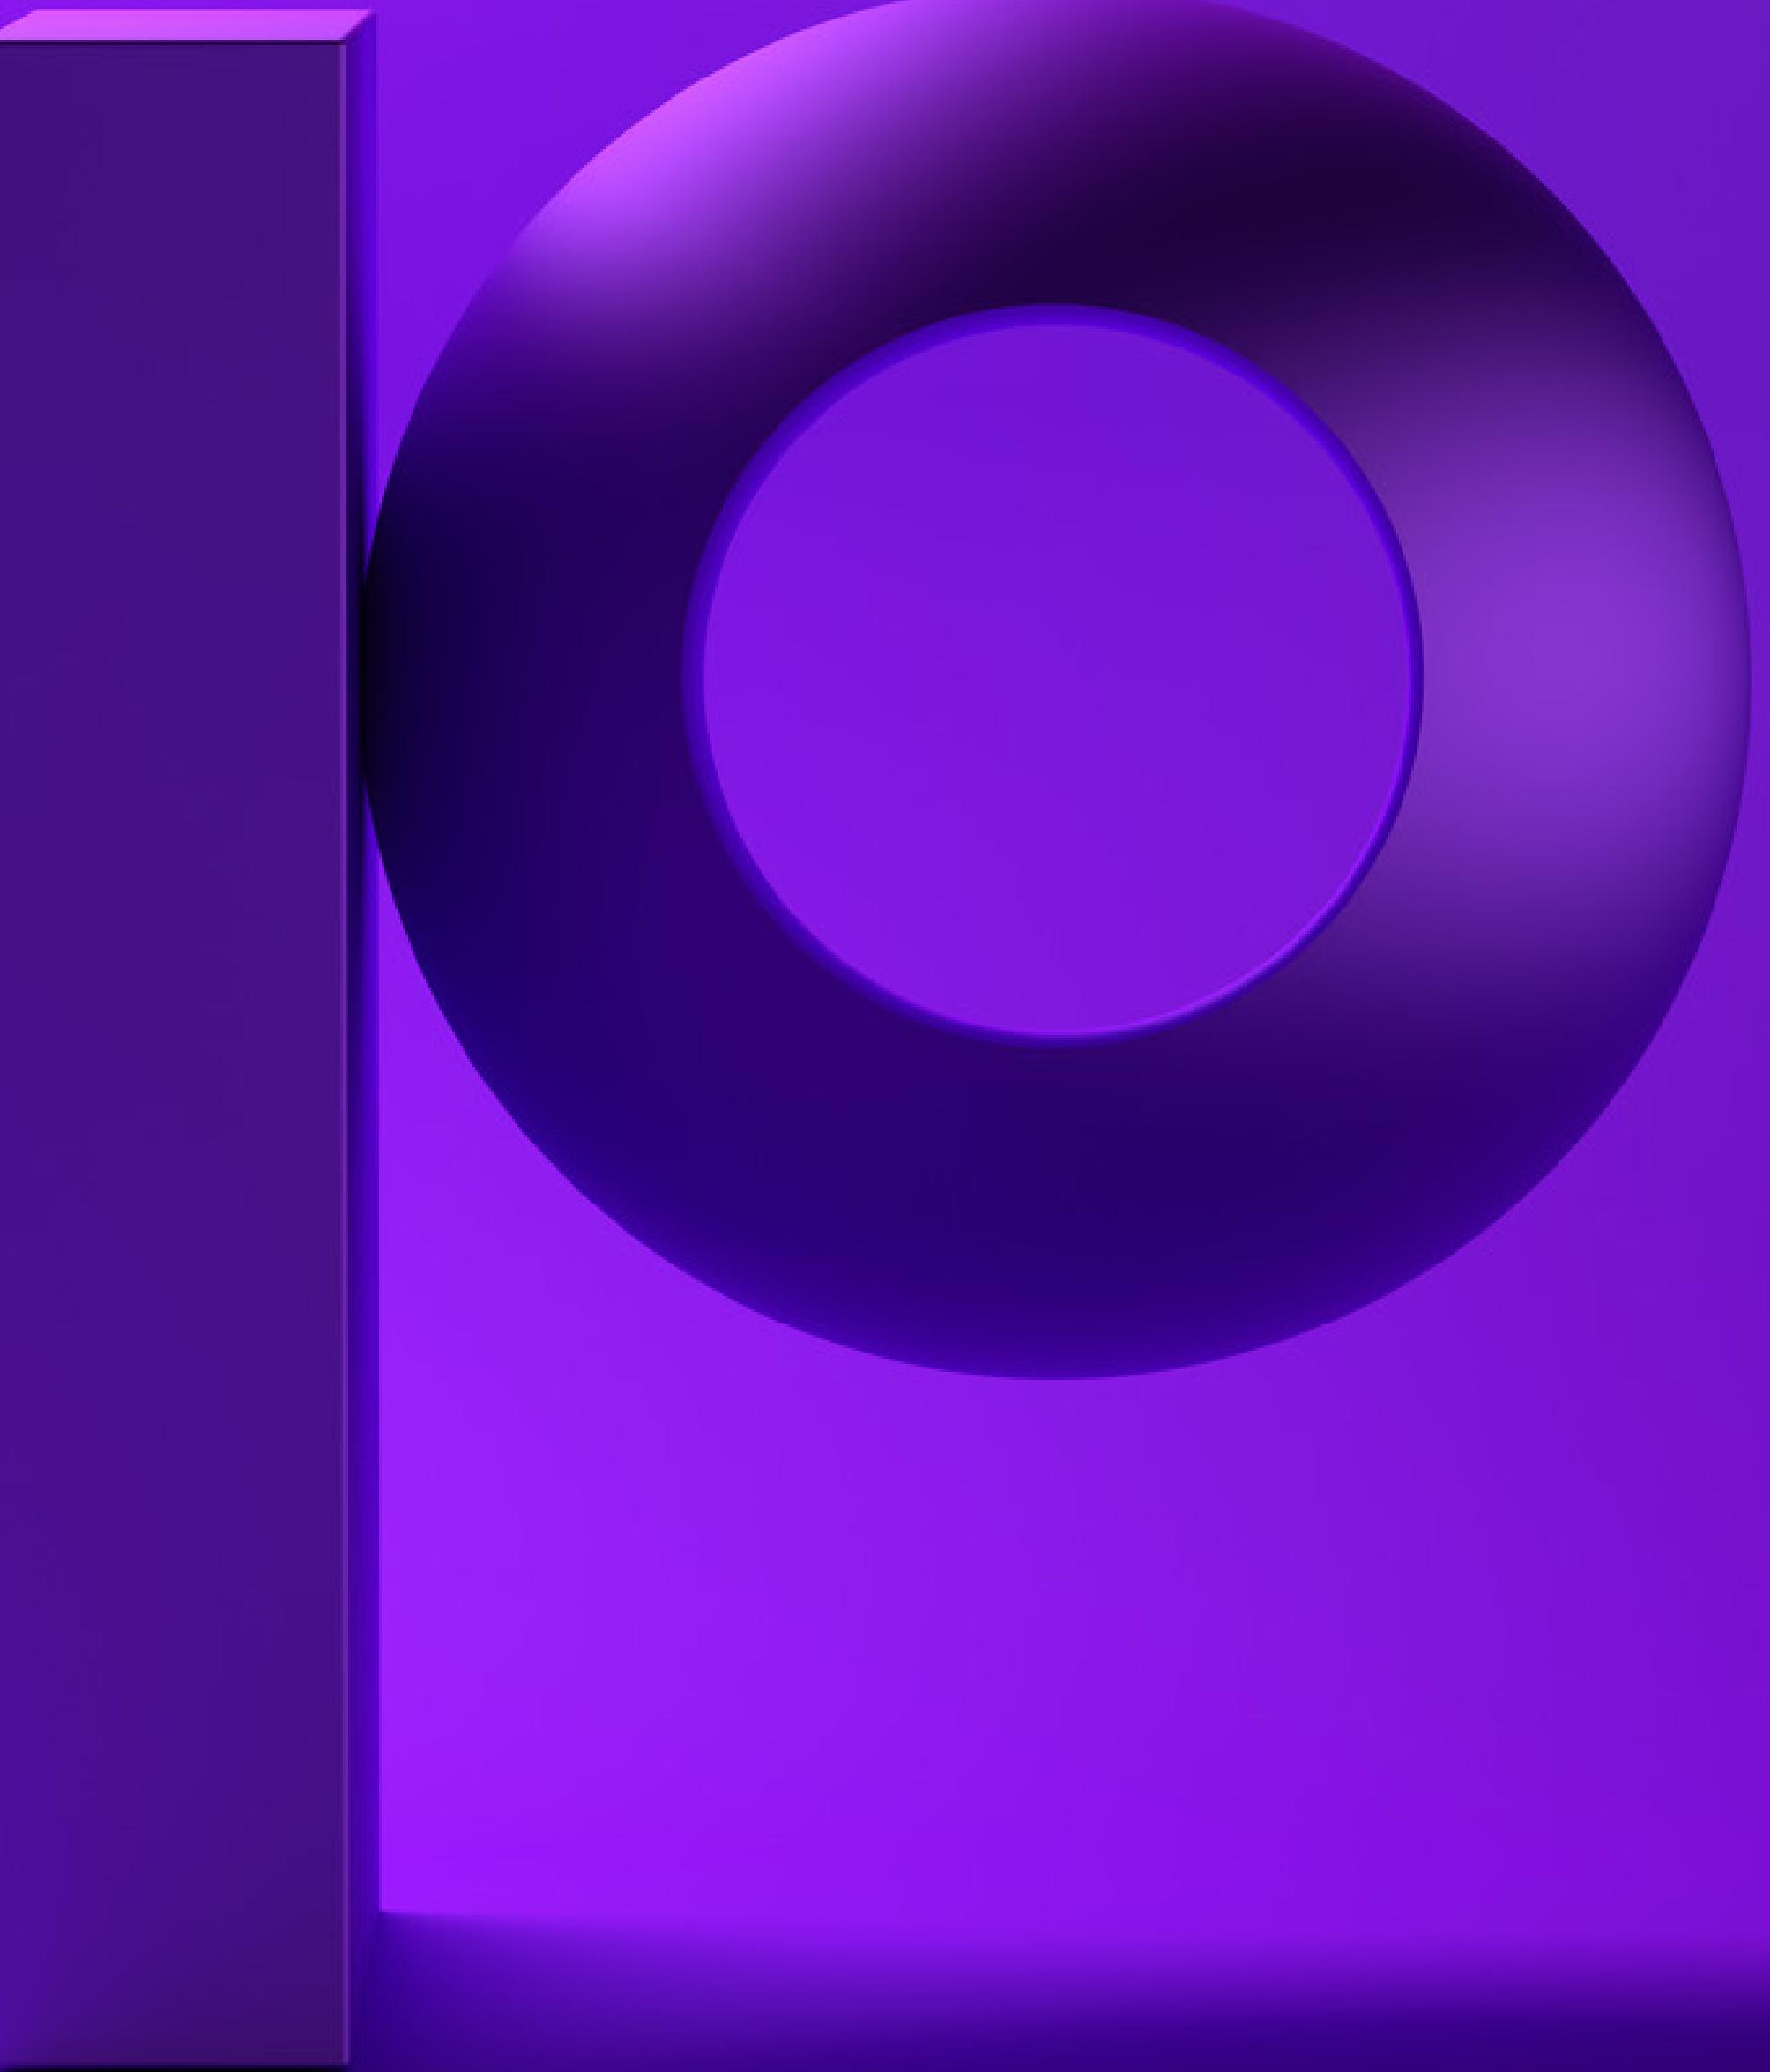


precisely

AI and Data Integration:

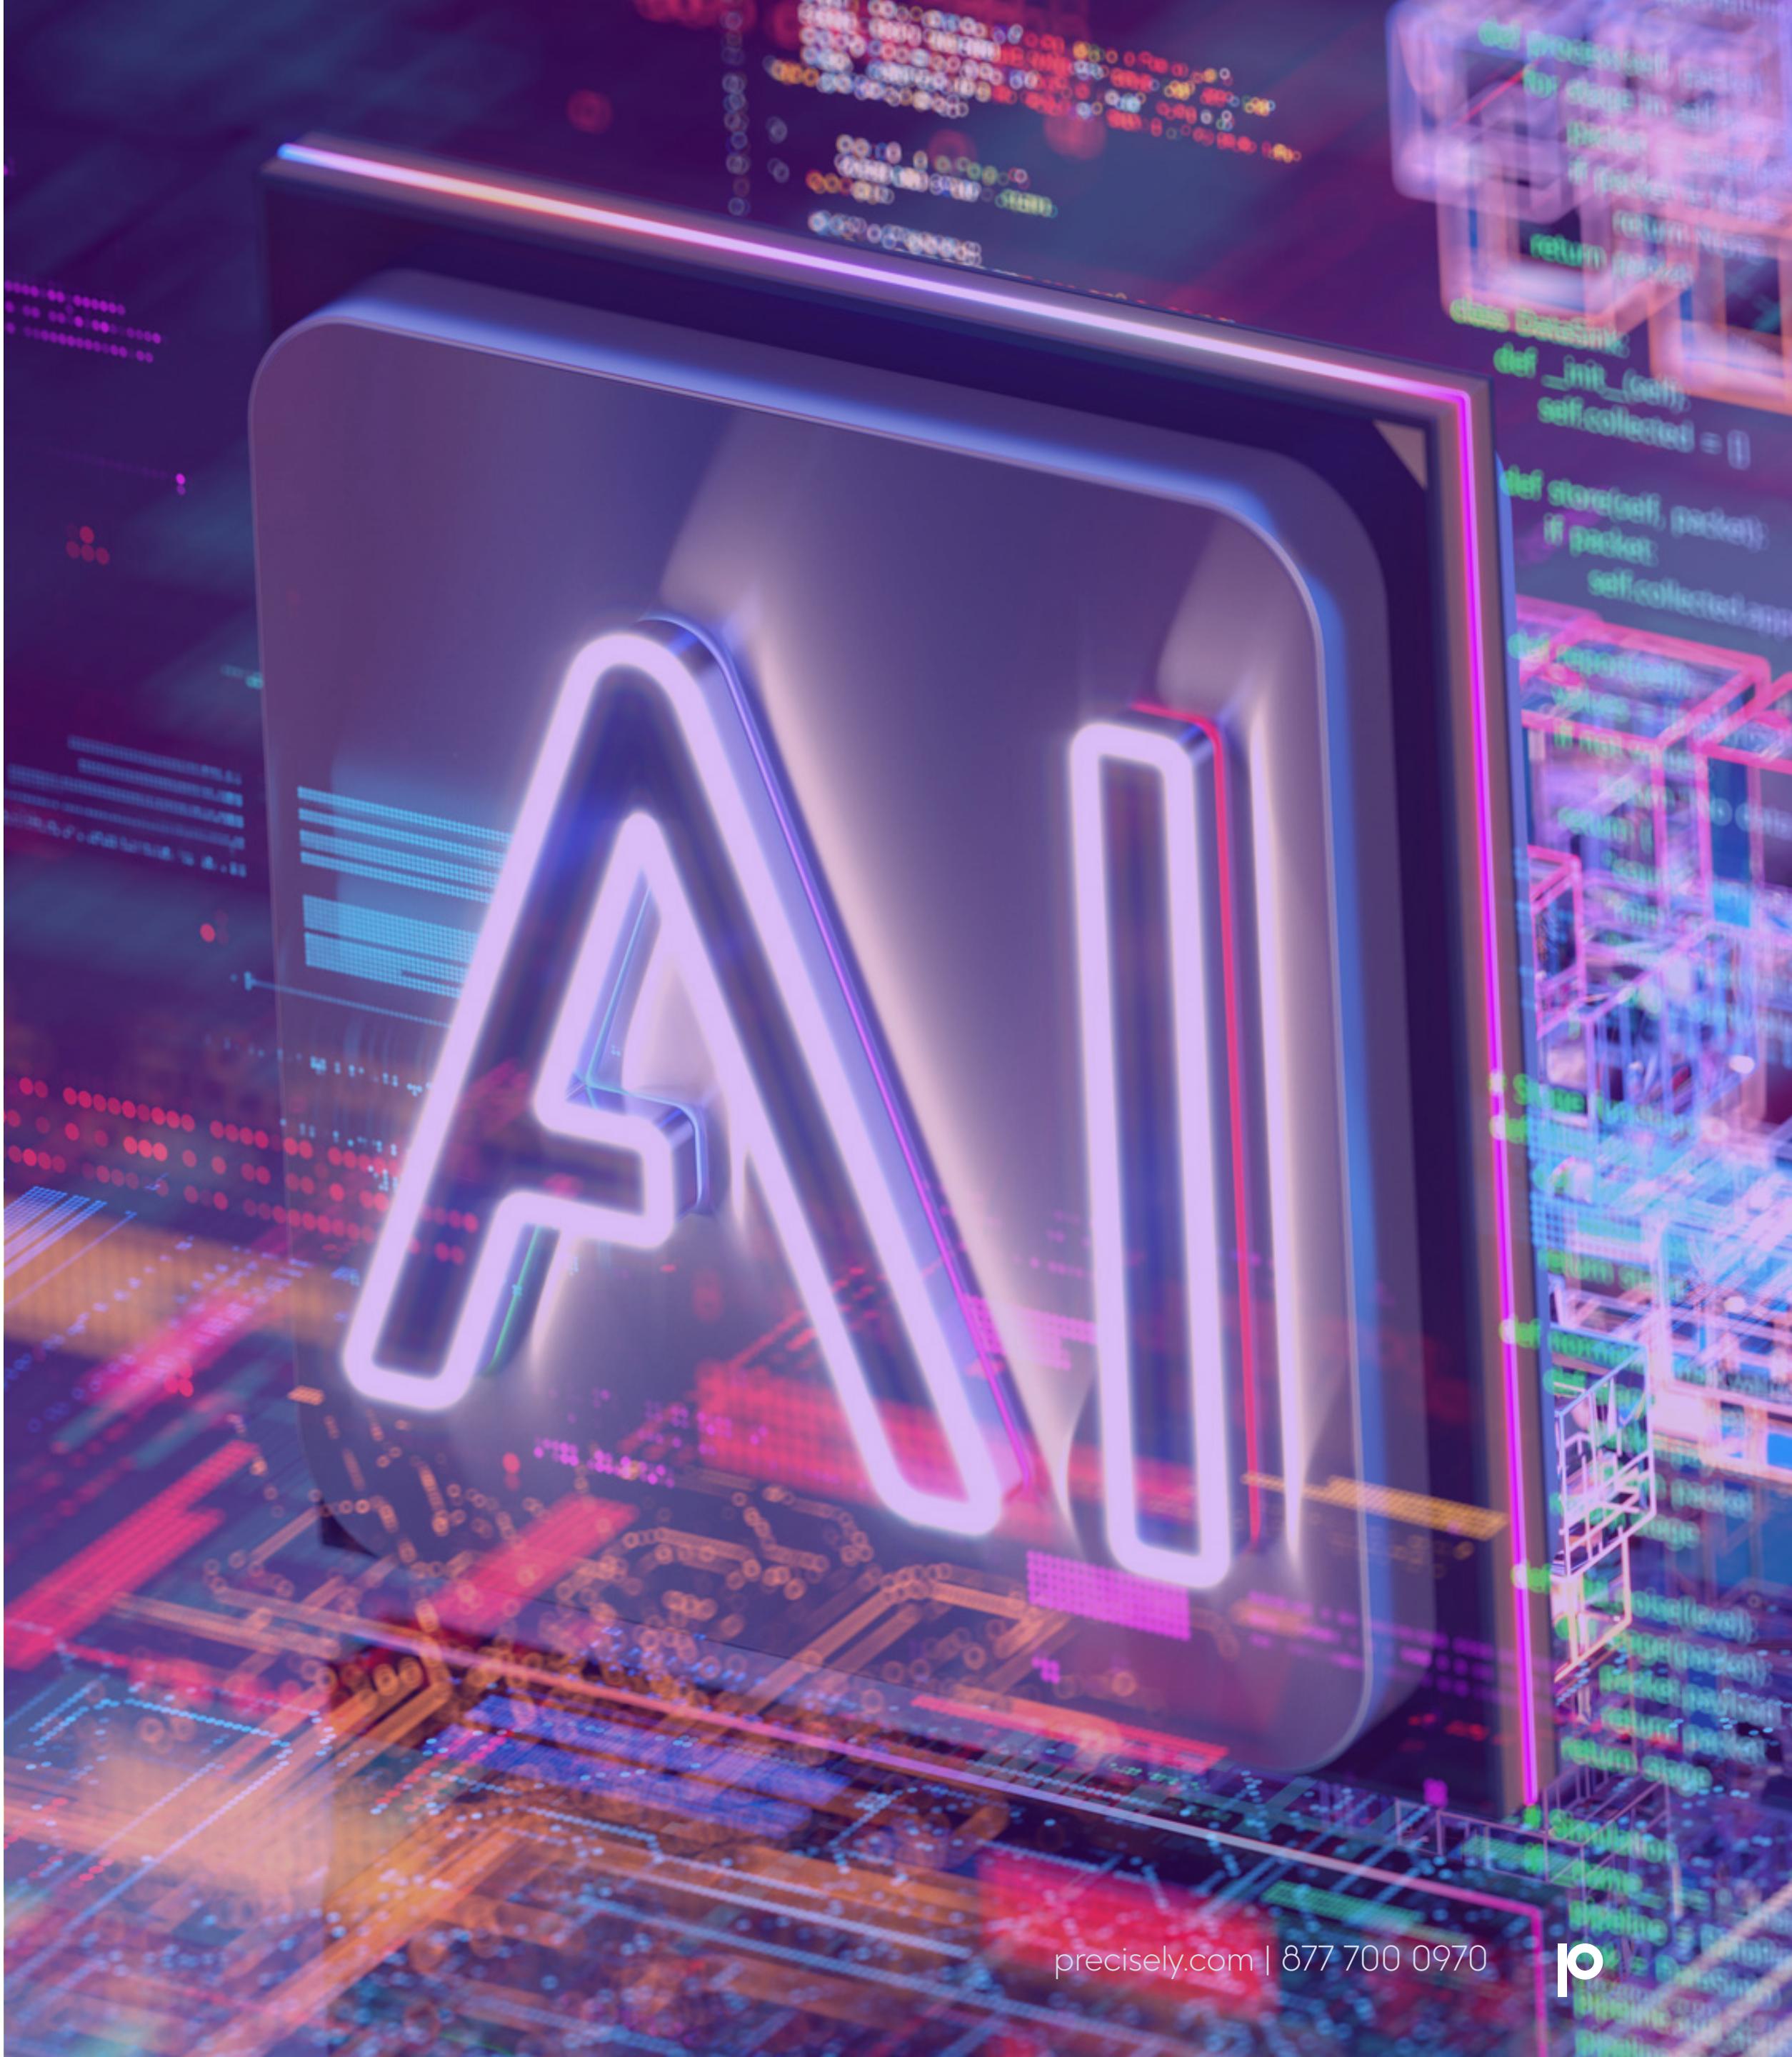
Seven Essential Principles for Success



Introduction

Artificial intelligence (AI) is reshaping data integration in ways we couldn't have imagined just a few years ago. The promise is compelling: faster insights, better decisions, and automated processes that free teams and individual contributors for higher-value work. Yet in the rush to embrace AI, organizations often skip critical groundwork and pay for it later with failed projects and wasted investments.

