

+ Framework



“ *Artificial Intelligence, deep learning, machine learning — whatever you’re doing if you don’t understand it — learn it. Because otherwise, you’re going to be a dinosaur ...*

—Mark Cuban, American billionaire entrepreneur [1]

Artificial Intelligence in SAFe

Artificial Intelligence (AI) is a term used to describe a wide range of smart machines capable of performing tasks that typically required human intelligence. AI can be applied at all levels of SAFe to build intelligent customer solutions, automate value stream activities, and improve customer insights.

It is a technology that can revolutionize solutions developed by SAFe organizations and has the potential to dramatically influence the operational and business models of enterprises as well.

The Growing Influence of AI

Every day, more and more aspects of our personal lives are aided and supported by intelligent systems. Applying for a loan, flying on an airplane, shopping in an online store, and scheduling a doctor's visit are just a few limitless examples of daily activities using Artificial Intelligence (AI). Further, consumers are not the only beneficiaries of AI-powered solutions. Enterprises use AI to understand their customers better and create better products and services. Banks use AI to detect fraud and money laundering. Governments provide services to their citizens using AI. Even military and national security systems have incorporated this emerging technology.

The trend is clear:

- A startling 72% of executives believe that AI will be the *most* significant business advantage of the future [2].
- According to a PwC survey performed in 2020 [3], 86% of respondents said AI would become a "mainstream technology" at their company in 2021.
- By 2025, 70% of organizations will have operationalized AI architectures [4].

Critical factors have accelerated the rapid growth of these ever-more intelligent technologies. Significant advancements in general-purpose hardware, cloud technology, and optimized infrastructure, such as Graphics Processing Units (GPUs) [5], provide unprecedented computing power that supports a different magnitude of AI capabilities. Concurrently, new learning mechanisms and AI architectures have emerged that expand the frontier of business problems solvable by AI. In addition, the COVID-19 pandemic forced many enterprises to accelerate their multi-year digital transformation plans into weeks or months, thus bringing a robust big data foundation to a broad range of AI-powered business opportunities. As a result, according to interviews with Gartner analysts, over 80% of Global 2000 companies have launched AI initiatives. Many of those solutions are in production and on the market today.

It is increasingly evident that every organization must seriously evaluate and leverage AI capabilities in their next generation of products and services to stay competitive.

This article provides a foundational context for applying AI in the enterprise and how SAFe can accelerate the successful adoption of this advanced technology. Key topics covered include:

- **AI as a Competitive Advantage** highlights the nature of AI's business value and what is required to build a sustained advantage over the competition.
- **Applying AI to Achieve Better Business Results** outlines typical application opportunities for AI within operational and development value streams and capabilities that enterprises can directly embed in the products and services delivered to the end user.
- **Understanding the Fundamental Types of AI** describes common AI concepts and architectures that technology and business leaders must understand to apply AI within their enterprise.
- **Success Factors for AI Initiatives with SAFe** describes essential practices and a decision framework for successful AI initiatives with SAFe.

AI as a Competitive Advantage

AI opens various possibilities to extend existing solutions and make them more valuable and scalable. More importantly, it opens a frontier for new solutions and qualitatively distinct capabilities that benefit the customer and the business in a new way.

But where exactly is AI most helpful?

'Conventional' software solutions are good at addressing well-understood problems; such solutions execute a limited number of scenarios based on a predefined list of rules. However, many organizational workflows and customer scenarios involve parameters that can't be accounted for via conventional, preprogrammed means. AI addresses many of these most complicated scenarios and turns them into viable business opportunities. The landscape of potential AI applications within an enterprise is broad and diverse (Figure 1).

