CompTIA

ARTIFICIAL INTELLIGENCE IN BUSINESS:

Top Considerations Before Implementing AI

Business leaders and AI practitioners must ask the right questions before embarking on an AI project.

Artificial Intelligence (AI) is poised to transform many industries in the coming years, including augmenting human intelligence, powering automation, enabling optimization, offering decision support, paving the way for hyper-personalization and making possible natural interfaces to many business applications. Companies are actively exploring, experimenting and deploying AI-infused solutions in their business processes. Chatbots in customer support scenarios, doctors' assistants in hospitals, legal research assistants in the legal domain, marketing manager assistants in marketing, and face detection applications in the security domain are some early use cases of AI in enterprise.

TOP AI SOLUTIONS

The CompTIA AI Advisory Council identifies the solutions that offer the most opportunity for achieving success in the artificial intelligence market.

Many things must come together to build and manage AI-infused applications. Data scientists who build machine learning models need infrastructure, training data, model lifecycle management tools and frameworks, libraries, and visualizations. Similarly, an IT administrator who manages the AI-infused applications in production needs tools to ensure that models are accurate, robust, fair, transparent, explainable, continuously and consistently learning, and auditable. This requires new tools, platforms, training and even new job roles. AI-infused applications should be consumable in the cloud (public or private) or within your existing datacenter or in a hybrid landscape. All this can be overwhelming for companies trying to deploy AI-infused applications.

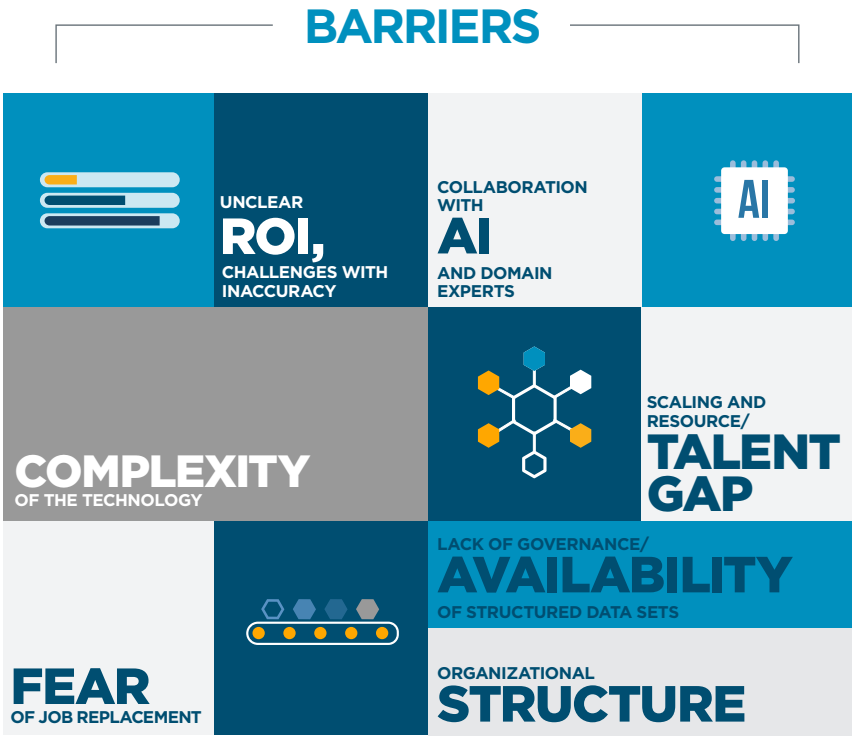

To help determine if AI is the right choice for your company and your business problem, CompTIA's Artificial Intelligence Advisory Council developed questions and answers that business decision-makers and AI practitioners should consider before investing in AI. While both decision-makers and practitioners have their own points to consider, it's recommended that they work in tandem to make the best, most appropriate decision for their respective environments.

