



Finding the Right AI Solution: A Guide for IT Management



Contents

Key elements to operationalize successful AI applications at scale.....	2
Fostering confidence in AI - address trust, transparency and interpretability head-on.....	5
Plan for the process implications of AI implementation.....	8
Align open source and commercial assets into a cohesive AI infrastructure.....	9
Conclusion.....	11

A growing trend in organizations across all industries is the urgent call for IT management to implement artificial intelligence (AI) capabilities as a means to solve key business issues – and deliver measurable value from them in record time. Sometimes this stems from the lines of business having bought into the starry-eyed promise of AI, and they are anxious to put these capabilities to work in meeting their goals. In other cases, AI implementation may start in pilot projects scattered throughout the organization, and the business stakeholders may not have a full view of all the possibilities with AI, how much it will cost, who is responsible for its implementation and maintenance, and the amount of time required to deliver real business value. In all cases, the interest in adopting AI has grown, and there are real implications for IT management that need to be closely examined.

Even amid this momentum, IT leaders may not be operating with a robust road map for successfully implementing, operating and managing AI, because AI applications often have only recently been fully developed for mainstream usage. [Recent research](#) shows that some IT and business leaders have made headway with real, practical progress in implementing AI and are starting to demonstrate value from AI implementation. We zeroed in on their experiences and recommendations to create a strategy for generating real value through AI innovation.

Through this initiative, we have tapped into the observations and insights of global technology executives at SAS, Cisco and Intel who have firsthand knowledge to share about enabling AI adoption. The goal is to share these insights so that you can be successful in finding and implementing AI solutions across your organization.



Key elements to operationalize successful AI applications at scale

Some difficulties are to be expected in making the transition from the proof of concept phase to broader deployment, or even discrete “live” AI projects. This is fundamentally part of the friction that may come with *operationalization* - of building on proven models and processes to facilitate decision making and provide tangible value. It is a predictable challenge given all the ways in which AI capabilities must be connected to data and systems throughout the organization, and all the work that can be required to ensure that the data is ready for use. There are some downsides to this challenge because the potential for AI has been so hyped that organizations will be looking for signs of momentum. Barring that, they may lose interest, and it may be more difficult to find support for other AI applications.

The upsides, however, seem to result from four key elements of operationalizing AI.



Selecting the right business challenges

Jennifer Nenadic, Director of Data and Analytics at SAS, says that based on her experience, “the most important ingredient in AI success is having a business problem to solve.” Nenadic has seen her share of AI projects undertaken primarily to experiment - AI for AI’s sake. “Those projects are going to either fail or just hit a dead end and be forgotten,” says Nenadic. She recommends starting AI projects by looking at the business plans. *Does our HR team need a way to halve the time required to sort through thousands of résumés? Is our customer service team routinely fielding calls from customers who have been misrouted?* When solving real problems like these is the starting point, it becomes easier to expand out from there because the parameters for the implementation become framed in terms of real-world issues, from infrastructure and security to system interoperability. Per Nenadic, “That’s why ‘proof of concept’ isn’t really in our vocabulary.”



Expanding access to AI tools

For some leaders, successful operationalization hinges on their ability to provide a wide range of teams and individuals with intuitive, ready-to-use tools for experimenting with - and ultimately implementing - AI on their own. For example, while the technology organization may be focusing on a handful of strategic initiatives, other groups may be able to sustain their own AI momentum, with some help getting started. “We’re trying to build the ‘long tail’ of AI within our organization,” says Itay Yogev, Director of Advanced Analytics at Intel. “We can’t afford to fund startups in every part of the business, but that doesn’t mean we can’t help them by giving them something - ideally, it would be within the context of an AI technology adoption and rollout strategy. For us, operationalizing includes saying, ‘We may not be able to go deep with you here, but we can give you some tools for elsewhere.’”

4 keys to operationalize AI

1. Select the right business challenges.
2. Expand access to AI tools.
3. Adopt ModelOps practices.
4. Choose the right platform.