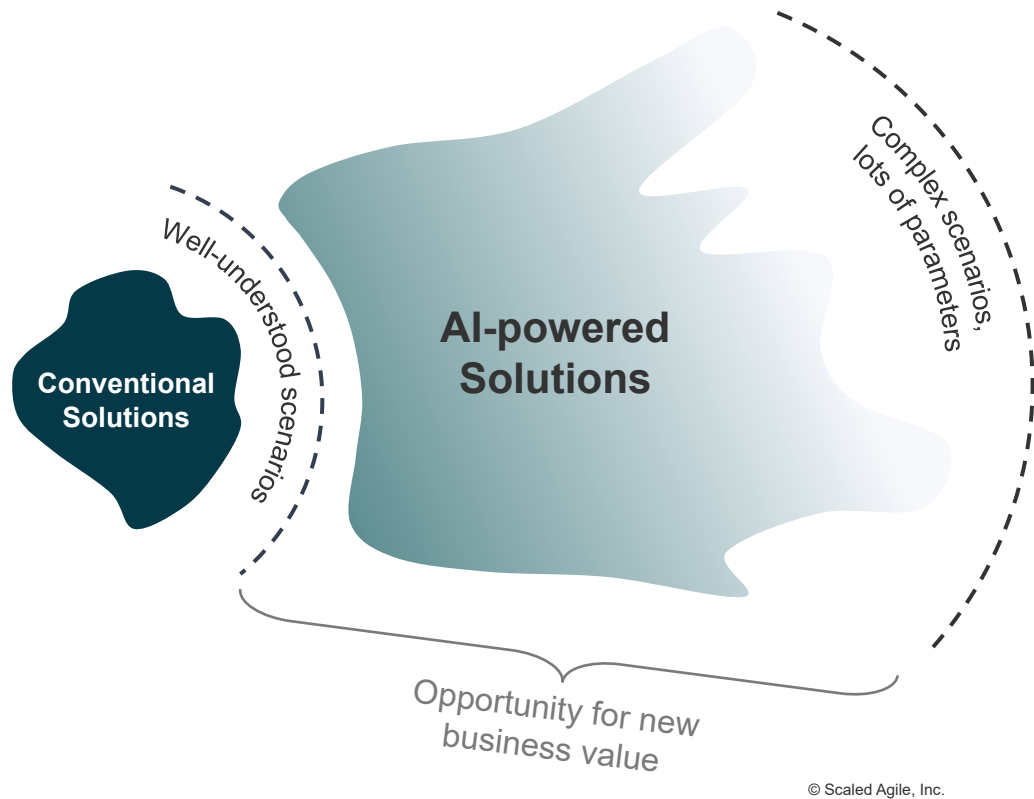


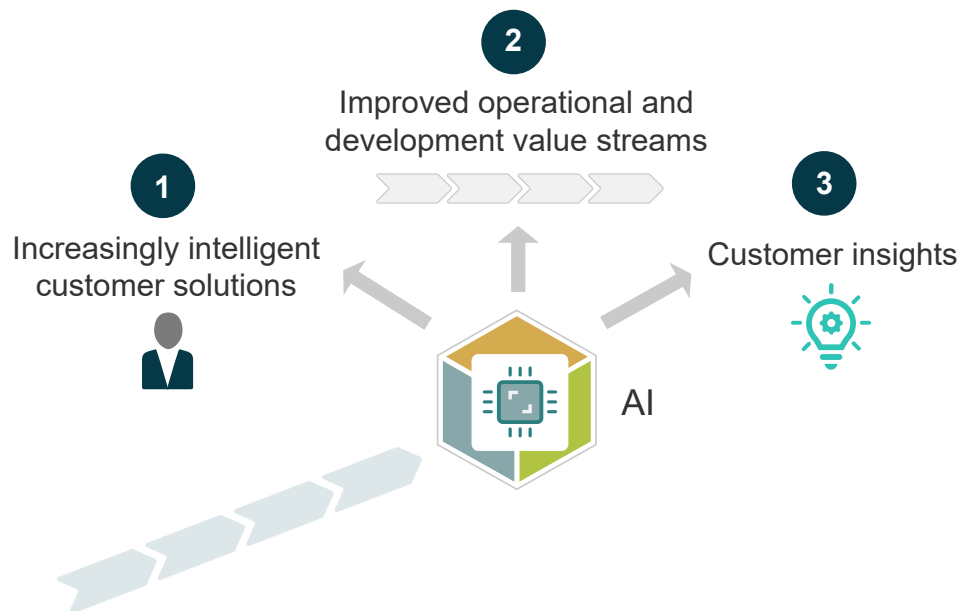
Figure 1. AI dramatically extends the range of business solutions

For example, a bank may utilize AI to automate customer service, identify suspicious account activity, understand customers' needs better and offer appropriate products, extract essential insights from customer feedback and social media, facilitate regulatory compliance, and so forth.

As specific AI applications become more commonplace, companies are moving beyond basic AI solutions to find new opportunities where complex problems exist that can be solved with the help of AI. The ultimate goal is to create innovative business value ahead of the competition [6]. But creating successful AI solutions is no easy task. While AI is still software, the nature of machine learning software development is fundamentally different. Many organizations don't know how to leverage an AI opportunity effectively, even when they devote significant resources to the problem. SAFe provides an operating model that helps identify opportunities, validate new capabilities, and translate them into valuable, working solutions that benefit the customer. The first step is understanding how AI helps improve an organization's outcomes.

Applying AI to Achieve Better Business Results

Generally, an organization's opportunities for leveraging AI lie in three areas (Figure 2):



© Scaled Agile, Inc.

Figure 2. The three primary ways to apply AI in an organization

1. Increasingly Intelligent Customer Solutions. More and more intelligence is embedded in the products and services we access every day. Common examples include self-driving vehicles and driver-assisted capabilities, facial recognition on our phones, personalized product recommendations, and smarter devices powering the Internet of Things (IoT). In addition, many AI-powered functions are now directly embedded in our mobile and desktop computer software applications.

