in AI, assess your AI-readiness

Before AI can achieve what you expect it to, your enterprise content must be AI-enabled. Here's the AI-readiness framework to consider.

Switching to Hyland

Smarter, stronger, faster growth is waiting for you. It's time to elevate how you and your teams connect, innovate and interact with the world.

Your agentic enterprise: The new reality of AI and enterprise content

With Hyland's bold approach to weaving AI into your content universe, you can move faster than ever and realize incredible outcomes.