Adopting ModelOps practices

Operationalizing AI can also be achieved by embracing ModelOps principles (based on longstanding DevOps principles), in which models are cycled from AI and data science teams to the IT production team in a regular cadence of deployment and updates. ModelOps allows IT leaders to move models from the lab to validation, testing and production as quickly as possible while also ensuring quality results. Leaders can manage and scale models to meet demand and continuously monitor them to spot and fix early signs of degradation.

ModelOps addresses several common obstacles to successful AI operationalization, including:

- **Data quality:** Subtle shifts or changes in data may go unnoticed or have a diminished effect in traditional analytics processes, but can have an exponentially larger impact on AI and machine learning model accuracy. ModelOps processes are designed to address these issues head-on.
- **Model creation:** Many organizations develop portfolios of AI models created by different teams for differing purposes. Such portfolios can function as a valuable marketplace of AI solutions for teams across the organization, which can repurpose the solutions to meet their own challenges and opportunities. This is a critical aspect of operationalization, allowing organizations to accelerate their efforts by building on the work of other teams while avoiding rework.
- **Time to deployment:** Because the model development/deployment cycle can be long, you first assess how long that cycle is for your organization, then set benchmarks to measure improvement. Break down your process into discrete steps, then measure and compare projects to identify best and worst practices. Also consider model management software that can help automate some activities.
- **Model monitoring:** Because AI capabilities are constantly changing and evolving based on new learning, constant monitoring is important to ensure against degradation or unintended outcomes. Analysts can then retrain the models based on insights from monitoring.
- **Degradation:** Drift and bias can affect even the most rigorously designed models. The answer to these problems is creating a strong approach to model stewardship in your organization. If everyone from model developers to business users takes ownership for the health of your models, these problems can be addressed before they affect your bottom line.
- **Governance:** AI engages a wide range of data sources and types, which can make it difficult and time-consuming to trace data lineage and preserve other "track-back" information. This is particularly important in the context of audit compliance. ModelOps approaches introduce proven governance techniques for managing this type of complexity.

Among technology leaders implementing AI strategies, ModelOps remains an under-utilized approach. But those who have used it have found it to be a critical tool for scaling their AI projects, which is the foundation of operationalization. Analyst firm Gartner recognized SAS® Viya® in its Magic Quadrant for Data Science and Machine Learning Platforms as offering one of the best model operationalization and management platforms. “It includes performance monitoring for model decay, automated retraining of models when thresholds are exceeded, governance via a centralized model repository with templates and version control, and lineage for both SAS and open source models.”



Choosing the right platform

The choice of platform plays an important role in operationalizing AI at scale, enabling self-sustained, ongoing AI service in production. There are a host of reasons that leveraging these platforms works to the advantage of IT leaders, including:

- Making it easy to quickly deploy models into production.
- Bringing relevant data to the AI algorithms - or vice versa with streaming and batch data.
- Enabling closed feedback loops at scale, in a timely manner.
- Sustaining model integrity over time through model management, retraining models, deploying logs and indicators, and contributing to stronger manageability.
- Enabling greater transparency and “explainability” - which is especially important in heavily regulated industries requiring documentation of decision-making inputs and rationale.
- Bolstering privacy and security across a wide range of in-house, commercial and open source solutions.

Depending on how mature your organization is with using analytics, the way you implement your AI platform can impact how effectively you can achieve your strategic goals. In line with what SAS’ Jennifer Nenadic recommends, the key is for each application to be built around a clear business case. “Tell me what you’re trying to do, what are the constraints and other variables describing the issue, what types of insights do you need, and what kind of data are you working with,” says Intel’s Chandhu Yalla. “Using that information, we can channel those use cases into the most effective AI application.”

Fostering confidence in AI – address trust, transparency and interpretability head-on

In an AI environment, the business lines are being asked to trust algorithms to perform important tasks for them, and regardless of the size or importance of the task, they initially may be hesitant to do so. The experience of early AI adopters has shown that the willingness to trust AI to take on tasks will likely vary depending on the type of work AI is doing. When AI is replacing automation of basic noncritical tasks, this is generally viewed as a welcome improvement – the work is not changing, only the technology being used to take on the work. But when the work itself has farther-reaching impacts – as is often the case with more advanced applications – trust, transparency and interpretability are the elephants in the room standing in the way of full AI adoption. In short, that’s when the difficult questions start.