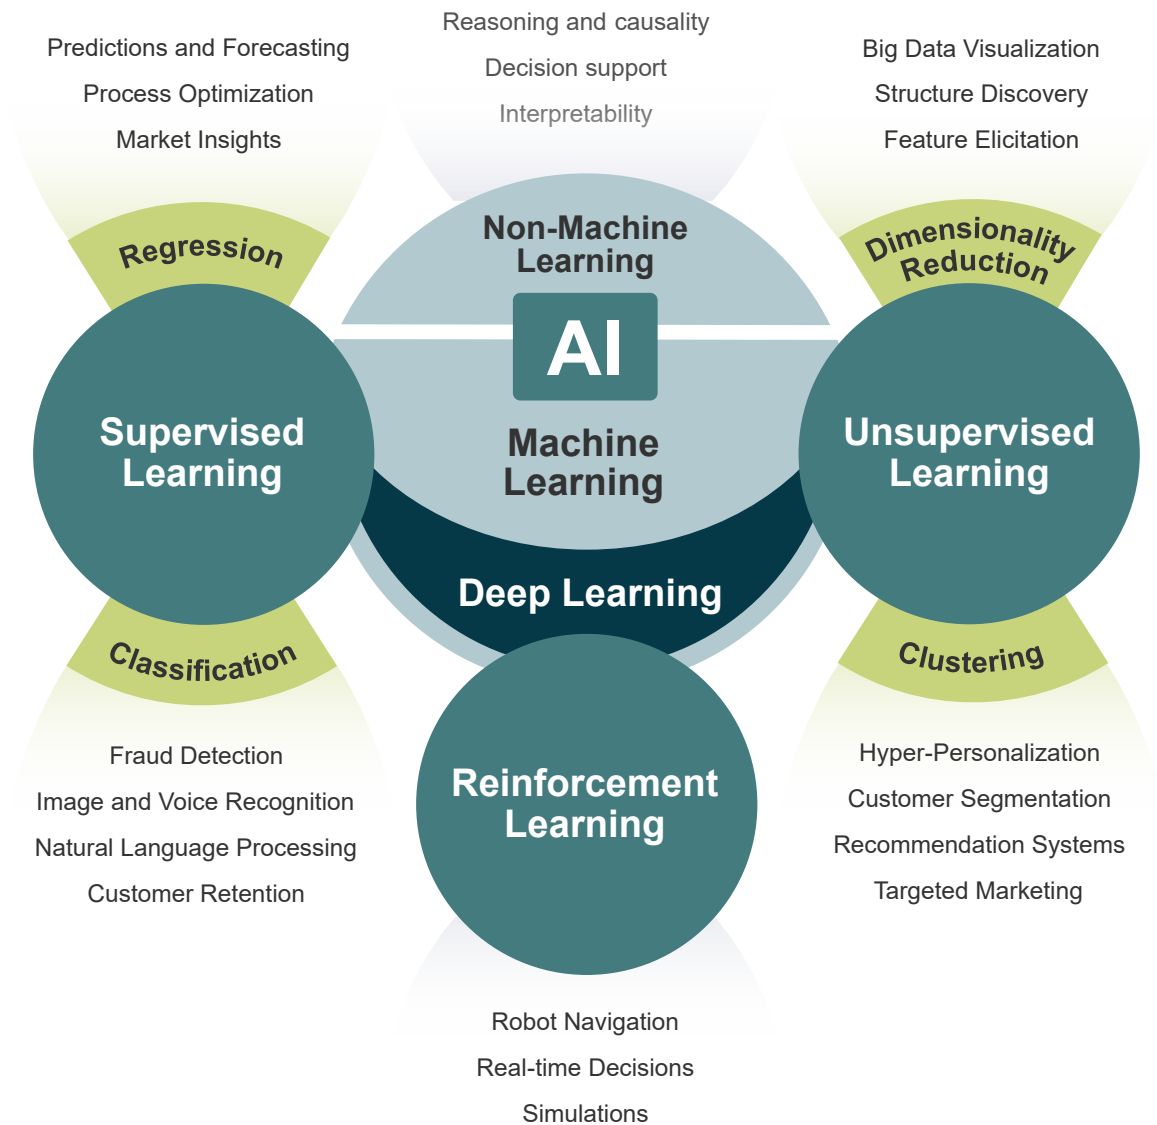
2. Improved Operational and Development Value Streams. Organizational processes are ripe for AI applications to make an internal enterprise process more productive or introduce new capabilities. AI solutions can be applied to both [Operational](#) and [Development Value Streams](#). Inventory and demand management, personalized consumer experiences, and fraud detection are examples of how AI can power operational workflows. In the case of development value streams, companies can use AI to navigate the solution context, analyze production data, identify optimal parameters for user scenarios, and facilitate effective testing.

3. Customer Insights. AI can help organizations identify new business opportunities, learn more about the customer, and extract market insights that create entirely new offerings on an even broader scale. In this latter case, AI supports launching new

initiatives via the BAVS (the [Business Agility Value Stream](#)) that were otherwise undiscoverable.

Understanding the Fundamental Types of AI

Artificial intelligence (AI) is a wide range of smart machines capable of performing tasks that typically require human intelligence. The potential applications represented by AI are extensive and affect almost every facet of business and consumer life. Many of today's AI systems are based on Machine Learning (ML). ML-based solutions are designed to improve based on experiences and data. However, some AI architectures do not involve Machine Learning and instead are based on a comprehensive set of static rules that encode some complex reasoning. Figure 3 provides a typology of various AI and machine learning approaches. It also illustrates some of the capabilities that these technologies enable. However, some illustrated capabilities in Figure 3 can be built with more than one AI approach, and a combination of methods may also be applied.



© Scaled Agile, Inc.

Figure 3. The landscape of Artificial Intelligence

As shown in this graphic, there are three major categories of ML, differentiated by how the *learning* is achieved. All three approaches involve three primary components: the data, the learning algorithm, and the learning model, as Figure 4 illustrates.