This guide brings together hard-won wisdom from data integration experts who've seen both successes and failures. Their insights distill into seven principles that consistently separate initiatives that deliver value from those that don't.



1. FOUNDATION FIRST:

Data Management is Non-Negotiable

Integration, quality, governance; the elements of foundational data management have always mattered, but generative AI (Gen AI) makes them absolutely critical. The problem is that most organizations don't realize their data infrastructure isn't ready for AI. They review their databases and systems, assume their data is fine (it usually isn't) realize something is wrong, then try to fix it without knowing where to start.

Consider real-world examples: two companies merge, each with different customer numbering systems; a manufacturing system updates part numbers from seven digits to ten; someone makes a mistake in the conversion; or a data migration goes wrong, creating duplicate records or missing information. It happens all the time.

For AI to work, your data needs three things:

- **Accuracy** (no duplicate records)
- **Completeness** (no gaps that create blind spots in decision-making)
- **Context** (the business meaning behind the numbers, so AI can make informed decisions rather than just pattern-match)

This is where data management becomes essential. Imagine you have a customer with both a home address and a PO Box in your system. Which one is correct? Proper data management helps establish the rules for these decisions, determining which sources are authoritative for which types of data and creating a consistent foundation that AI can actually rely on.

But having clean, trusted data is only the beginning. You also need systems in place to keep it that way, which underscores the critical importance of validation and governance.



2. QUALITY OVER SPEED:

Validation, Governance, and Strategic Thinking

High quality data doesn't happen by accident; it has to be built into every step of your data workflow. This means automated tests that catch problems before humans ever see them, validation rules that act as guardrails, and unit testing that verifies data integrity at each stage. Every data source needs validation, no exceptions.

The traditional approach to validation has made this relatively straightforward. For example, with batch processing, you could run comprehensive quality checks overnight before anyone uses the data. But real-time data changes the game entirely. Now you must check quality as data is in flight, validating that it fits your standards, matches expected patterns, and catches problems before they reach decision-makers or customers. This is where data observability becomes your early warning system, spotting data quality issues before they cascade into bad AI decisions.

Generative AI adds new wrinkles to governance challenges that have existed for years. Yes, you still need comprehensive data visibility, proper access controls, and continuous quality monitoring. But Gen AI introduces two specific risks that demand attention: bias (prejudices inherited from training data or baked-in design assumptions) and hallucination (the phenomenon in which AI generates completely wrong information with absolute confidence).

The biggest pitfall involves sacrificing the quality of your data for the appearance of speed. Organizations increase risk when they move too fast without proper documentation or skip stakeholder alignment just to get started. The tech world celebrates “move fast and break things,” but that philosophy crashes hard when applied to AI and data.

For instance, when organizations rush to implement gen AI with unstructured data, they often face data reconciliation challenges. It becomes difficult to confirm whether information was lost during transformation because traditional validation methods (like comparing row counts) don't work with chunked and tokenized data. Break things here, and you're making decisions on broken, bad information, often without realizing it until significant damage is done.

With quality and governance principles established, the next question becomes: how do you decide what level of investment each type of data actually needs?

3. BALANCE AND PRAGMATISM: Right-Sizing Your Approach

When you pair real-time data with AI, something powerful happens: you eliminate the gap between an event occurring and your ability to respond to it. AI can monitor dozens or hundreds of data streams simultaneously, watching for patterns that matter and alerting humans only when something requires attention.

But not all data needs the same treatment. Some information genuinely requires real-time processing because immediate action creates value (think fraud detection or supply chain disruptions). Other data works perfectly fine in batch or micro-batch cycles updated every few hours. The key is matching the processing approach to actual business needs.

Most successful organizations run a mixed model: real-time where it matters, batch where it works, and micro-batch for everything in between. Invest in real-time infrastructure when the value justifies it. Continue with batch processing when timing isn't critical.

Where should you focus your AI investment?
Three areas consistently deliver value:

- Sentiment analysis for making natural language actionable
- Automated pipeline building to free people from rote data movement
- Vibe coding that helps developers write code faster while maintaining validation guardrails

Understanding these principles of balance and pragmatism sets the stage for actual implementation.

The question becomes: where do you start?

4. START SMALL, SCALE SMART:

A Practical Implementation Framework

The most common failure pattern in AI initiatives is trying to transform everything at once, which only works if you're fully modernized, and that modernization is likely happening at the same time you're experimenting with AI. It's overwhelming, expensive, and rarely works. The successful approach is nearly opposite: start with one small, achievable project with clear business value. Make it low-risk but high-impact. Get it right. Celebrate the win. Use that success to build momentum and justify the next investment.

There are several factors that can help make your implementation successful.

1. Build a data catalog.

This means inventorying what data you have, where it lives, what it means, and how reliable it is.

2. Document everything.

This includes definitions, construction methods, and quality measures. It's not glamorous work, but it's essential. You can't build on a data foundation that you don't understand.

3. Build a semantic layer.

This translates raw data structures into business concepts. For example, instead of table_customer_rev_2023_q4, your semantic layer calls it "quarterly customer revenue." This bridges the gap between how databases store information and how humans think about business problems.

4. Test with 10 questions.

Questions should be specific, such as "Which products had the highest return rate in the northeast region last quarter?" Make sure your data model has the correct dimensions to answer them. Tune the system until those 10 questions work perfectly, then expand.