“The majority of concerns we see are about the unsupervised learning aspect of AI,” says Ananth Sankaranarayanan, Director of Engineering in Intel’s Artificial Intelligence Products Group. “If the system has learned and is predicting some outcome, users want to know what processes the system went through in order to develop that intelligence.”

Compounding the challenges of gaining trust in AI, most of these systems rely heavily on cloud-based operations, which some people believe limits visibility for users, especially those accustomed to on-premises tools. “We often have users who say, ‘I want exactly the same visibility and control that we have on-prem,’” says Jon Heaton, IT architect for Cisco. “So we spend a lot of time working with those users to make sure they have what they need in order to feel comfortable that their information is secure.”

On some level, such skepticism is healthy – [recent research](#) has shown that AI systems require human oversight to perform at a high level. But if left unchecked, rampant skepticism of AI can also be crippling to the organization, which serves to undermine the significant investments your organization may have made in AI.

Here are five practical recommendations to cultivate more confidence in AI systems.



1. Make use of benchmarking

The best way to cultivate trust in AI systems is to prove that they are working as planned. When AI systems are proven to beat preexisting benchmarks using real business data, people begin to trust these systems. This can be easily accomplished in cases where AI is replacing existing automation capabilities because there are already benchmarks for performance and accuracy in place. This is more difficult when all-new tasks are being assigned to AI, where there are no benchmarks.

“This is a big issue for us,” says Intel’s Itay Yogev. “We find that people are very tolerant of their own mistakes compared to those that a machine might make, so if we create an imperfect algorithm that is making mistakes, people can have an emotional response to that. We have to remind our people frequently that we’re not assuming perfection is within reach, but instead that we can deploy systems that make fewer mistakes than humans. It’s all about managing expectations.”

5 ways to build confidence in AI

1. Make use of benchmarking.
2. Expand existing governance structures.
3. Focus on culture and organizational factors.
4. Enable practical collaboration.
5. Address IT talent and staffing implications.

In that regard, look to the metrics and standards already in use in your organization as the most effective benchmarks to work with because people already know them, and they provide the best link to organizational priorities.



2. Expand existing governance structures

In this environment, trust, transparency and interpretability can be built through formal, structural IT channels and facilitated by the use of an AI platform with these capabilities. AI is just one of the many endless processions of powerful new technologies - all of which are governed by existing IT policies, procedures and guidelines. It follows that governance tools will need to be updated in light of AI's advances. As is the case with benchmarking, the best starting point is to examine your existing IT policies, procedures and guidelines, update them accordingly, then verify their relevance to AI adoption.



3. Focus on culture and organizational factors

While some aspects of AI's impact have been overblown, [research has shown](#) its impact on your organization's culture will be real and immediate - at least if it's working properly. Because AI applications by definition involve adding learning to analytic processes, and then automating that learning, they often reach across traditional organizational boundaries and then can scale quickly. Those factors elevate the importance of cultivating transparency across the organization and enabling greater sharing between departments, business units and individuals. With AI deployment as a catalyst, AI has a unique and powerful role to play in helping multiple operating units align more tightly. The ability to foster greater alignment is an additional benefit of IT's broader mandate to drive innovation, given the increasing connection between technology and innovation in organizations today.



4. Enable practical collaboration

The ability for AI implementations to cross organizational boundaries logically suggests that having different teams working in isolation toward similar goals and potentially using similar AI technologies and processes may limit the impact of AI. When one group or team develops a successful AI capability, other teams elsewhere in the organization not only should be aware of it but also be able to borrow, modify and apply the capability to their own initiatives. At Intel, this type of collaboration is enabled at the AI team leader level. These leaders understand the importance of sharing and collaborating in an AI environment and meet regularly to discuss new developments. "We don't require these meetings; we just make sure our AI leaders have the ability to do so easily," says Intel's Itay Yogev. "We believe strongly in the free will of our teams and leaders - I don't want to trap them in any unnecessary processes. They're responsible. But we've found that given that freedom, our AI leaders gravitate to one another. They know they can't do it all. So when given the opportunity, they collaborate."