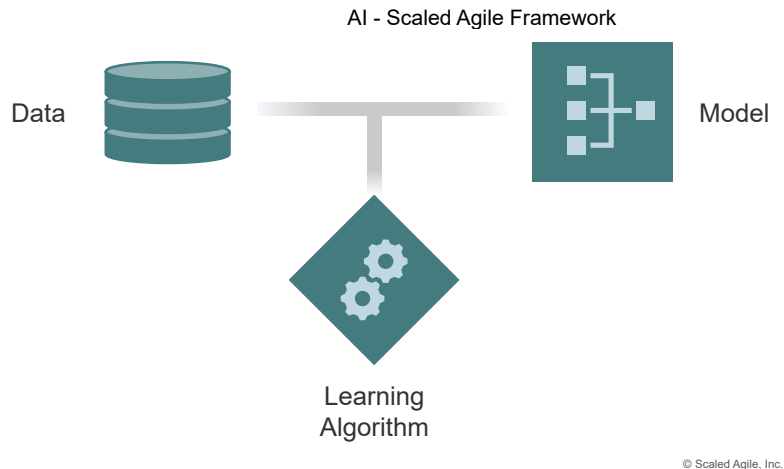


Figure 4. Three critical components of machine learning

Most generally, Machine Learning can be described as one of four types, as described below.

Supervised Learning

Supervised learning utilizes training data to *teach* the model how to produce the desired output (Figure 5). The training data must contain the inputs and the desired outputs as labels. The learning algorithm runs inputs through the model, compares them with the labels, and computes the model output.

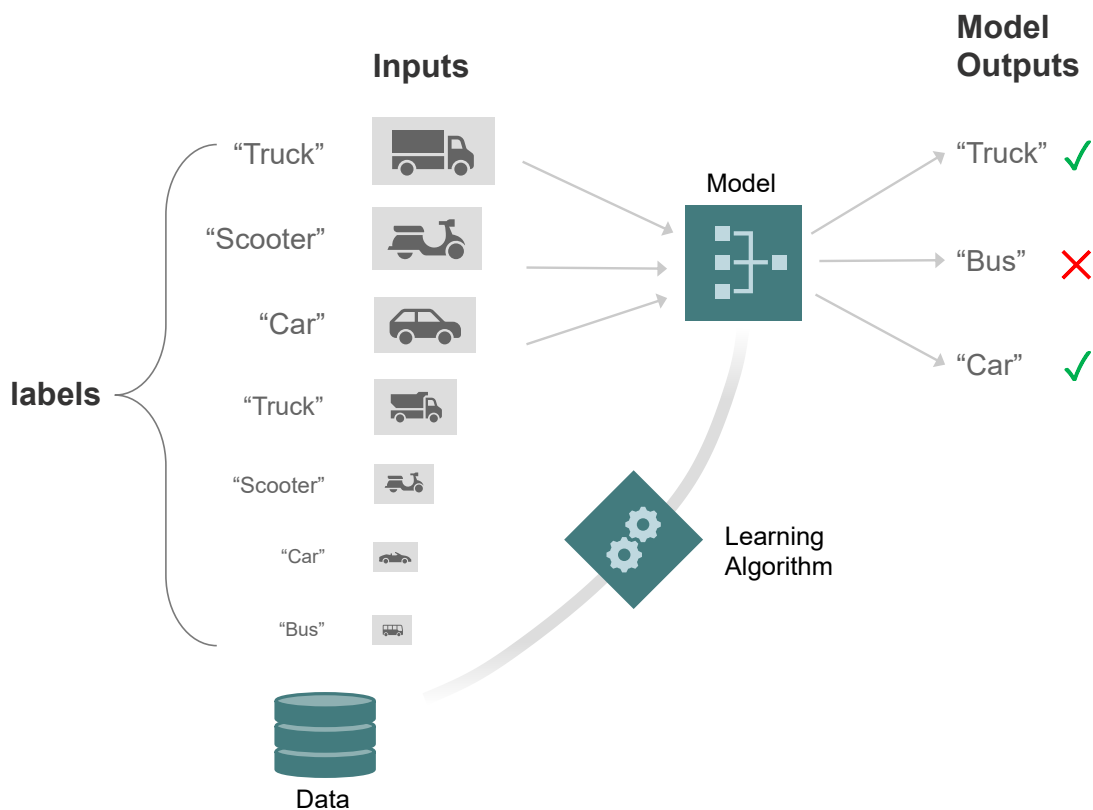



Figure 5. Supervised learning

The algorithm adjusts model parameters and repeats the process until it reaches a sufficiently low number of errors. It is called *supervised* because the desired outputs are supplied alongside the inputs and are used to 'supervise' or 'guide' the learning process. Unless the data initially includes both the inputs and the labels, it requires a 'labeling' process *before* training the model.

Supervised learning can help detect known patterns (fraudulent transactions, spam messages) and data categorization (image recognition, text sentiment analysis). In some instances, the output data may be readily available or easily attainable in an automated manner (such as a customer name alongside the profile photo for face recognition or a five-star rating score next to the product review text for sentiment detection; the situation often referred to as *self-supervised learning*). Identifying such facets of data opens excellent opportunities for applying supervised learning to organizational processes.

Unsupervised Learning

Unlike the previous approach, Unsupervised Learning does not utilize any feedback mechanism. Instead, it extracts valuable information merely by analyzing the internal structure of the data.



© Scaled Agile, Inc.

Figure 6. Unsupervised learning

Unsupervised learning has a significant advantage because the input data doesn't need to be labeled, allowing the learning algorithms to use vast volumes of data. This approach supports easier scaling of the unsupervised learning-enabled capabilities.