Questions for Business Decision Makers

1. Do we have executive sponsorship to infuse AI within existing business processes?
2. Have we clearly defined the business objectives and outcomes to be achieved using AI?
3. Have we set the right initial expectations about the potential benefits of AI?
4. Do we have the budget required to support near-term and long-term AI objectives?
5. Do we understand the timeline needed to successfully deploy an AI project within the organization?
6. Can we manage market or competitive pressures to accelerate AI infusion within the organization?
7. Do we have the required skillset/domain expertise within the organization to execute on an AI vision?
8. Do we have the right IT infrastructure needed to deploy my AI solution? Do I know what questions to ask vendors when evaluating solutions?
9. Can we trust AI and make the business decisions based on the model predictions? How do we course correct if an AI model or prediction goes wrong?
10. Can we market our value proposition or differentiate our organization from competition using AI-infused solutions?
11. Can the organization reconcile the impact of AI on staff resourcing and culture?

Questions for AI Practitioners

1. What is the goal and purpose of the AI model I'm asked to build? Do I have representative training and test data for building this AI model? If not, is there is a data manager I can work with to acquire or prepare the data?
2. Is my data architecture adequate for leveraging AI?
3. How can I ensure biased data won't skew results?
4. Will robotic process automation, or a cheaper, non-AI process deliver the same outcome?
5. What AI or machine learning (ML) methods will be relied upon and why?
6. How long will implementation take?
7. Do I understand the legal, privacy, compliance, security implications of building AI solutions at this company?
8. Will the solution offer explainable results and transparency into the decision-making process?
9. Will the solution provide an adequate level of accuracy, given the application?
10. How can I make AI models more trustworthy for business decision makers? What are the tradeoffs to consider when building more transparent and explainable AI models?
11. What resources already exist instead of starting from scratch? What's the best tool available for automating the modeling that we need to do?
12. What kind of expenses related to infrastructure, data acquisition, crowd/SME annotation and others should my organization think about to implement AI best practices?
13. How will the AI function when it encounters a previously unseen situation or data point?
14. What governance policies and controls have I prepared for preventing and dealing with "AI breaches"?
15. How will I protect my models and data from adversarial attacks and malicious input?

Considerations Before Implementing AI: Questions for Business Decision Makers

When determining whether your company should implement an artificial intelligence (AI) project, decision makers within an organization will need to factor in a number of considerations. Use the questions below to get the process started and help determine if AI is right for your organization right now.

Artificial Intelligence Project Planning

Do I have executive sponsorship to infuse AI within existing business processes?

Artificial intelligence (AI), when implemented properly, has the potential to optimize and automate business processes and provide decision support to humans. AI might even lead to new business opportunities via new business models. To experience the true benefits, organizations will need to be flexible and adapt. Executive-level sponsorship is critical before embarking on a new AI project. Executive buy-in must include financial and human resources support as well as cultural changes to business processes to scale your AI project from pilot to production.

Have I clearly defined the business objectives and outcomes to be achieved using AI?

Defining milestones for an AI project upfront will help you determine the level of completion or maturity in your AI implementation journey. The milestones should be in line with the expected return on investment and business outcomes.

For example, companies may choose to start with using AI as a chatbot application answering frequently asked customer support questions. In this case, the initial objective for the AI-powered chatbot could be to improve the productivity of customer support agents by freeing up their time to answer complex questions. A milestone would be a checkpoint at the end of a proof-of-concept (PoC) period to measure how many questions the chatbot is able to answer accurately in that timeframe. Once the quality of AI is established, it can be expanded to other use cases.

Have I set the right initial expectations about the potential benefits of AI?

AI involves multiple tools and techniques to leverage underlying data and make predictions. Many AI models are statistical in nature and may not be 100% accurate in their predictions. Business stakeholders must be prepared to accept a range of outcomes (say 60%-99% accuracy) while the models learn and improve. It is critical to set expectations early on about what is achievable and the journey to improvements to avoid surprises and disappointments.

Do I have the budget required to support near-term and long-term AI objectives?