Look at industries getting this right. Software companies started small with AI-powered customer support, such as chatbots that handle common questions before escalating to humans. Airlines and retail organizations began to rely on real-time inventory and pricing insights. Insurance and financial services firms picked focused pilots despite facing massive complexity, heavy regulation, and high compliance stakes. Notice the pattern: nobody successfully transformed everything at once. They all started small, proved value, and then expanded.

But even the best implementation framework isn't enough if you're trying to master everything on your own. The pace of change in AI demands a different approach to building capability.

5. PARTNERSHIP IS STRATEGIC: You Can't Go It Alone

The AI landscape is changing so quickly that by the time you master what exists today, it's already shifting beneath your feet. New models emerge monthly. Best practices evolve weekly. Yesterday's cutting-edge approach becomes today's baseline. It's not just difficult, it's impossible.

This makes partnership with reliable vendors essential, not optional. You need allies who are specialists in different aspects of your technology stack. When considering vendors, ensure they have integration expertise in cloud and legacy technologies and can provide comprehensive governance offerings.

Partnership mitigates risk. Partners bring technical expertise in rapidly evolving technologies. Working with proven partners increases confidence with business leaders and provides credibility when requesting a budget. It accelerates your team's learning curve through knowledge transfer and dramatically reduces the likelihood of expensive mistakes.

With the right partnerships in place, you have access to technical expertise and proven approaches. But technology and expertise alone won't transform your organization. Success requires attention to dimensions that have nothing to do with algorithms or infrastructure.



6. HOLISTIC APPROACH:

People, Process, Culture, and Tools

Here's what organizations consistently get wrong: they treat AI as a technology problem. They buy tools, hire data scientists, and wonder why nothing changes. The reality is that technology is only one piece of the puzzle. Success requires getting three other things right: your people, your processes, and your organizational culture.

Consider what's been happening over the past 10 to 15 years.

The people who work with data have become steadily less technical. Marketing managers query databases. Sales leaders build their own reports. Product managers explore customer data directly. Everyone wants access to information for ad hoc questions, not just static dashboards updated weekly.

This democratization of data is accelerating, and AI is the catalyst.

Natural language interfaces and AI agents are eliminating the learning curve. Instead of learning SQL or mastering a business intelligence tool, you'll simply ask questions in plain English. "Which products had the highest return rate in the northeast last quarter?" The system figures out how to access the data, aggregate it correctly, and present the answer. Technical users experience an even more dramatic shift. In the AI-enabled future, agents handle the initial diagnosis and resolution, bringing in humans only when needed.

This pattern transforms integration operations fundamentally. Nobody wants to spend time debugging data pipelines or tweaking configurations. They want pipelines that are resilient, quick to set up, and self-healing when something does go wrong. AI agents will learn from past performance, understand best practices for everyday issues, and improve through feedback loops. The system gets smarter over time instead of requiring constant human intervention.

Building this future requires intentional cultural work. Celebrate wins publicly; they prove the approach works and build momentum. Use successes to justify additional investment. Create a culture of continuous improvement in which teams expect systems to improve over time. Foster experimentation within appropriate guardrails. Let people try new approaches, but with safety nets that prevent catastrophic failures.

The core promise underlying all of this is simple and powerful: jobs should get easier for everyone involved in integration and data operations over the next few years. Not eliminated - easier. Technical staff do more valuable work. Business users get answers without having to go through technical gatekeepers. Organizations make decisions based on current information instead of stale reports.

With these organizational elements in place, one final principle ties everything together: a way of thinking that should precede every AI initiative you undertake.