This type of AI algorithm is applied to data clustering, anomaly detection, association mining, and latent variable extraction tasks. These processes partition the data by similarity and establish existing relationships within data to be used by other solution capabilities or functions. Some common use cases of such tasks are customer or product segmentation, similarity detection, and recommendation systems. Unsupervised learning can also be leveraged as a link in a broader chain of a supervised learning process to extend data labeling to unlabeled datasets.

Reinforcement Learning

Reinforcement Learning is similar to supervised learning because it also involves a feedback mechanism that verifies the model. However, the feedback does not rely on labeled data in this case. Instead, the system *acts* in a particular *environment* and is supplied with a *reward function* that helps the model learn what action leads to successful outcomes. Therefore, the learning algorithm generates exploratory activity and selects scenarios that lead to the highest reward.




Figure 7. Reinforcement learning

Reinforcement learning finds applications in robotics, gaming, decision support systems, personalized recommendations, bidding and advertising, and other contexts where simulated exploratory behaviors can be evaluated in terms of their value.

Deep Learning

Deep learning is the label for machine learning models based on artificial neural networks (ANN). Deep learning can be effectively applied to supervised, unsupervised, and reinforcement learning and, in many practical tasks, has produced results comparable to or surpassing human expert performance. An artificial neural network is loosely modeled after the structure of neurons in the brain. An ANN has *inputs*, *outputs*, and consists of a connected set of *neurons*. An example of such a model could be a neural network that accepts pixel colors in an image as an input and determines what type of object is in that image as an output.




Figure 8. A deep neural network applied to pattern recognition

Every connection has a specific *weight* that either strengthens or inhibits the signal. When all the connectors leading to a particular neuron convey a sufficiently strong cumulative signal, the neuron *activates* and transmits the signal to other neurons further downstream.

A neural network with multiple hidden layers is called a *deep neural network* and is the foundational architecture for deep learning.

Generative AI

Generative AI is a subfield of artificial intelligence that focuses on creating new content and experiences through machine learning algorithms. It has the potential to revolutionize the way businesses operate and create value. Both businesses and the general public became aware of generative AI virtually overnight with the release of ChatGPT in late November 2022. Initially offered as a free tool, ChatGPT gained over a million users in its first five days and exceeded 100 million users by the end of January 2023.

Generative AI differs from other types of AI in its focus on creating new content and experiences. Unlike other AI applications, such as supervised learning and reinforcement learning, generative AI algorithms are trained to generate original output, such as images, videos, or text. This makes generative AI a powerful tool for businesses looking to automate creative tasks, generate digital assets, and drive innovation. While other AI applications are designed to recognize patterns in existing data, generative AI focuses on producing new, unique content not seen in the training data.

One major application of generative AI is to generate realistic product images or write compelling ad copy. This can significantly reduce the time and resources required for businesses to produce high-quality digital assets. Another area where generative AI can have a significant impact is in automating repetitive tasks and processes. For example, generative AI algorithms can be trained to perform tasks such as invoicing, data entry, and report generation, freeing up employees to focus on higher-value tasks. Generative AI also has the potential to drive innovation and new product development. By combining generative AI algorithms with other technologies, such as simulation and computer-aided design, businesses can quickly generate and test new product concepts, helping to speed up the development process.

However, it's important for businesses to be aware of the potential ethical implications of generative AI, such as the creation of fake, misleading, or biased content. It's critical for organizations to have strong governance and ethical frameworks in place to ensure that their use of generative AI aligns with their values and mission. Precautions must also be taken to avoid exposing customer data, personally identifiable information (PII), or sensitive company data to public generative AI tools.

(Note: The preceding description of generative AI was generated by ChatGPT! [7])

Success Factors for AI Initiatives in SAFe

Developing and delivering successful AI solutions is a task challenging to many organizations. The following factors are critical in establishing a productive solution development process for AI systems.

Apply a Clear AI Decision-Making Framework

Many AI initiatives fail to deliver the better results promised when advocates propose investments in this technology. This failure to meet expectations is often caused by poor decision-making about how and why AI will be used. Organizations often want to 'do AI' because 'everyone else does,' without understanding the effort required to operationalize and scale this technology, its impact on the organization, or even if it will produce the intended benefits.

SAFe organizations already have powerful tools to make better decisions on the appropriate use of AI. A few of these are highlighted in Figure 9.