Embarking on a new AI project or expanding from an initial pilot phase will require a separate budget to cover new tools and technologies as well as expertise that may not be present within your organization. Consider these tips when determining the scope of your budget:

- Review your projected business value from the AI project over a 12- to 36-month period.
- Consider the total budget including internal headcount, contract resources, and IT infrastructure (including application and cloud licenses) to calculate your total investment on an AI project.
- Review cost of data acquisition (sourced internally or externally) as part of the total budget.

Do I understand the timeline needed to successfully deploy an AI project within my organization?

AI projects typically take anywhere from three to 36 months depending on the scope and complexity of the use case. Often, business decision makers underestimate the time it takes to do "data prep" before a data science engineer or analyst can build an AI algorithm. There are certain open source tools and libraries as well as machine learning automation software that can help accelerate this cycle.

Once you have identified a project or a business challenge, you can begin planning for a proof of concept (PoC), which will include data sources, technology platforms, tools and libraries to train the AI models leading to predictions and business outcomes. Depending on the use case and data available, it may take multiple iterations to achieve the levels of accuracy desired to deploy AI models in production. However, that should not deter companies from deploying AI models in an incremental manner. Error analysis, user feedback incorporation, continuous learning/training should be integral parts of AI model lifecycle management.

Lastly, nearly 80% of the AI projects typically don't scale beyond a PoC or lab environment. Businesses often face challenges in standardizing model building, training, deployment and monitoring processes. You will need to leverage industry tools that can help operationalize your AI process—known as ML Ops in the industry.

Can I manage market or competitive pressures to accelerate AI infusion within my organization?

Despite the hype, in McKinsey's Global State of AI report, just 16% of respondents say their companies have taken deep learning beyond the piloting stage. While many enterprises are at some level of AI experimentation—including your competition—do not be compelled to race to the finish line. Every organization's needs and rationale for deploying AI will vary depending on factors such as fit, stakeholder engagement, budget, expertise, data available, technology involved, timeline, etc. No two projects between organizations will look similar.

IT Infrastructure and Data Management

Can I access the data that exists within my organization to meet my project goals?

Data often resides in multiple silos within an organization in multiple structured (i.e., sales, CRM, ERP, HRM, marketing, finance, etc.) or unstructured (i.e., email, text messages, voice messages, videos, etc.) platforms. Depending on the size and scope of your project, you may need to access multiple data sources simultaneously within the organization while taking data governance and data privacy into consideration. Additionally, you may need to tap into new, external data sources (such as data in the public domain). Expanding your data universe and making it accessible to your practitioners will be key in building robust artificial intelligence (AI) models.

Have I considered data governance issues as related to compliance and privacy?

Data preparation for training AI takes the most amount of time in any AI solution development. This can account for up to 80% of the time spent from start to deploy to production. Data is the longest pole in the tent. Data in companies tends to be available in organization silos, with many privacy and governance controls. Some data maybe subject to legal and regulatory controls such as GDPR or HIPAA compliance. Having a solid strategy and plan for collecting, organizing, analyzing, governing and leveraging data must be a top priority.

Large organizations may have a centralized data or analytics group, but an important activity is to map out the data ownership by organizational groups. There are new roles and titles such as data steward that help organizations understand the governance and discipline required to enable a data-driven culture. Additionally, as AI taps into multiple new and old data sources, you will need to ensure that privacy and compliance guidelines are adhered to during the model development and training or inference process (i.e., facial recognition data that can be tracked to an individual in retail or banking or patient data protected by HIPAA).

Do we have the required skillset/domain expertise within the organization to execute on an AI vision?

Infusing AI into business processes requires roles such as data engineers, data scientists, and machine learning engineers, among others. Organizations should consider their current team and then determine a people strategy, which could include reusing or repurposing existing resources, upskilling and training current staff, hiring, or working with outside consultants or contractors. Some organizations might need to contract with a third-party IT service partner to provide supplementary, needed IT skills to model data or implement the software.

As the organization matures, there are several new roles to be considered in a data-driven culture. Depending on the size of the organization and its needs new groups may need to be formed to enable the data-driven culture. Examples include an AI center of excellence or a cross-functional automation team. Take into consideration the end-to-end requirements during your planning phase as getting the right skillset—whether it is building your own or utilizing outside expertise like consultants—will take time and impact your project delivery timelines.