7. THINK BEFORE IMPLEMENTING: Ask the Critical Question

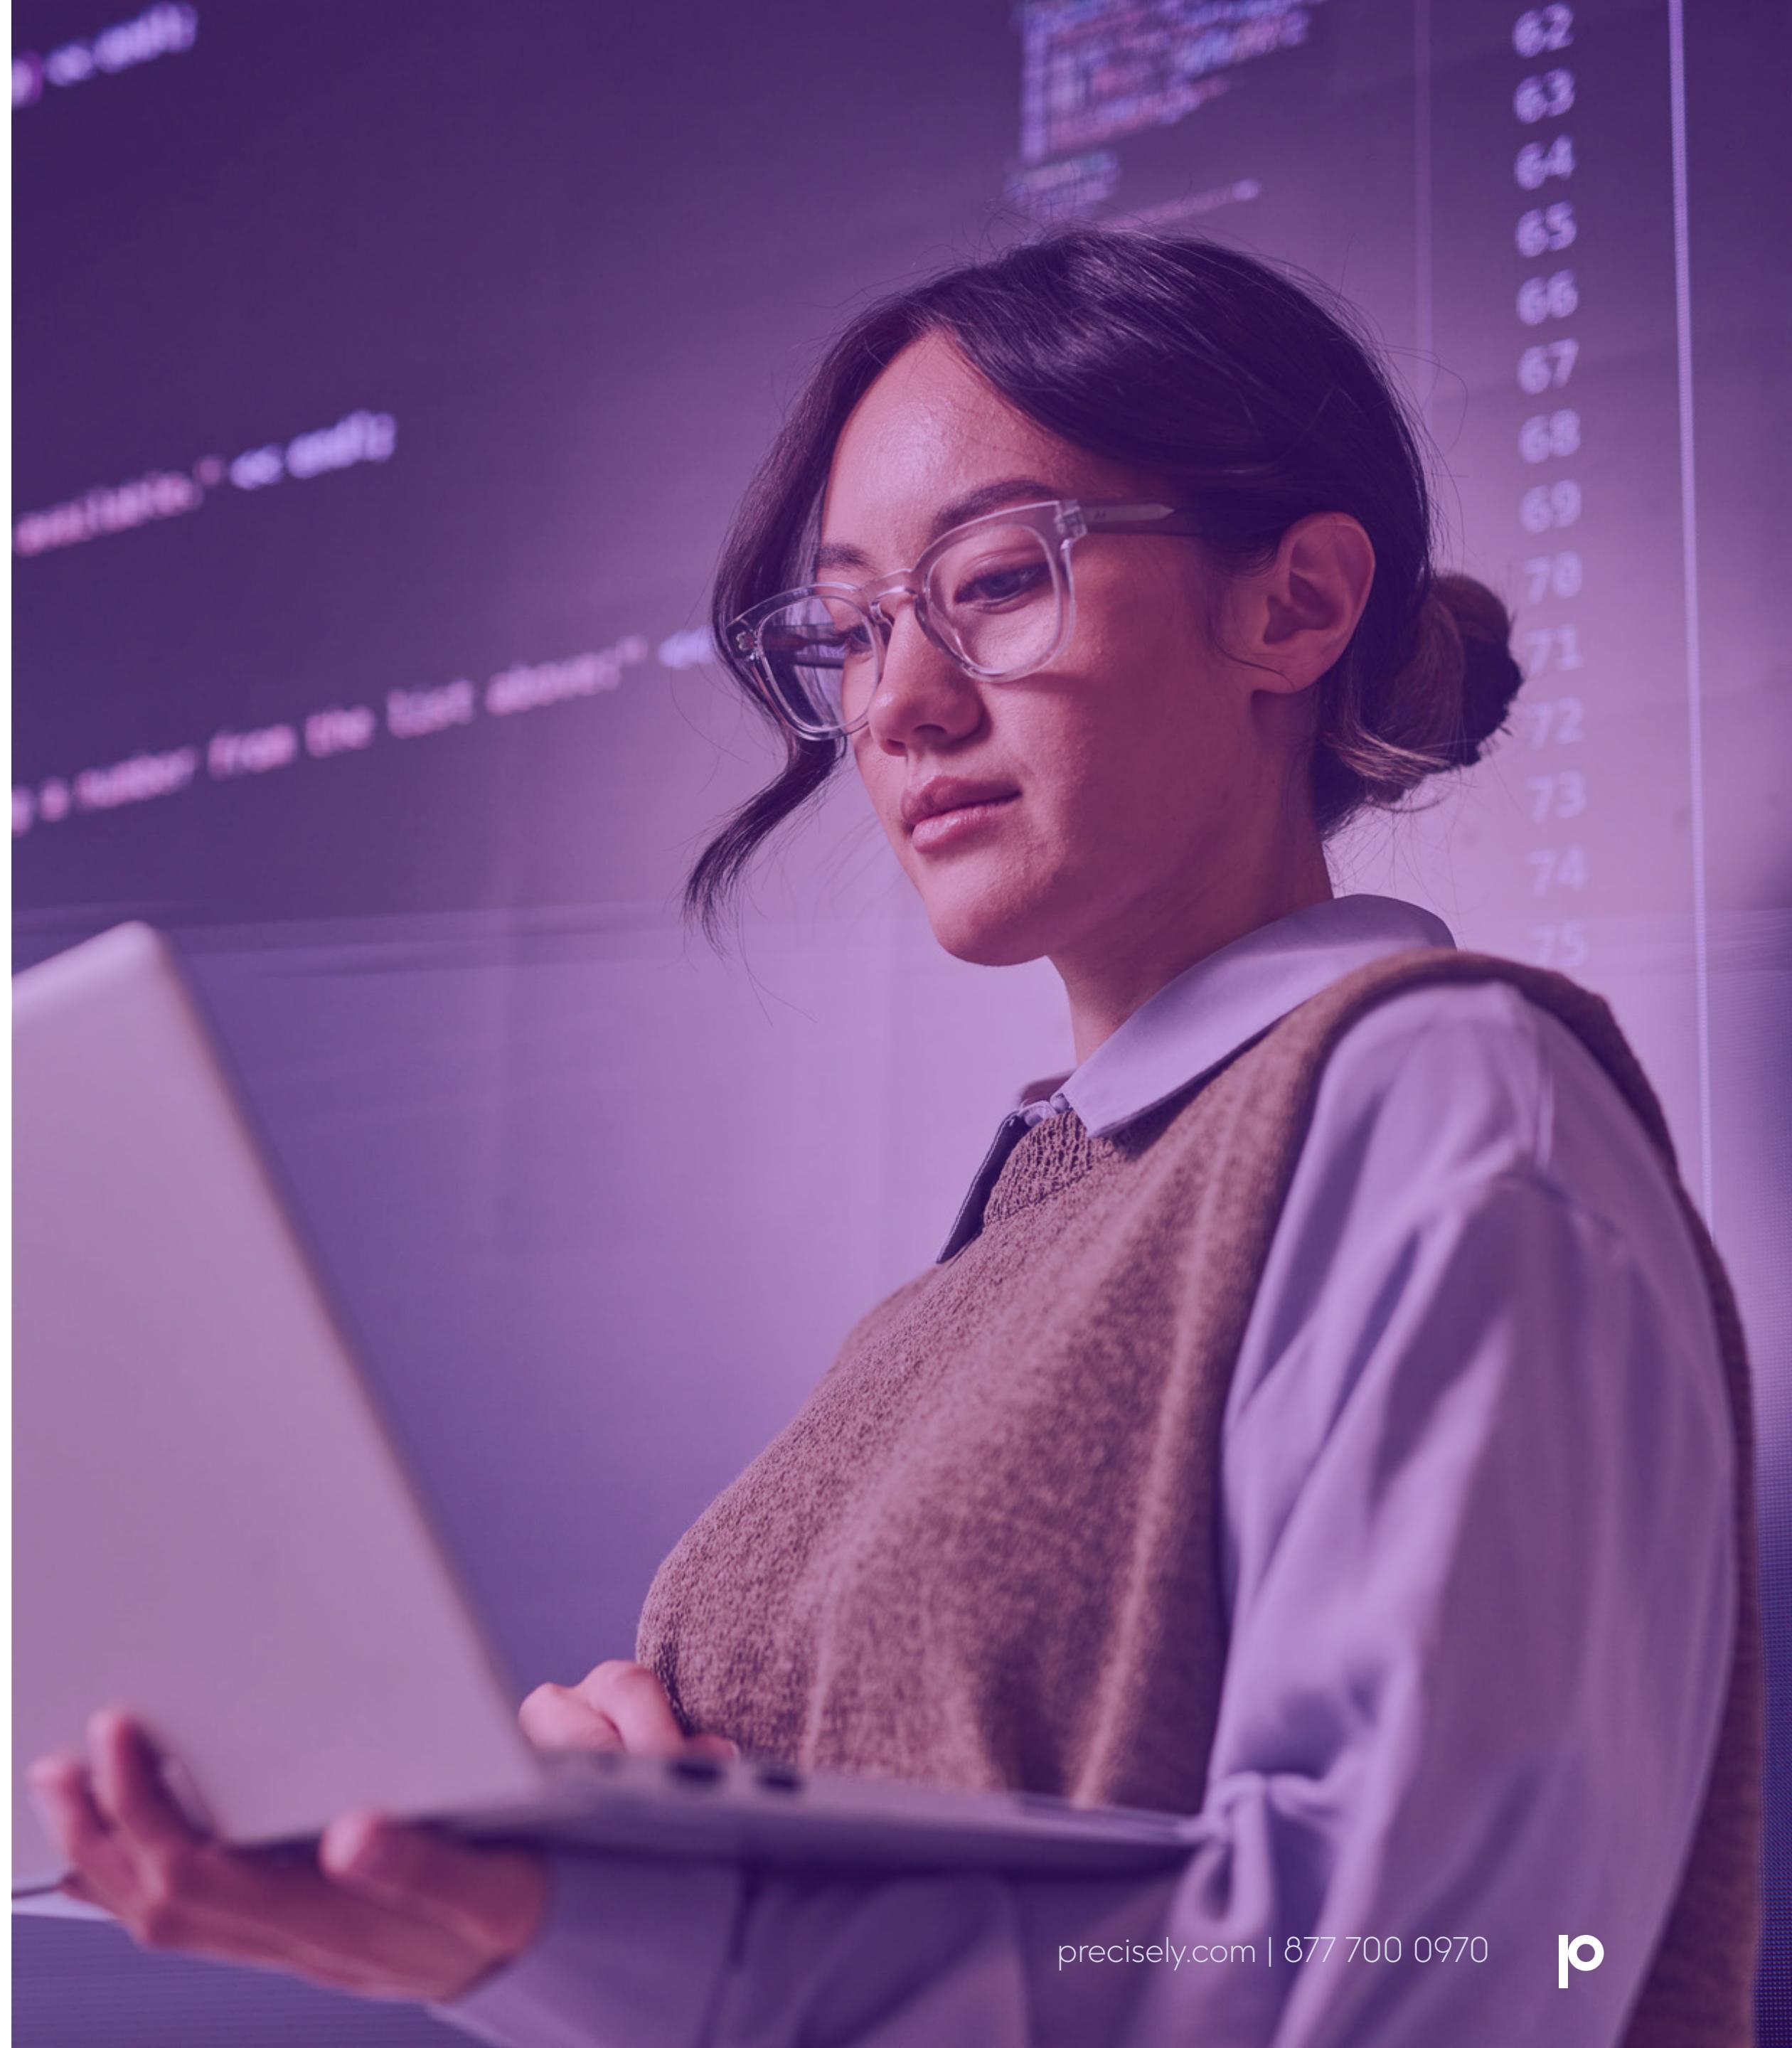
Before implementing any AI solution, ask yourself one essential question: How would I do this without AI? If you can't describe the manual process with specific signals and clear logic, AI won't magically solve it for you.