5. Address IT talent and staffing implications

IT staffing models and strategies will play a large role in the success of any AI initiative. Data scientists (typically professionals who have experience interpreting data, have robust coding skills, and have earned advanced degrees) are particularly important resources, and they are not only difficult to find and hire, but also have a high rate of turnover due to the strong demand for their skills. Given the difficulty of securing that type of talent, chief information officers (CIOs) need to identify additional talent profiles to supplement scarce data scientists.

Intel, for example, has identified what it calls “modern data analysts” as a key pillar of its AI talent strategy. According to Bradley Schwartz, Senior Director of Enterprise Analytics at Intel, “these are people who have statistics and data modeling skills, along with a passion for learning new technologies and techniques.” By broadening the definition of desirable AI talent to include modern data analysts, Intel can draw from a larger talent pool. “This has taken us from targeting hundreds of data scientists to thousands of analysts,” says Schwartz. “From there, the question becomes, ‘How can we best enable this valuable talent?’”

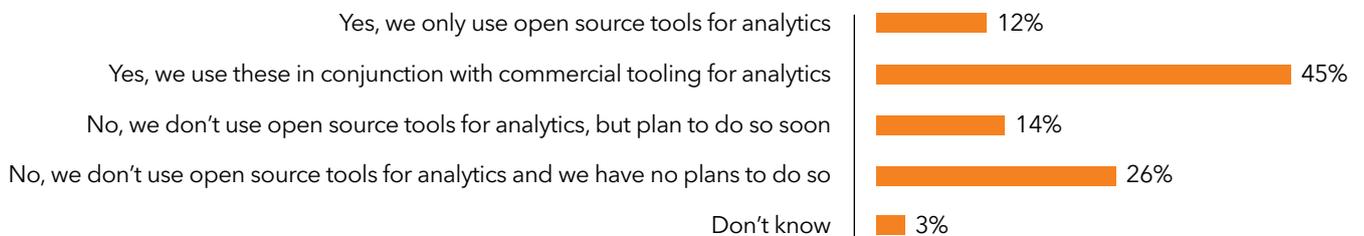
Other organizations turn to established AI and analytics platforms, such as [SAS® Viya®](#), or leverage opportunities such as the [SAS Academy for Data Science](#) to bolster the skills of existing team members.



Plan for the process implications of AI implementation

If you are concerned about the process implications of AI implementation, that may be because you still have the scars to show from recent large-scale ERP or CRM implementations. Such sweeping implementations likely required business process reengineering before the technology implementation began and have understandably made CIOs wary of process-intensive technology implementations.

Do you use open source tools for analytics?



Source: TDWI Pulse Report: Using a Hybrid Open Source and Commercial Analytics Ecosystem

While AI implementations certainly have process implications, in those cases where they are initially rolled out as a limited-scope pilot, disruptions to operations can be similarly limited. "We've found that yes, there are some amount of business process improvements required," says Intel's Chandhu Yalla, "but there are a lot of drag-and-drop applications, lots of customer modifications in these technologies that allow you to mold the AI systems to the processes, rather than the other way around." By the same token, that does not negate the need to perform a thorough analysis of the business processes affected by the introduction of AI - because AI is a disruptive technology, and it's reasonable to expect ripple effects in processes.

Align open source and commercial assets into a cohesive AI infrastructure

Today, your organization is likely relying on a mix of commercially available solutions, homegrown technology tools, and open source capabilities to support a wide range of business and technology initiatives. That's just a fact of today's advanced technology organizations - one that is reflected in a recent [TDWI Pulse Report: Using a Hybrid Open Source and Commercial Analytics Ecosystem](#). In the survey, only 26% indicated that they did not use open source tools for analytics and had no plans to do so. Meanwhile, 45% indicated they are already using open source tools in conjunction with commercial tools for analytics. Your AI plans may deepen your reliance on open source tools - and there's nothing wrong with that. They offer a level of flexibility that can be instrumental in executing your AI strategies. At the same time, open source capabilities open the door to choices on everything from strategy to implementation. As the TDWI report says, "Organizations using open source for operationalizing models will often need to piece together tools to do so, whereas commercial products may provide unified functionality" - a solid rationale for deploying both commercial and open source solutions together. Here are some important considerations for getting started with open source models in the context of AI strategy.