Do I have the right IT infrastructure needed to deploy my AI solution? Do I know what questions to ask vendors when evaluating solutions?

Myth: I can run my AI project with open source tools and in the cloud today.

Reality: This is partially true only if you are running a small proof of concept (PoC) for kickstarting your AI project. AI projects by definition are complex; however, their implementation doesn't have to be. There are industry vendors available today to assist you in your AI journey. However, you need to be cautious in your selection of the right vendor(s). There are hundreds of startups and emerging vendors that may not have the resources or investment capital to sustain in the long run. Keep in mind that AI solutions are not cookie cutter. While most solutions available today may meet 80% of your requirements, you will still need to work on customizing the remaining 20%.

Begin by researching use cases and white papers available in the public domain. These documents often mention the types of tools and platforms that have been used to deliver the end results. Explore your current internal IT vendors to see if they have offerings for AI solutions within their portfolio (often, it's easier to extend your footprint with an incumbent solution vendor vs. introducing a new vendor). Once you build a shortlist, feel free to invite these vendors (via an RFI or another process) to propose solutions to meet your business challenges. Based on the feedback, you can begin evaluating and prioritizing your vendor list.

Checklist: Questions to Tech Vendors When Evaluating AI Solutions

- ☐ How long have you been offering AI solutions?
- ☐ Do you have any use cases/examples you can share that align closely with my industry or organization's needs?
- ☐ Do you have alliances with multiple vendors and an ecosystem to deliver me a complete solution?
- ☐ Do you have the resource bench to assist me in deploying this solution across the enterprise and across geographies (probe for roles like data scientists, data engineers, machine learning architects, service engineers, etc.)?
- ☐ Can you integrate the AI solution within my existing IT footprint?

Implementation and Ongoing Management

Can we trust AI and make the business decisions based on the model predictions? How do I course correct if an AI model or prediction goes wrong?

Most artificial intelligence (AI) models will make prediction mistakes. No AI model, be it a statistical machine learning model or a natural language processing model, will be perfect on day one of deployment. Therefore, it is imperative that the overall AI solution provide mechanisms for subject matter experts to provide feedback to the model. AI models must be retrained often with the feedback provided for correcting and improving. Carefully analyzing and categorizing errors goes a long way in determining where improvements are needed.

A mature error analysis process should enable data scientists to systemically analyze a large number of "unseen" errors and develop an in-depth understanding of the types of errors, distribution of errors, and sources of errors in the model. A mature error analysis process should be able to validate and correct mislabeled data during testing. Compared with traditional methods such as confusion matrix, a mature process for an organization should provide deeper insights into when an AI model fails, how it fails and why. Creating a user-defined taxonomy of errors and prioritizing them based not only on the severity of errors but also on the business value of fixing those errors is critical to maximizing time and resources spent in improving AI models. It is important to select vendors who can offer the full AI model lifecycle management capabilities as opposed to just a model that can make initial predictions but are incapable of taking feedback or learning from feedback and self-reflection via error analysis.

Can we market our value proposition or differentiate our organization from competition using AI-infused solutions?

As a decision maker/influencer for implementing an AI solution, you will grapple with demonstrating ROI within your organization or to your management. However, if you plan the AI infusion carefully with a strategic vision backed by tactical execution milestones in collaboration with the key business stakeholders and end users, you will see a faster adoption of AI across the organization.

AI value translates into business value which is near and dear to all CxOs—demonstrating how any AI project will yield better business outcomes will alleviate concerns they may have.

Can the organization reconcile the impact of AI on staff resourcing and culture?

The goal of AI is to either optimize, automate, or offer decision support. AI is meant to bring cost reductions, productivity gains and in some cases even pave the way for new products and revenue channels. All this may result in change to the way people do their work. In some cases, people's time will be freed up to perform more high-value tasks. In some cases, more people may be required to serve the new opportunities opened up by AI and in some other cases, due to automation, fewer workers may be needed to achieve the same outcomes. Companies should analyze the expected outcomes carefully and make plans to adjust their work force skills, priorities, goals, and jobs accordingly. Managing AI models requires new type of skills that may or may not exist in current organizations. Companies have to be prepared to make the necessary culture and people job role adjustments to get full value out of AI.