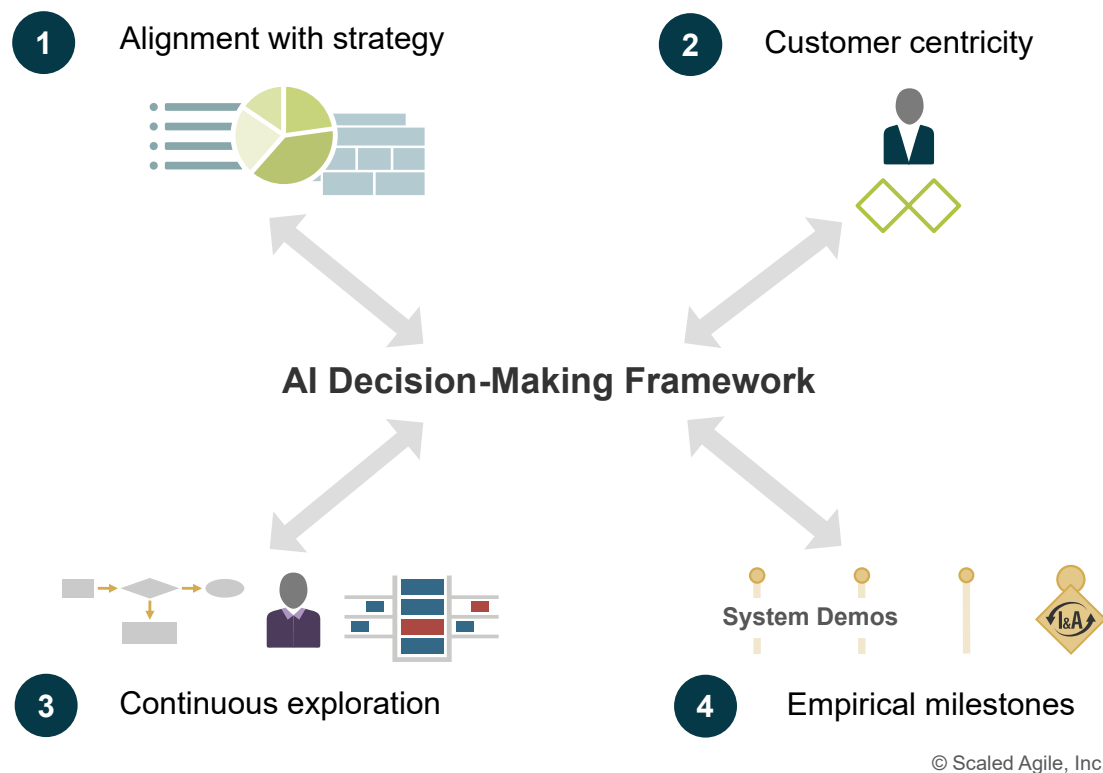


Figure 9. AI decision-making framework enabled by SAFe

1. Alignment with strategy ensures that AI initiatives pursue beneficial outcomes for the business. Aligning the AI roadmap with the [Portfolio](#) strategy is an essential

step in this direction. Some AI initiatives may require more financial support; others may need to be repurposed or canceled based on a progressively assessed economic viability. Productively managing the spending requires [Lean Budgeting](#). Using the [Portfolio Kanban System](#) and a [Lean Business Case](#) helps establish better alignment with strategy. Additionally, [PI Planning](#) provides the foundation for recurrent alignment of AI strategy with the actual implementation.

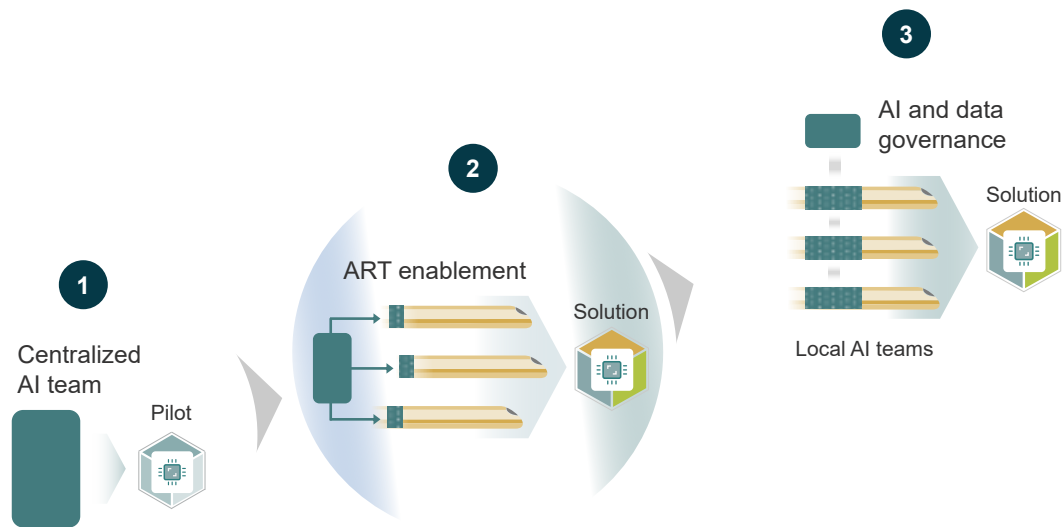
2. Customer centricity is critical to ensuring that an AI initiative is actually solving a customer problem. For example, intrinsic AI capabilities (such as image recognition or natural language processing) must be appropriately integrated into a favorable customer scenario. Explicitly defining the customer problem is an important step and hugely benefits from applying [Design Thinking](#) to AI capabilities.

3. Continuous exploration paves the pathway to a successful AI-powered solution. Solution development always contains a significant degree of uncertainty. In the case of AI, however, the level of uncertainty is exceptionally high in creating the right solution and implementing it correctly. This is where the [SAFe Lean Startup Cycle](#) is very useful. Creating a clear business hypothesis, building an AI MVP, and validating it against suitable measures are at the core of successful exploration.

4. Empirical milestones guide the development of a successful AI solution. AI capabilities must be continuously integrated with the rest of the solution throughout the incremental development process. Solution increments are used to elicit important customer feedback.

Organize to Deliver AI Solutions

Organizing around value is a critical enabler of flow and is one of the SAFe operating [principles](#). In the case of AI-powered solutions, it has specific implications for SAFe organizations. We've observed three stages of maturity and organization evolution, as Figure 10 illustrates.