Right-sizing your open source model

Where does open source fit in your AI strategy? Today, the question is less "Should we be using open source capabilities?" and more "Which open source models might come into play with this initiative?" "It generally makes sense to find capabilities that may not be as readily available commercially," says Intel's Itay Yogev. "If it's a mission-critical initiative, we'll usually go with a hybrid model - we'll go with open source tools but also find commercial offerings that may offer greater support because if something goes wrong, we know they'll be there for us."

Securing support commitments from all sources at the outset

Yogev's colleague Chandhu Yalla, Head of Big Data and AI for Intel, emphasizes the importance of considering potential support resources. "Doing that will give you more options to resolve potential issues if your environment has both commercial and open source solutions deployed," says Yalla. "And, of course, the level of support varies among open source and commercial providers - do your homework."

Managing costs

Especially for organizations managing multiple AI initiatives at the same time, the total cost of ownership (TCO) may be the most effective tool for deciding how to budget for support. "We're constantly looking for ways to optimize our TCO," says Yalla. With a cohesive AI infrastructure of both open source softwares and commercial platforms, such as SAS Viya, you can unify disparate toolsets and analytic assets into a streamlined, collaborative environment that fosters productivity, business agility and tangible results. Further validating this point, a recent IDC MarketScape Worldwide General-Purpose Artificial Intelligence Software Platforms Vendor Assessment reported, "46% of the respondents using SAS Viya indicated that they chose the platform due to lower cost."

Leaning on enterprise-level infrastructure

Every IT environment today is dynamic, constantly shifting as new capabilities emerge and are adopted. For CIOs, the practical impact is immense. The organization may be moving from on-premises to cloud architectures, for example, or to a hybrid environment. Some CIOs will need to manage their infrastructure, while others may outsource it. In this context, advancing AI often serves to increase the level of complexity, given the dynamic nature of AI applications and their ability to scale quickly. Just as importantly, AI capabilities rely on an ecosystem of other technology assets and data sources, underscoring the practical advantages of implementing capabilities within a single, flexible analytics platform, such as SAS Viya.

Another key element of that enabling technology ecosystem includes reliable hardware supporting an operating environment that meshes cloud-deployed software and data with traditional data center-based operations. A good example of that is the [Cisco Unified Computer System™](#), which integrates high-velocity delivery of information in real time to boost operating efficiencies and scalability while reducing data center costs.

A third critical area to consider is the processing speeds and capacities ingrained in the hardware and the software so that it can keep up with the dynamic nature of AI algorithms and the likelihood of quickly scaling. SAS continues to embed powerful AI capabilities into Viya, while Intel is driving major advances in the processing capabilities required to make sense of all that valuable data. The result? Integrated processing power and world-class artificial intelligence capabilities for smarter organizations and stronger outcomes.

Conclusion

One constant in AI applications today is that they are underpinned by core capabilities, such as machine learning, deep learning, computer vision, natural language processing, and forecasting and optimization. The power of these technologies comes from their ability to uncover patterns, get better at doing so, and then automating those learning processes accurately, repeatedly and reliably.

A CIO's role was once narrowly focused on operational IT responsibilities, and more recently those responsibilities frequently include a mandate to support innovation across the organization. And since innovation goals often are strongly tied to business goals, the upshot for IT management is increased pressure to move the organization more quickly from analysis to action. It's within this dynamic context that AI applications continue to provide some of the brightest opportunities for IT management to foster innovation.

As these IT leaders at SAS, Cisco and Intel attest, real success with AI increasingly involves a multipronged approach across four dimensions that involves operationalizing successful AI applications at scale, fostering confidence in AI, planning carefully for the process implications of AI implementation, and successfully aligning open source and commercial assets into a cohesive AI infrastructure. Together, SAS, Cisco and Intel have the right combinations of technology and support to equip IT management to succeed across all four of those dimensions.

Learn more about how partnering with each of these technology leaders can become one of the most important steps to take for succeeding with AI.

Find more details at:

sas.com/AI-CIO

sas.com/cisco

sas.com/intel