Considerations Before Implementing AI: Questions for Practitioners

It's critical for the practitioners of artificial intelligence (AI) solutions—those using and supporting the solutions and analyzing the data—to have a different but no less important understanding of the technology and its benefits and challenges. The following are some questions practitioners should ask during the AI consideration, planning, implementation and go-live processes.

What is the goal and purpose of the AI model I'm asked to build? Do I have representative training and test data for building this AI model? If not, is there is a data manager I can work with to acquire or prepare the data?

AI models must be built upon representative data sets that have been properly labeled or annotated for the business case at hand. Attempting to infuse AI into a business model without the proper infrastructure and architecture in place is counterproductive. Training data for AI is most likely available within the enterprise unless the AI models that are being built are general purpose models for speech recognition, natural language understanding and image recognition. If it is the former case, much of the effort to be done is cleaning and preparing the data for AI model training. In latter, some datasets can be purchased from external vendors or obtaining from open source foundations with proper licensing terms.

Is my data architecture adequate for leveraging AI?

A company's data architecture must be scalable and able to support the influx of data that AI initiatives bring with it. Data architecture involves instrumenting business processes, applications and infrastructure to collect data, building data connectors to observe, collect and store data, data lake strategy and implementation, data versioning, lineage management, data bias checking and normalization. If data is required to be collected or purchased from external sources, proper data management function is needed to ensure data is procured legally and that compliance standards are met in data storage, including GDPR-type of compliance management.

How can I ensure biased data won't skew results?

Biased training data has the potential to create not only unexpected drawbacks but also lead to perverse results, completely countering the goal of the business application. To avoid data-induced bias, it is critically important to ensure balanced label representation in the training data. In addition, the purpose and goals for the AI models have to be clear so proper test datasets can be created to test the models for biases. Several bias-detection and debiasing techniques exist in the open source domain. Also, vendor products have capabilities to help you detect biases in your data and AI models. Leverage those capabilities.

Will robotic process automation, or a cheaper, non-AI process deliver the same outcome?

AI has its time and place. Not every automation needs to be solved with AI. Some automations can likely be achieved with simpler, less costly and less resource-intensive solutions, such as robotic process automation. However, if a solution to the problem needs AI, then it makes sense to bring AI to deliver intelligent process automation.

What AI or machine learning (ML) methods will be relied upon and why?

AI and ML cover a wide breadth of predictive frameworks and analytical approaches, all offering a spectrum of advantages and disadvantages depending on the application. It is essential to understand which approaches are the best fit for a particular business case and why.

How long will implementation take?

Understanding the timeline for implementation, potential bottlenecks, and threats to execution are vital in any cost/benefit analysis. Most AI practitioners will say that it takes anywhere from 3-36 months to roll out AI models with full scalability support. Data acquisition, preparation and ensuring proper representation, and ground truth preparation for training and testing takes the most amount of time. The next aspect that takes the most amount of time in building scalable and consumable AI models is the containerization, packaging and deployment of the AI model in production.

Do I understand the legal, privacy, compliance, security implications of building AI solutions at this company?

Different industries and jurisdictions impose varying regulatory burdens and compliance hurdles on companies using emerging technologies. With AI initiatives and large datasets often going hand-in-hand, regulations that relate to privacy and security will also need to be considered. Data lake strategy has to be designed with data privacy and compliance in mind. Many vendor tools offer these capabilities. Companies must make decisions about and understand the tradeoffs with building these capabilities in-house or working with external vendors.

Will the solution offer explainable results and transparency into the decision-making process?

Consumers, regulators, business owners, and investors may all seek to understand the process by which an organization's AI engine makes decisions, especially if those decisions can impact the quality of human lives. Black box architectures often do not allow for this, requiring developers to give proper forethought to explainability. Data scientists must make tradeoffs in the choice of algorithms to achieve transparency and explainability.