© Scaled Agile, Inc.

Figure 10. Observed organizational models for AI maturity in SAFe

Stage 1. The first step is to build a critical mass of AI expertise. This step includes hiring and developing skills in AI and ML, and the accompanying skills in data science, data analytics, and data engineering. This centralized team engages with subject matter experts. It may take on building and deploying some initial proof of concept, pilot projects, or even an MVP for an anticipated AI initiative. This stage seeds the enterprise with AI thinking and shows the potential for extending the range and types of AI applications.

Stage 2. It quickly becomes apparent that these new capabilities need to be developed and integrated into the enterprise's product offerings. Integrating AI capabilities into existing products requires the engagement of the ARTs responsible for enterprise solutions. As a result, the focus of the central AI team moves from developing AI capabilities to enabling ARTs to do so. At this stage, the team may embed some subject matter experts on board each train to streamline the process. Typically, the centralized function also continues to develop certain types of AI capabilities, particularly those that cross value streams and cause the creation of more mature, enterprise-level data.

Stage 3. With time, more AI functions get embedded as a part of the Agile Release Trains. Including these skill sets on the ART helps manage integration risks, establishes a more robust delivery pipeline, and fosters critical knowledge. The decentralized effort, however, typically benefits from ongoing centralized expertise. A central AI architecture team can continue to explore new AI technologies, implement portfolio and enterprise cross-cutting initiatives, and provide

organizational governance that ensures proper management of customer data and AI's ethical and business implications in the marketplace.

Across all three stages, the AI group must involve engineering and business professionals. This is vital for the organization to identify and leverage productive opportunities for AI-enabled solutions.

Proactively Manage Data

Data is a critical link in most AI implementations and, when handled poorly, leads to a failed AI initiative. Every organization must establish an effective process around data management.

Data is critical to AI in a couple of ways. First, most AI systems require a *training dataset* that must be large enough—and thus referred to as Big Data—for the AI model to learn to perform its task correctly. For example, a fraud detection system may require tens or even hundreds of thousands of transactions until it learns to recognize the fraudulent ones amongst them reliably. To validate whether the system was successfully trained, a whole separate dataset is needed – the *testing dataset* that must be readily available.

But the great challenge comes not only from the size of Big Data. An even more significant challenge is that data seldom exists in a ready-to-go, consumable form. Data is often dispersed across multiple, disjointed sources within and outside the enterprise. Additionally, the actual format of the data may be unconsumable. It may lack some critical attributes, without which training is impossible.

Any AI approach will depend on a portfolio strategy that recognizes that a systematic approach to acquiring and normalizing Big Data is virtually always the precedent for an AI application. This includes validating assumptions about data as a first order of the business in every AI initiative. Figure 11 illustrates some central elements of how SAFe guides enterprises on this parallel journey.

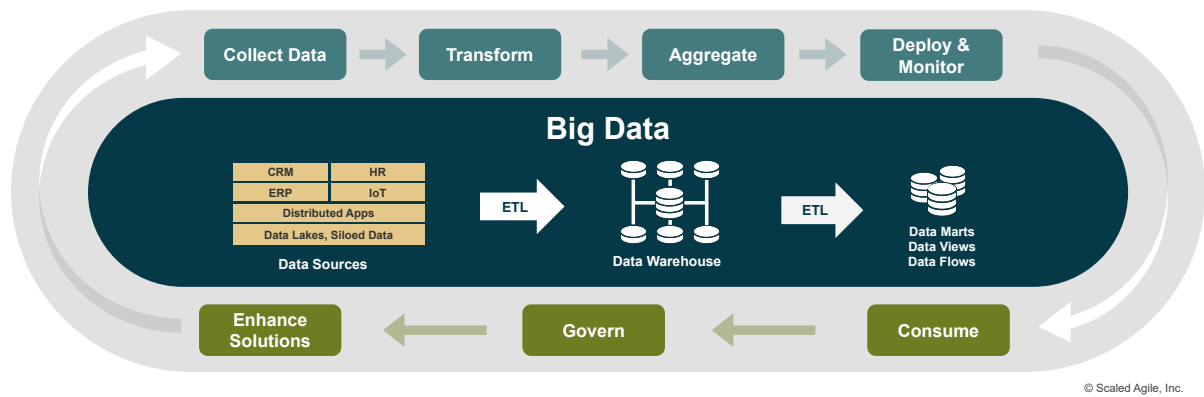


Figure 11. The Big Data cycle

Build Organizational Competency Around AI

Artificial Intelligence represents a novel type of technology to many organizations. And even those companies with some AI initiatives in flight don't have the right expertise applied in the right place. AI is a complex technology that requires a good understanding of AI capabilities and potential business applications' landscape. Navigating critical opportunities requires business and technology people in an organization to work collaboratively to develop AI competency.

SAFe emphasizes the importance of the constant interaction of business and technology professionals. This interaction is achieved through multiple means, including strategic effort around the portfolio kanban system, events like PI planning, [System](#) and [Solution Demos](#), and [Inspect & Adapt](#). Additionally, various continuous exploration activities, design thinking, and work definitions leverage the overlap between business and technology representatives. In the case of AI, all these techniques need to be utilized even more rigorously, as disconnects may be very costly.