Walk through the problem as if you were solving it yourself. What signals would you look for? Document the decision-making logic step-by-step. Identify edge cases and exceptions that always come up. Recognize where human judgment is genuinely essential versus where it's just habit.

AI excels at specific tasks: automating repetitive tasks you understand, monitoring conditions you define, scaling approaches that already work manually, and providing consistent execution. Its superpower is handling volume that would overwhelm humans. This frees humans for higher-value work that requires creativity, nuance, or complex judgment.

But AI has clear limits. It can't define unclear business requirements, create a strategy from an ambiguous goal, replace understanding of your business processes, or compensate for poor data you wouldn't trust if you were making decisions manually.

If you can't walk through the manual process, you're not ready for AI. Close that gap first, and AI becomes a powerful accelerator. Skip that step, and AI becomes an expensive way to automate confusion at scale.



Conclusion

The gap between AI's promise and its actual delivery comes down to fundamentals. These seven principles aren't theoretical ideals; they're patterns observed across hundreds of implementations, distilled from both successes that transformed organizations and failures that wasted millions. The organizations thriving with AI didn't get there by implementing everything at once or treating it purely as a technology problem. They started small, proved value, and expanded methodically. They invested in data quality before algorithms, brought partners in early, and recognized that culture, process, and people matter as much as technology itself. Most importantly, they asked the hard questions before implementation, not after deployment.

AI will continue evolving at a breathtaking pace, but these seven principles—integration, quality, pragmatism, strategic implementation, partnership, holistic thinking, and thoughtful preparation will remain relevant because they address the human and organizational challenges that don't change with technology cycles. The real question isn't whether AI will transform data integration. It already is. The question is whether your organization will capture that value or squander it by skipping the foundational work that makes transformation possible.





About Precisely

As a global leader in data integrity, Precisely ensures that your data is accurate, consistent, and contextual. Our portfolio, including the Precisely Data Integrity Suite, helps integrate your data, improve data quality, govern data usage, geocode and analyze location data, and enrich it with complementary datasets for confident business decisions. Over 12,000 organizations in more than 100 countries, including 95 of the Fortune 100, trust Precisely software, data, and strategy services to power AI, automation, and analytics initiatives. Learn more at www.precisely.com.

www.precisely.com