Will the solution provide an adequate level of accuracy, given the application?

Depending on the use case, varying degrees of accuracy and precision will be needed, sometimes as dictated by regulation. Understanding the threshold performance level required to add value is an important step in considering an AI initiative. In some cases, precision and recall tradeoffs might have to be made. For example, when analyzing sentiment in social media context, precision might be more important than recall whereas in a security scenario recall is just important as precision since you may not want to miss out on any security violation incidents.

How can I make AI models more trustworthy for business decision makers? What are the tradeoffs to consider when building more transparent and explainable AI models?

AI continues to represent an intimidating, jargon-laden concept for many non-technical stakeholders and decision makers. Gaining buy-in from all relevant parties may require ensuring a degree of trustworthiness and explainability embedded into the models. User experience plays a critical role in simplifying the management of AI model life cycles.

What resources already exist instead of starting from scratch? What's the best tool available for automating the modeling that we need to do?

Large cost savings can often be derived from finding existing resources that provide building blocks and test cases for AI projects. There are many open source AI platforms and vendor products that are built on these platforms. Evaluating those and selecting those might be cost efficient.

What kind of expenses related to infrastructure, data acquisition, crowd/SME annotation and others should my organization think about to implement AI best practices?

AI initiatives require might require medium-to-large budgets or not depending on the nature of the problem being tackled. Therefore, it needs to be properly scoped. AI strategy requires significant investments in data, cloud platforms, and AI platform for model life cycle management. Each initiative could vary greatly in cost depending on the scope, desired outcome, and complexity.

How will the AI function when it encounters a previously unseen situation or data point?

Over a long enough period of time, AI systems will encounter situations for which they have not been supplied training examples. It is prudent to anticipate and account for these eventualities. It may involve falling back on humans to guide AI or for humans to perform that function till AI can get enough data samples to learn from.

What governance policies and controls have I prepared for preventing and dealing with "AI breaches"?

AI may act or be manipulated in unintended and undesired ways. It is vital that proper precautions and protocols be put in place to prevent and respond to breaches. This includes incorporating proper robustness into the model development process via various techniques including Generative Adversarial Networks (GANs). GANs simulate adversarial samples and make the models more robust in the process during model building process itself.

How will I protect my models and data from adversarial attacks and malicious input?

Stakeholders with nefarious goals can strategically supply malicious input to AI models, compromising their output in potentially dangerous ways. It is critical to anticipate and simulate such attacks and keep a system robust against adversaries. As noted earlier, incorporating proper robustness into the model development process via various techniques including Generative Adversarial Networks (GANs) is critical to increasing the robustness of the AI models. GANs simulate adversarial samples and make the models more robust in the process during model building process itself.

Artificial Intelligence Resources

There are multiple data sources and experts available in the industry including the ***CompTIA AI Advisory Council***.

Analyst reports and materials on artificial intelligence (AI) business case from sources like Gartner, Forrester, IDC, McKinsey, etc., could be a good source of information. Gartner and Forrester publish quadrant matrices ranking the leaders/followers in AI infusion in specific industries. Descriptions of those leaders/followers can give a sense of the strengths and weaknesses of the vendors. This helps in knowing what to look for from a business case perspective. Most vendors publish client case studies and success stories. Read them—with a pinch of salt—as they can be overselling, but still helpful.

CompTIA Resources

Practical Insights on AI

Emerging Business Opportunities in AI

CompTIA Use Case Library

Top AI Solutions

Accelerators and Barriers to AI Business Growth

CompTIA
Artificial Intelligence Advisory Council

About the CompTIA AI Advisory Council

CompTIA's AI Advisory Council brings together thought leaders and innovators to identify business opportunities and develop innovative content to accelerate adoption of artificial intelligence and machine learning technologies. The council is committed to building the strategies and resources necessary to help companies leverage AI to be more successful and collaborates with CompTIA's other industry advisory councils to further study the IT channel, blockchain, drones, business applications and internet of things.