Establish an AI Solution Path

Driven by the flow of [Business Agility](#), an AI solution advances through the following logical steps:



Figure 12. Example AI solution path

Define – This step involves identifying the business opportunity and understanding the customer's need. The initial assessment determines how suitable AI architectures could support the business's needs. Both technology and business stakeholders are closely involved in this step.

Pilot – The step usually involves building a [Minimum Viable Product](#) (MVP) for AI capabilities and proceeding further based on the learnings from the MVP. It is critical to ensure that the MVP solves the business problem rather than simply demonstrating technical feasibility. The basic level of AI capability monitoring and measurement is applied at this step. Pilots should be designed to get to an MVP that can test the hypothesis of the AI capability as quickly as possible.

Operationalize – At this stage, the AI capability—if proven viable in the previous step—gets designed and implemented to allow full integration with the existing solution ecosystem and provides full-fledged support for required business scenarios. Operationalization is expensive in solution development and may lead to significant waste if the MVP does not validate the viability of integration with existing systems. The organization must also plan for the investments required to monitor, adjust, and retrain AI components continuously.

Scale – An AI capability is successful if it grows to support increased volumes of customers. In the case of AI—quite unlike conventional solutions—scaling introduces some unique challenges. So, with increasing processing volume, the actual business parameters, the data, and the AI model may grow out of alignment (something called “drift” in AI terminology). Progressive adjustment of the model, the learning algorithm, and data processing is often required to sustain the model through scaling and beyond. Additionally, as a result of scaling, some qualitatively new scenarios may emerge or be subsumed by the solution. This may sometimes lead to a significant update to the AI capability or even a drastic change to the ML model or the learning algorithm.

Govern – Enterprises are often presented with multiple opportunities to apply AI and can do it in different areas. Over time, an ecosystem of intelligent capabilities emerges across various solutions operating in the same Big Data pool that establishes potential interoperability to create higher-level enterprise value. This requires a governance process that spans multiple Solution Trains in the portfolio or even [multiple portfolios](#) in the enterprise. This governance process can be a part of the general portfolio and enterprise governance. Still, it must especially account for some areas that grow far more important with the advent of AI. Such areas include effective data management, privacy and security, computing power, the bias

problem, and traceability. Ethical implementation of AI is a dominant topic in this domain and requires serious investment by organizations leveraging this technology.

Due to the inherent complexity of AI solutions and the lack of adequately established practices supporting AI development and operation, many organizations experience serious challenges driving their AI solutions through operationalization, scaling, and proper governance. This significantly reduces the number of initiatives that produce any business value. That is why an organization must establish a productive solution path that fosters rapid feedback and effectively uncovers and manages technology and business risks.

Summary

Succeeding with AI is critical to surviving and thriving in the Digital Age. SAFe provides various tools that enable the successful development and delivery of AI-powered solutions. By organizing the teams and trains to incorporate AI capabilities, using practical, strategically sound decisions, and proactively building up the necessary cloud and Big Data capabilities, enterprises create the foundation for a productive AI effort. It is essential that technology and business professionals collaboratively develop competency around AI and understand where Artificial Intelligence can help the organization seize critical business opportunities.

Learn More

[1] Cuban, Mark. Upfront Summit, 2017.

[2] *The Macroeconomic Impact of Artificial Intelligence*. PwC, 2018.

<https://www.pwc.co.uk/economic-services/assets/macroeconomic-impact-of-ai-technical-report-feb-18.pdf>

[3] *AI Predictions*. PwC, 2022. Retrieved October 10, 2023, from

<https://www.pwc.com/us/en/tech-effect/ai-analytics/ai-predictions.html>

[4] *The 4 Trends That Prevail on the Gartner Hype Cycle for AI, 2021*. Gartner, September 22, 2021. Retrieved October 10, 2023, from

<https://www.gartner.com/en/articles/the-4-trends-that-prevail-on-the-gartner-hype-cycle-for-ai-2021>

[5] *Using GPUs for Training Models in the Cloud*. Google. Retrieved October 10, 2023, from <https://cloud.google.com/ai-platform/training/docs/using-gpus>

[6] Iansiti, Marco, and Karim R. Lakhani. *Competing in the Age of AI: Strategy and Leadership When Algorithms and Networks Run the World*. Harvard Business Review Press, 2020.

[7] OpenAI. ChatGPT, 2021. OpenAI API, accessed February 5, 2023.

Last update: 10 October 2023

The information on this page is © 2010-2023 Scaled Agile, Inc. and is protected by US and International copyright laws. Neither images nor text can be copied from this site without the express written permission of the copyright holder. Scaled Agile Framework and SAFe are registered trademarks of Scaled Agile, Inc. Please visit [Permissions FAQs](#) and [contact us](#) for permissions.

Framework

[Download SAFe Posters & Graphics](#)

[Watch and download SAFe videos and presentations](#)

[Blog](#)

Training

[Course Calendar](#)

[About Certification](#)

[Become a Trainer](#)

Content & Trademarks

[FAQs on how to use SAFe content and trademarks](#)

[Permissions Form](#)

[Usage and Permissions](#)

Scaled Agile, Inc

Contact Us

5400 Airport Blvd., Suite 300
Boulder, CO 80301 USA

Business Hours

Weekdays: 9am to 5pm

Weekends: CLOSED

© 2023 Scaled Agile, Inc.

[FAQs on how to use SAFe content and trademarks](#)

[Permissions Form](#)

[Usage and Permissions](#)

